

Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/141402/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Fussey, Pete, Davies, Bethan ORCID: <https://orcid.org/0000-0003-4997-8432>
and Innes, Martin ORCID: <https://orcid.org/0000-0002-8950-8147> 2021.
'Assisted' facial recognition and the reinvention of suspicion and discretion in digital policing. *British Journal of Criminology* 61 (2) , pp. 325-344.
10.1093/bjc/azaa068 file

Publishers page: <https://academic.oup.com/bjc/article/61/2/325/5921...>
<<https://academic.oup.com/bjc/article/61/2/325/5921789>>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies.
See

<http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



‘ASSISTED’ FACIAL RECOGNITION AND THE REINVENTION OF SUSPICION AND DISCRETION IN DIGITAL POLICING

PETE FUSSEY, BETHAN DAVIES and MARTIN INNES*

Automated facial recognition (AFR) has emerged as one of the most controversial policing innovations of recent years. Drawing on empirical data collected during the United Kingdom’s two major police trials of AFR deployments—and building on insights from the sociology of policing, surveillance studies and science and technology studies—this article advances several arguments. Tracing a lineage from early sociologies of policing that accented the importance of police discretion and suspicion formation, the analysis illuminates how technological capability is conditioned by police discretion, but police discretion itself is also contingent on affordances brought by the operational and technical environment. These, in turn, frame and ‘legitimate’ subjects of a reinvented and digitally mediated ‘bureaucratic suspicion’.

Key Words: facial recognition, digital policing, suspicion, discretion

Introduction

Automated facial recognition (AFR) surveillance has emerged as a particularly controversial technology among the growing armoury of digital policing tools. This reflects how AFR engages several longstanding dilemmas affecting the execution of social control: those of identification, suspicion and discretion. In particular, questions of identification and linking individuals to bureaucratic traces held about them and their past behaviours have constituted defining and elemental challenges. As Cole (2002) attests, accomplishing this linkage has been a significant ‘driver’ of innovation in the conduct of surveillance and forensic science. For example, fingerprinting has its origins in 19th century India, as colonial institutions sought to individuate members of the population they were attempting to administer. More recently, Innes *et al.* (2020) describe how certain neighbourhood policing officers acquire ‘cop cultural capital’ through their ability to identify ‘on sight’ members of street gangs and to have knowledge of their patterns of behaviour and association.

This identification imperative underpins the development trajectories and operational adoption of some of the most significant and controversial social control technologies. Included here are the rapid growth of biometrics, such as the UK National DNA Database, as well as the proliferation of visual surveillance technologies (*inter alia* Marx 2016). Advancing this momentum have been recent manoeuvres to integrate AFR technologies into the routines and processes of policing as part of the wider and growing movements in ‘digital policing’ and its ‘biometrification’. AFR is especially intriguing because of how it blends and amalgamates the principles of both visual and

*Pete Fussey, Department of Sociology, University of Essex, Wivenhoe Park, Colchester, Essex CO4 3SQ, UK; pfussey@essex.ac.uk; Bethan Davies and Martin Innes, Crime and Security Research Institute, School of Social Sciences, Cardiff University, Cardiff, CF10 3AT, UK

biometric surveillance. As a consequence, it has proven highly controversial, with significant political and legal challenges mounted with respect to several ‘early adopter’ policing applications.

An aspect of these challenges resides in deep-rooted normative formations of suspicion that reflect how surveillance is routinely steered towards segments of the population that Reiner (1992) dubbed ‘police property’. Indeed, it was arguably the seminal finding of the early sociological ethnographies of street policing that the police ‘gaze’, and decisions to intervene via the enactment of criminal law, were not uniformly distributed (Dixon 1997). The tenor of these arguments were synthesized in Matza’s (1969) differentiation between ‘bureaucratic’ and ‘incidental’ suspicion. The latter pivots around the archetypal crimefighter myth of police work, whereby the skilled investigator determines ‘whodunit’ by linking aspects of the incident to a suspect with the means, motive and opportunity to commit the offence (Innes 2003). Whilst this may be the favoured fictional representation of police crime investigation work, according to research, it is relatively infrequent. Far more routine is ‘bureaucratic suspicion’ involving rounding up the ‘usual suspects’ (Gill 2000), where police search strategies are based upon knowledge of individuals previously displaying forms of criminal behaviour similar to the incident under investigation. This schema was notably elaborated by Marx’s (1988) notion of ‘categorical suspicion’, whereby suspect identification becomes based upon specific socio-demographic characteristics. He identifies this as especially important to remote surveillance camera practices, a view enriched through detailed work by Norris and Armstrong (1999) who concluded that operators possessed considerable discretion in determining ‘suitable’ subjects for the surveillant gaze.

Discretion thus constitutes another concept of fundamental concern here that, along with suspicion, was similarly excavated by the ground-breaking police ethnographies of the 1960s and 1970s. For example, Skolnick (1966) concluded that officers on the street almost irrevocably possessed determinative decision-making power in terms of against whom, when, how and why criminal law was enforced. Collectively, these insights and evidence convey that processes of police surveillance and identification are influenced by multiple influences.

These concerns remain central to current debates over efforts to harness digital AFR technologies for contemporary policing. Supporters of AFR extol its virtues in aiding austerity-afflicted law enforcement agencies identify wanted suspects. To date, two UK police forces—South Wales Police (SWP) and the London Metropolitan Police Service (MPS)—have undertaken the world’s first long-term public multisite trials of this technology. In response, these initiatives have induced a chorus of opposition that approached a crescendo during 2019. In 2018, the UN Special Rapporteur for the Right to Privacy criticized SWP’s use of the technology on the grounds of necessity and proportionality (UN OHCHR 2018). During July 2019, the House of Commons Science and Technology Committee (2019) called for authorities to halt its use, pending the creation of specific regulations. Summer 2019 also saw the first major legal challenge against AFR. Whilst the judicial review upheld that police use was lawful, the Court of Appeal overturned parts of this judgement in August 2020, ruling that the current legal basis for AFR was insufficient. Crucially for the analysis presented in this article, the court highlighted the presence of excessive police discretion over *who* becomes targeted for surveillance and *where* the camera systems are deployed. This emphasis

on discretion is central to the empirically informed analysis presented below, wherein the article identifies a wider range and scope of discretionary practices than acknowledged by the court.

In this article, we explicitly set aside these normative and ethical debates to focus upon how AFR is operationalized in support of street policing. We draw upon data derived from unprecedented access to these two trial studies to interrogate police uses of AFR. Central to our account is how these operational deployments pivot around a digital reinvention of bureaucratic suspicion and reformulation of discretion invoked by the system and its users. This is important given it contests a more technologically deterministic discourse surrounding AFR promoted through some press and political commentary. The picture emerging from our high-resolution case studies is that, rather than AFR being a purely algorithmically driven machine, its use constitutes a socio-technical assemblage that both shapes police practices yet is also profoundly shaped by forms of police suspicion and discretion. Consequently, we argue that AFR should be reframed as 'assisted' rather than 'automated' facial recognition.

The next section outlines AFR technology, situating this in wider literatures on policing, surveillance and technology. This is followed by an account of how the data were collected by the two projects and then synthesized for purposes of analysis and reporting. Discussion of the findings starts with an overview of the outcomes achieved in order to set out the roles AFR performs in a policing context. These results are then interpreted and unpacked by introducing a series of 'operator', 'organisational' and 'system' issues that account for these outcomes. To help make sense of these data, they are next brought alongside the concepts of discretion and suspicion outlined above. The concluding section reprises key tenets of the argument, considering their wider applicability to the conduct of digital policing and the work of contemporary social control technologies.

Situating Facial Recognition Technologies

Facial recognition technology involves real-time biometric¹ processing of captured video images for the purposes of matching to a database and identifying individuals.² For police uses analysed herein, faces of members of the public are scanned as they pass fixed and mobile camera points. These images are then analysed by a facial recognition algorithm that deciphers a subject's facial dimensions, typically focusing on the size and relationship between individual features (such as shape of the eye socket, distance between eyes, shape of the jaw, and so on). More advanced systems use digital modelling to compensate for low light, lens distortion or faces tilted away from cameras. This computational code is then matched to a database containing similarly processed images (colloquially dubbed a 'watchlist') that also contains text-based data, such as

¹The EU General Data Protection Regulation (GDPR) and companion EU Law Enforcement Directive both define facial images as biometric data.

²Facial recognition technologies have many forms and applications. This paper focuses on technology that monitors the public as they traverse observed areas and compares their faces to a database of known suspects (sometimes called 'live facial recognition'). In the United States, face recognition technology adopts a more investigative role, where surveillance camera, social media or other video footage is analysed after an offence has occurred. Here, suspects are unknown to police in advance. The other widespread use is for 'verification' (as opposed to 'identification'), such as that used at borders to confirm traveller identities. Each application holds different uses, rights implications and demands on human discretion.

ethnicity code and warrant information. If the algorithm detects a ‘possible match’,³ an alert is sent to a human operator for review. If the alert is deemed credible, this stimulates an attempt to engage the matched individual. Overall, while AFR is demonstrably technological in character, human interpretation and intervention constitute key parts of the policing process.

While it has become increasingly familiar through applications, such as passport control and Apple’s integration within their iPhone, facial recognition technology actually has a surprisingly long history. For Gates (2011), NEC’s first public demonstration of the technology at the 1970 World’s Fair Osaka marked a key moment in its biography.⁴ While the technology had greater success in capturing the public’s imagination than their faces, computational limitations have not prevented experimentation with AFR in civic spaces across the world. Several small-scale experiments with AFR technology arose in the United Kingdom and United States, including deployments in the London Borough of Newham during the 1990s, reportedly in response to perceived resurgent threats from Irish Republican terrorism (Fussey 2010). Demonstrating the linkages between large sporting events and expansions in advanced surveillance architecture, a prototype version was trialled at Manchester City’s former Maine Road stadium during the mid-1990s (Davies 1996) and in Tampa, FL, during the 2001 Superbowl (Stanley and Steinhardt 2002). Notable about all of these attempts is that the technology was subsequently abandoned due to poor performance. Such failings also contributed to the technology being rejected by 2002 Salt Lake City Winter Olympics planners at the height of post-9/11 security anxieties (Fussey and Coaffee 2012).

Accordingly, police engagements with AFR have, until recently, remained small scale and episodic. This has left the technology relatively underexamined by surveillance scholars, particularly when compared to the attention lavished on other instruments, such as CCTV and the bulk monitoring of online activity. Notwithstanding academic commentary on the wider topic (*inter alia* Introna and Wood 2004; Gates 2011), police operational uses of AFR technology have yet to be studied empirically. This article addresses this gap whilst aspiring to avoid the risk of technological determinism facing the study of surveillance technology (Lyon 2001) by accenting the role of human interpretation and intervention.⁵ To achieve this, three intersecting literatures are considered particularly important for understanding the socio-technical practices surrounding AFR use: sociological studies of police suspicion and discretion; analyses of surveillance and society and insights from science and technology studies. These are briefly considered in turn.

The ‘discovery’ of discretion in police decision-making during the 1960s and 1970s opened important questions over how police powers are enacted (Dixon 1997). Delivering formal social control was revealed as grounded in finely grained officer judgements influenced by myriad factors, including, e.g., the values and precepts of ‘cop culture’ and various heuristic ‘easing behaviours’. Related scholarship was brought together most completely in Davis’ (1971) analysis of injustices and inconsistencies arising from police discretion. Although Davis’ key contribution emphasizes harms caused by

³This terminology has been introduced by SWP in the period since the Cardiff University evaluation.

⁴In a point of historical continuity, NEC developed the algorithms supporting the face recognition technology used in both police trials analysed in this paper.

⁵Recent legislation has established a requirement for retaining meaningful human decision-making amid algorithmic policing processes (e.g. The Data Protection Act 2018; EU GDPR and companion Law Enforcement Directive).

licentious uses of discretion, particularly intriguing is the analysis of how such judgements relate to the legal and regulatory structures governing policing. Such structures not only frame and constrain subjective action but also create spaces for new forms of police discretion to flourish. While Davis' formulation has proved durable, social scientists have challenged the excessive significance given to the law in this and related socio-legal theories (e.g. [Dworkin 1977](#)), pointing to a range of different shaping and structuring processes affecting the application of police discretion ([Campbell 1999](#)). As detailed below, to these we add the role of technology in shaping and framing police discretion.

If the concept of discretion addressed how police interacted with citizens when deciding whether to operationalize legal powers, the invocation of suspicion framed who was subject to such encounters and why, appended to which was work on the organizational construction of suspicion rehearsed above (see also [Manning 1978](#)). More recent scholarship has provided particular nuance to this debate. Drawing inspiration from [Goffman's \(1972\) *microstudies in public order*](#), [Quinton \(2011\)](#) highlights the role of tacit knowledge in structuring suspicion. Although often deemed unquantifiable by its possessors, Quinton's analysis identifies a discernible bearing of contextual cues, 'backstage' stereotyping and practical rules of thumb on the formulation of suspicion.

AFR technology nuances suspicion in important ways. It draws elements of *recognition* into deliberations of suspicion. To some extent, this reflects established legal grounds for suspicion, such as those coded in the *Police and Criminal Evidence Act 1984*, whereby recognition may constitute grounds for stop and search. Yet, an important difference arises when AFR technology is inserted into the process. Namely, *initial* recognition does not originate from the officer exercising discretion. Instead, officers act more akin to intermediaries, interpreting and then acting upon a (computer instigated) suggestion originating outside of, and prior to, their own intuition. The technology thus performs a framing and priming role in how suspicion is generated.

One point of convergence in this wider literature is the linkages between surveillance, risk and discretion drawn by [Ericson and Haggerty \(1997\)](#). Particularly relevant for this study of advanced AFR surveillance is Ericson and Haggerty's emphasis upon how routine policing practices encode, communicate and give visibility to specific renderings of risk. Several empirical questions arise concerning operational uses of AFR. These include how operators' risk judgements become framed by information brought to their attention by the computer, the impact on ensuing operational judgements, the dynamic nature of human-computer interaction and how this affects levels of risk appetite and, thus, willingness to intervene. Important here is how appetites escalate and de-escalate according to numbers of AFR alerts, themselves an artefact of how the machine is calibrated and configured.

The conduct of surveillance has been a significant and rapidly expanding sub-theme in the general literature on social control over the past two decades. One theoretical touchstone has been [Foucault's \(1977\)](#) panoptic frame as a particular manifestation of intricate intertwinings of power and knowledge. Work in this tradition has captured how the surveillance net has both widened and deepened through the incorporation of increasingly sophisticated technological apparatus for (self-)monitoring facets of identity and behaviour. Critical scholars have variously censured the over-literal and often incorrect application of the panoptic metaphor (e.g. [Bauman and Lyon 2013](#); [Marx 2016](#)). The panopticon metaphor also emphasizes the self-regulatory experiences of

the surveilled rather than the activities of surveyors that constitute the focus herein. However, less acknowledged is the continued relevance of disciplinary power in the context of Foucault's other works, notably emphasis on multiscale articulations of power during his major project on biopolitics (Foucault 2008). Additionally, and as a preparatory staging post for his wider project on governmentality, Foucault (2007) problematized the issue of circulation, a concept that may find resonance for contemporary biometric surveillance strategies as they variously permit and prohibit movements of populations through spaces.

Recent years have seen Haggerty and Ericson's (2000) concept of the *surveillant assemblage*—itself drawn from Deleuze and Guattari's (1987) formulation—gain influence in this field. The assemblage concept accents the positive magnetism drawing different surveillance systems together into powerful upscaled ensembles of observation techniques. Such accounts have been important in detailing the nuanced ways surveillance technologies interrelate with other technological systems to accumulate greater potency. Informed by such theoretical contributions, advanced surveillance systems cannot be seen to function in isolation or, even, at one specific register of operation. Instead, it has become increasingly necessary to recognize the interoperability of different technological architectures, such as the linkages of biometric scanning with databases of individuals and their analysis through complex algorithmic techniques. Yet, such digital configurations beg important questions over what remains of human agency and intervention.

In engaging with such matters, this article takes a cue from a tradition long instituted in science and technology studies concerning the agentic qualities of technology. Stated simply, this holds that, while practitioners shape and condition the application and potential of their technological instruments, these practices, forms of action and ways of thinking are simultaneously shaped and conditioned by these technologies and the affordances they bring (*inter alia* Latour 1987). This approach induced a field of study investigating how the form, meaning and outcomes of a technology are shaped as it passes through social settings (Bijker and Law 1992). One theoretical advance emerging from this field particularly germane to this paper is Hutchby's (2001) novel application of Gibson's (1977) 'affordance theory'. Affordances emphasize how objects *invite* certain (inter)actions and generate particular conditions of possibility. As explained by MacKenzie *et al.* (2017: 736): 'the vital quality of affordances is the opportunity for action'.

Applied to AFR use, the concept of affordances illuminates how its outcomes are not determined by specific technological properties but the opportunities for action it provides (and denies) users. How officers engage with these opportunities constitutes the focus for this paper. It is also here, the way AFR technology frames and invites opportunities for action, that we make a link to bureaucratic suspicion. As noted above, traditional sociologies of policing recognize how bureaucratic suspicion structures officer decision-making processes in specific ways. This paper interrogates the complex ways affordances of AFR technology scaffold police suspicion in the information age. Face scanning processing power operates at rates exceeding human capability and induces new conditions of policing possibility. These, in turn, assert hitherto unexplored influences of technology on discretion and suspicion formation within policing.

Data Collection and Analysis

Following the established traditions of the sociology of policing outlined above, the concepts and insights developed here are informed by detailed 'high resolution' fieldwork. Findings emerge from two independent academic reviews of police AFR technology in the United Kingdom. Researched AFR deployments spanned 12 events in the SWP force area between June 2017 and March 2018 and a more episodic ten deployments by the MPS between August 2016 and February 2019.⁶ The two evaluative studies were conducted separately and independently. It was only after completing the bulk of data collection that the authors were able to compare and contrast their results. The fact that the two separate studies illuminated similar issues and themes enhanced the evidential strength of the principal claims and justified the attempt to blend the findings. This article is the result.

On the grounds that different kinds of data would afford deeper insights into the nuances and complexities of AFR in policing contexts, both studies collected empirical data using multiple-method research designs. Ethnographic observations and interviews were the principal methods at both sites. This involved being stationed in surveillance vans observing how police operators engaged with AFR technology, with most emphasis placed on how officers interpreted and responded to computer-issued matches of suspects. Detailed and extensive conversations with police staff also took place during these long observation periods. Observations were carried out at seven SWP deployments, totalling 35 hours. As SWP usually deployed AFR in multiple sites, researchers moved around the relevant locations while observing. For the MPS component, observations were conducted at six live operations (single location each time), totalling over 50 hours, with additional observations of pre-operational briefings, post-operational de-briefings and key operational planning meetings.

Data on numbers of AFR-generated alerts (or 'possible matches') were analysed at both research sites. In Wales, a quantitative analysis of system outputs (CSV files) and operator logs, provided by police, was carried out. There were fewer computer-generated alerts in London. Here, outcomes were manually logged by researchers *in situ* and verified with police evaluation teams in follow-up meetings. These data indicate how many correct (and incorrect) alerts were generated by the technology, how human operators viewed the credibility of computational alerts and the outcomes of such deliberations.⁷ Taken together, these engagements with the two policing agencies trialling AFR systems offered a rare opportunity to trace the deployment of technologies through sensitive operational environments.

AFR-equipped surveillance vans were typically stationed at key busy locations, such as the main entrance gate to the Principality Stadium in Cardiff or a thoroughfare leading to Leicester Square in London. Two operators were normally allocated to each van, along with a street-based intervention team usually comprising several police officers. While the MPS used one vehicle, three vans were usually deployed at larger events in Wales, such as the 2017 Champions League final in Cardiff.

⁶Other smaller-scale UK police AFR activities include a joint one-off MPS–Humberside Police trial (utilizing MPS technical capability and counted by them as one of 'their' trials) and Greater Manchester Police's short-lived collaboration with the Trafford Centre shopping mall in Manchester. Both occurred during summer 2018.

⁷Due to data quality issues in operator logbooks used in Wales, slight differences exist between the outcome data of SWP and the evaluators. SWP logbooks have now been modified to account for this and are now specifically designed for AFR use.

TABLE 1 *Results from SWP deployments from 31 May 2017 to 7 March 2018*

Event	Date	Watchlist size	Alerts	Number of AFR matches deemed credible by human observers (%)
Champions League	31 May 2017–3 June 2017	1,049 (average)	2,632	78 (3%)
Elvis Festival (Porthcawl)	22 September 2017–23 September 2017	472	18	11 (61%)
Operation Fulcrum	18 October 2017	616	20	9 (45%)
Anthony Joshua Boxing	28 October 2017	609	60	5 (9%)
Autumn Rugby Internationals (four dates)	11 November 2017–2 December 2017	628–1,262	91	13 (14%)
Kasabian concert	4 December 2017	442	7	3 (43%)
Liam Gallagher concert	13 December 2017	41	6	6 (100%)
Operation Fulcrum	21 December 2017	920	10	8 (80%)
Operation Malacite	22 December 2017	923	3	3 (100%)
Royal visit	18 January 2018	9	0	–
Six Nations Rugby (three dates)	3 February 2018–17 March 2018	801–873	48	22 (46%)
Stereophonics concert	7 March 2018	576	5	0 (0%)

Regarding data analysis, qualitative thematic analysis of the observational and interview data revealed extensive conceptual themes. This included those of interest for the purposes of the original evaluations, as well as those dissected below. Regarding what is reported herein, thematic and data selections were based upon intensive discussions between the authors over a period of several months. This iterative process progressively distilled the focus of this article.

AFR and its Outcomes

Political and media-based discussions of AFR have been largely pre-occupied with outcomes and ‘if it works’. But this position fails to define what appropriate measures of success should be (i.e. the number of convictions, arrests or accurate identifications or minimizing the volume of inaccurate ‘matches’). For AFR critics, a key issue concerns purported high numbers of ‘false positives’ generated by the system. This refers to when the AFR algorithm suggests a ‘possible match’ that is inaccurate. This is counterpointed by the ‘true positive’ category (a correct match) and ‘false negative’ (when the system fails to identify a suspect).

Tables 1 and 2 below list quantitative results from deployments of AFR by SWP and the MPS. The columns provide a sense of how the systems were deployed, the size of each watchlist and numbers of computer-generated alerts alongside ‘true positive’ outcomes. Crucially, these outcomes were significantly influenced by the activities of police operators. The first table includes data on officer judgements of computer-generated AFR matches in Wales.⁸

While officers may judge algorithmically generated alerts as credible, a further question arises over whether these individuals were right to agree with the computer. In

⁸The Champions League operation was disproportionate in terms of the number of false positives for reasons explored in more detail below.

TABLE 2 Results from MPS deployments from 28 June 2018 to 14 February 2019

Event	Date	Watchlist size	Alerts	Number of AFR matches deemed credible by human observers (%)	Number of total matches proved correct after subject ID checks (%)
Stratford (1)	28 June 2018	489	4	4 (100%)	0 (0%)
Stratford (2)	26 July 2018	306	1	1 (100%)	0 (0%)
Soho (1)	17 December 2018	2,226	5	3 (60%)	1 (20%)
Soho (2)	18 December 2018	2,226	9	3 (33%)	1 (11%)
Romford (1)	31 January 2019	2,401	9	7 (78%)	4 (44%)
Romford (2)	14 February 2019	1,996	15	9 (60%)	3 (20%)

other words, did officers cede discretion to a computer judgement that turned out to be wrong? It was possible to explore this issue further in London, where AFR was deployed in single locations and evaluators could be present at all times. Below, instances of officer agreement with computer-generated matches are additionally compared against data on street-based checks to confirm a subject's identity. For example, at Soho (1), *three* possible matches from the AFR system were deemed correct by operators, but when IDs were checked, *one* proved to be correct.

Across both sets of data, results follow no particular pattern. They are mixed in terms of their accuracy, with some deployments yielding higher error rates in possible matches than others. SWP uses of AFR showed some system improvement over time where, proportionally, the number of correct matches improved (from the original 3 per cent). In London, one striking finding is the tendency of officers to agree with (cede discretion to) the algorithm and the high chance that computer-generated matches would not be verifiably correct.⁹

However, depicting the policing outcomes of AFR solely in such terms is of limited utility. Such figures, and their representation through terminology of 'true' and 'false positives' and 'negatives' appropriated from computer science, lack the context needed to interpret their implications, particularly for suspicion. For example, algorithms were upgraded in the middle of the South Wales trial significantly improving system accuracy, something invisible when considering the figures in isolation. The following passages unpack some of the complexities and nuances that shape decision-making around AFR practices. Specifically, analysis attends to the multitude of factors impacting and interacting with AFR use and their implications for police officers' constructions of suspicion and discretion. Broadly speaking, we delineate three key groups of factors originally highlighted in the South Wales study: 'organizational', 'system' and 'operational'. It is argued that how AFR functions in practice, the ways it is used and the outcomes produced are shaped by interactions within and between these three core sets of influences.

Organizational factors are defined as the policing routines, policies, strategic choices and standard operating procedures that directly impact the deployment of AFR. System-based concerns relate more explicitly to technical aspects of AFR, including

⁹This does not mean that the computer is always incorrect on other occasions. For example, if an individual is matched yet becomes lost in the crowd or if a possible match is *judged* incorrect by a human operator, it remains possible that the computer made a correct match. However, because the identity of matched individuals are not checked, it cannot be known if the judgement was correct (hence 'not verifiably correct').

hardware and software issues, such as: the specific facial recognition algorithms; setting parameters for how the technology operates and constructing ‘watchlist’ databases of ‘suspects’ images for the camera technology to scan. Operational factors cover human–machine interactions shaping how the outcomes of AFR are produced and used. Of particular salience are officer judgements about the credibility of algorithmically suggested suspect identities and types of human deference to computational outputs. Analysis of the role of discretion active in each of these three spheres of action elucidates the myriad ways suspicion becomes parameterized and primed through the use of digital facial recognition technology.

Organizational Factors

An immediate result of adopting researching contrasting deployments in two police force areas is the clarity it brings about how distinctive forces implement AFR differently. Key systems and processes across SWP and the MPS were significantly different from each other. Thus, organisational decisions and choices shaped the outcomes of the AFR technology. Among these were active selections regarding where and when to deploy cameras and the extent to which their presence was signalled to the public. Often such decisions were based on the availability of large numbers of people on which to test the technology and police ‘beliefs’ that multiple ‘suspected’ individuals are likely to be present. Such considerations were germane to the uses of AFR at the Notting Hill Carnival in London (twice, in 2016 and 2017) and the UEFA Champions League Final in Cardiff (2017).

At a more granular level, decisions about the specific positioning of cameras significantly influenced any outcomes generated, including: the ideal location for street-based intervention teams; risks to the public; risks to police and, ultimately, judgements over optimizing AFR performance. For example, in London, local intelligence briefings were used to justify the selection of Stratford as a site for one AFR trial. Borough-level crime statistics attributing it the fourth highest rate of violence and fifth highest rate of robbery in London seemed to support this decision. Yet, the actual positioning of cameras, at Westfield shopping mall, almost half a mile away from the highest crime ‘hotspot’, was in part motivated by environmental factors that eased the temporary installation of AFR technology.¹⁰ At other sites, such as MPS operations in Romford in 2019, preference was given to a place allowing optimal positioning of on-the-ground intervention teams (to increase their chances of intercepting subjects matched by the AFR technology) rather than the location with highest concentrations of offences. Thus, the arraying of physical bodies in relation to the aptitude and affordances of AFR technology took primacy over the optimizing of policing aims.

In South Wales, decisions over deployment locations were informed by intelligence, as well as strategic decisions relating to numbers of physical bodies. For example, the positioning of a van at the Principality Stadium during large sporting events was an obvious choice as it both guaranteed large numbers of people passing the cameras and offered a promising location for apprehending ticket touts and ‘known’ troublemakers. Similarly, the main shopping street in Cardiff also promised heavy footfall, including

¹⁰Notes from observation of police planning meeting, 23 July 2018.

people from across the whole force area. Situational factors that arose while 'on location' nonetheless had to be managed appropriately. For instance, flag sellers also operate for sporting events in Cardiff and the following field note excerpt illuminates how such exogenous environmental factors impact the implementation of AFR:

A flag seller arrives on scene to set up stall. There is a large Italy flag which at times obscures the camera view as it waves in the wind. Operator discussion about whether to ask nearby PCSOs to intervene and ask to move on. Over time, it appears the stall might be a bonus because people tend to look up as they approach it. (fieldnote, SWP, 31 May 2017).

Organizational considerations also had a temporal dimension. AFR deployments in Leicester Square, London, were justified on the grounds of addressing crime and disorder closely linked to the West End's night-time economy. However, due to limitations in low-light camera performance, AFR operations took place during the day. Such decisions drifted from their initial rooting in intelligence-based considerations to those predicated on affordances arising from AFR technological capabilities. Similar problems were observed with low-light clouded deployments in South Wales, many of which took place during winter evenings—such as the Autumn Internationals and Six Nations rugby events.

System Factors

Essential to AFR performance, and pivotal in the construction and resolution of suspicion, are watchlists; the database of subject images against which captured video images are compared. There has also been considerable speculation and concern over 'watchlist' composition (see [Information Commissioner's Office 2019](#)).

Bespoke watchlists were created for each deployment, with both SWP and the MPS investing considerable resources in constructing them. Watchlists ranged in size between 400 and 1,200 individuals in South Wales and up to 2,226 individuals in London.¹¹ Significantly, criteria for inclusion on watchlists is a highly discretionary practice. Both forces included 'persons of interest' in their watchlists. For SWP, this included individuals wanted on warrant or as suspects for offences, as well as potentially vulnerable individuals (i.e. missing persons). Some deployments focused on specific crime types.

MPS documentation included individuals as 'wanted by police and the courts' that later became refined as 'wanted for serious violent offences'. MPS Data Protection Impact Assessments repeatedly referenced non-specific categories, such as: '[t]o support ongoing policing activity with regards to a specific problem or location' and '[t]o assist police in identifying individuals who may be at risk or vulnerable'.¹² Such broad categories offer significant latitude for interpretation, creating a space for officer discretion with regards to who was enrolled and excluded from such databases. Confirming the previous assertion about the importance of organizational factors, SWP used colour-coded watchlists to denote the status of individuals (e.g. pink for warrants and

¹¹During the planning stages for one London trial, some officers expressed a desire to construct a watchlist of 'all wanted people in London'. According to officers using police databases to construct these watchlists (interviewed 20 September 2018), this would have implicated more than 23,700 individuals. Notwithstanding debates over what was meant by 'wanted', it is likely that this approach would have failed basic human rights tests of necessity and proportionality.

¹²Data Protection Impact Assessment for the Use of Live Facial Recognition Within the MPS', 12 December 2018, p. 2.

amber for suspects), which were displayed within alerts to operators. Though colours were not offence specific, they did denote categories of suspicion and relative urgency for apprehension. For example, during the Champions League deployment, the ethnographic data recorded operators' concern when they could not locate a 'red' nominal, whose watchlist colour indicated risk of danger. Across both research sites, the degree to which individuals graded from 'of interest' to 'wanted' was not always clearly defined. Thus, in an expression of techno-social integration, each algorithmically derived 'possible match' reflected a more generalized sense of operator suspicion and one further framed by the computer's issuance of an 'alert' when someone resembling these 'suspects' passed a camera.

The quality of input data used to configure watchlists directly influences AFR performance. Yet, the quality of images inputted to the system was also crucial, as South Wales Police quickly learned during their initial pilots. Images used for the first deployment at the Champion's League final varied greatly in quality. This was partly due to the international nature of the event but also the force's lack of understanding around the importance of high-quality images. Prior to the deployment, images were sourced from British, Spanish and Italian police, as well as other European agencies. The majority were 'custody' (controlled environment) images, but they differed greatly in standard and resolution. In addition, images taken from other surveillance operations were also used for some individuals—these had no 'quality control' and many were taken under varied lighting. Because they were 'naturalistic' with people, sometimes smiling or squinting or tilting their heads at an angle, this affected the shape of their faces and, therefore, the algorithm's ability to analyse them effectively. Having encountered difficulties with poor-quality source images during the Champion's League operation, SWP initiated a major force-wide programme to enhance the quality of all their new custody images. Such developments reveal further affordances at play: the quality of source images—captured in an entirely different context and often before the possibilities of AFR had been imagined—became key to determining if a passing 'suspect' would be identified.

In the context of AFR identification, watchlists structure the police gaze, interjecting a technologically inflected element of 'bureaucratic suspicion'. By nature, in creating watchlists comprised of police-held custody images, AFR specifically targets police attention towards individuals already known to the authorities (the 'usual suspects'). In one SWP operation, individuals suspected by intelligence units of pickpocketing at live music events were included on watchlists as a preventative measure. Similarly, in Soho, London, attempts were made to enrol local 'characters' known for engaging in antisocial behaviour into watchlists. Deployments of this nature highlight how suspicion can be directed and further enacted by this technology. AFR systems induce a technologically framed bureaucratic suspicion in digital policing.

Digital parameters

Multiple additional technical components shape AFR outputs and outcomes. Fundamentally, AFR technology is not 'plug and play' both in the sense that it requires organizations to accommodate its idiosyncrasies and also because the system itself has

to be adjusted and calibrated to operate effectively. Such adjustments require continued discretion to negotiate a series of competing aims.

For example, the AFR systems analysed here use an adjustable 'similarity score' ranging between 0 and 1—broadly analogous to a correlation coefficient—when determining whether faces constitute possible matches. In theory, the higher the similarity score, the higher degree of confidence the computer 'has' in the match. Users set the threshold of similarity needed before the machine issues an alert. Over time, SWP raised their threshold score from 0.55 (as recommended by the commercial supplier) to 0.59 (during the Six Nations rugby internationals) in order to improve results (the MPS used the recommended 0.55 score throughout). Adjusting the threshold is a deliberate intervention intended to influence the performance of AFR suspect identification in specific ways. Setting a higher threshold increases the likelihood of a computer match being correct but reduces the number of overall matches made. It also risks incorrectly discarding more borderline cases that may be correct matches (false negatives). Lowering the threshold generates more possible matches but raises the likelihood that some of these will be inaccurate (false positives). In both force areas, system accuracy and the availability of resources to intercept flagged individuals were key considerations when deliberating over similarity score threshold levels and where to set these.

Paradoxically, whilst human adjustment of threshold settings shapes the operation of AFR, a common consequence of this decision is the ceding of human agency to the technology. The more the threshold increases, the more the algorithm becomes the decision-maker. During SWP training sessions, operators were instructed to ignore similarity scores and instead focus their attention on the images presented to them. Despite this instruction, observational data suggested that these scores regularly appeared to influence operators' sense of suspicion when adjudicating AFR alerts. Another common practice was for operators to read scores aloud, noting when they were particularly high. In several instances in London, a high similarity score proved pivotal in tilting a (human) decision towards engaging with an individual matched by the computer, thus exerting decisive influence on the framing of suspicion.

Other technological trade-offs involved operator judgements about appropriate resolution, pixilation and processing power. Images with higher resolution can be captured at a greater distance, enabling cameras to adopt a wider angle and, hence, analyse more people in each frame. However, the higher the image resolution, the harder the system has to work to analyse it. A similar issue affects decisions over how images were framed and cropped (the 'zone of recognition'). A wider angle captured more faces yet required more processing power and reduced system performance, where a tightly cropped view reduced the numbers of faces available for processing. Configuring these settings, therefore, involved striking a balance between the time taken to process an image and the amount of a scene deemed acceptable to lose. Such decisions rendered some faces visible to the machine yet, in doing so, rendered others invisible.

This issue played out during the final stop of the Leicester Square deployment. An alert was generated as an individual passed the very bottom of the camera's field of vision. A clear image of the face was displayed yet the suspect's body was outside the frame, making it impossible to provide street-based officers with a meaningful description of his appearance beyond reference to ethnicity and hair colour. A similar problem was observed in Cardiff as an individual rode a bicycle down the main shopping street.

In the captured image, the bicycle was out of frame, meaning operators were completely unaware of the bicycle and unable to relay the full description to the intervention teams, resulting in widespread confusion over how the suspect disappeared so quickly. While much discussion of AFR technology rightly emphasizes biometric signatures of an individual's face, intervention teams rely heavily on more traditional signals for identification, such as the description of a suspect's clothing. Such considerations demonstrate the continually negotiated nature of such technologies, where optimal performance in some areas is predicated on reduced capability in others.

In Wales, this issue connected with the level of habituation users had to the AFR technology. For example, ethnographic observations revealed how, lacking familiarity with the system, operators reverted to what they already knew, which, in many cases, meant using the system in a similar way to CCTV. For example, one officer experienced in CCTV use zoomed the cameras out in order to capture a much wider scene, assuming that, because the system looked similar to CCTV, it functioned accordingly. However, because this added load to computer processors, the system then stopped detecting faces.

One notable issue concerns levels of trust and the extent to which discretionary agency is delegated to technology, despite awareness of potential technological limitations. During the two MPS trials in Newham, London, e.g., pre-operational briefings consistently highlighted a belief in the near infallibility of AFR. A representative quote was:

'it is 100% effective in spotting those uploaded into the system' (MPS Officer 28 June 2018), and 'the technology is very accurate, despite misinformation' (MPS Officer 24 July 2018).

Throughout the MPS trials, a commonly articulated and prevailing view was one of faith in AFR systems to enhance policing, but the challenge being in proving its worth externally.

Operator Factors

The final group of factors relate to the interpretations and decisions made by individual AFR operators. AFR as a socio-technical system reformats key aspects of police decision-making processes and brings elements of bureaucratic suspicion to the fore. Potential drawbacks of this were highlighted by Matza (1969) when he noted how bureaucratic suspicion has a tendency to frame persons of interest as different to others, regardless of any available evidence. Individuals listed on watchlists and databases are cast as warranting suspicion and the AFR surveillant gaze is specifically oriented towards them. But, in so doing, the social biases of police activity that disproportionately focuses on young people and members of African Caribbean and other minority ethnic groups (*inter alia* The Lammy Review 2017) are further inflected by alleged technological biases deriving from how technical accuracy recedes for subjects who are older, female and for some people of colour (Buolamwini and Gebru 2018; National Institute of Standards and Technology 2019). As such, it has been suggested that the algorithm could be more likely to be incorrect and generate 'false positives' when 'identifying' people from these groups. Such issues notwithstanding, discretion remains a fundamental aspect of AFR policing and underpins the array of decisions that must be made by officers to render it operationalizable.

Arguably though, the principal instance of police discretion during AFR use is in the moment of 'adjudication'. When assessing the credibility of an AFR match, the related activities of discretion and negotiation took place across three main arenas: between human operators and the computer; between different operators assessing the same image and between system operators and street-based intervention teams.

Human-computer interaction

The potential of AFR becomes realized when computer-generated matches are resolved through human activity. In most instances, this either involves an initial decision to disregard a match or, conversely, a suspect being engaged by street-based intervention teams tasked with conducting additional identity checks with AFR-matched individuals. In practice, this 'second stage' of activity was an arena of contestation and negotiation.

Human operators, therefore, constitute an essential component of the AFR process and play the primary role in adjudication. Two officers usually carried out this role. Many received formal or informal training prior to deployment, and some occupied non-operational roles, meaning AFR was a novel experience for them. Variances in operator capability were evident across both research sites and these disparities mirrored those encountered in other forms of biometric policework, such as that identified in DNA typing activities (Cole 2002).

Such considerations shaped the deference some officers gave to the algorithm and, conversely, why others were more sceptical of its performance. During one SWP deployment, one operator was visibly frustrated with a lack of correct alerts being generated in their van, while, on the same day, operators conducting surveillance elsewhere in the city centre had succeeded in locating and arresting multiple 'persons of interest'. As a result, this operator became less trusting of alerts generated by the system. Despite habituation to the system, the technology thus *reduced* the sense of suspicion he experienced. Similarly, in London, once an AFR match had first been deemed incorrect by operators (on the third day observed), the overall rate of disconfirmed alerts increased slightly. Such incidences demonstrate the varied responses among human operators of AFR. However, while deference to suggestions generated by algorithmic decision-making was largely habitual—and with 26 of 42 computer-generated alerts considered suitably credible to intercept a matched individual in London—it is important to acknowledge the important role of some officers' (techno)scepticism.

Roles and interactions between adjudicating officers undertaking this duty varied considerably. Sometimes, one operator would be looking for the person in the crowd while the other was describing them aloud from the image captured on the screen: operators reported using key facial features, such as eyes, nose, mouth, jawline and hairline to inform their decisions. While not relevant to a subject's appearance, some officers also recruited background information (e.g. offence type) for their deliberations. At other times, contrasting approaches occurred within the same operational team. In London, any disagreements over whether to launch an intervention were always resolved in the affirmative, though this was not the case during SWP deployments.

Technical difficulties sometimes limited the role of AFR. During mobile deployments in central London, radios continually failed to work inside the AFR van. The corollary effect of these network fractures is illustrated by the following field note:

Fourth AFR match of the first Soho deployment (0.57 threshold). The officer adjudicating images attempts to radio a request to intercept a suspect. Responding to failures of both the radio and mobile tablets he lent out of the van and tried to radio again. When this failed, he took off after the suspect on foot. By this time the individual had crossed almost the length of Leicester Square. Limited adjudication time. Decision-making was near instant. With reflection significant differences were apparent between the probe and gallery images. The gallery image had moles on the suspects face, the probe image had none. While difficult to ascertain at first, most tellingly they had different colour eyes. A false positive (MPS, 17 December 2018, 14:22 pm).

Compensating for technical difficulties, therefore, not only limited AFR capability but also compressed the time available for discretionary adjudication. During South Wales' initial deployments for the Champions League Final when the system was still being configured, it was slow and often produced 'lag'. For example, 90 seconds elapsed between the camera timestamp and real alert time, 'which was especially evident where a potential match was brought up by the system' (field note, SWP, 31 May 2017). This relationship between different components of human-technical networks also reflects a critique among accounts of surveillance informed by assemblage theories rehearsed above: while surveillance practices involve intricate relationships between different forms of technology, they are not necessarily enhanced by such unions. Single points of failure inhibit the network and reduce overall surveillance possibilities.

Technical and environmental influences

Additional to the role of functioning and malfunctioning technology, formulations of suspicion were shaped by amalgams of environmental, technical and 'human' influences. The spatial dynamics of areas into which AFR cameras were deployed were particularly crucial in shaping the outcomes of this technology. The Stratford, London, deployments revealed a trade-off between the risk of losing suspects in the crowd versus the need to preserve sufficient time for human deliberation of AFR matches. Here, in one example, intervention teams were stationed close to the cameras. After a day of inactivity, a late afternoon alert matched a female to a watchlist entry listing serious firearms offences. The ensuing flurry of activity revealed that the intervening officers were too close to the cameras and the suspect had already merged into the adjacent busy shopping mall. Eventually discovered, and discounted as a false positive, this sequence of events demonstrates important linkages between algorithmic decision-making, police discretion and the spatial dynamics in which an intervention is performed. Spatial juxtapositions of cameras to intervention teams assert a significant constraint on discretion and the formulation of suspicion: physical proximities compress decision-making and, it can be argued, invite an intervention.

Depleted human resources, specifically the limited availability of intervention teams, also affected the capability of AFR technology. At the October 2017 Anthony Joshua boxing fight in Cardiff, near simultaneous possible matches issued to different intelligence units left intervention teams unavailable to deal with both. Constraints imposed by the available human resources to service the 'demand' created by the AFR system moderates simplistic claims that such technologies can address austerity restrictions by replacing policing functions.

Other ergonomic, 'human factor' and, crucially, affective influences shaped AFR use. Significant concerns were raised in Wales over the potential for 'face blindness' and fatigue caused by looking at faces on laptop/desktop screens for long shifts in a confined space. At the 2017 Champions League final, e.g., operators were observed being bored and frustrated as they were responding to large numbers of 'false positives' with low similarity scores. A state of ennui was also palpable during some London deployments, albeit for different reasons. For the most part of several AFR deployments, very little happened and long days of watching prosaic surveillance feeds were only occasionally punctuated with moments of drama. Such incidences connect with longstanding observations documenting boredom, weariness and disengagement among surveillance operatives (Norris and Armstrong 1999). In response, AFR alerts rapidly transform a setting of boredom to one of excitement. Moreover, prior states of tedium might shift the weight of decision-making towards the algorithm as a rare alert offers relief through an invitation to action.

Conclusion

Facial recognition technologies provide the capability for police to identify suspects they probably would not be otherwise able to. That said, numerous caveats to such claims require acknowledgement. AFR does not directly replicate extant policing approaches to identifying wanted individuals (who are more commonly sought by more directed and often covert tactics). Moreover, the evidence from across the two originating studies reveals the considerable investment of effort, and willingness to navigate within the parameters of the framework of standard police operating procedures, required to implement AFR with a reasonable degree of efficacy and efficiency.

Given how it functions in support of street policing, facial recognition technologies are better understood as providing 'assisted facial recognition' rather than anything solely 'automated'. This is important in terms of acknowledging the complex socio-technical mediations that exist in operational AFR environments. The role of 'assistance', rather than 'automation', opens a discretionary space in which agency is enacted and techno-social interactions become negotiated. Particularly significant here, therefore, are operator decision-making activities involving discretionary and suspicious judgements over who should be stopped once a possible identification has been articulated by the algorithm. Ultimately then, this article has interrogated how technological capability is conditioned by police discretion, but police discretion itself is also contingent on the operational and technical environment.

The 'discovery' of police discretion and allied operational constructions of police suspicion in the 1960s was responsible for painting a more complex and nuanced set of relations between police behaviour and law. Multiple street-level ethnographies documented the distinction between 'law in books' and 'law in action'. This cast police as far more than simple law 'enforcers', as active decision-makers negotiating social order and against whom, when and why statutory instruments would be applied or not.

In their recent judgement, the Court of Appeal articulated two concerns over 'two impermissibly wide areas of discretion' afforded to police during AFR operations: (1) who is included on a watchlist and (2) where systems are deployed. Both issues are illuminated by our analysis. However, what the court's formulation seems to miss is that

discretionary decision-making infiltrates several other components of how AFR is used in street policing (i.e. confirming suspected matches). Thus, positioned in a more conceptual register, what our analysis demonstrates is how, although officer decision-making in digital policing systems may be algorithmically framed, it still draws upon forms of discretion redolent of the working practices familiar to previous generations of police.

The assertion that police discretion persists and recurs in terms of how leading-edge digital technologies are operationally deployed similarly guides our attention to a more supple and negotiated understanding of digital policing and the contributions of technologies such as AFR. For what has been illuminated by the empirical data is a complex series of layered interactions between operators, organizations and technical systems. These are vital in shaping what systems such as AFR can and cannot do.

There is a clear analogy here with how law in policing came to be understood over half a century ago. AFR algorithms steer and guide officer decision-making, but they do not wholly determine it. The rules encoded within the algorithms are not ‘unbending’ and inflexible but configured and constructed via a range of policing influences. Officers take active decisions about inclusion criterion for watchlists and how to configure the sensitivity of the system and, thus, which matches it is enabled to make or not. This, in turn, frames who becomes the ‘legitimate’ subjects of a reinvented and digitally mediated ‘bureaucratic suspicion’. The outcomes of such interactions are myriad affordances shaping the operation of AFR. It is for these reasons that we prefer the term ‘assisted’ rather than ‘automated’ facial recognition. For whilst the application of facial recognition technologies for verification purposes in border environments may be automated, this is not an accurate descriptor of its current functioning in street policing.

Thus, at the same time as claiming a recurrence of police discretion and suspicion, we are also suggesting a degree of reinvention and reconfiguration. There is a process of algorithmically mediated co-constitution that, in subtle ways, alters relations between police and their suspects. Accenting this facet of facial recognition technologies in policing does not, and should not, blind us to the fact that much concern has been expressed about it. There is, however, a difference between the ‘surveillance imaginaries’ and the ‘surveillance actualities’ of this and other digital policing technologies. Such interpretations have the potential to inform a more sophisticated and supple understanding of the configuration of Foucauldian biopolitics and the capillaries of power/knowledge in the information age.

Thus, in addition to making a substantive contribution to understanding AFR’s use in policing, we conceive of the evidence and insights presented as simultaneously helpful in relation to unpacking the broader domain of digital policing. Herein, we have explored the interactions between police operators and algorithmic processes to elucidate how they are mutually and recursively shaping and structuring one another. In much popular and political commentary on artificial intelligence and ‘big data’, there is a tendency for ‘algorithmic reification’ to be induced. The technologies concerned are attributed an almost mystical power, derived from their capacity to ingest and process vast quantities of raw data beyond what would be possible for any single human. They are after all, as Pasquale (2015) notes, ‘black boxes’ where it is difficult to determine how they are arriving at their conclusions. AFR simultaneously expands and constrains spaces for police discretion and suspicion. These varied articulations challenge more deterministic readings of police technology and also the kinds of discursive frames promoted by the commercial developers and suppliers of these systems.

Funding

Fussey's work was supported by the Economic and Social Research Council grant [ES/M010236/1]. Davies and Innes's study was funded via South Wales Police.

REFERENCES

- BAUMAN, A. and LYON, D. (2013), *Liquid Surveillance*. Polity.
- BIJKER, W. AND LAW, J. (eds) (1992), *Shaping Technology/Building Society: Studies in Sociotechnical Change*. MIT Press.
- BUOLAMWINI, J. and GEBRU, T. (2018), 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification', *Proceedings of Machine Learning Research*, 81: 1–15.
- CAMPBELL, E. (1999), 'Towards a Sociological Theory of Discretion', *International Journal of the Sociology of Law*, 27: 79–101.
- COLE, S. (2002), *Suspect Identities: A History of Fingerprinting and Criminal Identification*. Harvard University Press.
- DAVIES, S. (1996), *Big Brother*. Pan
- DAVIS, K. (1971), *Discretionary Justice: A Preliminary Inquiry*. Louisiana State University Press.
- DELEUZE, G. and GUATTARI, F. (1987), *A Thousand Plateaus*. Continuum.
- DIXON, D. (1997), *Law in Policing: Legal Regulation and Police Practices*. Clarendon.
- DWORKIN, R. M. (1977), *Taking Rights Seriously*. Harvard University Press.
- ERICSON, R. and HAGGERTY, K. (1997), *Policing the Risk Society*. Clarendon Press.
- FOUCAULT, M. (1977), *Discipline and Punish: The Birth of the Prison*. Penguin.
- (2007) *Security, Territory, Population: Lectures at the Collège de France, 1977–1978*. Palgrave Macmillan.
- (2008) *The Birth of Biopolitics: Lectures at the Collège de France, 1978–1979*. Palgrave Macmillan.
- FUSSEY, P. (2010) 'Protecting Britain's Crowded Spaces From Terrorist Attacks: Key Criminological Reflections', in A. Silke, ed., *Psychology, Terrorism and Counterterrorism*, 164–85. Routledge.
- FUSSEY, P. and COAFFEE, J. (2012), 'Balancing Local and Global Security Leitmotifs: Counter-terrorism and the Spectacle of Sporting Mega-Events', *International Review for the Sociology of Sport*, 47: 268–85.
- GATES, K. (2011), *Our Biometric Future: Facial Recognition Technology and the Culture of Surveillance*. NYU Press.
- GIBSON, J. 1977. 'The Theory of Affordances', in R. Shaw and J. Bransford, eds, *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*, 67–82. Erlbaum.
- GILL, P. (2000), *Rounding Up The Usual Suspects? Developments in Contemporary Law Enforcement Intelligence*. Ashgate.
- GOFFMAN, E. (1972), *Relations in Public: Microstudies in Public Order*. Basic Books.
- HAGGERTY, K. and ERICSON, R. (2000), 'The Surveillant Assemblage', *British Journal of Sociology*, 51: 605–22.
- HOUSE OF COMMONS SCIENCE AND TECHNOLOGY COMMITTEE (2019), *House of Commons Science and Technology Committee: The Work of the Biometrics Commissioner and the Forensic Science Regulator, Nineteenth Report of Session 2017–19*. House of Commons.
- HUTCHBY (2001), 'Technology, Texts and Affordances', *Sociology*, 35: 441–56

- INFORMATION COMMISSIONER'S OFFICE (2019), 'ICO Investigation Into How the Police Use Facial Recognition Technology in Public Places', available online at <https://ico.org.uk/media/about-the-ico/documents/2616185/live-frt-law-enforcement-report-20191031.pdf> (accessed on 1 March 2020).
- INNES, M. (2003), *Investigating Murder: Detective Work and the Police Response to Criminal Homicide*. Oxford University Press.
- INNES, M., ROBERTS, R., LOWE, T. and INNES, H. (2020), *Neighbourhood Policing: The Rise and Fall of a Policing Model*. Clarendon Press.
- INTRONA, L. D. and WOOD, D. (2004), 'Picturing Algorithmic Surveillance: The Politics of Facial Recognition Systems', *Surveillance and Society*, 2: 177–98.
- LATOUR, B. (1987), *Science in Action: How to Follow Scientists and Engineers Through Society*. Harvard University Press
- LYON, D. (2001), *Surveillance Society: Monitoring Everyday Life*. Open University Press.
- MACKENZIE, R., MARKS, A., and MORGAN, K. (2017), 'Technology, Affordances and Occupational Identity Amongst Older Telecommunications Engineers: From Living Machines to Black-Boxes', *Sociology*, 51: 732–48.
- MANNING, P. (1978), 'Rules, Colleagues and Situationally Justified Actions', in P. Manning and J. Van Maanen, eds, *Policing: A View from the Streets*, 71–90. Random House.
- MARX, G. (1988), *Undercover: Police Surveillance in America*. University of California Press.
- (2016), *Windows Into the Soul: Surveillance and Society in an Age of High Technology*. University of Chicago Press.
- MATZA, D. (1969), *Becoming Deviant*. Prentice-Hall.
- NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY (2019), *Ongoing Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects*. US Department of Commerce.
- NORRIS, C. and ARMSTRONG, G. (1999), *The Maximum Surveillance Society: The Rise of CCTV*. Berg
- PASQUALE, F. (2015), *The Black Box Society: The Secret Algorithms That Control Money and Society*. Harvard University Press
- QUINTON, P. (2011), 'The Formation of Suspicions: Police Stop and Search Practices in England and Wales', *Policing and Society*, 21: 357–68
- REINER, R. (1992), *The Politics of the Police (2e)*. Harvester Wheatsheaf.
- SKOLNICK, J. (1966), *Justice Without Trial: Law Enforcement in Democratic Society*. Wiley
- STANLEY, J. and STEINHARDT, B. (2002), 'Drawing a Blank: The Failure of Facial Recognition Technology in Florida, an ACLU Special Report', available online at <https://www.aclu.org/report/drawing-blank-report-tampa-police-records-reveals-poor-performance-face-recognition> (accessed on 9 January 2020).
- THE LAMMY REVIEW (2017), 'An Independent Review Into the Treatment of, and Outcomes for, Black, Asian and Minority Ethnic Individuals in the Criminal Justice System', available online at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/643001/lammy-review-final-report.pdf (accessed on 1 March 2020).
- UN OHCHR (2018), 'End of Mission Statement of the Special Rapporteur on the Right to Privacy at the Conclusion Of his Mission to the United Kingdom of Great Britain and Northern Ireland', available online at <https://www.ohchr.org/EN/NewsEvents/Pages/DisplayNews.aspx?NewsID=23296&LangID=E>.