GRESHAM COLLEGE

5th May 2020

Is Robocop now a reality?

Professor Richard Harvey FBCS

In the UK a population of around 70million people made a staggering 35million calls to the emergency services in 2017. Of these, the majority were to the police or ambulance services[1]. And of those calls, roughly one third to a quarter were real emergencies. There is whole thesis to be written about why there are so many false alarms, but it is generally acknowledged that the emergency services are being treated as the government's first responder service. The UK is not alone in feeling this pressure and there is therefore a pressing need to consider automation and digitisation of routine police activities and, more ambitiously, the criminal justice system. This lecture examines some of the issues associated with this digitisation and cherry-picks a few examples that might show us if Robocop is a realistic prospect for the future.

The UK National Police Chiefs' Council identifies three aspects to digital policing that are of current interest. Firstly, what they term "Digital First" which is the automation and systemisation of the criminal justice system; second is "Digital Intelligence and Investigation," which is the exciting part always covered in the media and TV dramas, and thirdly "Digital Public Contact" which is making the interface to the public a bit more modern than turning up the local Police Station or ringing in via that peculiar British invention the Police Box.

Digital First sounds marvellous but like the health service[2] the criminal justice system has a fiendish complexity, illustrated, for the UK, in Figure 1[3].

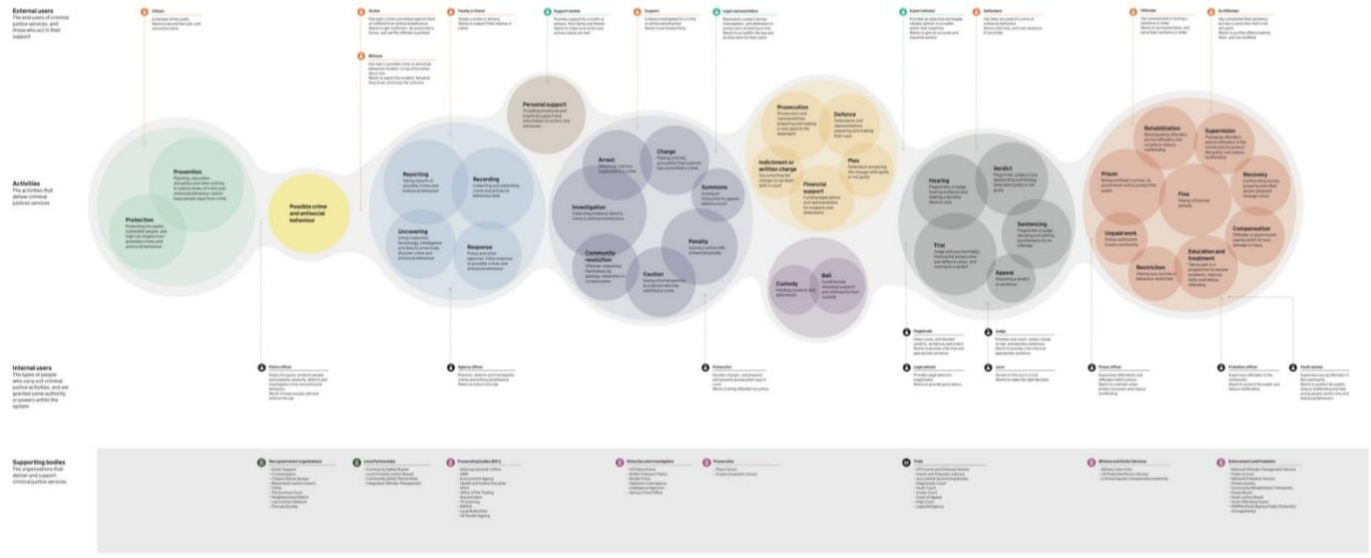

Figure 1: *Simplified diagram of the UK Criminal Justice system taken from [1].*

[1] Readers might be surprised by the small, and decreasing, volume of calls the fire service but, notwithstanding, recent disasters, the risks of fire have considerably reduced through good design and building control. Indeed, many parts of the world now merge the fire service into the police service and either used multi-skilled officers or share back-office services.

[2] "Digital healthcare: will the robot see you now?", Gresham Lecture, Richard Harvey 2020 series.

[3] I'm afraid you will need to zoom in to see the detail.

The Government Digital Service, part of the UK Cabinet Office, acknowledge that the complexity of Figure 1 is essentially an admission that the UK criminal justice system is not a system at all. It is an accretion of custom and practice with hundreds of corporate stakeholders, thousands of vested interests with very few of them representing the citizen who is both a user of the system and the funder of the whole enterprise. Thus, the phrase "Digital First" expresses the hope that we can move from "Digital Last" or perhaps "Digital second from last[4]".

The basis of good system design in IT is to first understand the system you intend to replace or augment. This is sometimes called the requirements-capture phase in software engineering. Figure 1 makes one think that for the criminal justice system the system itself is too complicated to be properly captured. There are several approaches to such complexity. The first is to pick a manageable part of the system, say The Courts, and to digitise it. This has been the approach of the UK government up till this point and, in the case of the Courts has been a spectacular failure – the Digital Courts system has currently cost £1.2Bn is over-budget and late. Another approach, commonly used in large military and civil systems, is to design common interfaces so that a multiplicity of different systems can interface with each other. This is possible, but even a large country such as the UK, does not have enough users to make it commercially worthwhile. However, it might be viable if the common interfaces could embrace multiple justice systems[5] or the taxpayer might fund work on common legal interfaces in the hope that future systems would be simpler which they would be.

A more plausible alternative would be to design a parallel digital justice system that works. Thus, if people wish to have a lengthy and expensive experience, they can indeed opt for the legacy system or they can opt for a more streamlined digital version. To my mind there are many advantages to such a system, not least of which is the opportunity to remove senseless variation between systems[6]. Variations between systems lead to hopeless confusion within the citizenry, a laddering of cost and frequent accusations that the severity of the punishment is a Kafkaesque lottery based on where the crime was committed. We can see the early signs of such systems in online payment for minor offences such as illegal parking – when confronted with a digital photo from a parking official one tends to pay-up rather than risk the time commitment of a tribunal.

Unfortunately, for the time-being neither of the redesign options (common interfaces or digital justice) seem very widespread so IT Professionals are engaged on Sisyphean tasks associated with trying to capture a system that is not properly understood even by people who operate it. As with Health Services therefore, sensible IT professionals avoid it since as large projects are fraught with risk and may even be ruled out by the BCS Code of Conduct[7]. To my mind, this is a very profound issue for much of government – many of the proposed e-government projects are so unmanageable that competent people will avoid them if they possibly can, instead they focus on smaller projects which are manageable and provide shorter-term benefit[8]. With this in mind, we turn to techniques related to digital intelligence and investigation.

Digital intelligence and investigation is the area that captures the imagination of the press and dramatists – visions of nerdy scientists with wild hair (old version) or punky tattoos (new version) are hunched over flat screens zooming pixelated CCTV images into full-resolution images at the touch of button[9]. It also includes the examination of the new activity of cyber-crime. Cyber-crime really deserves a lecture in its own right so I would

4 A complex system with many vested interests clearly has the taxpayer last!
5 Although experts on comparative legal systems tend to emphasize the differences between systems, there are also many many commonalities.
6 Is there are any rational reason why people in England should be tried using a different system to those in Scotland, or Ireland or the Netherlands?
7 The British Computer Society (BCS) Code of Conduct states we should "only undertake to do work or provide a service that is within your professional competence" This would appear to rule out work which is of an unimaginable scale.
8 Such systems often fall into the category of "cool stuff". They are often popular within the service as they do not demand people work differently and they can often, in a relatively modest way, improve the quality of the service. Needless to say they usually cost more money than the old way and thus the cost of the public service grows in an unsustainable way.
9 This trope is particular bugbear of mine. It is very difficult to restore pixels which were not captured in the first place. There are two approaches one uses a model of motion and the other uses a very detailed model of the object. Neither works very well and ceeainly nowhere near as well as on the movies.

wish to emphasise, that while cyber-crime might catch the attention, the majority of crime is old crime with a new cyber dimension. It is nowadays impossible to imagine an allegation of rape where the police did not examine the personal communication devices of both the accuser and the accused. This can be an aggravatingly time-consuming process[10] – imagine dealing with a locked iPhone where the accused as "forgotten" the password and has been using a secure communication system such as WhatsApp or Wickr to communicate. Not only does your forensic evidence suite have to handle a bewildering array of communication options, many of them of end-to-end secure so can only be read with access to the phone. This is the business of digital forensics and it's growing rapidly with specialist systems for certain criminal activities. One that I have previously helped developed systems for is the analysis of child pornography. Child pornography is particularly challenging as the volume of material can be daunting and there is general agreement that it is damaging for police officers to be exposed to high doses of child pornography during the collection and summarisation of evidence. Furthermore, if officers identify new material then there is clearly a child at risk so it's very important to distinguish between legacy images and those which might be decades old. The solution is to build automatic systems that can scan confiscated hardware, summarise the contents, analyse images and classify them into categories.

Another area of growth, as with Medicine, are decision support systems. In the lecture I devote some time to Predpol [2] which is a decision support tool for what is called "predictive policing". The idea that crime is highly imhomgenous is quite ancient and certainly in the 1800s Quetelet was gathering and publishing statistics on crime by geographical area [3]. Predpol devises hotspots which are then passed to local officers for action. The system is said to have been evaluated as effective [4] although one aspect of an early variant of the algorithm can lead to dangerous feedback loops in which officers are repeatedly dispatched to hotspots which then become reinforced as hotspots. When these areas do not have a balanced mix of ethnicities there is the risk that policing is seen as racially biased. Such loops can be easily corrected however by resampling the data to force a more representative ethnicity sample [5]. There are persistent stories in the media that Predpol targets individuals but, as far as I can ascertain, this is not true – it is system that is trained on anonymised crime statistics and reports.

Another example of a decision support system is the decision to grant probation [6]. A system was first trailed in Philadephia and concerned the problem of recidivism: the aim being to predict, before probation began, whether a candidate was "high risk"[11], "moderate risk" or "low risk." Of course, it is tricky experiment as the law enforcement agencies cannot run an experiment in which high risk criminals are let free. This system has been deployed in the UK by Durham Constabulary where it is known as HART, the Harm Assessment Risk Tool. The context is slightly different, as the US and UK legal systems vary[12] but the principle of operation is the same: the system is trained on several years' worth of offender data. HART uses 34 predictor values of which the majority (29) come from the suspect's offending history. These behavioural predictors are combined with age, gender, two forms of residential postcode, and a count of police intelligence reports relating to the offender. A critique of the system is provided in [7]. The principal observation is that the evaluation of algorithms in policing can itself be described by an algorithm, a human-implemented algorithm which the authors describe by the acronym ALGO-CORE. There is not space here to describe all the aspects but I'd like to take one to illustrate an important point. The "R" in ALGO-CORE stands for "Responsible: would the operation of the algorithm be considered fair? Is the use of the algorithm transparent (taking account of the context of its use), accountable and placed under review alongside other developments in policing? Would it be considered to be for the public interest and ethical?"[7]. On the face it these are very reasonable points and indeed few readers would wish to disagree with them I suspect. However, a very important aspect is missing.

[10] In preparing this lecture, I spoke to a number of senior police officers who lamented the time it took to gather cyber evidence.

[11] Which they defined as the offender was predicted to commit at least one serious offense – defined as murder, attempted murder, aggravated assault, robbery, or a sexual crime – during the first two years after their case start date;

[12] In the lecture I refer to all these decisions, US and UK, as probation decisions which is an over-simplification which will probably irritate lawyers. This is partly provocative because, the over-complexity of the legal complexity is a major problem as I want to emphasize in this lecture and partly because I'm interested not in the precise legal definition of the decisions, Durham is actually referring potential recidivists to a pathway in a programme known as "Checkpoint", but rather in the concept of automated decision-making in the criminal justice system.

When considering the operation of an algorithm, it is very important to consider the efficacy of the human system it replaces. Existing humans might well be biased, capricious, tired, poorly informed and have all those lovely human failings with which are so familiar. The key question with any algorithm, is to identify ways in which the algorithm is superior or inferior to the human. If the algorithm is superior to the human, then sensible people ought to argue that it is dangerous to continue with human operation. The word "superior" turns out to be quite nuanced – many algorithms are superior to a buffoon – so the usual approach, developed in Medicine, is to compare the algorithm to the balanced judgement of a mixture of experts. That might not be appropriate in the legal system – if the police didn't make mistakes then there might be enough work for the lawyers! But it is a critical consideration. If we are constantly asking perfection of computer systems they will never be deployed and we will be left with the highly imperfect system that is described in Figure 1.

Of course, the views of civil libertarians are useful and one area where they are hottest under the collar is digital intelligence – the gathering of information that might help prevent crime. One certainly has every sympathy with subjects that are detained or suffer inconvenience due to a poorly designed IT system although again I would issue a plea that such negative consequences are balanced against the current situation. That said, there a number of interesting developments in the area of intelligence of which behaviour recognition has much potential. Former FBI officer Joe Navarro tells an entertaining story online about how he detected a Russian interloper by observing that the suspect exited a florist shop with the bloom of flowers held stems upwards.[13] Understanding the semiotics of flowers might be beyond current automatic systems which currently fall into two categories. In the first, we learn a detailed version of normal. The computer vision system consumes hours and hours of footage from a particular camera and learns spatio-temporal signatures. In essence which colours neighbour others in both space and time. When a motion doesn't fit "normal" it is flagged for analysis by a human operator. Such systems make the workload on video surveillance manageable[14] or, rather, allow the great number of unattended cameras to be used when something valuable happens but otherwise remain dormant. The second approach is more "model-based". There is now quite a body of work related to tracking people. In the second category are computer systems that have "hard-wired" models of interesting objects such as people. To a computer vision system, a person is rectangular box with relatively static colour distribution that moves with a small range of motions. Tracking people allows for detection of a variety of anomalous behaviours such as loitering around one end of a railway platform (a known precursor to suicide); casing a high-risk terrorist target (schools are a popular target for modern terrorists because they often have poor security and are open) or assault[15]. Here I have picked only a couple of Digital Intelligence and Investigation tools, but it is a highly active research area with a great number of experimental systems and a few deployed systems.

The third area identified by the NPCC is Digital Public Contact. This work is concerned with making the criminal justice system accessible to people via digital means. There are several advantages of digital interfaces with the public among which we might mention the ease of multilingualism, accessibility to persons with disabilities, anonymity if desired and responsiveness. Those are advantages for the public, but for the police or courts, a digital communication system will usually provide better tracking, useful statistics and the potential for better localisation – working out to which Police force one should route an emergency call is not that easy. In post Covid times, a robot receptionist in an unmanned Police station might be highly desirable as, in between visits, the robot can dock be irradiated with UV light which kills the virus and return to duty. Less imaginatively most Police forces operate a social media team and will follow-up reports from a variety of sources. Maybe our future will see the return of Mackenzie-Trench style Police boxes[16] as unmanned police stations.

1. *Mapping New Ideas for the Criminal Justice System*, Mike Bracken, Government Digital Service Blog, 2015, https://gds.blog.gov.uk/2015/08/18/mapping-new-ideas-for-the-digital-justice-system-2/
2. https://www.predpol.com
3. *A treatise on man and the development of his faculties*, M A Quetelet, English translation ,1842, in French 1835

[13] He asserts that Westeners invariable hold bunches of flowers with the stem downwards. I don't but then maybe I spent too long hanging around in the wrong places.
[14] There are rough estimates of around 0.6M CCTV cameras operating in London. Operators can handle around 10 simultaneous monitors so 0.6M cameras implies a workforce of 60,000 CCTV operators in 24 hours.
[15] Assault is the example I used in the lecture – to a computer vision system fighting is two rectangular boxes that are too close to each and other and have high velocities.
[16] AKA "The Tardis" for Doctor Who fans.

4. *Data and evidence challenges facing place-based policing*, Hutt, O., Bowers, K., Johnson, S. and Davies, T. (2018), , *Policing: An International Journal*, Vol. 41 No. 3, pp. 339-351. https://doi.org/10.1108/PIJPSM-09-2017-0117
5. *A Penalized Likelihood Method for Balancing Accuracy and Fairness in Predictive Policing*, G. Mohler, R. Raje, J. Carter, M. Valasik and J. Brantingham, 2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC), Miyazaki, Japan, 2018, pp. 2454-2459.
6. *Classifying Adult Probationers by Forecasting Future Offending*, Final Technical Report ,Geoffrey Barnes and Jordan M. Hyatt, March 2012 https://www.ncjrs.gov/pdffiles1/nij/grants/238082.pdf.
7. *Algorithmic risk assessment policing models: lessons from the Durham HART model and 'Experimental' proportionality*, Marion Oswald, Jamie Grace, Sheena Urwin & Geoffrey C. Barnes (2018), Information & Communications Technology Law, 27:2, 223-250, DOI: 10.1080/13600834.2018.1458455