

police analysed blood from the car they found its DNA matched a sample from a rape carried out in another multi-storey car park two years earlier. That attack had taken place some 100 kilometres further south, in Nottingham. And, after an appeal on the BBC *Crimewatch* programme, police also managed to link the case to three other incidents a decade before in Bradford, Leeds and Leicester.²

But tracking down this particular serial rapist was not going to be easy. Together, these crimes spanned an area of 7,046 square kilometres – an enormous stretch of the country. They also presented the police with a staggering number of potential suspects – 33,628 in total – each one of whom would have to be eliminated from their enquiries or investigated.³

An enormous search would have to be made, and not for the first time. The attacks ten years earlier had led to a massive man-hunt; but despite knocking on 14,153 front doors, and collecting numerous swabs, hair samples and all sorts of other evidence, the police investigation had eventually led nowhere. There was a serious risk that the latest search would follow the same path, until a Canadian ex-cop, Kim Rossmo, and his newly developed algorithm were brought in to help.⁴

Rossmo had a bold idea. Rather than taking into account the vast amount of evidence already collected, his algorithm would ignore virtually everything. Instead, it would focus its attention exclusively on a single factor: geography.

Perhaps, said Rossmo, a perpetrator doesn't randomly choose *where* they target their victims. Perhaps their choice of location isn't an entirely free or conscious decision. Even though these attacks had taken place up and down the country, Rossmo wondered if there could be an unintended pattern hiding in the geography of the crimes – a pattern simple enough to be exploited. There was a chance, he believed, that the locations at which crimes took place

could betray where the criminal actually came from. The case of the serial rapist was a chance to put his theory to the test.

Operation Lynx and the lawn sprinkler

Rossmo wasn't the first person to suggest that criminals unwittingly create geographical patterns. His ideas have a lineage that dates back to the 1820s, when André-Michel Guerry, a lawyer-turned-statistician who worked for the French Ministry of Justice, started collecting records of the rapes, murders and robberies that occurred in the various regions of France.⁵

Although collecting these kinds of numbers seems a fairly standard thing to do now, at the time maths and statistics had only ever been applied to the hard sciences, where equations are used to elegantly describe the physical laws of the universe: tracing the path of a planet across a sky, calculating the forces within a steam engine – that sort of thing. No one had bothered to collect crime data before. No one had any idea what to count, how to count or how often they should count it. And anyway – people thought at the time – what was the point? Man was strong, independent in nature and wandering around acting according to his own free will. His behaviour couldn't possibly be captured by the paltry practice of statistics.⁶

But Guerry's analysis of his national census of criminals suggested otherwise. No matter where you were in France, he found, recognizable patterns appeared in what crimes were committed, how – and by whom. Young people committed more crimes than old, men more than women, poor more than rich. Intriguingly, it soon became clear that these patterns didn't change over time. Each region had its own set of crime statistics that would barely change year on year. With an almost terrifying exactitude, the numbers of robberies, rapes and murders would repeat themselves from one year to the next. And even the methods used by the murderers

were predictable. This meant that Guerry and his colleagues could pick an area and tell you in advance exactly how many murders by knife, sword, stone, strangulation or drowning you could expect in a given year.⁷

So maybe it wasn't a question of the criminal's free will after all. Crime is not random; people are predictable. And it was precisely that predictability that, almost two centuries after Guerry's discovery, Kim Rossmo wanted to exploit.

Guerry's work focused on the patterns found at the country and regional levels, but even at the individual level, it turns out that people committing crime still create reliable geographical patterns. Just like the rest of us, criminals tend to stick to areas they are familiar with. They operate locally. That means that even the most serious of crimes will probably be carried out close to where the offender lives. And, as you move further and further away from the scene of the crime, the chance of finding your perpetrator's home slowly drops away,⁸ an effect known to criminologists as 'distance decay'.

On the other hand, serial offenders are unlikely to target victims who live *very* close by, to avoid unnecessary police attention on their doorsteps or being recognized by neighbours. The result is known as a 'buffer zone' which encircles the offender's home, a region in which there'll be a very low chance of their committing a crime.⁹

These two key patterns – distance decay and the buffer zone – hidden among the geography of the most serious crimes, were at the heart of Rossmo's algorithm. Starting with a crime scene pinned on to a map, Rossmo realized he could mathematically balance these two factors and sketch out a picture of where the perpetrator might live.

That picture isn't especially helpful when only one crime has been committed. Without enough information to go on, the so-called *geoprofiling algorithm* won't tell you much more than good old-fashioned common sense. But, as more crimes are added, the

picture starts to sharpen, slowly bringing into focus a map of the city that highlights areas in which you're most likely to catch your culprit.

It's as if the serial offender is a rotating lawn sprinkler. Just as it would be difficult to predict where the very next drop of water is going to fall, you can't foresee where your criminal will attack next. But once the water has been spraying for a while and many drops have fallen, it's relatively easy to observe from the pattern of the drops where the lawn sprinkler is likely to be situated.

And so it was with Rossmo's algorithm for Operation Lynx – the hunt for the serial rapist. The team now had the locations of five separate crimes, plus several places where a stolen credit card had been used by the attacker to buy alcohol, cigarettes and a video game. On the basis of just those locations, the algorithm highlighted two key areas in which it believed the perpetrator was likely to live: Millgarth and Killingbeck, both in the suburbs of Leeds.¹⁰

Back in the incident room, police had one other key piece of evidence to go on: a partial fingerprint left by the attacker at the scene of an earlier crime. It was too small a sample for an automatic fingerprint recognition system to be able to whizz through a database of convicted criminals' prints looking for a match, so any comparisons would need to be made meticulously by an expert with a magnifying glass, painstakingly examining one suspect at a time. By now the operation was almost three years old and – despite the best efforts of 180 different officers from five different forces – it was beginning to run out of steam. Every lead resulted in just another dead end.

Officers decided to manually check all the fingerprints recorded in the two places the algorithm had highlighted. First up was Millgarth: but a search through the prints stored in the local police database returned nothing. Then came Killingbeck – and after 940 hours of sifting through the records here, the police finally came up with a name: Clive Barwell.

Barwell was a 42-year-old married man and father of four, who had been in jail for armed robbery during the hiatus in the attacks. He now worked as a lorry driver and would regularly make long trips up and down the country in the course of his job; but he lived in Killingbeck and would often visit his mother in Millgarth, the two areas highlighted by the algorithm.¹¹ The partial print on its own hadn't been enough to identify him conclusively, but a subsequent DNA test proved that it was he who had committed these horrendous crimes. The police had their man. Barwell pleaded guilty in court in October 1999. The judge sentenced him to eight back-to-back life sentences.¹²

Once it was all over, Rossmo had the chance to take stock of how well the algorithm had performed. It had never actually pinpointed Barwell by name, but it did highlight on a map the areas where the police should focus their attention. If the police had used the algorithm to prioritize their list of suspects on the basis of where each of them lived – checking the fingerprints and taking DNA swabs of each one in turn – there would have been no need to trouble anywhere near as many innocent people. They would have found Clive Barwell after searching only 3 per cent of the area.¹³

This algorithm had certainly been proved effective. And it brought other positives, too. As it only prioritizes your existing list of suspects, it doesn't suffer from bias of the kind we met in the 'Justice' chapter. Also, it can't override good detective work, only make an investigation more efficient; so there's little chance of people putting too much trust in it.

It is also incredibly flexible. Since Operation Lynx, it has been used by more than 350 crime-fighting agencies around the world, including the US Federal Bureau of Investigation and the Royal Canadian Mounted Police. And the insights it offers extend beyond crime: the algorithm has been used to identify stagnant water pools that mosquitoes use as breeding grounds, based on the locations

of cases of malaria in Egypt.¹⁴ A PhD student of mine at University College London is currently using the algorithm in an attempt to predict the sites of bomb factories on the basis of the locations where improvised explosive devices are used. And one group of mathematicians in London have even used it to try to track down Banksy,¹⁵ the reclusive street artist, on the basis of where his paintings have been found.

The kinds of crimes for which geoprofiling works best – serial rapes, murders and violent attacks – are, fortunately, rare. In reality the vast majority of infractions don't warrant the kind of man-hunt that the Clive Barwell case demanded. If algorithms were to make a difference in tackling crime beyond these extreme cases, they'd need a different geographical pattern to go on. One that could be applied to a city as a whole. One that could capture the patterns and rhythms of a street or a corner that every beat officer knows instinctively. Thankfully, Jack Maple had just the thing.

Charts of the Future

A lot of people would think twice about riding the New York City subway in the 1980s. It wasn't a nice place. Graffiti covered every surface, the cars stank of stale urine, and the platforms were rife with drug use, theft and robbery. Around 20 innocent people were murdered below ground every year, making it practically as dangerous a place as any in the world.

It was against this backdrop that Jack Maple worked as a police officer. He had recently earned himself a promotion to transit lieutenant, and in his years in the force he'd grown tired of only ever responding to crime, rather than fighting to reduce it. Out of that frustration was born a brilliant idea.

'On 55 feet of wall space, I mapped every train station in New York City and every train,' Maple told an interviewer in 1999. 'Then

I used crayons to mark every violent crime, robbery and grand larceny that occurred. I mapped the solved versus the unsolved.’¹⁶

They might not sound like much, but his maps, scrawled out on brown paper with crayons, became known as the ‘Charts of the Future’ and were, at the time, revolutionary. It had never occurred to anyone to look at crimes in that way before. But the moment Maple stepped back to view an entire city’s worth of crime all at once, he realized he was seeing everything from a completely new perspective.

‘It poses the question, Why?’ he said. ‘What are the underlying causes of why there is a certain cluster of crime in a particular place?’ The problem was, at the time, every call to the police was being treated as an isolated incident. If you rang to report an aggressive group of drug dealers who were hanging out on a corner, but they ducked out of sight as soon as the cops arrived, nothing would be recorded that could link your complaint to other emergency calls made once the gang retook their positions. By contrast, Maple’s maps meant he could pinpoint precisely where crime was a chronic problem, and that meant he could start to pick apart the causes. ‘Is there a shopping centre here? Is that why we have a lot of pickpockets and robberies? Is there a school here? Is that why we have a problem at three o’clock? Is there an abandoned house nearby? Is that why there is crack-dealing on the corner?’¹⁷

Being able to answer these questions was the first step towards tackling the city’s problems. So in 1990, when the open-minded Bill Bratton became head of the New York Transit Police, Maple showed him his Charts of the Future. Together they used them to try to make the subway a safer place for everyone.¹⁸

Bratton had a clever idea of his own. He knew that people begging, urinating and jumping the turnstiles were a big problem on the subway. He decided to focus police attention on addressing those minor misdemeanours rather than the much more serious

crimes of robbery and murder that were also at epidemic levels below ground.

The logic was twofold. First, by being tough on any anti-social behaviour at the crime hotspots, you could send a strong signal that criminal activity was not acceptable in any form, and so hopefully start changing what people saw as ‘normal’. Second, the people evading fares were disproportionately likely to be criminals who would then go on to commit bigger crimes. If they were arrested for fare evasion they wouldn’t have that chance. ‘By cracking down on fare evasion, we have been able to stop serious criminals carrying weapons at the turnstiles before they get on the subways and wreak havoc,’ Bratton told *Newsday* in 1991.¹⁹

The strategy worked. As the policing got smarter, the subways got safer. Between 1990 and 1992, Maple’s maps and Bratton’s tactics cut felonies on the subway by 27 per cent and robberies by a third.²⁰

When Bratton became Police Commissioner for the New York Police Department, he decided to bring Maple’s Charts of the Future with him. There, they were developed and refined to become CompStat, a data-tracking tool that is now used by many police departments, both in the United States and abroad. At its core remains Jack Maple’s simple principle – record where crimes have taken place to highlight where the worst hotspots are in a city.

Those hotspots tend to be quite focused. In Boston, for instance, a study that took place over a 28-year period found that 66 per cent of all street robberies happened on just 8 per cent of streets.²¹ Another study mapped 300,000 emergency calls to police in Minneapolis: half of them came from just 3 per cent of the city area.²²

But those hotspots don’t stay the same over time. They constantly move around, subtly shifting and morphing like drops of oil on water, even from one day to the next. And when Bratton moved to Los Angeles in 2002, he began to wonder if there were other patterns

that could tell you *when* crime was going to occur as well as *where*. Was there a way to look beyond crimes that had already happened? Rather than just responding to crime, the source of Maple's frustration, or fighting it as it happened, could you also predict it?

The flag and the boost

When it comes to predicting crime, burglary is a good place to start. Because burglaries take place at an address, you know precisely where they happened – unlike pickpocketing, say. After all, it's quite possible that a victim wouldn't notice their phone missing until they got home. Most people will report if they've been burgled, too, so we have really good, rich datasets, which are much harder to gather for drug-related offences, for instance. Plus, people often have a good idea of when their homes were burgled (perhaps when they were at work, or out for the evening), information you won't have for crimes such as vandalism.

Burglars also have something in common with the serial murderers and rapists that Rossmo studied: they tend to prefer sticking to areas they're familiar with. We now know you're more likely to be burgled if you live on a street that a burglar regularly uses, say on their way to work or school.²³ We* also know that there's a sweet spot for how busy burglars like a street to be: they tend to avoid roads with lots of traffic, yet home in on streets with lots of carefree non-locals walking around to act as cover (as long as there aren't also lots of nosey locals hanging about acting as guardians).²⁴

But that's only the first of two components to the appeal of your home to a burglar. Yes, there are factors that don't change over time, like where you live or how busy your road is, that will 'flag' the steady

* On this occasion, when I say 'we' I actually mean it. This particular study was a paper I worked on with my marvellous PhD student Michael Frith.

appeal of your property as a target. But before you rush off to sell your house and move to a quiet cul-de-sac with a great Neighbourhood Watch scheme, you should also be aware that crime hotspots don't stay still. The second component of your home's appeal is arguably the more important. This factor depends on what exactly is going on right now in your immediate local neighbourhood. It's known as the 'boost'.

If you've ever been broken into twice within a short space of time, you'll be all too familiar with the boost effect. As police will tell you after you've first been victimized, criminals tend to repeatedly target the same location – and that means that no matter where you live, you're most at risk in the days right after you've just been burgled. In fact, your chances of being targeted can increase twelvefold at this time.²⁵

There are a few reasons why burglars might decide to return to your house. Perhaps they've got to know its layout, or where you keep your valuables (things like TVs and computers are often replaced pretty quickly, too), or the locks on your doors, or the local escape routes; or it might be that they spotted a big-ticket item they couldn't carry the first time round. Whatever the reason, this boost effect doesn't just apply to you. Researchers have found that the chances of your neighbours being burgled immediately after you will also be boosted, as will those of your neighbours' neighbours, and your neighbours' neighbours' neighbours, and so on all the way down the street.

You can imagine these boosts springing up and igniting across the city like a fireworks display. As you get further away from the original spark, the boost gets fainter and fainter; and it fades away over time, too, until after two months – unless a new crime re-boosts the same street – it will have disappeared entirely.²⁶

The flag and the boost in crime actually have a cousin in a natural phenomenon: earthquakes. True, you can't precisely predict where

and when the first quake is going to hit (although you know that some places are more prone to them than others). But as soon as the first tremors start, you can talk quite sensibly about where and how often you expect the aftershocks to occur, with a risk that is highest at the site of the original quake, lessens as you get further away, and fades as time passes.

It was under Bratton's direction that the connection between earthquake patterns and burglary was first made. Keen to find a way to forecast crimes, the Los Angeles Police Department set up a partnership with a group of mathematicians at the University of California, Los Angeles, and let them dig through all the data the cops could lay their hands on: 13 million crime incidents over 80 years. Although criminologists had known about the lag and the boost for a few years by this point, in searching through the patterns in the data the UCLA group became the first to realize that the mathematical equations which so beautifully predicted the risk of seismic shocks and aftershocks could also be used to predict crimes and 'aftercrimes'. And it didn't just work for burglary. Here was a way to forecast everything from car theft to violence and vandalism.

The implications were indeed seismic. Rather than being able to say that a recently victimized area of the city was vaguely 'more at risk', with these equations, you could quantify exactly what that risk was, down to the level of a single street. And knowing, probabilistically speaking, that a particular area of the city would be the focus of burglars on a given night, it was easy to write an algorithm that could tell police where to target their attention.

And so PredPol (or PREDictive POLicing) was born.

The crime tipster

You might well have come across PredPol already. It's been the subject of thousands of news articles since its launch in 2011, usually

under a headline referencing the Tom Cruise film *Minority Report*. It's become like the Kim Kardashian of algorithms: extremely famous, heavily criticized in the media, but without anyone really understanding what it does.

So before you fill your mind with images of seers lying in pools of water screaming out their premonitions, let me just manage your expectations slightly. PredPol doesn't track down people before they commit crimes. It can't target individual people at all, only geography. And I know I've been throwing the word 'prediction' around, but the algorithm can't actually tell the future. It's not a crystal ball. It can only predict the risk of future events, not the events themselves – and that's a subtle but important difference.

Think of the algorithm as something like a bookie. If a big group of police officers are crowded around a map of the city, placing bets on where crime will happen that night, PredPol calculates the odds. It acts like a tipster, highlighting the streets and areas that are that evening's 'favourites' in the form of little red squares on a map.

The key question is whether following the tipster's favourite can pay off. To test the algorithm's prowess,²⁷ it was pitted against the very best expert human crime analysts, in two separate experiments; one in Kent in southern England, the other in the Southwest Division of Los Angeles. The test was a straight head-to-head challenge. All the algorithm and the expert had to do was place 20 squares, each representing an area of 150 square metres, on a map to indicate where they thought most crime would happen in the next 12 hours.

Before we get to the results, it's important to emphasize just how tricky this is. If you or I were given the same task, assuming we don't have extensive knowledge of the Kentish or Californian crime landscapes, we'd probably do no better than chucking the squares on to the map at random. They'd cover pathetically little of it, mind – just one-thousandth of the total area in the case of Kent²⁸ – and every

12 hours you'd have to clear your previous guesses and start all over again. With this random scattering, we could expect to successfully 'predict' less than one in a hundred crimes.²⁹

The experts did a lot better than that. In LA, the analyst managed to correctly predict where 2.1 per cent of crime would happen,³⁰ and the UK experts did better still with an average of 5.4 per cent,³¹ a score that was especially impressive when you consider that their map was roughly ten times the size of the LA one.

But the algorithm eclipsed everyone. In LA it correctly forecast more than double the number of crimes that the humans had managed to predict, and at one stage in the UK test, almost one in five crimes occurred within the red squares laid down by the mathematics.³² PredPol isn't a crystal ball, but nothing in history has been able to see into the future of crime so successfully.

Putting predictions into practice

But there's a problem. While the algorithm is relatively great at *predicting* where crime will happen over the next 12 hours, the police themselves have a subtly different objective: *reducing* crime over the next 12 hours. Once the algorithm has given you its predictions, it's not entirely clear what should be done next.

There are a few options, of course. In the case of burglary, you could set up CCTV cameras or undercover police officers and catch your criminals in the act. But perhaps it would be better for everyone if your efforts went into preventing the crime before it happened. After all, what would you prefer? Being the victim of a crime where the perpetrator gets caught? Or not being the victim of a crime in the first place?

You could warn local residents that their properties are at risk, maybe offer to improve the locks on their doors, maybe install burglar alarms or timers on their light switches to trick any dodgy people

passing by into thinking there's someone at home. That's what one study did in Manchester in 2012,³³ where they managed to reduce the number of burglaries by more than a quarter. Small downside, though: the researchers calculated that this tactic of so-called 'target hardening' costs about £3,925 per burglary it prevents.³⁴ Try selling that to the Los Angeles Police Department, which deals with over 15,000 burglaries a year.³⁵

Another option, which deviates as little as possible from traditional policing, is a tactic known as 'cops on the dots'.

'In the olden days,' Steyve Colgan, a retired Metropolitan Police officer, told me, '[patrols] were just geographic, you got the map, cut it up into chunks, and divvied it up. You're on that beat, you're on that beat. As simple as that.' Problem was, as one UK study calculated, a police officer patrolling their randomly assigned beat on foot could expect to come within a hundred yards of a burglary just once in every eight years.³⁶

With cops-on-the-dots, you simply send your patrols to the hotspots highlighted by the algorithm instead. (It should be called cops-on-the-hotspots, really. I guess that wasn't as catchy.) The idea being, of course, that if the police are as visible as possible and in the right place at the right time, they're more likely to stop crime from happening, or at least respond quickly in the immediate aftermath.

This is exactly what happened in Kent. During the second phase of the study, as the evening shift started, the sergeant printed off the maps with red squares highlighting the at-risk areas for that evening. Whenever the police on patrol had a quieter moment, they were to go to the nearest red box, get out of their cars and have a walk around.

On one particular evening, in an area they would never normally go to, police found an east European woman and her child in the street. It turned out that the woman was in an abusive relationship and, just minutes before, the child had been sexually assaulted. The sergeant on duty that night confirmed that 'they found these people

because they were in a PredPol box.³⁷ The suspect was arrested nearby later that night.

That mother and her child weren't the only people helped by the algorithm during the cops-on-the-dots trial: crime in Kent overall went down by 4 per cent. Similar studies in the United States (conducted by PredPol itself) report even bigger drops in crime. In the Foothill area of Los Angeles, crime dropped by 13 per cent in the first four months of using the algorithm, despite an increase of 0.4 per cent in the rest of the city, where they were relying on more traditional policing methods. And Alhambra, a city in California not far from LA, reported an extraordinary 32 per cent drop in burglaries and a 20 per cent drop in vehicle theft after deploying the algorithm in January 2013.³⁸

These numbers are impressive, but it's actually difficult to know for sure whether PredPol can take the credit. Toby Davies, a mathematician and crime scientist from UCL, told me: 'It's possible that merely encouraging policing officers to go to places and get out of their cars and walk around, regardless of where, actually could lead to reductions [in crime] anyway.'

And there's another issue here. If, the harder you look for crime, the more likely you are to find it, then the act of sending police out could actually change the crime records themselves: 'When police are in a place,' Davies told me, 'they detect more crime than they would have done otherwise. Even if an equal value of crime is happening in two places, the police will detect more in the place they were than the one that they weren't.'

That means there is one very big potential downside of using a cops-on-the-dots tactic. By sending police into an area to fight crime on the back of the algorithm's predictions, you can risk getting into a feedback loop.

If, say, a poorer neighbourhood had a high level of crime in the first instance, the algorithm may well predict that more crime will

happen there in future. As a result, officers are sent to the neighbourhood, which means they detect more crime. Thus, the algorithm predicts more still, more officers are sent there, and so on it goes. These feedback loops are more likely to be a problem for crimes that are linked to poorer areas such as begging, vagrancy and low-level drug use.

In the UK, where some sections of society regularly complain about a lack of police presence on the streets, focusing police attention on certain areas might not immediately seem unfair. But not everyone has a positive relationship with the police. 'It is legitimate for people who see a police officer walking in front of their house every day to feel oppressed by that, even if no one's doing any crimes, even if that police officer is literally just walking up and down,' Davies told me. 'You almost have a right not to be constantly under pressure, under the eye of the police.'

I'm rather inclined to agree.

Now, a well-tuned algorithm *should* be built so that it can take account of the tactics being used by the police. There are ways, theoretically at least, to ensure that the algorithm doesn't disproportionately target particular neighbourhoods – like randomly sending police to medium-risk areas as well as high-risk ones. But, unfortunately, there's no way to know for sure whether PredPol is managing to avoid these feedback loops entirely, or indeed whether it is operating fairly more generally, because PredPol is a proprietary algorithm, so the code isn't available to the public and no one knows exactly how it works.

PredPol is not the only software on the market. One competitor is HunchLab, which works by combining all sorts of statistics about an area: reported crimes, emergency calls, census data (as well as more eyebrow-raising metrics like moon phases). HunchLab doesn't have an underlying theory. It doesn't attempt to establish *why* crime occurs in some areas more than others; it simply reports

on patterns it finds in the data. As a result, it can reliably predict more types of crime than PredPol (which has at its heart theories about how criminals create geographical patterns) – but, because HunchLab too is protected as intellectual property, it is virtually impossible from the outside to ensure it isn't inadvertently discriminating against certain groups of people.³⁹

Another opaque predictive algorithm is the Strategic Subject List used by the Chicago Police Department.⁴⁰ This algorithm takes an entirely different approach from the others. Rather than focusing on geography, it tries to predict which individuals will be involved in gun crime. Using a variety of factors, it creates a 'heat list' of people it deems most likely to be involved in gun violence in the near future, either doing the shooting or being shot. The theory is sound: today's victims are often tomorrow's perpetrators. And the programme is well intentioned: officers visit people on the watch list to offer access to intervention programmes and help to turn their lives around.

But there are concerns that the Strategic Subject List might not be living up to its promise. One recent investigation by the non-profit RAND Corporation concluded that appearing on it actually made no difference to an individual's likelihood of being involved in a shooting.⁴¹ It did, however, mean they were more likely to be arrested. Perhaps – the report concluded – this was because officers were simply treating the watch list as a list of suspects whenever a shooting occurred.

Predictive policing algorithms undoubtedly show promise, and the people responsible for creating them are undoubtedly doing so in good faith, with good intentions. But the concerns raised around bias and discrimination are legitimate. And for me, these questions are too fundamental to a just society for us simply to accept assurances that law enforcement agencies will use them in a fair way. It's one of many examples of how badly we need independent experts

and a regulatory body to ensure that the good an algorithm does outweighs the harm.

And the potential harms go beyond prediction. As we have already seen in a variety of other examples, there is a real danger that algorithms can add an air of authority to an incorrect result. And the consequences here can be dramatic. Just because the computer says something doesn't make it so.

Who do you think you are?

Steve Talley was asleep at home in South Denver in 2014 when he heard a knock at the door.⁴² He opened it to find a man apologizing for accidentally hitting his car. The stranger asked Talley to step outside and take a look. He obliged. As he crouched down to assess the damage to his driver's side door,⁴³ a flash grenade went off. Three men dressed in black jackets and helmets appeared and knocked him to the ground. One man stood on his face. Another restrained his arms while another started repeatedly hitting him with the butt of a gun.

Talley's injuries would be extensive. By the end of the evening he had sustained nerve damage, blood clots and a broken penis.⁴⁴ 'I didn't even know you could break a penis,' he later told a journalist at *The Intercept*. 'At one point I was actually screaming for the police. Then I realized these were cops who were beating me up.'⁴⁵

Steve Talley was being arrested for two local bank robberies. During the second robbery a police officer had been assaulted, which is why, Talley thinks, he was treated so brutally during his arrest. 'I told them they were crazy,' he remembers shouting at the officers, 'You've got the wrong guy!'

Talley wasn't lying. His arrest was the result of his striking resemblance to the right guy – the real robber.

Although it was a maintenance man working in Talley's building who initially tipped off the police after seeing photos on the



local news, it would eventually be an FBI expert using facial recognition software⁴⁶ who later examined the CCTV footage and concluded that 'the questioned individual depicted appears to be Talley'.⁴⁷

Talley had a cast-iron alibi, but thanks to the FBI expert's testimony, it would still take over a year to clear his name entirely. In that time he was held in a maximum-security pod for almost two months until enough evidence surfaced to release him. As a result, he was unable to work, and by the time his ordeal was over he had lost his job, his home and access to his children. All as a direct result of that false identification.

Seeing double

Facial recognition algorithms are becoming commonplace in modern policing. These algorithms, presented with a photograph, footage or snapshot from a 3D camera, will detect a face, measure its

characteristics and compare them to a database of known faces with the aim of determining the identity of the person pictured.

In Berlin, facial recognition algorithms capable of identifying known terrorism suspects are trained on the crowds that pass through railway stations.⁴⁸ In the United States, these algorithms have led to more than four thousand arrests since 2010 *just* for fraud and identity theft in the state of New York alone.⁴⁹ And in the UK, cameras mounted on vehicles that look like souped-up Google StreetView cars now drive around automatically cross-checking our likenesses with a database of wanted people.⁵⁰ These vans scored their first success in June 2017 after one drove past a man in south Wales where police had a warrant out for his arrest.⁵¹

Our safety and security often depend on our ability to identify and recognize faces. But leaving that task in the hands of humans can be risky. Take passport officers, for instance. In one recent study, set to mimic an airport security environment, these professional face recognizers failed to spot a person carrying the wrong ID a staggering 14 per cent of the time – and incorrectly rejected 6 per cent of perfectly valid matches.⁵² I don't know about you, but I find those figures more than a little disconcerting when you consider the number of people passing through Heathrow every day.

As we shall see, facial recognition algorithms can certainly do better at the task than humans. But as they're applied to hunt for criminals, where the consequences of misidentification are so serious, their use raises an important question. Just how easily could one person's identity be confused with another's? How many of us have a Steve Talley-style lookalike lurking out there somewhere?

One study from 2015 seems to suggest the chances of you having your own real-life doppelgänger (whether they're a bank robber or otherwise) are vanishingly small. Teghan Lucas at the University of Adelaide painstakingly took eight facial measurements from photographs of four thousand people and failed to find a single match

among them, leading her to conclude that the chances of two people having exactly the same face were less than one in a trillion.⁵³ By that calculation, Talley wasn't just 'a bit' unlucky. Taking into account that his particular one-in-a-trillion evil twin also lived nearby *and* happened to be a criminal, we could expect it to be tens of thousands of years before another ill-fated soul fell foul of the same miserable experience.

And yet there are reasons to suspect that those numbers don't quite add up. While it's certainly difficult to imagine meeting someone with the same face as yourself, anecdotal evidence of unrelated twin strangers does appear to be much more common than Lucas's research might suggest.

Take Neil Douglas, who was boarding a plane to Ireland when he realized his double was sitting in his seat. The selfie they took, with a plane-load of passengers laughing along in the background, quickly went viral, and soon redheads with beards from across the world were sending in photos of their own to demonstrate that they



too shared the likeness. 'I think there was a small army of us at some point,' Neil told the BBC.⁵⁴

I even have my own story to add to the pile. When I was 22, a friend showed me a photo they'd seen on a local band's Myspace page. It was a collage of pictures taken at a gig that I hadn't attended, showing a number of people all enjoying themselves, one of whom looked eerily familiar. Just to be sure I hadn't unwittingly blacked out one night and wandered off to a party I now had no recollection of attending, I emailed the lead singer in the band, who confirmed what I suspected: my synth-pop-loving doppelgänger had a better social life than me.

So that's Talley, Douglas and me who each have at least one doppelgänger of our own, possibly more. We're up to three in a population of 7.5 billion and we haven't even started counting in earnest – and we're already way above Lucas's estimate of one in a trillion.

There is a reason for the discrepancy. It all comes down to the researcher's definition of 'identical'. Lucas's study required that two people's measurements must match one another exactly. Even though Neil and his lookalike are incredibly similar, if one nostril or one earlobe were out by so much as a millimetre, they wouldn't strictly count as doppelgängers according to her criteria.

But even when you're comparing two images of the same person, exact measurements won't reflect how each one of us is continually changing, through ageing, illness, tiredness, the expressions we're pulling or how our faces are distorted by a camera angle. Try to capture the essence of a face in millimetres and you'll find as much variation in one person's face as you will between people. Put simply, measurements alone can't distinguish one face from another.

Although they might not be perfectly identical, I can none the less easily imagine mixing up Neil and his twin-stranger in the photograph. Likewise in the Talley case – poor Steve didn't even look *that* similar to the real robber, and yet the images were misinterpreted

by FBI experts to the point where he was charged with a crime he didn't commit and thrown into a maximum-security cell.

As the passport officers demonstrated, it's astonishingly easy to confuse unfamiliar faces, even when they bear only a passing resemblance. It turns out that humans are astonishingly bad at recognizing strangers. It's the reason why a friend of mine claimed she could barely sit through Christopher Nolan's beautifully made film *Dunkirk* – because she struggled to distinguish between the actors. It's why teenagers find it worthwhile to 'borrow' an older friend's ID to buy alcohol. And it's why the Innocence Project, a non-profit legal organization in the United States, estimates that eyewitness misidentification plays a role in more than 70 per cent of wrongful convictions.⁵⁵

And yet, while an eyewitness might easily confuse Neil with his travel companion, his mother would surely have no problem picking out her son in the photo. When it comes to people we know, we are tremendously good at recognizing faces – even when it comes to real-life doppelgängers: a set of identical twins might be easily confused if they are only your acquaintances, but just as easily distinguished once you know them properly.

And herein lies a critical point: similarity is in the eye of the beholder. With no strict definition of similarity, you can't measure *how* different two faces are and there is no threshold at which we can say that two faces are identical. You can't define what it means to be a doppelgänger, or say how common a particular face is; nor – crucially – can you state a probability that two images were taken from the same individual.

This means that facial recognition, as a method of identification, is not like DNA, which sits proudly on a robust statistical platform. When DNA testing is used in forensics, the profiling focuses on particular chunks of the genome that are known to be highly variable between humans. The extent of that variation is key: if the DNA

sequence in a sample of body tissue found at the scene of a crime matches the sequence in a swab from a suspect, it means you can calculate the probability that both came from the same individual. It also means you can state the exact chance that some unlucky soul just happened to have an identical DNA sequence at those points.⁵⁶ The more markers you use, the lower your chances of a mismatch, and so, by choosing the number of markers to test, every judicial system in the world has complete power to decide on the threshold of doubt they're willing to tolerate.⁵⁷

Even though our faces feel so intrinsically linked to who we are, without knowing the variation across humans the practice of identifying felons by their faces isn't supported by rigorous science. When it comes to identifying people from photos – to quote a presentation given by an FBI forensics unit – 'Lack of statistics means: conclusions are ultimately opinion based.'⁵⁸

Unfortunately, using algorithms to do our facial recognition for us does not solve this conundrum, which is one very good reason to exercise caution when using them to pinpoint criminals. Resemblance and identity are not the same thing and never will be, however accurate the algorithms become.

And there's another good reason to tread carefully with face-recognition algorithms. They're not quite as good at recognizing faces as you might think.

One in a million?

The algorithms themselves work using one of two main approaches. The first kind builds a 3D model of your face, either by combining a series of 2D images or by scanning you using a special infrared camera. This is the method adopted by the Face ID system that Apple uses in its iPhones. These algorithms have worked out a way to get around the issues of different facial expressions and ageing by

focusing on areas of the face that have rigid tissue and bone, like the curve of your eye socket or the ridge of your nose.

Apple has claimed that the chance of a random person being able to unlock your phone with Face ID is one in a million, but the algorithm is not flawless. It can be fooled by twins,⁵⁹ siblings,⁶⁰ and children on their parents' phones. (Soon after the launch of Face ID, a video appeared of a ten-year-old boy who could hoodwink the facial recognition on his mother's iPhone. She now deletes her texts if there is something she doesn't want her son to look at.)⁶¹ There have also been reports that the algorithm can be tricked by a specially built 3D printed mask, with infrared images glued on for the eyes.⁶² All this means that while the algorithm might be good enough to unlock your phone, it probably isn't yet reliable enough to be used to grant access to your bank accounts.

Nor are these 3D algorithms much use for scanning passport photos or CCTV footage. For that you need the second kind of algorithm, which sticks to 2D images and uses a statistical approach. These algorithms don't directly concern themselves with landmarks that you or I could recognize as distinguishing features, but instead build a statistical description of the patterns of light and dark across the image. Like the algorithms built to recognize dogs in the 'Medicine' chapter, researchers realized recently that, rather than having to rely on humans to decide which patterns will work best, you can get the algorithm to learn the best combinations for itself, by using trial and error on a vast dataset of faces. Typically, it's done using neural networks. This kind of algorithm is where the big recent leaps forward in performance and accuracy have come in. That performance, though, comes with a cost. It isn't always clear precisely *how* the algorithm decides whether one face is like another.

That means these state-of-the-art algorithms can also be pretty easily fooled. Since they work by detecting a statistical description



of the patterns of light and dark on a face, you can trick them just by wearing funky glasses with a disruptive pattern printed on them. Even better, by designing the specific disruptive pattern to signal someone else's face, you can actually make the algorithm think you are that person – as the chap in the image above is doing, wearing glasses that make him look 'like' actress Milla Jovovich.⁶³ Using glasses as a disguise? Turns out Clark Kent was on to something.

But, targeted attacks with funky glasses aside, the recognition abilities of these statistical algorithms have prompted many admiring headlines, like those that greeted Google's FaceNet. To test its recognition skills, FaceNet was asked to identify five thousand images of celebrities' faces. Human recognizers had previously attempted the same task and done exceptionally well, scoring 97.5 per cent correct identifications (unsurprisingly, since these celebrity faces would have been familiar to the participants).⁶⁴ But FaceNet did even better, scoring a phenomenal 99.6 per cent correct.

On the surface, this looks as if the machines have mastered superhuman recognition skills. It sounds like a great result, arguably good enough to justify the algorithms being used to identify criminals. But there's a catch. Five thousand faces is, in fact, a pathetically

small number to test your algorithm on. If it's going to be put to work fighting crime, it's going to need to find one face among millions, not thousands.

That's because the UK police now hold a database of 19 million images of our faces, created from all those photos taken of individuals arrested on suspicion of having committed a crime. The FBI, meanwhile, has a database of 411 million images, in which half of all American adults are reportedly pictured.⁶⁵ And in China, where the ID card database gives easy access to billions of faces, the authorities have already invested heavily in facial recognition. There are cameras installed in streets, subways and airports that will supposedly spot everything from wanted criminals to jay-walkers as they travel through the country's cities.⁶⁶ (There's even a suggestion that a citizen's minor misdemeanours in the physical world, like littering, will form part of their Sesame Credit score – attracting all of the associated punishments that we uncovered in the 'Data' chapter.)

Here's the problem: the chances of misidentification multiply dramatically with the number of faces in the pile. The more faces the algorithm searches through, the more chance it has of finding two faces that look similar. So, once you try using these same algorithms on bigger catalogues of faces, their accuracy plummets.

It would be a bit like getting me to match ID cards to ten strangers and – when I got full marks – claiming that I was capable of correctly identifying faces 100 per cent of the time, then letting me wander off into the centre of New York to identify known criminals. It's inevitable that my accuracy would drop.

It's just the same with the algorithms. In 2015, the University of Washington set up the so-called MegaFace challenge, in which people from around the world were invited to test their recognition algorithms on a database of 1 million faces.⁶⁷ Still substantially smaller than the catalogues held by some government authorities,

but getting closer. Even so, the algorithms didn't handle the challenge well.

Google's FaceNet – which had been close to perfect on the celebrities – could suddenly manage only a 75 per cent* identification rate.⁶⁸ Other algorithms came in at a frankly pathetic 10 per cent success rate. At the time of writing, the world's best is a Chinese offering called Tencent YouTu Lab, which can manage an 83.29 per cent recognition rate.⁶⁹

To put that another way, if you're searching for a particular criminal in a digital line-up of millions, based on those numbers, the *best-case scenario* is that you won't find the right person one in six times.

Now, I should add that progress in this area is happening quickly. Accuracy rates are increasing steadily, and no one can say for certain what will happen in the coming years or months. But I can tell you that differences in lighting, pose, image quality and general appearance make accurately and reliably recognizing faces a very tricky problem indeed. We're some way away from getting perfect accuracy on databases of 411 million faces, or being able to find that one-in-a-trillion doppelgänger match.

Striking a balance

These are sobering facts, but not necessarily deal-breakers. There are algorithms good enough to be used in some situations. In Ontario, Canada, for instance, people with a gambling addiction can voluntarily place themselves on a list that bars them from entering a casino. If their resolve wavers, their face will be flagged by

* This number relates to how many true matching faces the algorithm missed when it was tuned to avoid misidentification. There's more on the kinds of errors an algorithm can make in the 'Justice' chapter, and the different ways to measure accuracy in the 'Medicine' chapter.

recognition algorithms, prompting casino staff to politely ask them to leave.⁷⁰ The system is certainly unfair on all those mistakenly prevented from a fun night on the roulette table, but I'd argue that's a price worth paying if it means helping a recovering gambling addict resist the temptation of their old ways.

Likewise in retail. In-store security guards used to have offices plastered with Polaroids of shoplifters; now algorithms can cross-reference your face with a database of known thieves as soon as you pass the threshold of the store. If your face matches that of a well-known culprit, an alert is sent to the smartphones of the guards on duty, who can then hunt you out among the aisles.

There's good reason for stores to want to use this kind of technology. An estimated 3.6 million offences of retail crime are committed every year in the UK alone, costing retailers a staggering £660 million.⁷¹ And, when you consider that in 2016 there were 91 violent deaths of shoplifting suspects at retail locations in the United States,⁷² there is an argument that a method of preventing persistent offenders from entering a store before a situation escalates would be good for everyone.

But this high-tech solution to shoplifting comes with downsides: privacy, for one thing (FaceFirst, one of the leading suppliers of this kind of security software, claims it doesn't store the images of regular customers, but shops are certainly using facial recognition to track our spending habits). And then there's the question of who ends up on the digital blacklist. How do you know that everyone on the list is on there for the right reasons? What about innocent until proven guilty? What about people who end up on the list accidentally: how do they get themselves off it? Plus again there's the potential for misidentification by an algorithm that can never be perfectly accurate.

The question is whether the pros outweigh the cons. There's no easy answer. Even retailers don't agree. Some are enthusiastically adopting the technology, while others are moving away from

it – including Walmart, which cancelled a FaceFirst trial in their stores after it failed to offer the return on investment the company were hoping for.⁷³

But in the case of crime the balance of harm and good feels a lot more clear cut. True, these algorithms aren't alone in their slightly shaky statistical foundations. Fingerprinting has no known error rate either,⁷⁴ nor do bite mark analysis, blood spatter patterning⁷⁵ or ballistics.⁷⁶ In fact, according to a 2009 paper by the US National Academy of Sciences, none of the techniques of forensic science apart from DNA testing can 'demonstrate a connection between evidence and a specific individual or source'.⁷⁷ None the less, no one can deny that they have all proved to be incredibly valuable police tools – just as long as the evidence they generate isn't relied on too heavily. But the accuracy rates of even the most sophisticated facial recognition algorithms leave a lot to be desired. There's an argument that if there is even a slight risk of more cases like Steve Talley, then a technology that isn't perfect shouldn't be used to assist in robbing someone of their freedom. The only problem is that stories like Talley's don't quite paint the entire picture. Because, while there are enormous downsides to using facial recognition to catch criminals, there are also gigantic upsides.

The tricky trade-off

In May 2015, a man ran through the streets of Manhattan randomly attacking passers-by with a black claw hammer. First, he ran up to a group of people near the Empire State Building and smashed a 20-year-old man in the back of the head. Six hours later he headed south to Union Square and, using the same hammer, attacked a woman quietly sitting on a park bench on the side of the head. Just a few minutes later he appeared again, this time targeting a 33-year-old woman walking down the street outside the park.⁷⁸ Using

surveillance footage from the attacks, a facial recognition algorithm was able to identify him as David Baril, a man who, months before the attacks, had posted a picture on Instagram of a hammer dripping with blood.⁷⁹ He pleaded guilty to the charges stemming from the attacks and was sentenced to 22 years in prison.

Cold cases, too, are being re-ignited by facial recognition breakthroughs. In 2014, an algorithm brought to justice an American man who had been living as a fugitive under a fake name for 15 years. Neil Stammer had absconded while on bail for charges including child sex abuse and kidnapping; he was re-arrested when his FBI 'Wanted' poster was checked against a database of passports and found to match a person living in Nepal whose passport photo carried a different name.⁸⁰

After the summer of 2017, when eight people died in a terrorist attack on London Bridge, I can appreciate how helpful a system that used such an algorithm might be. Youssef Zaghba was one of three men who drove a van into pedestrians before launching into a stabbing spree in neighbouring Borough Market. He was on a watch list for terrorist suspects in Italy, and could have been automatically identified by a facial recognition algorithm before he entered the country.

But how do you decide on that trade-off between privacy and protection, fairness and safety? How many Steve Talleys are we willing to accept in exchange for quickly identifying people like David Baril and Youssef Zaghba?

Take a look at the statistics provided by the NYPD. In 2015, it reported successfully identifying 1,700 suspects leading to 900 arrests, while mismatching five individuals.⁸¹ Troubling as each and every one of those five is, the question remains: is that an acceptable ratio? Is that a price we're willing to pay to reduce crime?

As it turns out, algorithms without downsides, like Kim Rossmo's geoprofiling, discussed at the beginning of the chapter, are the

exception rather than the rule. When it comes to fighting crime, every way you turn you'll find algorithms that show great promise in one regard, but can be deeply worrying in another. PredPol, HunchLab, Strategic Subject Lists and facial recognition – all promising to solve all our problems, all creating new ones along the way.

To my mind, the urgent need for algorithmic regulation is never louder or clearer than in the case of crime, where the very existence of these systems raises serious questions without easy answers. Somehow, we're going to have to confront these difficult dilemmas. Should we insist on only accepting algorithms that we can understand or look inside, knowing that taking them out of the hands of their proprietors might mean they're less effective (and crime rates rise)? Do we dismiss any mathematical system with built-in biases, or proven capability of error, knowing that in doing so we'd be holding our algorithms to a higher standard than the human system we're left with? And how biased is too biased? At what point do you prioritize the victims of preventable crimes over the victims of the algorithm?

In part, this comes down to deciding, as a society, what we think success looks like. What is our priority? Is it keeping crime as low as possible? Or preserving the freedom of the innocent above all else? How much of one would you sacrifice for the sake of the other?

Gary Marx, professor of sociology at MIT, put the dilemma well in an interview he gave to the *Guardian*: 'The Soviet Union had remarkably little street crime when they were at their worst of their totalitarian, authoritarian controls. But, my God, at what price?'⁸²

It may well be that, in the end, we decide that there should be some limits to the algorithm's reach. That some things should not be analysed and calculated. That might well be a sentiment that eventually applies beyond the world of crime. Not, perhaps, for lack of trying by the algorithms themselves. But because – just maybe – there are some things that lie beyond the scope of the dispassionate machine.