



ARTICLE

# Criminal justice profiling and EU data protection law: precarious protection from predictive policing

Orla Lynskey\*

Law Department, London School of Economics

\*Corresponding author. E-mail: [O.Lynskey@lse.ac.uk](mailto:O.Lynskey@lse.ac.uk)

## Abstract

This paper examines the application of the latest iterations of EU data protection law – in the General Data Protection Regulation, the Law Enforcement Directive and the jurisprudence of the Court of Justice of the EU – to the use of predictive policing technologies. It suggests that the protection offered by this legal framework to those impacted by predictive policing technologies is, at best, precarious. Whether predictive policing technologies fall within the scope of the data protection rules is uncertain, even in light of the expansive interpretation of these rules by the Court of Justice of the EU. Such a determination would require a context-specific assessment that individuals will be ill-placed to conduct. Moreover, even should the rules apply, the substantive protection offered by the prohibition against automated decision-making can be easily sidestepped and is subject to significant caveats. Again, this points to the conclusion that the protection offered by this framework may be more illusory than real. This being so, there are some fundamental questions to be answered – including the question of whether we should be building predictive policing technologies at all.

**Keywords:** data protection; personal data; profiling; automated decision-making; predictive policing; law enforcement

## 1 Introduction

All four vectors of change analysed in this Special Issue – from technological assistance to replacement, from non-automated to automated processes, from humans being in the loop to being out of the loop and from *ex post* to *ex ante* criminal justice – are present when law enforcement authorities deploy technologies to facilitate ‘predictive policing’. Predictive policing is defined as ‘any policing strategy or tactic that develops and uses information and advanced analysis to inform forward-thinking crime prevention’ (Uchida, 2009, p. 1). The automated analysis of data for predictive purposes is a foundational concern of EU data protection law, reflected in the prescient inclusion in the 1995 Data Protection Directive<sup>1</sup> of a provision offering individuals a right not to be subjected to automated decision-making. However, doctrinal scholars have tended to focus on the commercial deployment of profiling and prediction techniques, in particular in the context of online behavioural advertising. Normative scholars have, for their part, identified the challenges that such a shift to automated decision-making will entail, with an entire field of interdisciplinary work coalescing around the notion of fair, accountable and transparent machine learning, or what could perhaps cumulatively be referred to as algorithmic due process (Keats Citron and Pasquale, 2014).

The EU’s General Data Protection Regulation (GDPR),<sup>2</sup> which entered into force in May 2018, provided a legal response to these normative challenges and, as such, attracted significant academic

<sup>1</sup>European Parliament and Council Directive 95/46/EC of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. [1995] OJ L281/23.

<sup>2</sup>Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). [2016] OJ L 119/1.

attention regarding the relevance of its provisions to automated decision-making (the lively disagreement between, *inter alia*, Wachter, Mittelstadt and Floridi, and Selbst and Powles in 2017 editions of *International Data Privacy Law* being one case in point). Although, as mentioned above, the 1995 Directive provided for similar protection, the GDPR's provisions have captured the *Zeitgeist*, in particular in relation to how they apply to the provision of 'free' personalised digital services where the harvesting of personal data acts as the *quid pro quo* for access to the service.

Thus, while the use of profiling techniques by the public sector, including in the context of criminal justice, has attracted the attention of normative scholars, there has been comparatively little scholarship on such criminal justice usage from a doctrinal perspective (with some exceptions, most recently Garstka (2018)). Indeed, it is the lesser-known and discussed legislative reform accompanying the GDPR, the Law Enforcement Directive (LED),<sup>3</sup> that will, ordinarily, govern automated decision-making in the context of law enforcement activities (the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties). This contribution therefore focuses on these predictive policing technologies through the underutilised doctrinal lens. It emerges from this scrutiny that the law in this area is both underdeveloped and of potentially little assistance to those affected by the use of predictive policing technologies. It is underdeveloped as fundamental questions, such as whether the processing of personal data required by predictive policing technologies falls within the scope of data protection law, are difficult to answer with certainty. It may be of little substantive assistance to those impacted by these technologies, as the substantive protection it offers in the event of automated decision-making is subject to a number of powerful caveats and provisos.

Before elaborating on this reasoning, it is worth noting that this contribution does not claim to be exhaustive in its application of the data protection rules to the predictive policing context. As such, important questions such as who is responsible for data processing (i.e. the data controller) and how the data protection safeguards (Article 4 LED) apply in the predictive policing context are not addressed and merit separate future consideration.

## 2 Identifying the relevant legal framework

Prior to the entry into force of the EU's Lisbon Treaty in 2009, the EU legally comprised three distinct pillars: one pillar each respectively for the European Communities; Common Foreign and Security Policy; and Police and Judicial Co-operation in Criminal Matters. This structure assured that the more extensive judicial oversight and more stringent legislative procedures that were applied in the context of first-pillar internal market measures, such as the Data Protection Directive, were not applied to the second and third pillars, the latter being areas more directly linked to the sovereignty of Member States. Despite the collapse of this pillar system by the Lisbon Treaty, and the introduction of an explicit legal basis for all EU data protection laws in Article 16 TFEU, the EU legal framework for data protection remains differentiated along these historic lines.

The data protection reform package that culminated in the entry into force of the GDPR and the LED in 2018 continues to separate data processing for law enforcement purposes from 'general' data-processing operations – a possibility that was foreseen in Declaration No. 21, which accompanied the Lisbon Treaty. This section shall therefore provide a brief overview of the relevant legal framework applicable to data processing in the criminal justice context.

### 2.1 The legislative framework

The 2018 legislative package incorporates two legal instruments: the GDPR and the LED. The GDPR, as its title suggests, sets out a generally applicable framework that regulates the processing of personal

<sup>3</sup>Directive (EU) 2016/680 of the European Parliament and Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties. [2016] OJ L119/89.

data (Art. 2(2) GDPR). This regulation is generous in scope, as discussed below, with the core concepts of ‘personal data’ and ‘processing’ being expansively defined and interpreted. Once within the regulation’s scope, individual ‘data subjects’ are the beneficiaries of certain rights, the cornerstone of which – a right to access their personal data – enables the exercise of a number of corollary rights (Arts 12–22 GDPR), such as a right to rectify or erase. Data controllers – those who determine the purposes and means of personal data processing – are accountable for this processing pursuant to the GDPR (Art. 5 (2) GDPR). Controllers must respect the principles, or safeguards, relating to personal data processing and ensure the lawfulness of that processing (Arts 5 and 6 GDPR).

The GDPR allows Member States to enact limitations to specified Chapters or provisions in certain contexts, notably where necessary to reconcile data protection rights with freedom of information and expression (Art. 85) or for archiving and related purposes (Art. 89). It also allows Member States to introduce legislation restricting the application of the aforementioned data protection principles and rights in order to pursue specified purposes, including national security, defence, public security and law enforcement purposes (Art. 23), provided this restriction is necessary and proportionate in a democratic society and respects the essence of fundamental rights and freedoms. However, beyond these possible restrictions to specific provisions, Article 2(2) GDPR also excludes from its scope entirely certain other policies, including: policies on border checks, asylum and immigration, and police and judicial co-operation in criminal matters (Art. 2(2)(b) GDPR); data processing in the course of an activity that falls outside the scope of Union law (Art. 2(2)(a) GDPR); and data processing by competent authorities for the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, including the safeguarding against and the prevention of threats to public security (Art. 2(2)(d) GDPR).

Article 2(2)(d) GDPR therefore paves the way for the LED, which applies to precisely this activity: namely data processing by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, including the safeguarding against and the prevention of threats to public security. This explicit exclusion of ‘law-enforcement data processing’ from the GDPR may give the initial misleading impression that the division of labour between the two legislative instruments is clear. However, the interplay between these two potentially relevant instruments is, in fact, complicated.

First, it is worth emphasising that, in order to fall within the scope of the LED, the data processing must be undertaken by a ‘competent authority’. A competent authority is defined in Article 3(7)(a) and (b) LED as

‘any public authority competent for the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, including the safeguarding against and prevention of threats to public security; or any other body or entity entrusted by Member State law to exercise public authority and public powers for these purposes.’

A private actor conducting data processing for law enforcement purposes therefore needs to be entrusted with this role by Member State law before it is removed from the material scope of the GDPR and the relevant provisions of the LED apply. It follows that, in the absence of such a legislative enactment, the provisions of the GDPR continue to apply to private entities processing personal data for law enforcement purposes. This is implicitly confirmed by Article 23 GDPR: this provision allows Union or Member State law to restrict the obligations stemming from specified GDPR provisions where such a restriction is necessary for law enforcement purposes, thus acknowledging that the GDPR would otherwise apply to such processing in certain circumstances. A lot therefore turns on the designation of a ‘competent authority’ for LED purposes and when a private entity could be said to be entrusted by law with such status. Nevertheless, one might query, for instance, whether ‘entrustment’ requires a distinct legislative instrument (Garstka, 2018).

Second, even when data processing is undertaken by a ‘competent authority’, whether such processing falls within the scope of the GDPR or the LED will depend on the purpose of the processing. Thus,

when it comes to data sharing by competent authorities, the act of transmitting data from a competent authority to a non-competent authority for law enforcement purposes is within the scope of the LED (e.g. a data transfer from the police to a private predictive policing software provider) while a transfer for non-law enforcement purposes (e.g. a data transfer to medical or social services) would be covered by the GDPR (Recital 34, LED). Equally, where a competent authority initially collects personal data for law enforcement purposes but these data are then processed for alternative, non-law enforcement purposes, the GDPR applies (Art. 9(1) and Recital 11 LED). Garstka characterises the switch between legal frameworks, from the LED to GDPR, as a ‘downgrade’, although, as shall be discussed below, the LED is not necessarily more protective than the GDPR when it comes to automated decision-making (Garstka, 2018).

Perhaps most critically, it may be that neither legal instrument applies to certain personal data-processing operations. Like the GDPR, the LED does not apply to personal data processing for national security purposes (Art. 2(3) and Recital 14 LED). The LED does, however, incorporate data processing by competent authorities for law enforcement purposes, ‘including the safeguarding against and prevention of security threats’ (Art. 1(1) read in conjunction with Art. 2(1) LED). Nevertheless, the distinction between law enforcement activity, public security and national security activity is a blurred one. For example, Recital 12 states that the LED applies to police activities before the police have knowledge of whether an incident is criminal in nature and that such activities include coercive measures at demonstrations and riots. It might equally be argued by a Member State that state oversight and enforcement at a political rally or protest engages national security concerns, with the recent protestations of the *Gilets Jaunes* in France providing a good example (BBC, 2018). Yet, the LED reserves to the Court of Justice of the EU (CJEU) the controversial task of delineating the dividing line between criminal enforcement activity and national security measures through the ostensibly innocuous assertion that the notion of ‘criminal offence’, within the meaning of the LED, is ‘an autonomous concept of EU law’ (Recital 13 LED). One could imagine that any attempt by the EU to define as ‘criminal’ conduct that would otherwise be classified as a ‘national security’ threat will be highly contested given that national security is listed amongst the ‘essential state functions’ that the Union is required to respect pursuant to Article 4(2) TEU. Indeed, Article 4(2) TEU specifically emphasises that ‘national security remains the sole responsibility of each Member State’.

## 2.2 Article 8 ECHR and the EU Charter

Challenges to the use of data by the state and law enforcement authorities for criminal justice purposes have traditionally been anchored in Article 8 ECHR, which provides for the right to respect for private life (e.g. *S and Marper v. UK*<sup>4</sup>). This is for several reasons. Most evidently, the EU Charter of Fundamental Rights only acquired binding status in December 2009, by which point a considerable body of Article 8 ECHR jurisprudence had developed. Yet, even since the Charter’s entry into force, it has been of limited utility in the policing context, as the Charter applies to Member States ‘only when they are implementing Union law’ (Art. 51(1) EU Charter). Before the enactment of the LED, no Union law applied to the processing of personal data by police authorities for domestic purposes, thus excluding the application of the Charter to such processing. Now that such legislation does exist, the Charter’s application in this context raises queries.

Two issues in particular are likely to be contentious. First, as mentioned above, by reserving for the CJEU the competence to define ‘criminal offence’ for the purposes of the LED, the LED de facto enables the CJEU to police the outer limits of the concept of ‘national security’. A broad interpretation of ‘criminal offence’ could encroach upon what Member States perceive to be their exclusive competence in the field of national security and would allow the Charter to apply in this new territory. Although Article 6(1) TEU states that the provisions of the Charter shall not extend in any way the competences of the Union as defined in the Treaties, the CJEU has not hesitated to deploy

<sup>4</sup>*S and Marper v. United Kingdom* (2009) 48 EHRR 50.

bold readings of legislative provisions that ensure the Charter's application in contested contexts. For instance, in *Tele2 Sverige and Watson*,<sup>5</sup> following on from *Digital Rights Ireland*,<sup>6</sup> the court examined the compatibility of national legislative provisions providing for bulk communications meta-data retention for law enforcement purposes in light of the EU Charter rights to data protection and privacy. What was notable about this situation was that, at that time, there was no provision of EU law regulating data retention by law enforcement authorities or access to retained data by them. The court therefore justified the Charter's application in a contestable manner. It reasoned that the EU E-Privacy Directive provides for a right to confidentiality of communications; that Article 15(1) E-Privacy Directive provides for exceptions to this principle, including for public security and law enforcement purposes; and that such exceptions must comply with general principles of EU law, including fundamental rights. The court therefore held that national legislative measures adopted pursuant to Article 15(1) E-Privacy Directive must be interpreted in light of the Charter (*Tele2 Sverige and Watson*, para. 91). The court proposed this interpretation despite the fact that Article 1(3) E-Privacy Directive states that it shall not apply to, inter alia, the activities of the state in areas of criminal law. The court's application of the Charter provisions to communications meta-data retention, and in particular its potential encroachment into the field of national security, has been contested at the national level (see e.g. the pending preliminary reference from the UK: *Privacy International*<sup>7</sup>). Second, the level of protection offered by the LED is not identical to that offered by the GDPR. One may therefore wonder whether such differentiated protection, such as the omission of certain 'due-process' safeguards in the context of automated decision-making in the LED (as discussed below), is in fact compatible with the Charter right to data protection.

### 3 The deployment of intrusive profiling technologies

According to *Big Brother Watch* – a civil society organisation in the UK – there is an increasing trend for police forces in the UK to acquire, develop and operationally deploy technologies that are intrusive, untested and of questionable compatibility with fundamental rights (*Big Brother Watch*, 2018). One such range of technologies are those used for 'predictive policing' purposes. Predictive policing is, as mentioned above, any policing strategy or tactic that develops and uses information and advanced analysis to inform forward-thinking crime prevention. Degeling and Berendt specify that such policing encompasses the 'variety of techniques used by police departments to generate and act on crime probabilities, often referred to as *predictions*' (Degeling and Berendt, 2018, emphasis in original). It is thus the ostensible capability to predict future criminal outcomes based on big data analytics that is the decisive characteristic of these predictive policing technologies. Numerous categories of predictive policing have been identified in the literature. For instance, Oswald *et al.* identify three contexts in which predictive policing occurs: it can be incorporated into strategic planning and prioritising on a macro-level; it can be used to link operational intelligence; and it can be used to make decisions or risk assessments in relation to individuals (Oswald *et al.*, 2018). Broadly speaking, therefore, predictive policing can be used to make both systemic decisions (perhaps relating to times and places of crimes) as well as identification decisions (e.g. predicting the identity of alleged criminal offenders or victims of crime) (van Brackel, 2016). Before analysing whether and how data protection law constrains such policing predictions, a brief description of an application illustrative of systemic and identification decision-making respectively will be provided.

#### 3.1 Applications of predictive policing

Systemic decisions in this context are aggregate predictions regarding future criminal activity. This form of predictive policing has been the most widely adopted. It seeks to undertake what van

<sup>5</sup>Joined Cases C-203/15 and C-698/15, *Tele2 Sverige, Watson and Others* EU:C:2016:970.

<sup>6</sup>Joined Case C-293/12 and 594/12 *Digital Rights Ireland Ltd and Seitlinger and Others* EU:C:2014:238.

<sup>7</sup>Case C-623/17 *Privacy International* [2018] OJ C22/29 (pending).

Brakel labels ‘predictive mapping’ (or geo-spatial crime prediction): identifying when and where crimes may take place based on an aggregate-level analysis. The most well-known commercial example of this predictive mapping is the PredPol application. According to its website, PredPol sells a web application to police forces making predictions about possible future crime hotspots. PredPol was developed by researchers at UCLA and is now used by over sixty US police departments. In the UK, Kent Police uses the software (Dencik *et al.*, 2018). PredPol is sold as a ‘Software as a Service’ package, which means that it is run on centrally controlled servers rather than locally by each police force using it. The application can then be accessed using an ordinary web browser.

PredPol was inspired by the use of seismology algorithms. In seismology, aftershocks often occur after a first earthquake: it is therefore possible to assume that, after a first earthquake, the likelihood of a repeat event increases. PredPol, using a ‘near repeat theory’ model, extends this logic to the criminal context where certain crimes, most clearly burglary, occur in quick succession in close proximity to the initial crime (although, as Degeling and Berendt note, only a fraction of crimes can be considered ‘near repeat’). In predicting ‘crime hotspots’, PredPol makes use of three data variables: crime type, crime location and timing data (date and time).

Predictive mapping, or geo-spatial crime prediction, technologies have been developed over time in conjunction with risk terrain modelling, with the application commercialised by Hunchlab being one example of this. While this model is still rooted in geo-spatial crime prediction, it entails the creation of a ‘composite “risk terrain” map with values that account for all risk factors at every place throughout the geography’ (Caplan *et al.*, 2015, p. 8). Degeling and Berendt distinguish between the two by highlighting that, while near repeat theory focuses on internal/endogenous factors (such as repetitive criminal behaviour), risk terrain modelling incorporates exogenous factors into this assessment (e.g. the location of bars or parks). The environmental factors that potentially contribute to particular criminal activity (e.g. dimly lit areas in urban environments or potential drug-dealing in parks) can be factored into risk terrain modelling. This does, however, put an onus on those using this technology ‘to constantly update their dataset and inform the algorithms about any change of usage of buildings, construction sites or the location of events’ (Degeling and Berendt, 2018, p. 353).

The second broad category of predictive policing applications are those that are used to make identification decisions – that is, predictions regarding actual or potential offenders. Unlike systemic decisions, predictive identification focuses on identifying criminal behaviour at the individual or group level. In particular, such systems seek to calculate the likelihood that a given person will commit a crime or their likelihood of being prone to criminal behaviour. A well-known example of such identification profiling is used in Chicago, where the police department analyses networks of those arrested in order to calculate the likelihood of an individual being involved in a serious crime. In order to compile a ‘heat list’ of names of those likely to be involved in major crimes, various data points are analysed by the police. Individuals on this list are then provided with a ‘custom notification letter’ by the police that warns them, in advance, about any charges they may face should they engage in criminal activities. In the UK, the Durham Constabulary in conjunction with statistical researchers at the University of Cambridge developed a ‘Harm Assessment Risk Tool’ (HART) to assess risk of future offending and to provide those offenders classified as ‘moderate-risk’ with the opportunity to participate in a rehabilitation programme. HART is explicitly designed to aid custody officers in their decision-making regarding eligibility, rather than to make this decision for them (Oswald *et al.*, 2018). According to Oswald *et al.*, the HART model is built using approximately 104,000 ‘custody events’ over the five-year period between 2008 and 2012. Forecasts are determined based on thirty-four distinct predictors, twenty-nine of which focus upon the suspect’s offending history, with the remaining five being age, gender, two forms of residential postcode and ‘the count of existing police intelligence reports relating to the offender’ (Oswald *et al.*, 2018, p. 228). *Big Brother Watch* has indicated that the residential postcode data are based on a tool offered by Experian, a global data broker that it suggests uses 850 million data points, including information on child benefits and support and data scraped online, to create a ‘postcode stereotype’ (*Big Brother Watch*, 2018).

The HART model intentionally favours so-called ‘cautious error’ over ‘dangerous error’: a dangerous error would underestimate the offender’s risk of reoffending while a cautious error overestimates the risk of reoffending. While the accuracy rate of the HART model was assessed to be 62.8 percent in 2012, which reflected a drop from the initial construction estimate of 68.5 percent, Oswald *et al.* emphasise that the rates of the most dangerous form of error (low-risk offenders actually being high-risk) remained constant and it was cautious errors that were overestimated. This means, as they acknowledge, that, if the HART prediction was treated as decisive, one could argue that an unacceptable number of low- and medium-risk offenders are treated as high-risk.

## 4 The application of the underdeveloped data protection framework

### 4.1 The scope of application of EU data protection law

EU data protection provisions apply to the processing of personal data. One preliminary query may be whether the two forms of automated decision-making associated with systemic and identification predictive policing constitute personal data processing and thus fall within the scope of the provisions. ‘Processing’ is an all-encompassing concept defined in Article 3(2) LED and Article 4(2) GDPR as:

‘any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction.’

Each stage of the data handling outlined above would therefore constitute data ‘processing’. The more interesting question is thus whether it constitutes the processing of *personal* data. Personal data are defined as any information that relates to someone who is identified or identifiable on the basis of that data (Art. 3(1) LED; Art. 4(1) GDPR). Opining in 2007, the Article 29 Working Party, an advisory body on data protection matters constituted under the Data Protection Directive, proposed a very broad interpretation of ‘personal data’ (Article 29 Working Party, 2007). However, the CJEU has also had the opportunity to consider this concept in its jurisprudence in recent years, most notably in *YS and MS*,<sup>8</sup> *Breyer*<sup>9</sup> and *Nowak*,<sup>10</sup> and this jurisprudence has not always been unequivocal in endorsing such a broad approach to ‘personal data’. Therefore, as shall be discussed, there continues to be doctrinal debate and uncertainty regarding the appropriate reach of this notion (see e.g. Veale *et al.*, 2018).

#### 4.1.1 Parsing the concept of ‘personal data’

The definition of personal data has three constituent elements: personal data are (1) any information that (2) relates to (3) an identified or identifiable person. The Article 29 Working Party has advised that each of these three elements should be interpreted expansively. It thus suggests that ‘any information’ should incorporate objective and subjective information, information that is false and that the format in which the information is provided is not relevant. ‘Any information’ also includes information that would not be considered ‘private’ for the purposes of the right to respect for private life in Article 8 ECHR (this also follows e.g. from *Google Spain*<sup>11</sup>). The CJEU stated in *Nowak* that the expression ‘any information’ is used to reflect the legislature’s aim to ‘assign a wide scope to that concept’. The court continued by stating that it is not restricted to sensitive or private information and

<sup>8</sup>Joined Cases C-141/12 and 372/12 *YS v. Minister voor Immigratie, Integratie en Asiel and Minister voor Immigratie, Integratie en Asiel v. MS* EU:C:2014:2081.

<sup>9</sup>Case C-582/14, *Breyer* EU:C:2016:779.

<sup>10</sup>Case C-434/16, *Nowak v. Data Protection Commissioner* EU:C:2017:994.

<sup>11</sup>Case C-131/12, *Google Spain SL, Google Inc. v. Agencia Española de Protección de Datos and Mario Costeja González* EU:C:2014:317.

‘potentially encompasses all kinds of information’, giving the example of subjective information (para. 34). The court failed to endorse explicitly all of the examples provided by the Article 29 Working Party, while also qualifying its claim that any information encompasses all kinds of information (it *potentially* does so: emphasis added). There is therefore scope for the court to limit its interpretation of this definitional element in the future. Nevertheless, in the court’s existing jurisprudence, the more contested elements of ‘personal data’ are the requirements that this information ‘relates to’ an individual and that this individual is identifiable.

In its 2007 Opinion, the Article 29 Working Party suggested that information can relate to an individual in three ways, namely in terms of its content (when it is about a particular person); its purpose (when data are used with the purpose ‘to evaluate, treat in a certain way or influence the status or behaviour of an individual’) and its result (when it ‘is likely to have an impact on a certain person’s rights and interests, taking into account all the circumstances surrounding the precise case’). It further specified that it is not necessary that the information ‘focuses on someone in order to consider that it relates to him’.

The court was asked for guidance on this element of the definition in *YS*. *YS* concerned three non-EU nationals who applied for lawful residence in the Netherlands. During this process, the immigration official tasked with assessing the application prepares a draft decision (essentially, a recommendation on whether to grant the status requested) and a ‘minute’, which sets out the legal reasoning on which the draft decision is based. The minute also contains details of the relevant case officer, factual data regarding the applicant (name, nationality, etc.) as well as the documents submitted by the applicant and the procedural history of the claim. Each of the three applicants sought access to the minute relating to their application, claiming that, under data protection law, the minute constituted personal data and that they therefore had a right of access to such data. It was not contested before the court that the data contained in the minute about the applicant (what might be called the ‘content’ data, such as name, date of birth, gender, language, etc.) constituted personal data and the court confirmed this finding (para. 38). However, the court held that the legal analysis in the minute, which at most ‘is information about the assessment and application by the competent authority of the law to the applicant’s situation’, did not ‘relate to’ the applicant (para. 40). It was therefore not personal data (para. 39). The court claimed this conclusion was borne out by the objective and general scheme of the Data Protection Directive (para. 41). In particular, the court reasoned that the directive’s objectives of protecting the right to privacy are achieved by enabling data subjects to ensure that their data are correct and lawfully processed by exercising their right of access to those data (paras 42–44). It therefore concluded that, as the legal analysis in *YS* was not itself liable to such an accuracy check and subsequent rectification, extending access to the legal analysis in this case would not serve the directive’s purposes, but would rather constitute a right of access to administrative documents (paras 45–46).

It is suggested that this approach to the interpretation of ‘personal data’ is misguided. On the one hand, the court has frequently favoured a teleological or purposive approach to its interpretation of the EU data protection rules (as occurred in *Google Spain*). *YS* was therefore consistent in this regard. On the other hand, the application of data protection law could be thought of as a two-stage process. The first stage entails an assessment of whether the processing operation falls within the scope of the rules, with a consideration of which rights and responsibilities apply to the processing and how they apply taking place during the second stage. In *YS*, the court allowed its consideration of the second question to influence its adjudication on the first, logically prior issue by limiting the concept of ‘personal data’ to ensure an outcome where the right to access would not apply. The court had the opportunity to reconsider this approach in *Nowak*.

Mr Nowak had sat a number of professional examinations in Ireland. Having failed one of these examinations on a number of occasions and having unsuccessfully challenged the result of his fourth attempt, Mr Nowak then sought access to his personal data from the professional association. The professional association provided him with access to seventeen documents but refused to provide access to his examination script on the grounds that it did not constitute personal data. This refusal culminated



in a referral by the Irish Supreme Court to the CJEU querying whether the information recorded as answers in a professional examination constituted personal data and, if so, what factors are relevant in such a determination. When considering these questions, the court also assessed whether the comments made by an examiner in respect of those answers constitute the candidate's personal data.

The court's analysis focused primarily on when data 'relate to' an individual implicitly endorsing the Article 29 Working Party's claim that data relate to a data subject where the information by reason of its content, purpose or effect, is linked to a particular person (Article 29 Working Party, 2007). The court affirmed that the written answers submitted by an examination candidate are linked to that candidate (para. 36). It elaborated that the content of the answers reflects the candidate's knowledge and competence (para. 37); that the purpose of the data collection is to evaluate the candidate (para. 38); and that the use of that information in determining whether the candidate passed or failed is likely to have an effect on the candidate's rights (para. 39). With regard to the examiner's comments, the court held that the content of these comments reflects the opinion or assessment of the examiner of the individual's performance (para. 43). Moreover, the purpose of those comments is to record this evaluation and the comments are likely to have effects for the candidate (para. 43). The court therefore concluded that such comments could be both the personal data of the examination candidate and the personal data of the examiner (paras 44 and 45).

At this juncture, the court ostensibly departed from its prior reasoning in *YS* by disaggregating the scope of the rules from the ensuing rights and responsibilities (what it called 'classification' and 'consequences'). It noted that whether these answers and comments should be classified as personal data cannot be affected by the fact that the consequence of the classification is, in principle, that the candidate then has a right to access and rectification (para. 46). The court noted that, if data are not classified as 'personal data', they have the effect of entirely excluding that data from data protection's principles and safeguards and its rights (para. 49). It then went on to elaborate on how principles such as the accuracy and completeness of these data, as well as rights such as the right to rectification, may be relevant in the examination context. This led it to conclude that, in so far as these answers and comments are liable to be checked for accuracy and retention and to be subject to rectification or erasure, giving the candidate a right of access to both serves 'the purpose of the Directive of guaranteeing the protection of that candidate's right to privacy ... (see, *a contrario*, ..*YS and others*)' (para. 56).

The court's enigmatic reference to *YS* leaves room for disagreement as to when data 'relate to' an individual. It could be argued that *YS* is no longer good law following *Nowak* and that the court used the opportunity presented to endorse the broad Article 29 Working Party conception of personal data. However, an alternative interpretation is also possible: that *YS* still stands and can be distinguished from *Nowak* because, in the latter, it was possible for Mr Nowak to carry out the 'necessary checks' to ensure that his personal data were correct and lawfully processed and thus these constituted personal data. In sum, we can see that, as with the term 'any information', there remains scope for contestation regarding the meaning of the term 'relates to'.

According to the GDPR, an identifiable person is one 'who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity' (Art. 4(1) GDPR). Recital 26 specifies that, to determine whether a natural person is identifiable, 'account should be taken of all the means likely reasonably to be used either by the controller or by any other person to identify the said person'. In *Breyer*, Mr Breyer had contested the registration and storage of his dynamic Internet protocol (IP) address by the German state when he accessed several Internet sites run by the German federal institutions. The state also recorded the date on which the website was accessed.

The classification of the IP address as 'personal data' therefore turned on whether Mr Breyer is 'identifiable' on the basis of this address, given that the additional data necessary for the German state to identify the website user are held by the user's Internet service provider (ISP) (para. 39). This gave the court the opportunity to consider the meaning of 'indirect' identification involving 'all the means likely reasonably to be used' by either the controller or by 'any other person'. It found that this wording indicates that it is 'not required that all the information enabling the

identification of the data subject must be in the hands of one person' (para. 43). It thus examined whether the possibility for the German state to combine the dynamic IP address with the additional identifying information held by the ISP constituted a means likely reasonably to be used. It held that such identification would not be possible if it was prohibited by law or 'practically impossible, on account of the fact that it requires a disproportionate effort in terms of time, cost and man-power' (para. 46). The court observed that, while German law does not generally allow the transmission of such information between the ISP and the state, in the event of issues such as cyber-attacks, 'legal channels exist' for website providers to contact competent authorities who in turn can take the steps necessary to obtain this identifying information from ISPs (para. 47). On this basis, the court concluded that the website operators had the means 'which may likely reasonably be used' in order to identify the data subject, with the help of competent authorities and the ISP and thus that the dynamic IP address constituted personal data.

It can thus be seen that the concept of 'identifiability' has been interpreted extremely liberally, with no requirement that the information enabling identification must be in the hands of one person and with even very onerous steps constituting 'means reasonably likely to be used', rather than a 'disproportionate effort'. Indeed, the Article 29 Working Party had previously opined that a 'purely hypothetical possibility' of identification does not meet the standard of 'likely reasonably to be used' (Article 29 Working Party, 2007). It had suggested that a range of factors should be taken into consideration when assessing such likelihood including the purposes of the processing (when it 'only makes sense if it allows identification of specific individuals and treatment of them in a certain way'), the cost of identification and measures in place to prevent identification, amongst others. It is necessary now to consider how these jurisprudential developments might apply in the context presently being considered.

#### 4.1.2 Application to the predictive policing context

With regard first to systemic decisions, it might seem counter-intuitive to suggest that an application like PredPol, which, it is recalled, processes three types of input data (crime type, crime location and timing data) in order to predict potential sites of future criminal activity, processes 'personal data'. The functioning of the application might be thought of in three stages: the aggregation and inputting of data; the application of the PredPol algorithm to that data; and the recommendation. Yet, based on the jurisprudence outlined above, it is possible that all three stages entail personal data processing. First of all, as Purtova notes, the primary reason for information processing in the context of data-driven regulation is to treat people in a certain way or to influence human behaviour (Purtova, 2018, p. 55). The data processed at all three stages are therefore likely to 'relate to' a data subject by reason of their 'purpose' (to treat people in a certain way) or their 'effect' (to impact upon those in the 'hot-spots' identified). Yet, as outlined previously, a narrower interpretation is possible. Arguably, post *Nowak*, it remains possible that the interpretation of 'relating to' will be reverse-engineered by the court, with the court considering first whether key data protection safeguards and rights (such as accuracy and erasure) can be exercised in relation to data before classifying that data as personal data.

Whether these data could be linked to an identifiable person is equally difficult to establish. It could be argued that linking, for instance, crime location data to an individual is hypothetically possible but does not meet the 'likely reasonable' standard. However, the court in *Breyer* held that data could be linked to an identifiable person if, where there was a cyber-attack on a website, the website owner could liaise with a competent authority who would in turn liaise with authorities with identifying information (in that instance, an ISP). Similarly, in this context, if a criminal event did occur, it seems likely that a competent police authority could combine these input data with other data, for instance CCTV footage or mobile-phone meta-data, in order to identify individuals present. This same logic applies also to the other two stages of personal data processing.

The broad understanding of the concept of 'personal data' might therefore be welcomed in this context in so far as it includes such uses of technology within the protective scope of the data-protection rules. Prior to these judgments, it was clear that, while the identification of potential crime hotspots might have ethical implications, it was less clear whether identifiability required a

‘singling-out’ of particular individuals rather than a geographic area to be captured by data protection law. Indeed, Taylor had warned that, in order to operationalise the concept of data justice in the context of new technologies that tended ‘to sort, profile and inform action based on group rather than individual characteristics and behaviour ... it is inevitably going to be necessary to look beyond the individual level’ (Taylor, 2017, p. 8). Yet, whether intentionally or inadvertently, this case-law now seems to extend the application of data protection rules to almost all forms of data, irrespective of their proximity to the data subject. Purtova provocatively, yet quite rightly, queries whether even weather data collected as part of a smart-city project might be classified as personal data pursuant to this broad approach, eventually turning ‘data protection law into an uneconomical exercise of *regulating everything*’ (Purtova, 2018, p. 59).

Given that systemic decisions likely constitute personal data processing, it is unsurprising that identification decisions also entail such processing. HART provides a good example of this. The input data – the thirty-four predictors based on the offender’s personal history – are personal data, as their content concerns an identified person, while the output data – a suggested likelihood of reoffending – are personal data, as their purpose and effect are to influence an identified individual’s future prospects. It is perhaps less clear that the ‘middle’ processing stage – where the programme’s algorithm is applied to the input data – are personal data based on *YS* and *Nowak*. Based on *Nowak*, it could be argued that the application of the machine-learning model used by HART to the thirty-four predictors input into the system are personal data. However, if *Nowak* did not overrule *YS*, then it is perhaps necessary to consider whether – in defining them as personal data – the protections that would follow from that classification, such as the right to erase or rectify, are theoretically and practically available. If this was what the court was in fact proposing in *Nowak*, one must remark that it makes the applicability of data protection difficult to predict, particularly for individuals, the ostensible beneficiaries of the regime.

Thus, in conclusion, it can be seen that it is likely, yet not certain, that these two forms of predictive policing models will fall within the scope of application of the data protection framework. Key aspects of the concept of ‘personal data’ that determine what processing activities these rules cover are subject to contestation and will require further jurisprudential clarification. Furthermore, as shall now be discussed, even assuming these models fall within the scope of the rules, the protection they offer to individuals impacted by such processing is precarious.

#### 4.2 Substantive rights and predictive policing

The LED defines profiling in the same way as the GDPR. Profiling is

‘any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person’s performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements.’ (Art. 3(4) LED; Art. 4(4) GDPR)

If, as suggested above, predictive policing systems entail the processing of personal data, then it follows that those systems are ‘profiling’ systems in that they predict aspects relating to a natural person’s reliability and behaviour as well as potentially their location and movements. Profiling is itself thus a form of automated decision-making (Article 29 Working Party, 2017).

Article 11(1) LED prohibits decisions based solely on automated processing in certain circumstances. This provision is worded similarly to Article 22 GDPR leading the Article 29 Working Party to indicate that its guidelines on Article 22 GDPR are relevant to Article 11 LED, ‘albeit with important caveats and specifications’ (Article 29 Working Party, 2017, p. 11).

It is worth setting out Article 11 LED in full:

‘Member States shall provide for a decision based solely on automated processing, including profiling, which produces an adverse legal effect concerning the data subject or significantly affects him or her, to be prohibited unless authorised by Union or Member State law to which the controller is subject and which provides appropriate safeguards for the rights and freedoms of the data subject, at least the right to obtain human intervention on the part of the controller.’

One immediate observation about this provision is that, unlike Article 22 GDPR, which is framed as a right of the data subject, the ‘data subject shall have the *right* not to be subject to a decision based solely on automated processing’ (emphasis added), Article 11 LED is framed as a prohibition. On paper, Article 11 LED is therefore phrased in a more robust and protective way than its GDPR counterpart. This has led the Article 29 Working Party to recommend that Article 22 GDPR nevertheless be applied as a prohibition (Article 29 Working Party, 2018). As Kaminski has noted, interpreting Article 22 GDPR as a right would, perhaps counter-intuitively, limit the protection it offered by allowing data controllers to ‘regularly use algorithms in significant decision-making, adjusting their behaviour only if individuals actually invoke their rights’ (Kaminski, 2018, p. 4).

Yet, despite its tougher exterior, it is immediately apparent that the Article 11 LED prohibition is subject to a number of limitations. Most notably, like Article 22 GDPR, automated decision-making is permissible when provided for by Union or Member State law. By enacting legislation, Member States can therefore legitimise reliance by law enforcement authorities on fully automated decision-making to make systemic and/or individualised predictions regarding future criminal conduct. Very importantly, and in accordance with EU antidiscrimination law and the EU Charter (as per Recital 38 LED), the LED does however put in place a *per se* prohibition on profiling that leads to discrimination on the basis of sensitive data (Art. 11(3) LED). Moreover, unlike the GDPR, which allows automated decision-making if it is necessary for contractual purposes or is based on the data subject’s explicit consent, this is not possible under Article 11 LED. The obvious explanation for this is that ‘there is a clear imbalance of powers between the data subject and the controller’ in this context (Article 29 Working Party, 2017, p. 12). Indeed, the court has recognised the limitations of consent when citizens do not have the ability to object to data processing (in situ, the ability to object to fingerprints being used for a passport application).<sup>12</sup>

However, while the GDPR does allow automated decision-making in a wider range of circumstances, it also provides for a wider array of safeguards for individuals. The LED allows automated decision-making if authorised by Union or Member State law, subject to the condition that such law ‘provides appropriate safeguards for the rights and freedoms of data subjects, at least the right to obtain human intervention on the part of the controller’ (Article 11 LED). The non-binding recital further specifies that, in order to be significant, human intervention ‘must be carried out by someone who has the appropriate authority and capability to change the decision’ (Recital 38). In contrast, Article 22(3) GDPR sets out further safeguards in addition to this ‘human intervention’ criterion. It specifies that the data subject shall have the right to express their point of view and to contest the decision. This may indicate that these latter safeguards – the right to express their point of view and to contest the decision – are not required under the LED. The LED is not, however, a directive of maximum harmonisation and it explicitly states that Member States can provide for a higher level of protection: how this provision is transposed domestically will therefore be of real importance for the rights of data subjects (Recital 15; Art. 1(3) LED). Moreover, when the automated decision-making is based on sensitive data, Article 11(2) LED specifies that such decision-making shall not occur unless suitable measures to safeguard the data subject’s rights and freedoms and legitimate interests are in place. Member States are therefore left with a lot of responsibility – and discretion – when it comes to ensuring the rights of individuals, albeit that any national implementing legislation must be compatible with the EU Charter.

<sup>12</sup>Case C-291/12, *Michael Schwarz v. Stadt Bochum* EU:C:2013:670.

Beyond the significant possibility to sidestep the prohibition on automated decision-making through legislation, Article 11 LED is limited in other notable ways. Most importantly, the LED applies only to automated decisions based *solely* on automated processing. As a result, it is questionable, for instance, whether the recommendation generated by HART for relevant officers regarding rehabilitation prospects is one that is based 'solely' on automated processing. This depends on how the decision-making process occurs in practice. In this context, one would need to gauge to what extent the final decision entails the discretion and judgment of the officer making that decision. If, for instance, the relevant officers '(consciously or otherwise) prefer to abdicate responsibility for what are risky decisions to the algorithm', then, in addition to the deskilling and judgmental atrophy that might follow (Oswald *et al.*, 2018, p. 232), such decisions should also de facto be viewed as automated. However, if the decisive decision incorporates human judgment and is the final recommendation of the relevant officer, then this is not solely automated and Article 11 LED does not apply.

The next query would be whether this decision has a sufficient impact to trigger the application of Article 11 LED. To have such an impact, an automated decision must have an 'adverse legal effect' on the individual or 'significantly affect' them. The wording of Article 11 LED differs slightly in this regard from the GDPR, which requires the decision to produce 'legal effects' or to 'similarly significantly affect' the data subject. Under the LED, unlike the GDPR, the legal effect must be adverse. The Article 29 Working Party provides an example of a 'typical' adverse effect: namely the application of increased security measures or surveillance by competent authorities. A measure 'significantly affects' the individual where, for example, a passenger is refused access to transport as they are registered on a black list (Article 29 Working Party, 2017). For both Article 11 LED and Article 22 GDPR, automated decision-making with a 'trivial effect' would not be considered sufficient for the prohibition to apply. Again, there are some open questions when this stipulation is applied to the predictive policing context. For instance, in the context of an individualised predictive policing system like HART, it is necessarily the case for Article 11 LED to apply that the decision has involved human input. It could therefore be argued that it is not the HART recommendation (an interim decision, at most) that affects the individual, but rather it is the final (human) decision that is merely informed by this recommendation. This too would therefore require a close examination of how the decision-making procedure operates in practice. More controversially, one could query whether the collective adverse effect from the identification of 'crime hotspots' by systemic predictive policing technologies (such as stigmatisation) would be sufficient to trigger Article 11 LED or whether an individual would need to show adverse effects particular to them. Indeed, in *Tele2 Sverige and Watson*, the court adopted an uncritical approach to such geographic targeting when it suggested that the targeted retention of data pertaining to particular geographic areas would be a preferable alternative to blanket data retention (para. 108).

Furthermore, the extent to which automated decision-making via predictive policing applications might be transparent to those impacted by such applications is unclear. On the one hand, Article 24 LED provides that data controllers should maintain a record of all categories of processing activities under their responsibility, including, where applicable, the use of profiling. This requirement is not envisaged in such a general manner by the GDPR and the Article 29 Working Party has encouraged Member States to be particularly vigilant in enforcing it. On the other hand, unlike the GDPR, Article 13 LED does not explicitly indicate that the data subject must be provided with information regarding the existence of automated decision-making. However, such information may, as the Article 29 Working Party suggested, be provided pursuant to Article 13(2)(d) LED, which states that Member States shall provide by law for the controller to give further information to the data subject to exercise his or her rights. Indeed, the Article 29 Working Party highlights that providing appropriate information to the data subject regarding the existence of such automated decision-making including profiling and meaningful information about the logic involved are particularly relevant in respect of the fairness of data processing, with Article 4(1) providing that data should be processed lawfully and fairly. Once again, a lot will depend on how Member States implement and interpret the directive.

Finally, in practice, it may be difficult to disentangle the GDPR from the LED in this context. If one considers the application of Articles 11 LED and 22 GDPR to the HART model, one can see at once how complicated their application is likely to be in practice. The HART model includes thirty-four predictors that are input into the model. Some of these predictors (those relating to prior offences, for instance) are likely to have been initially collected by the ‘competent authority’ for law-enforcement purposes and thus the LED applies to their processing. This data processing may itself give rise to unique challenges: for instance, if data on spent convictions are included amongst the predictors. However, the inclusion of Experian data based on data scraped from various public sources complicates this picture. The initial processing of these personal data would fall within the scope of the GDPR. The creation of the Experian ‘postcode profile’ could therefore itself be subject to challenge on the ground that it constitutes an automated decision (a profile) under Article 22 GDPR. If such a profile has been created without a legal basis pursuant to Article 22(2) GDPR, it is a violation of that provision. Alternatively, should the profiling be legitimately based on consent, contract or law, it could be argued that the profiles themselves constitute unfair data processing, acting, as they do, as proxies for direct discrimination. *Big Brother Watch* has stated that the Experian profiles include categories such as ‘Asian Heritage’ and ‘Disconnected Youth’ and that these categories have ‘demographic characteristics’ attributed to them. Their report gives the example of the demographic characteristics associated with ‘Asian Heritage’, namely ‘extended families’ living in ‘inexpensive, close-packed Victorian terraces’, adding that ‘when people do have jobs, they are generally in low paid routine occupations in transport or food service’ (*Big Brother Watch*, 2018, p. 7). One could also query whether the subsequent inclusion of these data into the HART application is compatible with Article 11(2) and (3) LED: while it does not directly discriminate on the basis of protected characteristics, it does so at least indirectly.

In sum, the substantive protection offered by Article 11 LED may prove elusive for those impacted by predictive policing technologies. Member States can dodge the Article 11 LED prohibition by putting such technologies on a statutory footing. Furthermore, there is likely to be disagreement over the extent of the human intervention and the impact needed to avoid its application. Finally, the fluidity of data flows between actors involved in the predictive policing process will make even the identification of the appropriate legal regime a challenging task. This uncertainty will make the application of the law difficult for individuals to navigate, thereby enhancing the role and importance of national supervisory authorities and representative bodies (Articles 45 and 55 LED) in guaranteeing individual rights.

## 5 Conclusion

Predictive policing technologies are likely to gain increasing traction in the coming years, promising, as they do, more efficient policing at a lower cost. Yet, as this contribution has demonstrated, the legal framework applicable to these technologies is both unclear and unhelpful. When one takes stock, it is apparent that the application of data protection law to predictive policing technologies is uncertain and will depend on whether the CJEU confirms what Purtova labels its ‘law of everything’ approach. However, even if the data protection rules do apply, they may not be of much assistance to those whose fate is determined by predictive policing technologies. It can be seen that data protection law does not prevent the use of either systemic or individualised predictive policing applications if such applications are provided for by law while an array of provisos and conditions detract from the law’s certainty and may further limit its utility. For those concerned about the normative implications of this predictive capability, the existence of data protection law will be of little comfort. It is perhaps therefore little wonder that the normative narrative around automated decision-making is shifting: scholars have started to question the hunt for fair and transparent algorithmic decision-making and instead have encouraged us to query whether, in each instance, we should be building these systems at all (Powles and Nissenbaum, 2018).

## References

- Article 29 Data Protection Working Party** (2007) *Opinion 4/2007 on the Concept of Personal Data*, WP136, adopted on 20 June 2007.
- Article 29 Data Protection Working Party** (2017) *Opinion on Some Key Issues of the Law Enforcement Directive (EU 2016/680)*, WP258, adopted on 29 November 2017.
- Article 29 Data Protection Working Party** (2018) *Guidelines on Automated Individual Decision-making and Profiling for the Purposes of Regulation 2016/679*, WP251rev.01, adopted on 3 October 2017, as last revised and adopted on 6 February 2018.
- BBC** (2018) 'France fuel protests: Macron holds urgent security meeting. Available at <https://www.bbc.co.uk/news/world-europe-46417991> (accessed 21 February 2019).
- Big Brother Watch** (2018) *Home Affairs Select Committee: Policing for the Future Inquiry*.
- Caplan JM et al.** (2015) Risk terrain modeling for spatial risk assessment. *Cityscape: A Journal of Policy Development and Research* 17, 7–16.
- Degeling M and Berendt B** (2018) What is wrong with Robocops as consultants? A technology-centric critique of predictive policing. *AI & Society* 33, 347–356.
- Dencik L et al.** (2018) Data scores as governance: investigating uses of citizen scoring in public services. Project Report. Available at <https://datajustice.files.wordpress.com/2018/12/data-scores-as-governance-project-report2.pdf> (accessed 21 February 2019).
- Garstka K** (2018) Between security and data protection: searching for a model big data surveillance scheme within the European Union Data Protection Framework. Available at <https://48ba3m4eh2bf2sksp43rq8kk-wpengine.netdna-ssl.com/wp-content/uploads/2018/11/Garstka-Between-Security-and-Data-Protection-November-2018.pdf> (accessed 21 February 2019).
- Kaminski M** (2018) The right to explanation, explained. Available at SSRN: <https://ssrn.com/abstract=3196985> (accessed 21 February 2019).
- Keats Citron D and Pasquale F** (2014) The scored society: due process for automated predictions. *Washington Law Review* 89, 1–33.
- Oswald M et al.** (2018) Algorithmic risk assessment policing models: lessons from the Durham HART model and 'experimental' proportionality. *Information & Communications Technology Law* 27, 223–250.
- Powles J and Nissenbaum H** (2018) The seductive diversion of 'solving' bias in artificial intelligence, *The Medium*. Available at <https://medium.com/s/story/the-seductive-diversion-of-solving-bias-in-artificial-intelligence-890df5e5ef53> (accessed 21 February 2019).
- Purtova N** (2018) The law of everything: broad concept of personal data and the future of EU data protection law. *Law, Innovation and Technology* 10, 40–81.
- Selbst AD and Powles J** (2017) Meaningful information and the right to explanation. *International Data Privacy Law* 4, 233–242.
- Taylor L** (2017) What is data justice? The case for connecting digital rights and freedoms globally. Available at SSRN: <https://ssrn.com/abstract=2918779> (accessed 21 February 2019).
- Uchida CD** (2009) *A National Discussion on Predictive Policing: Defining our Terms and Mapping Successful Implementation Strategies*, California. NCJ 230404.
- Van Brakel R** (2016) Pre-emptive big data surveillance and its (dis)empowering consequences: the case of predictive policing. In van der Sloot B et al. (eds), *Exploring the Boundaries of Big Data*. Available at [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2772469](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2772469) (accessed 21 February 2019).
- Veale M, Binns R and Edwards L** (2018) Algorithms that remember: model inversion attacks and data protection law. *Philosophical Transactions of the Royal Society* 376. Available at <https://royalsocietypublishing.org/doi/10.1098/rsta.2018.0087> (accessed 21 February 2019).
- Wachter S, Mittelstadt B and Floridi L** (2017) Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law* 2, 76–99.