THE CENSORSHIP ISSUE

CENSORED

We were promised the internet would create a global commons for truly free speech. Instead, our online lives are increasingly ruled by unaccountable gatekeepers.
By **Luke Collins**

GATEKEEPERS OF OUR LIVES

WHAT'S WRONG WITH women's nipples? Or rather, why do some social-media platforms allow men to post topless pictures of themselves, but ban women from doing the same thing?

The answer to this question neatly encapsulates the key issue of online censorship today. We may think of the online world as a global commons, a public space for free expression, but it is increasingly a series of private spaces within which we are given limited permission to act. Our access to each other and to the resources available within these spaces is increasingly being controlled not just by law, but by gatekeepers whose motivations may include political advantage, social control, or simple profit.

For example, the Free The Nipple movement was launched in 2012 to highlight the differing ways in which society expects the bodies of women and men to be portrayed, as part of a campaign focused on "equality, empowerment and the freedom of all human beings".

When women started posting topless images of themselves using the #freethenipple hashtag on Instagram, their posts were banned for breaching the service's Community Guidelines. In a 2015 interview with *Business Insider*, Kevin Systrom, CEO of Instagram, explained that one reason for this was to meet the content guidelines of Apple's App Store. One gatekeeper (Instagram) bowed to another (Apple) to control how a campaign for equality was expressed online.

In a sense, this should be no surprise. When Steve Jobs launched the iPhone App Store in 2008, he said that apps would not be allowed to distribute pornography, hog bandwidth, breach user privacy, act maliciously, or otherwise violate rules set by the company. The motivation, Jobs said, was "to get a ton of apps out there".

This is fair enough in the context of furthering a business' interests. However, as the web, Facebook, YouTube, Instagram, Twitter, WhatsApp, Sina Weibo, Snapchat and more gather billions of users and become prime conduits for our communications, the way they are governed and controlled becomes increasingly important to the functioning of society.

The pressure group Access Now monitors and campaigns against internet shutdowns

How big an issue is online censorship?

According to estimates by Freedom House, a US watchdog backed by sponsors ranging from the US State Department and the Dutch Ministry of Foreign Affairs to BAE Systems and Google:

- 3.2 billion people have access to the internet
- 67 per cent live in countries where criticism of the government, military, or ruling family has been subject to censorship
- 60 per cent live in countries where ICT users were arrested or imprisoned for posting content on political, social and religious issues
- 49 per cent live in countries where individuals have been attacked or killed for their online activities since June 2015
- 47 per cent live in countries where insulting religion online can result in censorship or jail time
- 33 per cent live in countries where online discussion of LGBT+ issues can be repressed or punished
- 38 per cent live in countries where social media or messaging apps were blocked over the past year
- 27 per cent live in countries where users have been arrested for writing, sharing or even liking Facebook posts
- 38 per cent live under governments that disconnected internet or mobile phone access, often for political reasons.

In other words, the global commons is increasingly strictly policed, and offending its gatekeepers can lead to harsh punishments.

What gets censored?

Some online censorship is very broad, cutting users off from the internet for limited periods.

Access Now, a pressure group set up to defend and extend digital rights of users at risk around the world by influencing policy, advocacy, and direct technical support, runs a #keepiton campaign to counter internet shutdowns. It argues that such shutdowns harm everyone, from businesses through to human-rights activists. Its count of the number of shutdowns is rising: Access Now documented 15 in 2015, but 56 in 2016.

India's Software Freedom Law Centre maintains a website, internetshutdowns.in, that tracks instances of internet shutdowns in India. According to its data, there were 14 such shutdowns in 2015, 31 in 2016, and 47 so far in 2017. Many of the shutdowns were 'preventive', imposed in anticipation of a law-and-order issue, with the rest imposed as a reaction to such an issue.

This trend of increasing control over access to the internet doesn't sit well with the Indian government's Digital India initiative, announced in 2015 "to transform India into a digitally empowered society and knowledge economy".

Some censorship is capricious, focusing on protecting the pride of powerful people.

In July, for example, news agencies in Beijing reported that online searches for Winnie the Pooh, the honey-loving children's book character, were being blocked as >

CENSORED TOPICS BY COUNTRY

Censorship was reflected if state authorities blocked or ordered the removal of content, or detained or fined users for posting content on the topics considered. The chart does not consider extra-legal pressures such as violence, self-censorship or cyber attacks, even where the state is believed to be responsible.

Source: 'Freedom on the Net 2016', Freedom House

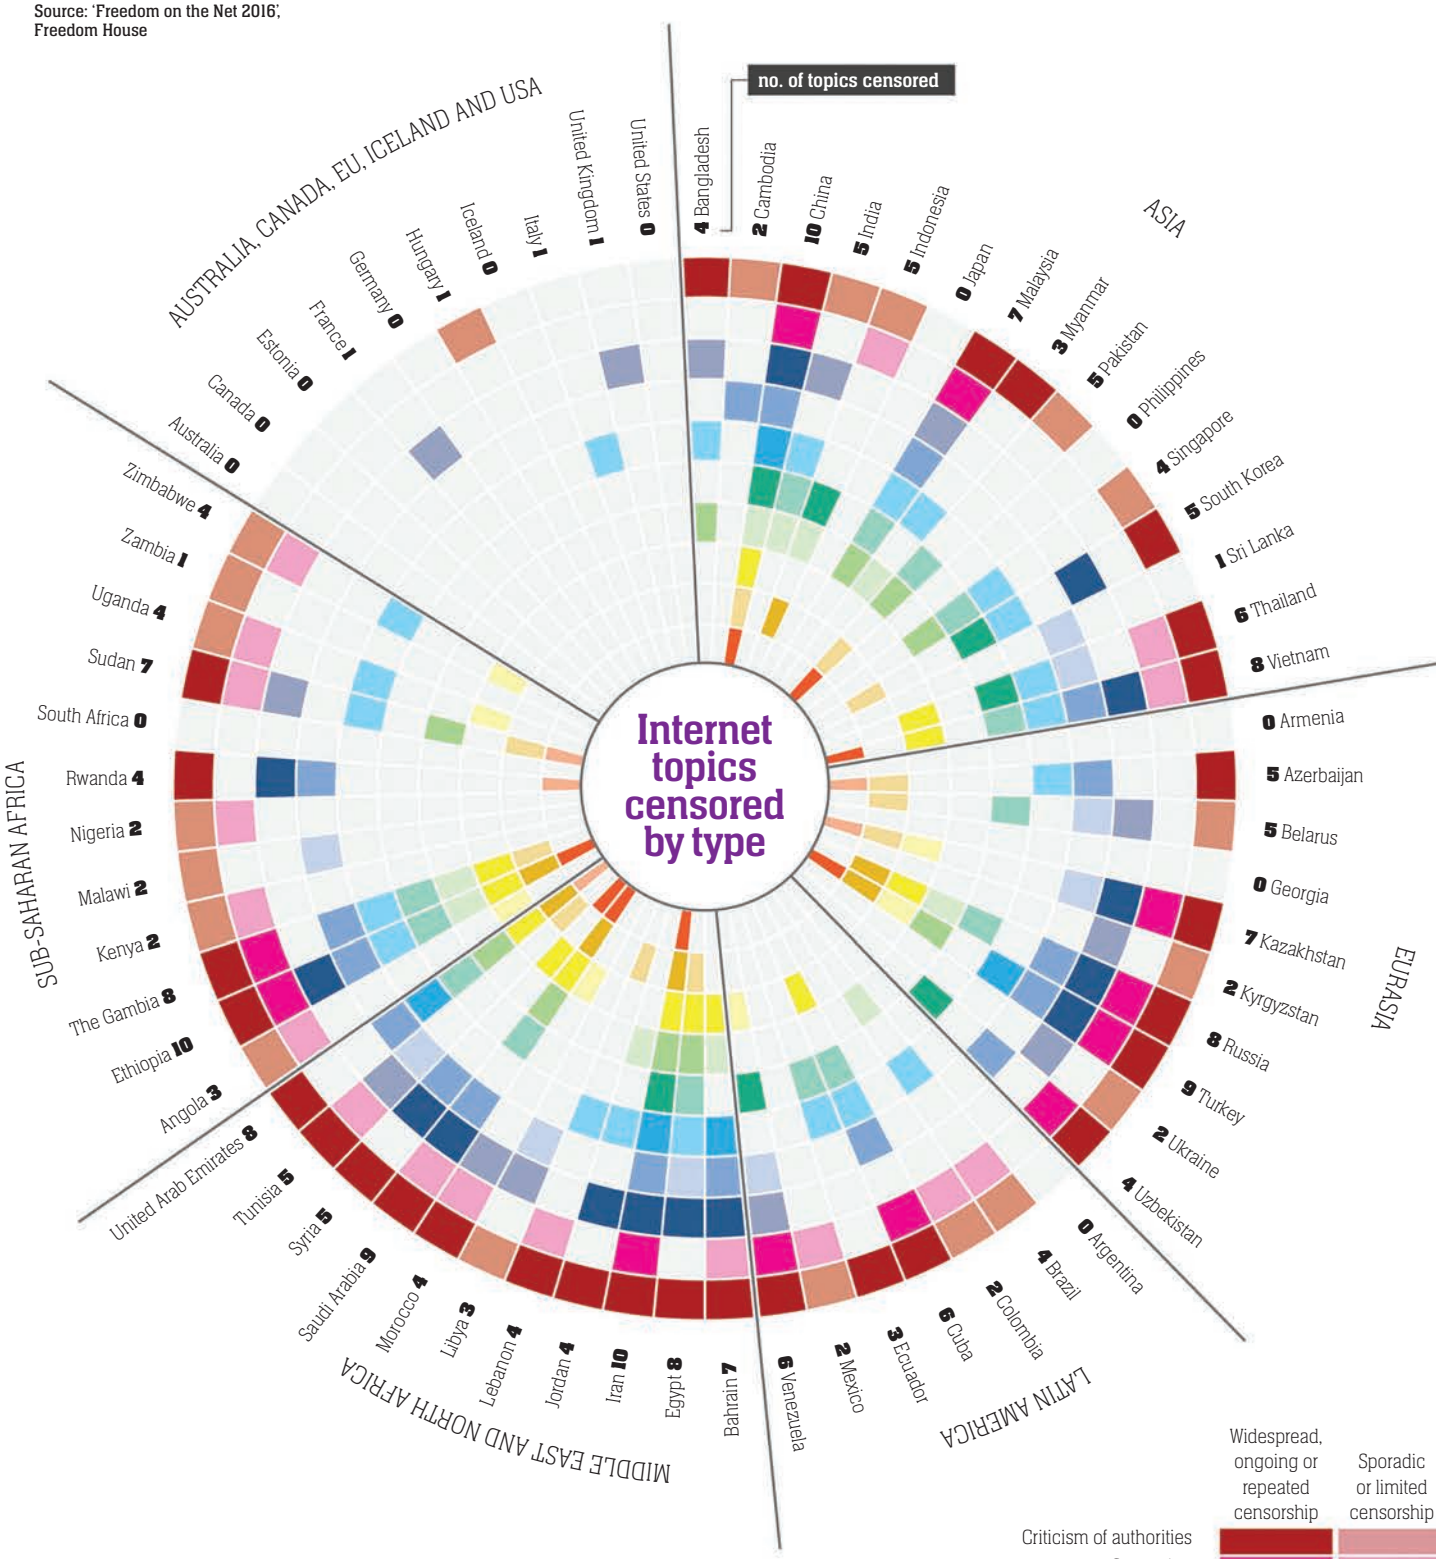

< 'illegal content', following a spate of unflattering comparisons between the portly bear and Xi Jinping, the Chinese president.

Similarly, a Turkish man was given a one-year suspended sentence for creating an image that juxtaposed Recep Tayyip Erdogan, Turkey's President, with Gollum from *Lord of the Rings*. An Egyptian student was jailed for three years for posting an image of his country's President with Mickey Mouse ears on Facebook.

Some censorship is about sustaining political control.

For example, in August the telecoms regulator of the Democratic Republic of Congo wrote to the telco Orange DRC, telling it to slow down access to social media sites, including Facebook, WhatsApp, Instagram and Twitter, in order to reduce users' ability to share "abusive messages". The order was widely seen as a reaction to growing opposition to President Joseph Kabila, who has said he will not step down from government when his term ends in December. Last year, eight other African governments – Chad, Ethiopia, Egypt, Gabon, Gambia, Uganda, Zambia and Zimbabwe – ordered similar internet disruptions in response to political events.

China has steadily increased its oversight of the internet and social media in the five years since President Xi Jinping took power, tending to exert particularly strong control when its political situation is changing rapidly, such as during party congresses.

In August, companies that host web servers in China were told to practise shutting down nominated web pages quickly, to deal with what China's Ministry of Public Security called "the problem of smaller websites illegally disseminating harmful information". It appears the drill was intended to prepare China's web hosts for the possibility of increased censorship during the Communist Party's forthcoming national congress.

How censorship is intensifying

China has also cracked down on the use of virtual private networks (VPNs) to circumvent its internet controls and surveillance. Apple recently had to remove a number of VPN apps from its China App Store, because they did not use the Chinese state's network infrastructure. Earlier in the year, Russia banned the use of VPNs and other anonymising services that allow users to access content it judges to be unlawful.

According to 'Freedom on the Net 2016', the most recent report on internet censorship from Freedom House, governments are now censoring a broader range of topics, including LGBT+ issues, digital activism, satire, and political opposition.

The level of control can be pervasive: Azerbeijan's national domain-name registrar refused to register website domains such as lgbt.az; in Indonesia, a messaging platform was asked to remove gay and lesbian-themed emojis from its service; South Korean regulators asked a website to exercise "restraint" after it linked to an online gay drama; the Turkish government blocked access to popular LGBT+ websites for a number of weeks in 2015.

Some censorship is a function of the collision between old business models and new technology, as seen in the suppression of file-sharing sites or access to VoIP services.

For example, in August some regions of India blocked access to the Wayback Machine because its URL was among a list of 'pirates' handed to the Madras High Court by an Indian film distributor trying to protect its copyright. As a result, users also lost access to the internet's key archive site for legitimate research purposes.

The UK has similarly been influenced by organisations such as the Premier League and the Motion Picture Association of America to block access to sites which help people pirate their content. Try to reach a leading BitTorrent-based file-sharing site from a UK broadband account and you're likely to be redirected by your internet service provider (ISP) to a web page which explains that access to the site has been blocked by order of the High Court.

Net neutrality also plays a role in the censorship debate: allowing ISPs to prioritise traffic from one source over another is effectively a powerful form of commercial censorship.

The Freedom of the Net report analyses the censorship situation in 65 countries, representing 88 per cent of the world's internet users. It scores these countries' 'internet freedom' from 0 (the most free) to 100 (the most restricted), based on three broad factors: obstacles to access, such as legal or technical barriers; limits on content, such as filtering and blocking; and violations of user rights, such as privacy, surveillance, and the repercussions of online activity.

Its latest analysis shows that 34 of the countries it surveyed have become less free online between 2015 and 2016. The worst performers were Uganda, Bangladesh, Cambodia, Ecuador and Libya, which have variously restricted social media platforms, put the entire telecommunications industry under government control, arrested people for their social-media posts, and even seen the murder of a blogger.

On the other hand, 14 countries registered modest improvements in internet freedom, in part due to regulatory changes and in part due to digital activism.

The threat of self-censorship

The internet's pervasive role as bearer of our communications is also causing people to censor themselves.

Social networks such as YouTube, Instagram and Snapchat have given people positive new ways to find and connect with each other. Yet it has also created pressure, especially among young people, to present an idealised version of their lives, in which the sun is always shining, friends are always smiling, and the food looks great. Such pressure narrows the range of publicly acceptable behaviours – such as the ordinary sadness and confusion of adolescents – with potentially worrying consequences. Indeed, a study released by the UK's Royal Society for Public Health in May this year argued that "social media may be fuelling a mental health crisis" among young people.

The extent to which individuals are surveilled on the internet is also giving people pause. For example, the UK's Investigatory Powers Act, also known as the 'Snooper's Charter', became law last November and provides the government with sweeping powers to intercept and collect communications data in bulk. Service providers must retain 'internet connection records', that is, a year-long history of the websites that each of their UK users has visited. The full consequences of this have yet to play out in public, but the notion that government can access a list of every website you visit, 'just in case', is likely to limit some people's behaviours.

Russia is taking a related, but perhaps even more pernicious, tack, according to an analysis of a 2014 survey of Russian citizens' concerns about internet and media usage.

Distribution of global internet users by country and FOTN status

The 65 countries covered in Freedom on the Net represent 88 per cent of the world's internet user population. Over 1.2 billion internet users, or 40 per cent of global users, live in three countries – China, India and the United States – that span the spectrum of internet freedom environments, from Free to Not Free.

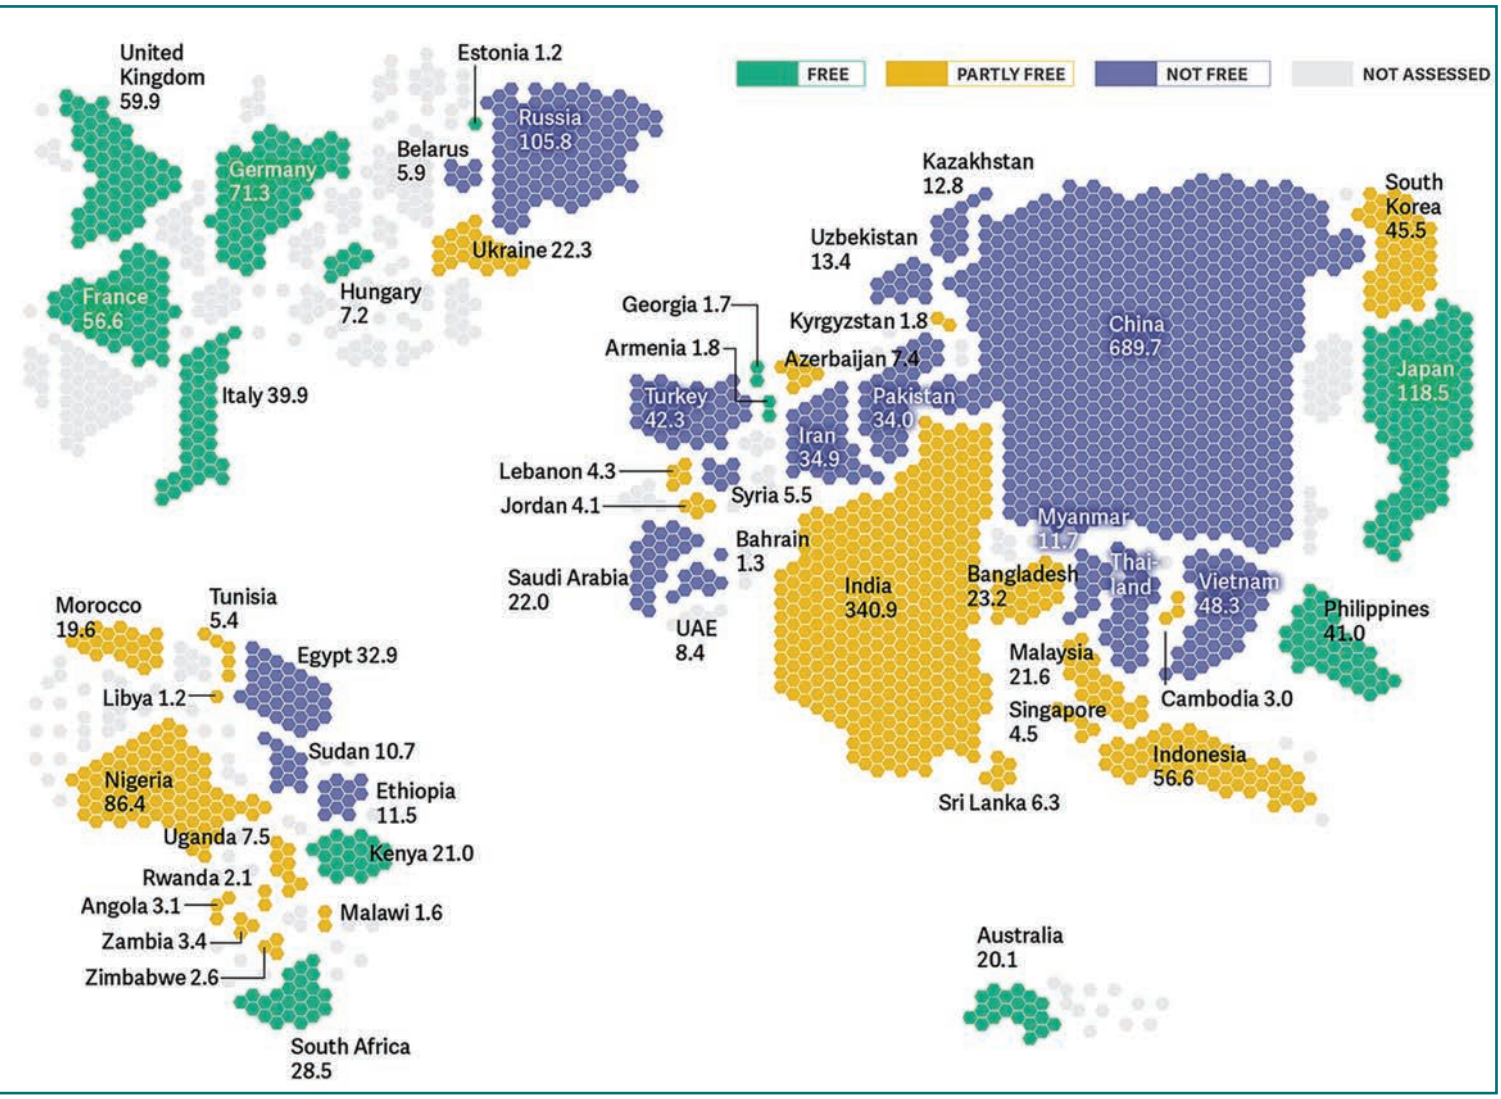

The study, carried out by researchers at Ohio State University, found Russians who consume relatively more state-backed media feel the internet is a greater national security threat than those who consume less. This has increased support for online political censorship, and even persuaded some users to censor their own internet usage.

"People report they don't go to certain websites 'because the government says it is bad for me'," says Erik Nisbet, co-author of the study and an associate professor of communication at The Ohio State University.

Fighting back

The vision of the online world as a global commons has been rapidly eroded by realities of its implementation, which has created myriad technical, regulatory and commercial opportunities for gatekeepers to surveil and control what we can access online.

Unsurprisingly, a number of technical responses have emerged.

Messaging apps such as WhatsApp and Apple's Messages are encrypted end to end, protecting users' conversations from all but the most determined surveillance systems. There are regular calls to ban encryption, or demand that service providers include 'back doors' in their platforms to enable government surveillance, but any such back doors would open the services up to hacking (and undermine the providers' business models) and so are strongly resisted.

On the web, a number of techniques and services have emerged to help users circumvent censorship.

VPNs create an encrypted tunnel between a user and another server, which can enable them to connect securely to servers that can access content that is otherwise unavailable in their own regions.

The Tor Project protects users' anonymity by routing communications through a globally distributed network of relays run by volunteers. This stops anyone who is watching a user's internet connection from learning which sites they are visiting, and also prevents the sites they visit from learning their physical location.

Lantern is a service for accessing blocked websites and apps, by encrypting the user's traffic and then using a variety of techniques to overcome firewalls and censorship strategies, such as dynamically allocated proxy servers. However, it is not a tool for anonymity like Tor.

Psiphon is another tool for overcoming internet censorship, using a combination of technologies such as VPNs, SSH (Secure Shell) and a global network of proxy servers to help users access the open internet, particularly from countries with heavy filtering and controls.

Increasing surveillance has also led to the growth of anonymity networks such as Freenet, a program to enable users to anonymously share files, browse and publish 'Freesites', websites that are only available through Freenet, and chat on private forums. Users give up part of their computers' storage and bandwidth to host Freenet nodes, which hold encrypted data and through which encrypted communications are routed. It's possible to add another layer of security to the service by only allowing it to connect to Freenet users that you trust.

Ceno is a peer-to-peer-based approach to circumventing censorship, which builds upon Freenet and adds the concept of caching any websites retrieved from the open internet within the Freenet network. This means the site is then available to all Ceno users without them having to make another request to the internet from within a censored zone.

Trust and censorship

Using any of the techniques discussed above takes a leap of faith – that a service is as secure as its developers believe it is, that other actors in private networks are only interested in their own anonymity rather than who their fellow network members are, and so on.

A lot of censorship is about trust, too – or lack of it: governments don't trust citizens to act responsibly; politicians don't trust the electorate's judgement of their performance; businesses don't trust their customers' reactions to unfamiliar situations.

If it is true that knowledge is power, then increasing asymmetry between those who control the flow of information and those who consume it is making us increasingly powerless.

The question at the heart of the censorship debate is therefore: who do we trust – ourselves, or an ever-expanding cohort of unaccountable gatekeepers? *

CENSORED

THE FACE OF FUTURE

CCTV-SHAUN DREYER, DREAMSTIME

A lack of clarity surrounding the legal aspects of biometric identification is holding back efforts to implement facial recognition technology. By **Josh Loeb**

"MEETING EYE TO LENS, the feeling of intrusion is unlike the ubiquitous CCTV we are usually subjected to – you know you are being measured, assessed, identified, invaded."

That is what Silkie Carlo, human rights group Liberty's technology policy officer, said about her experience of peering into the glassy artificial eye of an automated facial-recognition camera that was being trialled by the Metropolitan Police at the Notting Hill Carnival in London earlier this year.

At the same event, police deployed 140 'super recognisers'. These are officers with exceptionally high perception and memory skills, and they had been tasked with scanning crowds in search of the faces of known troublemakers.

This detail passed without comment from Carlo in a 1300-word blog post she wrote protesting against the Met's use of biometric technology, which she denounced as crude, wasteful and potentially racist.

She also did not mention the reams of footage of the annual street party that was being constantly recorded by spectators via smartphones and instantly uploaded to the internet, where it could be viewed by all –

SURVEILLANCE

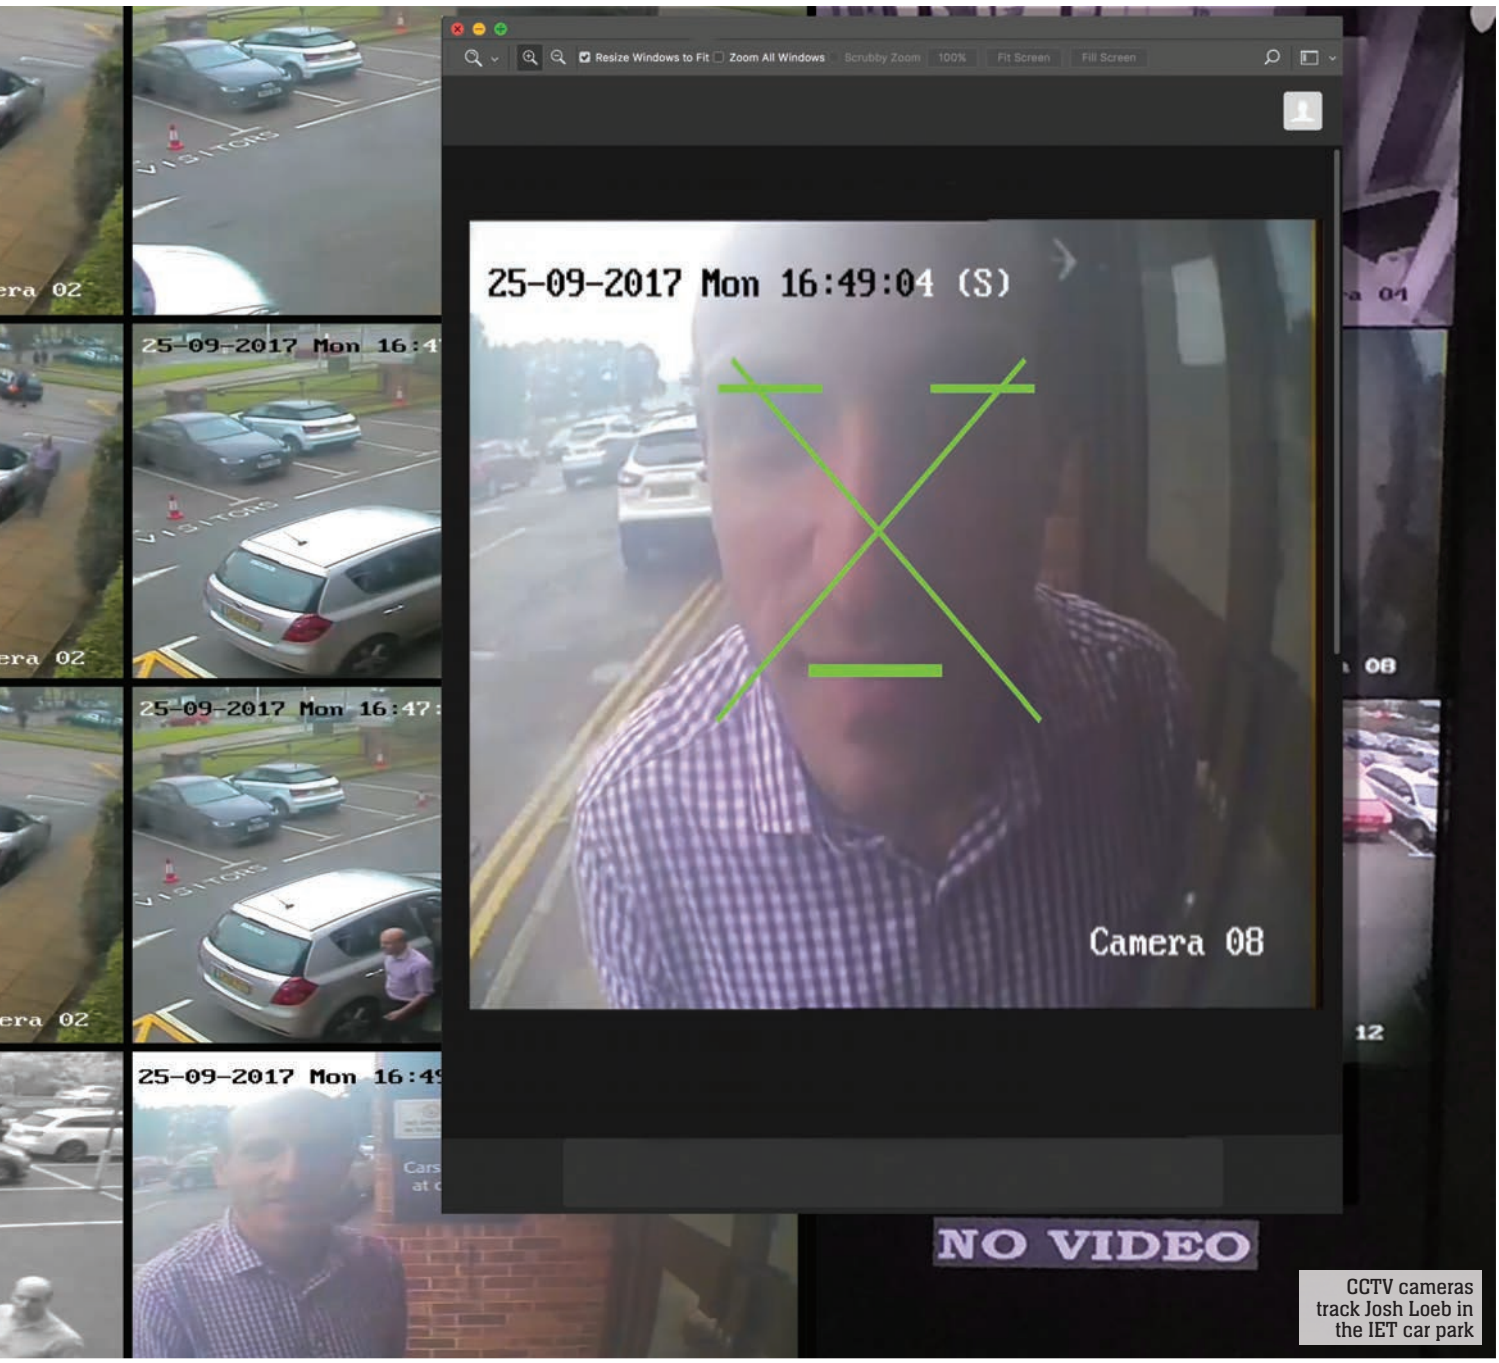

CCTV cameras track Josh Loeb in the IET car park

including the police.

These omissions are relevant because they show how new technology can spark particular kinds of fears, rational or otherwise, when used to help maintain law and order. Dystopian fiction – from 'Nineteen Eighty-Four' to 'Minority Report' – plays on these fears.

"It is the stuff of dystopian literature for a reason," Carlo wrote in her blog post. "In a society that has rejected ID cards, the prospect of biometric checkpoints overshadowing our public spaces is plainly unacceptable and frankly frightening."

Trust in governments was hit by Edward Snowden's revelations about the extent of surveillance in many Western countries, and offline as well as online anonymity is increasingly being sought by some concerned citizens. However, is biometric surveillance really so frightening in liberal, democratic societies where people are afforded legal and constitutional protections, and where daily life is quite different from that in totalitarian dictatorships?

"I think people worry about Big Brother far more than necessary," says Dr Josh Davis, an expert in super recognisers who has also carried out research into identification of unfamiliar faces via CCTV. "They've seen facial recognition in science fiction or 'Spooks' or something like that and they probably think it is better than it really is.

"It's right for people to be concerned that there may be a risk to their privacy or normal rights from computerised systems, I just don't think at the present time those systems are sufficiently good to cause those sorts of worries."

Regardless of whether the technology is any good, the clamouring of campaigners seems to have had an effect. Two weeks >

Digital Barriers says its SmartVis system could help trace missing young people

< after a media storm over the Met's use of facial biometrics it emerged that some other UK police forces had pulled back – for now, at least – from trialling some facial-recognition software in their jurisdictions, apparently for fear of an adverse public reaction.

Charlie Hedges, a former senior National Crime Agency officer who now runs a security consultancy, says trials meant to have taken place in several large shopping centres in conjunction with technology firm Facewatch had been scrapped at the eleventh hour. "The police initially were OK with it and then, just as we were about to make it go live, they pulled out. It's so frustrating," he told *E&T*.

"If there is a high-risk missing child, particularly in a shopping-centre scenario – a Jamie Bulger type of thing [a toddler kidnapped and murdered in 1993] – because in shopping centres you have such a huge number of CCTV cameras everywhere, the development of increasingly sophisticated facial-recognition technology means that if you have those extremely vulnerable missing children on a watch list that is linked to facial-recognition enabled cameras you've got a chance of being able to identify them more quickly."

UK trials of other facial-recognition products have also been slow to start. British security technology company Digital Barriers has not yet had its SmartVis software platform tested in real-world law-enforcement scenarios in the UK, despite offering it free of charge to police in cases involving missing young people.

Another company, NEC, did manage to have its facial-recognition service, NeoFace, road-tested by Leicestershire Police three years ago and by South Wales Police at the football Champions League final in Cardiff earlier this year, but there have been no UK trials of it announced since.

Another prime mover in the field of facial recognition, SeeQuestor – which works with the Met and British Transport Police – did not reply to an inquiry from *E&T* seeking information about any trials of its products.

Continuous live checking

Facial-recognition security software typically scans CCTV footage looking for matches with images on watch lists. It has been described as like running a constant Google-style search on a series of faces through CCTV footage in real time.

Privately some security insiders speak of reticence among police leaders about embracing the technology, saying a regulatory gap left because there is no official rule book around use of facial biometrics has made chief constables wary. They also point to the weaknesses of CCTV images or facial-recognition evidence when used in court prosecutions, and the potential for successful legal challenges against convictions reliant on it.

The UK nevertheless remains one of the most surveilled countries in the world. The Police National Database now holds 19 million facial images, many of people who were arrested but never charged or convicted. Earlier this year Mike Barton, the chief constable of Durham Constabulary, said his officers were compiling image databases of "villains" using footage from body-worn video cameras and were using this to study gait and mannerisms as well as facial features.

> 'Considering the amount of CCTV in the UK, we have never harnessed that information properly. We haven't trained police officers to use it.'
> **John Kennedy**
> Key Forensic Services

While videos are usually deleted after a month unless they are needed for prosecutions, they are reportedly retained for longer in Durham in the case of suspects with previous convictions.

The Biometrics Commissioner's recent assessment that images on police databases are being used in a way that goes "far beyond" custody purposes has given some in law enforcement pause for thought. Several years ago MPs on the Science and Technology Committee made the same point, but the lackadaisical attitude shows no sign of abating.

The government was supposed to have published a biometrics strategy three years ago, but this has been serially delayed. The Home Office promises only that the strategy will be released "in due course". Meanwhile, the Information Commissioner's Office and the forensics watchdog are understood to be probing the ramifications for the police of new data-protection laws.

Legally, DNA and fingerprint records cannot normally be retained for longer than six months in cases where no criminal charges are brought. However, there are no such safeguards around the retention of facial biometrics.

"Once you start using an unconvicted photograph for anything other than intelligence use, it's likely illegal," opines David Videcette, a former senior Scotland Yard detective who played a major role in the investigations into the July 2007 London Tube bombings.

The regulatory vacuum should, he says, be "sorted out" via a change in the law and a retrospective weeding of photographs of innocent people. He is also uncomfortable about situations where "outside contractors" – technology companies and private security staff at shopping centres, for example – might be allowed to leverage biometric data that police should arguably not have on file in the first place.

Restricted access

"Internally within the police, in counter-terrorism, we can look at anything," he says. "For borough-based officers dealing with high volume crime – car thefts, minor assaults, perhaps even a GBH [grievous bodily harm], it's 'No you're not allowed to look at this database' and 'No you're not allowed to look at that database'. You're not allowed to know what's there.

"We've got massive databases, huge databases, which we use for intelligence purposes only, but we can't justify putting them in other people's databases or going public with some of this stuff.

"Even around things like automatic number plate recognition, we've got an incredible amount of data which means you can virtually track a car from one point to another. We have these databases which will tell us exactly which car went where. Internally we say you can only access this database in the most extreme circumstances, and really it's only counter-terrorism. Even with murder, up to a point, we would say no."

Lord Harris of Haringey, a Labour peer who carried out a review for London Mayor Sadiq Khan into the city's preparedness for terror attacks, sees an irony in the fact that police surveillance of the public realm remains so controversial but smartphones and drone cameras have meant ordinary people can record so extensively in public, and sometimes private, spaces.

He also points out that Facebook allows users access to facial-recognition technology that can group photographs based on which

people have been automatically tagged in them. Because of different data-protection laws the service, called Moments, exists in different versions in the US and UK, though facial recognition is already used at passport control at many UK airports.

"This technology is advancing so rapidly that, in a sense, it's a bit silly for the police not to be making use of it," says Lord Harris. "Is this any different from the fact that you might have a briefing for 60 police officers showing them a picture of a suspect and telling them to go scan the crowd and if they see them, to go and arrest him or her? They've been using human spotters of people for a long time and I'm not sure what the difference really is apart from scale and numbers."

John Kennedy, head of digital forensics for the company Key Forensic Services and a pioneer in police use of video, argues that the extent of surveillance is less important than the ability to analyse and store images well – skills lacking in the UK.

"Historically, over the last two or three decades, the UK has led the world in surveillance, and there's an assumption that we are extremely good at using that evidence," he says. "But considering the amount of CCTV in the UK, we have never actually harnessed that information properly. The resources simply haven't been available. We haven't trained police officers to properly analyse [the images] or store them and create databases to use them for intelligence-gathering purposes.

China as a model?

"In China, they have actually gone about it properly. They have created proper databases which they have kept updated, adding to the metadata so that people can be flagged up on screen and you can see their name, date of birth and probably their antecedent history.

"We in the UK simply haven't done that. We haven't sat down and said, look, this is a huge evidential medium. How are we going to manage it? How are we going to use it to our best advantage? We're playing catch-up."

While privacy campaigners appear to view the UK as an Orwellian dystopia where the authorities have Stasi-style snooping powers, at the other end of the spectrum the securocrats claim they are fighting crime with one hand forever tied behind their backs. The truth is less clear-cut, though.

Technologically, the police have awesome surveillance capabilities. However, for various reasons, they are currently incapable of using them to anything like their full potential most of the time. Many images on databases have often been recompressed or have the wrong aspect ratio, so are utterly useless in practice. Police leaders either don't know how to use the technology or are reluctant to. In many cases there is too much data, of too poor quality, and it is not managed in a way that could ever make it useful. Moreover, cases are often unlikely to be successfully brought to court.

"The counter-terrorist agencies are obviously very, very good at this sort of thing, and they have the resources," says Kennedy. "But in your provincial forces it all depends on who the chief is and how technologically-minded they are and how much they believe in the value of CCTV as an investigative tool. There's no joined-up approach."

Unlike in China, where surveillance systems are centralised and all-embracing, and where the government's desire to rigidly control people is its incentive to pour resources into biometrics, the situation in the UK has been described as shoddy, fragmented, chaotic and lacking in direction.

That may or may not come as a relief to those worried about the power of the state, but it is frustrating for people wishing to make effective use of technology to enhance public safety.

In short, it's another typically British muddle. *

CENSORED

CONTAINING [redacted] CONTENT

The Terror of War, Nick Ut's iconic image showing a naked child fleeing a napalm attack during the Vietnam War, was inadvertently banned by Facebook in 2016

GETTY IMAGES

Are the social media companies finally getting their act together on filtering content?

By **Paul Dempsey**

IN JUNE, GOOGLE'S general counsel Kent Walker took to the influential op-ed pages of the *Financial Times* to announce that the company was instituting a more robust filtering strategy to remove or restrict extremist video content at its YouTube subsidiary.

By August, the company was forced to beat a partial retreat as several respected media outlets found that the new system had flagged and removed their content. Some reporters had had their accounts closed; others received warnings.

The deleted clips included footage of attacks during the Syrian civil war posted by independent investigative reporting group Bellingcat. After a very public spat with YouTube, Bellingcat's videos were restored, as were those from other sources.

A YouTube statement said: "With the >

< massive volume of videos on our site, sometimes we make the wrong call." Many of those affected wearily responded: "Twas ever thus."

This was hardly the first time that the media had found itself at odds with an internet giant over filtering and censorship. It is a decade since online censorship and abuse became public issues. The tensions behind it are proving slow and difficult to resolve.

Many were highlighted in a high-profile dispute during summer 2016. Facebook appeared to ban Nick Ut's iconic 'The Terror of War' image. It shows a naked young girl fleeing a napalm attack during the Vietnam War and had been posted by a journalist working for Norwegian newspaper *Aftenposten*.

Although some might challenge the ethical implications of sharing an image of a child in a distressing or vulnerable position, for Phan Thi Kim Phuc, the child in question, the photo has become what she calls "a path to peace".

In an interview with CNN in 2015, Phuc spoke of how she learned to accept the role that the photo has in demonstrating the horrors of war, and that her pain and terror have helped ensure that the past is not forgotten. "I realised that if I couldn't escape that picture, I wanted to go back to work with that picture for peace. And that is my choice," she said.

The image was only reinstated on Facebook after an aggressive front-page open letter to CEO Mark Zuckerberg from *Aftenposten* editor-in-chief Espen Egil Hansen. It followed private communication between the paper and the social media giant – and the deletion of a Facebook post that explained why the photograph had been used.

"Listen, Mark, this is serious," Hansen wrote. "First you create rules that don't distinguish between child pornography and famous war photographs. Then you practice these rules without allowing space for good judgement. Finally you even censor criticism against and a discussion about the decision – and you punish the person who dares to voice criticism."

Video filtering

Many of the same arguments applied during this year's confrontation between other outlets and YouTube. However, YouTube did bring a newer element into the delicate internet censorship equation: artificial intelligence.

In the *FT* article, Walker explained that AI was now one of four planks within the YouTube review process, alongside greater human analysis (Trusted Flaggers, including outside advice from members of respected NGOs), warning messages on extremist content and a mechanism to direct users away from content intended to radicalise them.

The problem is that, after another series of apparent errors, few active observers think YouTube is doing enough. Moreover, this kind of accidental deletion is only one side of the problem. What about truly offensive content?

The Fawcett Society, a leading UK campaigner for women's rights, collated a series of misogynistic posts made on Twitter and reported to its moderators around 14 August 2017 and then analysed what response there was a week later.

"By the morning of August 21, they were still up on the platform, despite the fact that they clearly violate Twitter's own community standards that do not allow direct or indirect threats or can be categorised as harassment or hateful content," the society said. "No response has been sent to the people who reported them, and no action had [sic] been taken against the users who posted them."

Again, Twitter is understood to be looking to AI to help it detect all forms of abuse, including such obnoxious techniques as dogpiling where a group of users coordinate an online attack on an individual woman. And again, the victims do not think they are doing enough.

At its root, the problem appears to have three fundamental components on both sides.

First, the internet companies do not consider themselves to be editorial organisations mandated to regulate comment in the way that a newspaper will, or consider its nuances.

Second, the internet companies do not staff their content review mechanisms sufficiently, giving human moderators sometimes scant time to evaluate posts.

Third, AI may be getting vaunted as an overarching panacea for the problem but it is not yet powerful enough. And maybe it never will be, though it could in time make things better.

Of the three, the editorial issue is about to become the most intriguing. Because, in US law at least, the internet companies were given a pass on this critical point more than two decades ago.

Section 230 (c) (1) of the US 1996 Communications Decency Act states: "No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider."

From that point of view, you or I would be legally liable – in the US – were we to go onto a social media platform and spread lies about someone else. But the platform itself would have effective immunity.

How widely this immunity is globally valid will, however, face its strongest test to date this month (October 2017). In Germany, a new law comes into force that places a responsibility on social media companies to remove "illegal" content within 24 hours of notification, and "criminal" content within seven days. Repeated failure to do so will expose the internet companies to fines of up to €50m (£44m).

The unresolved question here, though, is just how practical the law will prove in terms of enforcement. Still, it could be a wake-up call.

But there is then the question of resources. A major criticism levelled at social media giants is that they roll out content review systems based on commercial pressures (e.g. advertisers quitting if featured alongside hate speech) rather than the scale of the abuse.

This spring, Facebook found itself embroiled in further controversy over the streaming via its Live service of two murders – one showing a Thai man killing his 11-year-old daughter – and the rape of a 15-year-old girl.

Facebook CEO Mark Zuckerberg, unquestionably horrified, pledged to boost the human element within his company's content review process.

"Over the next year, we'll be adding 3,000 people to our community operations team around the world – on top of the 4,500 we have today – to review the millions of reports we get every week, and improve the process for doing it quickly," he said. "If we're going to build a safe community, we need to respond quickly."

But some were quick to ask whether even this would be enough. Last November, US National Public Radio undertook an investigation into the review volume facing Facebook. It found that the work was done largely by subcontractors, often in the Philippines or Poland.

"They are told to go fast – very fast; ... they're evaluated on speed; and ...on average, a worker makes a decision about a piece of flagged content once every 10 seconds," the report said.

The opportunities to properly evaluate, say, the Napalm Girl image as opposed to a piece of child pornography in such an environment are obviously limited. And even if Facebook were to double its review team, by this measure there would be just 20 seconds for each post.

Machine learning

So can AI come to the rescue, as the social media companies appear to hope? Not yet, but perhaps the good news is that the industry is climbing the learning curve.

Craig Fernsides, operations technical authority at web filtering and monitoring specialist Smoothwall, explains that while AI has – as YouTube's experience shows – still significant limitations, machine learning (ML) is already helping. Albeit for now with human help.

"ML allows someone to train a program using a known dataset. So you can give it a list of known pornographic sites and get it to learn what those look like. This means that when you show a new or unknown site to the pornographic ML program, it will be able to give you a true/false based on its reference material," he says.

"When we start connecting the ML programs together using something like a neural network, rather than a simple try/catch mechanism, then it gets interesting. That's the point that you can show anything to the neural network and if it doesn't know what something is, it will start compiling new categories and finding commonalities across multiple sites, spotting patterns or

Facebook CEO Mark Zuckerberg has pledged to improve his company's content review process

> 'Over the next year, we'll be adding 3,000 people to our community operations team around the world – on top of the 4,500 we have today – to review the millions of reports we get every week, and improve the process for doing it quickly.'
> **Mark Zuckerberg**

clues that no human could ever pick up on."

However, as Fernsides also recognises, there is no such thing as a universal filtering system. "Classifying an image or a video with the aim of monitoring or filtering access to it becomes a matter of opinion, and causing offence is such a personal definition that no system will ever be perfect," he says. "I'm not going to pretend I know what the answer is, but I know that our industry must try to provide tools and services that can adapt to the nuances of our many cultures and have faith that the tools will not be used for censorship or restricting freedom of information."

And thereby remains a likely eternal challenge though change is nevertheless coming.

Law beyond the US borders may force social media giants to rethink their roles, although the idea that difficult content may become more subject to lawyerly than editorial standards may itself raise serious issues.

More people are being hired to address the issue and though the numbers could fall short, it is at least a start.

Then, in terms of technology, maybe we just need to separate hype from reality and again monitor for gradual improvement.

But the one last thing to remember is that for all the charges flying around, and even though there is an astonishing amount of offensive, threatening and dangerous content reaching leading social media sites, this is a universe that has expanded more quickly than we have ever seen before.

At this stage in its development, perhaps the best we can ask is whether things are moving in the right direction. Some significant issues remain, but it looks as if there is progress – though a bit more haste would be welcome. *

CENSORED

[redacted] WATCHING WITHOUT [redacted] MOTHER

One important function of censorship is to keep unsuitable products and services away from those who are too young to be subjected to them. But how can technology determine how old a user is? The answer had better present itself soon, as a legal obligation is on the horizon.

By **Tim Fryer**

The development of Wi-Fi-connected smartphones and tablets has made discreet parental observation all the trickier when it comes to preventing access to unsuitable content

HOW MANY children under the age of 18 reach that point in a website when they are faced by that digital fortress that is the 'Are you over 18?' tick box, and then turn back to deal with more juvenile matters?

Not many.

This, though, is about to change. Following parental concerns that adult-orientated material was too easily accessible by children – by accident as well as through youthful curiosity – the Digital Economy Act was drawn up and passed by Parliament in April 2017. Its consequences are likely to be felt in the first half of next year when it becomes enforceable, probably at the beginning of May. Age verification will not only become compulsory; it will also be required to be legitimate. The 'age' tick box should become a thing of the past.

There are current ways of combatting age deception that can be perfectly adequate. Online ordering of goods that are to be delivered should theoretically be the easiest to deal with. Tobacco, alcohol, solvents, fireworks and knives – all require the purchaser to be 18 or over. Britain's largest retailer, Tesco, claims that the face-to-face contact between delivery person and the recipient for its home delivery service equates to the experience of buying over the counter. Unless the recipient is the person who ordered the goods and can produce proof of identity and age, the goods cannot be released. Having its own drivers gives Tesco the confidence that this practice is strictly adhered to, according to the company.

Many companies make deliveries through third-party services – couriers – which arguably do not have the same compunction to follow another company's delivery policy. There are anecdotal stories of cartons of cigarettes, bottles of spirits, fireworks or sharp knives being left with young children, rather than being put back in the van for redelivery. This is not quite the same as online ordering, but it's not a big step for a >

GETTY IMAGES, ALAMY, DREAMSTIME

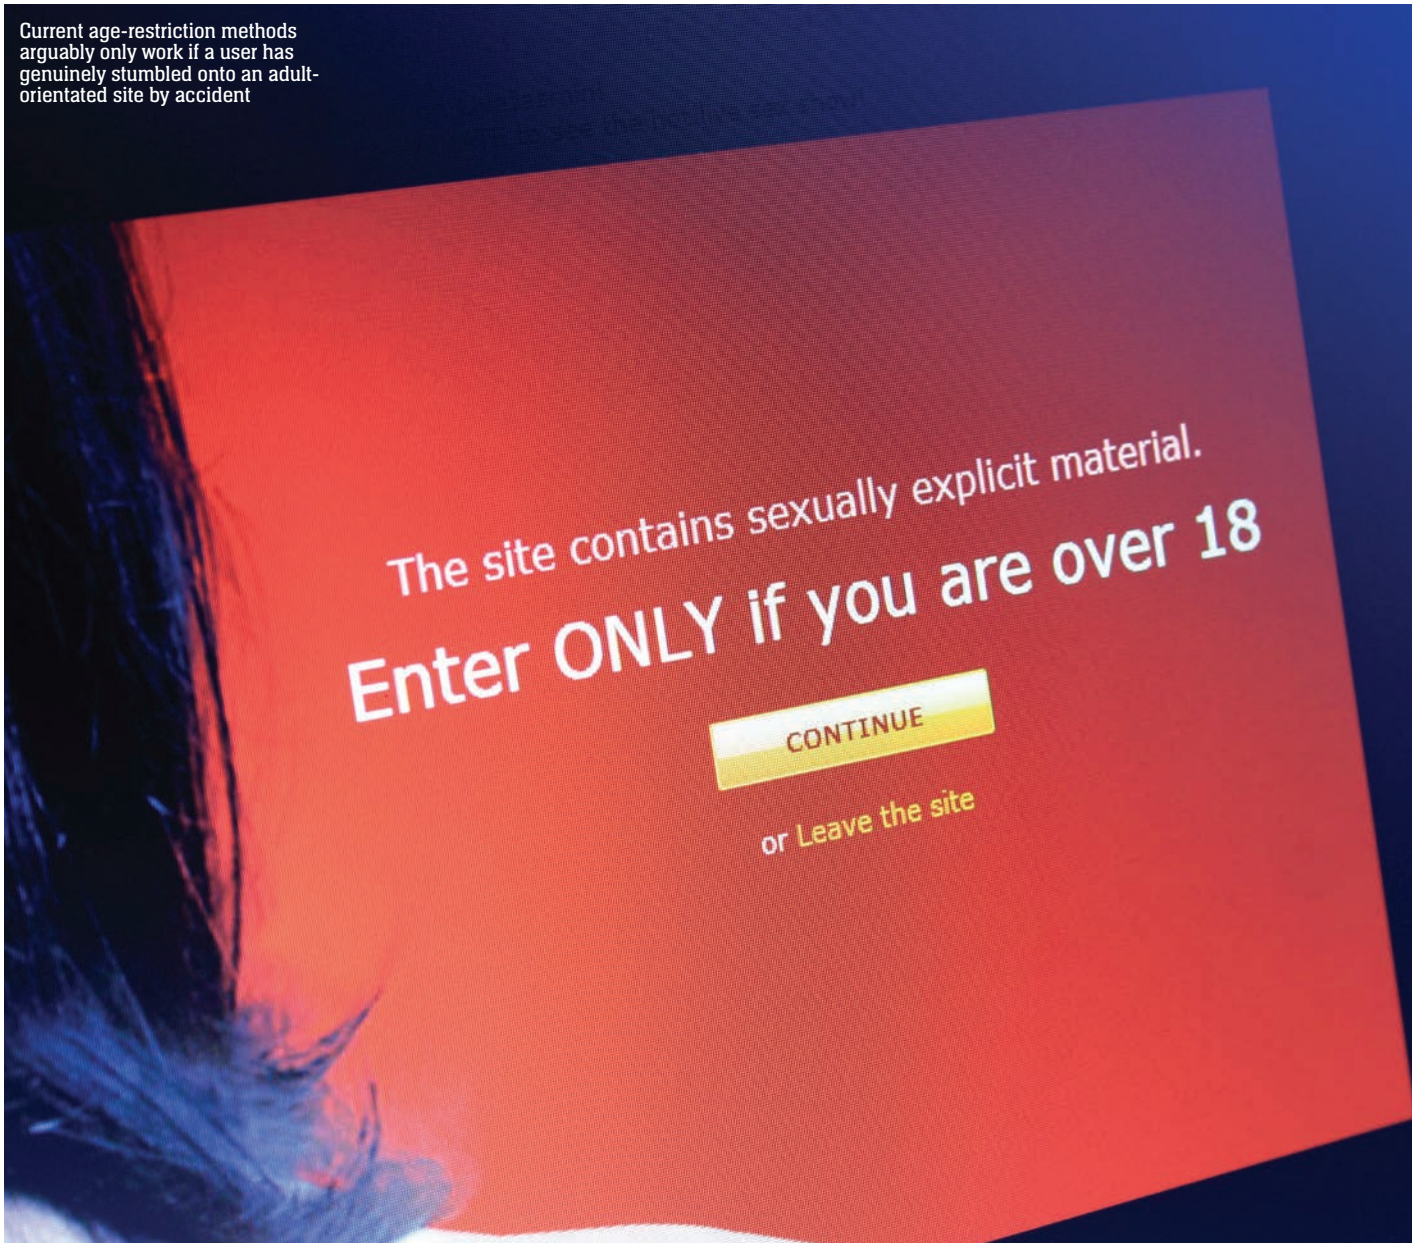

Current age-restriction methods arguably only work if a user has genuinely stumbled onto an adult-orientated site by accident

< wily teenager to get the hang of what will and what will not be delivered and then progress to making the purchases as well.

Amazon uses the qualifier: "By placing an order for one of these items you are declaring that you are 18 years of age or over. These items must be used responsibly and appropriately." That is hardly going to dissuade an enthusiastic teenager.

Having the ability to pay online is not a preventative measure either, as banks like Lloyds TSB and HSBC can offer junior accounts, including a use of a debit card, from the age of 11 onwards.

Assuming deliveries are not going to be abandoned in a porch or other 'safe place', however, a product has the advantage that it will be delivered with face-to-face contact, going some way to take away the underage concerns if done diligently. The opposite is true when it comes to services, which are more problematic.

Preventing children witnessing unsuitable content is at the heart of the issue. John Marsden is head of fraud and identity at Equifax, one of three UK credit reference agencies. He says: "We've had age checking over the counter for some time now, but there are issues with particularly sensitive bits of information that are not locked down to age at the moment, and the biggest one is pornography. The world of digital has turned it on its head."

It is difficult to determine the scale of the problem, as real numbers will only be in the domain of adult providers, who are understandably reluctant to offer the information. Alistair Graham, founder of age verification company Agechecked, says: "I would say when we are speaking to both adult providers and legislators about the size of the problem, around 25 per cent [of people watching adult entertainment being under age] is seen as a reasonable figure. So, it is in the millions."

But are companies providing pornography or other content deemed unsuitable for under-18s doing anything wrong? Netflix claims to have adequate parental controls built up to allow responsible parents to protect their offspring from anything inappropriate. This works as long as the adult remembers to set it up properly, is at least as adept technologically as the child, and in fact is an adult in the first place – it is only the universally ignored T&Cs that stipulate the customer must be over 18.

Graham says: "This is the crux of the problem – is it the parents' responsibility, the child or the provider? Realistically, if your child was given a bottle of whisky by an adult when you weren't aware, the finger instinctively points at the other adult. It is the same here – this is about policing the provider."

Digital Economy Act

The scope of the Digital Economy Act goes far beyond age verification – it is a wide-ranging act that covers many aspects of how digital information is distributed and used and how the businesses that use digital platforms do so fairly and responsibly.

One stated intention is that it is to "provide for restricting access to online pornography", and that is where age verification comes in. A policy group called the Digital Policy Alliance (which includes both Equifax and Agechecked) was set up to consult on how best to implement this, but Marsden observes that one particular step needs to be made: "To have a regulation you need a regulator who was going to be responsible for that online content. This type of business will not apply it unless it knows it is going to get policed. In policing the Bill we're talking about putting firewalls round the country to stop content coming in, and stopping providers in the UK with the original provision. Somebody will need to accept that regulatory body supervisory role for the deadline [next spring] to be applied."

That regulator, which will be appointed by the government's Department for Digital, Culture, Media and Sport (DCMS), is widely expected to be the British Board of Film Classification (BBFC). This would appear a natural extension of the BBFC's scope from television, film and physical media further into the online world.

Morally, claims Graham, "this is not taking it any further than we as customers have been used to for hundreds of years". It will signify an emphasis on everyone involved being responsible. "What is interesting in the legislation in the UK is that there are other service providers that also come under an umbrella of what is known as auxiliary service providers. If they are helping that distributor as part of their service to distribute adult content they are also obligated to ensure that merchant has age verification in place."

Payment providers would come under this umbrella, as could hosting services, VPNs and ISPs. Theoretically, no one will be able to view age-classified content without having their age verified, but it will take time for the industry to evolve to accommodate the new legislation. Payment providers could be an obvious frontline for age verification but then not all content is distributed in a way which immediately involves money.

"The majority of content online has no fee when the content is delivered," Graham says. "Payment on its own is not the sole solution. Potentially, hosting providers – people that are providing the hosting services for these organisations – would be responsible for making sure that whoever they are hosting has these measures in place."

He estimates that there are millions of relevant sites in the UK and the number of people that are likely to be verified is 20 to 30 million – all in a short space of time in 2018.

Anonymity (termed in the sector as 'pseudo anonymous') is the key to this verification, particularly for adult content sites. There is a sensitivity to such sites which means many of those, even if over 18 and legally allowed to view the content, do not want to leave any data trails that would associate them with it. The goal for much of that industry as well as the consumers of it, is to have a system where the user verifies once and thereafter just has to log in.

Moreover, the initial verification is not done with the adult content provider. "I think people really need to feel that they are giving their information to a site that doesn't do anything but verify them," says Graham. "For us, anonymity is absolutely critical to everything we do. In fact, once we verify a customer, we scrub any personal information from that verification so that person can still use their account to be able to login anonymously and we don't actually know who they are, which is fine because that is the service we provide. We don't do anything with any data other than age verification."

There are a number of ways of verifying age, cross-checking with the data held by the credit reference agencies (Equifax, Experian and Core Credit). One is the credit card, which cannot generally be issued in the UK until the age of 18. There is the added advantage here that 3D secure (such as 'MasterCard SecureCode' and 'Verified by Visa') can be used to determine that the person using the card is who they say they are and not someone using, for example, their parent's credit card.

Verification methods

And there are other methods. Any owner of a mobile phone in the UK that is viewing adult content must go through age verification to unlock their phones, so those numbers will have records against them as to whether they have been unlocked or not. Then there are also standard methods such as electoral roll look-ups, passports and driving licences.

The integrity of the test may vary according to requirements, as Graham explains: "Depending on which particular regulator you are talking to, certain checks will be robust enough for them and certain checks will be too light-touch. The goal is to move from a status where there is no proper effective age verification happening online to an area where it is ubiquitous; it works across the board. Some regulators may want more confidence that a person is a particular age and is definitely that person, and they might put more standards in place. But we are a long way from there.

"Currently the '18, yes or no' is the status quo, which is rubbish. This time next year, when everybody has got used to doing the light-touch verification that is proposed for the adult industry traffic, then we will be in a lot stronger place than we are today."

Unquestionably there is a huge issue, particularly with relation to the adult film sector, and it will be a challenge for the regulator to come up with a system that works for both the industry and its users. It will be the regulator who determines how the Digital Economy Act is enforced in its own sector.

"The current idea with the online adult content is there is a massive problem and fixing the vast majority of it with a light-touch age verification moves us a long, long way in the right direction," says Graham. As data increases and methods of using improve, it could be that the regulators start being more demanding, but as Graham adds: "The main thing, really, is solving the majority of the problem and not trying to make the perfect solution and becoming the enemy of the good."

There are 56 types of product that are age restricted in the UK, spanning 16 sectors. It is intended that age-verification technology will be demonstrably easy to use for both the service provider and the user and will become a familiar online feature next year. The most disturbing use of age deception is when an adult pretends to be a child. The technology is not there yet to provide protection for this, but development for underage misuse will hopefully lead to the technology for the reverse scenario as well.

EU LEGISLATION

IMMINENT CHANGES IN DATA LAW 'WILL HAVE A MASSIVE IMPACT'

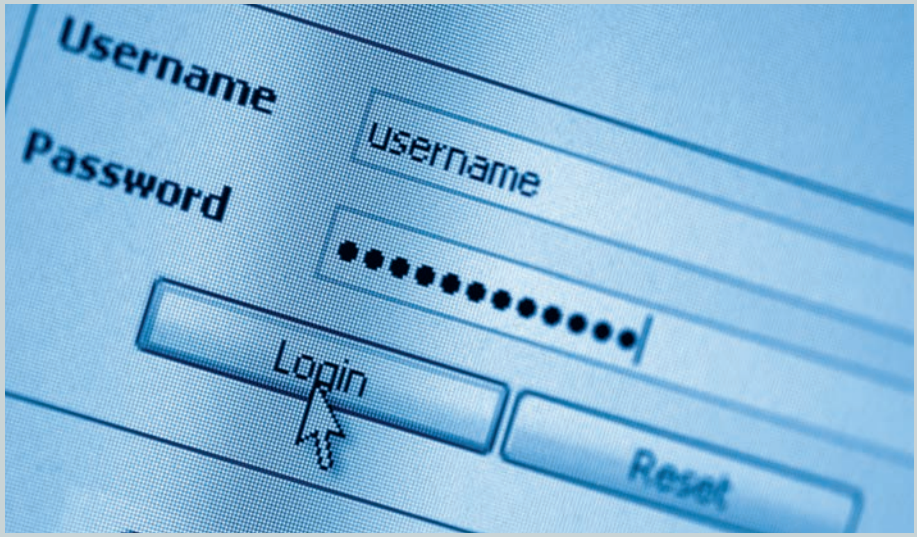

The General Data Protection Regulation needs to be complied with by 25 May 2018. It is EU legislation but the UK government has confirmed that Brexit will not affect British adoption. GDPR will be regulated by the Information Commissioner's Office (ICO).

The government is introducing measures related to this and wider data protection reforms in a Data Protection Bill, which needs to be in place before the May deadline for GDPR next year. The consequences are potentially huge. Data is only allowed to be used for the sole purpose that it was collected for and if permissions do not exist for historical data then they will need to be re-gathered. There will also be the right for an individual to examine any information that is held on them.

The upshot should be that details are more secure and unwanted use of it, for example sales calls and emails, will cease.

But the short-term impact, according to research from recruitment company Robert Half UK, is that 66 per cent of CIOs will hire additional, permanent employees to cope with the introduction of GDPR.

John Marsden from Equifax says: "We've probably put another 25 or 30 people on our staff just concentrating on how to audit data, explaining where it's from, deleting things if we think it's not going to comply and making sure that the terms and conditions that it was shared with us in the first place were correct and robust.

"So yes, it will have a massive impact on us."

CENSORED

IS RUSSIA [redacted] RIGGING GPS [redacted]?

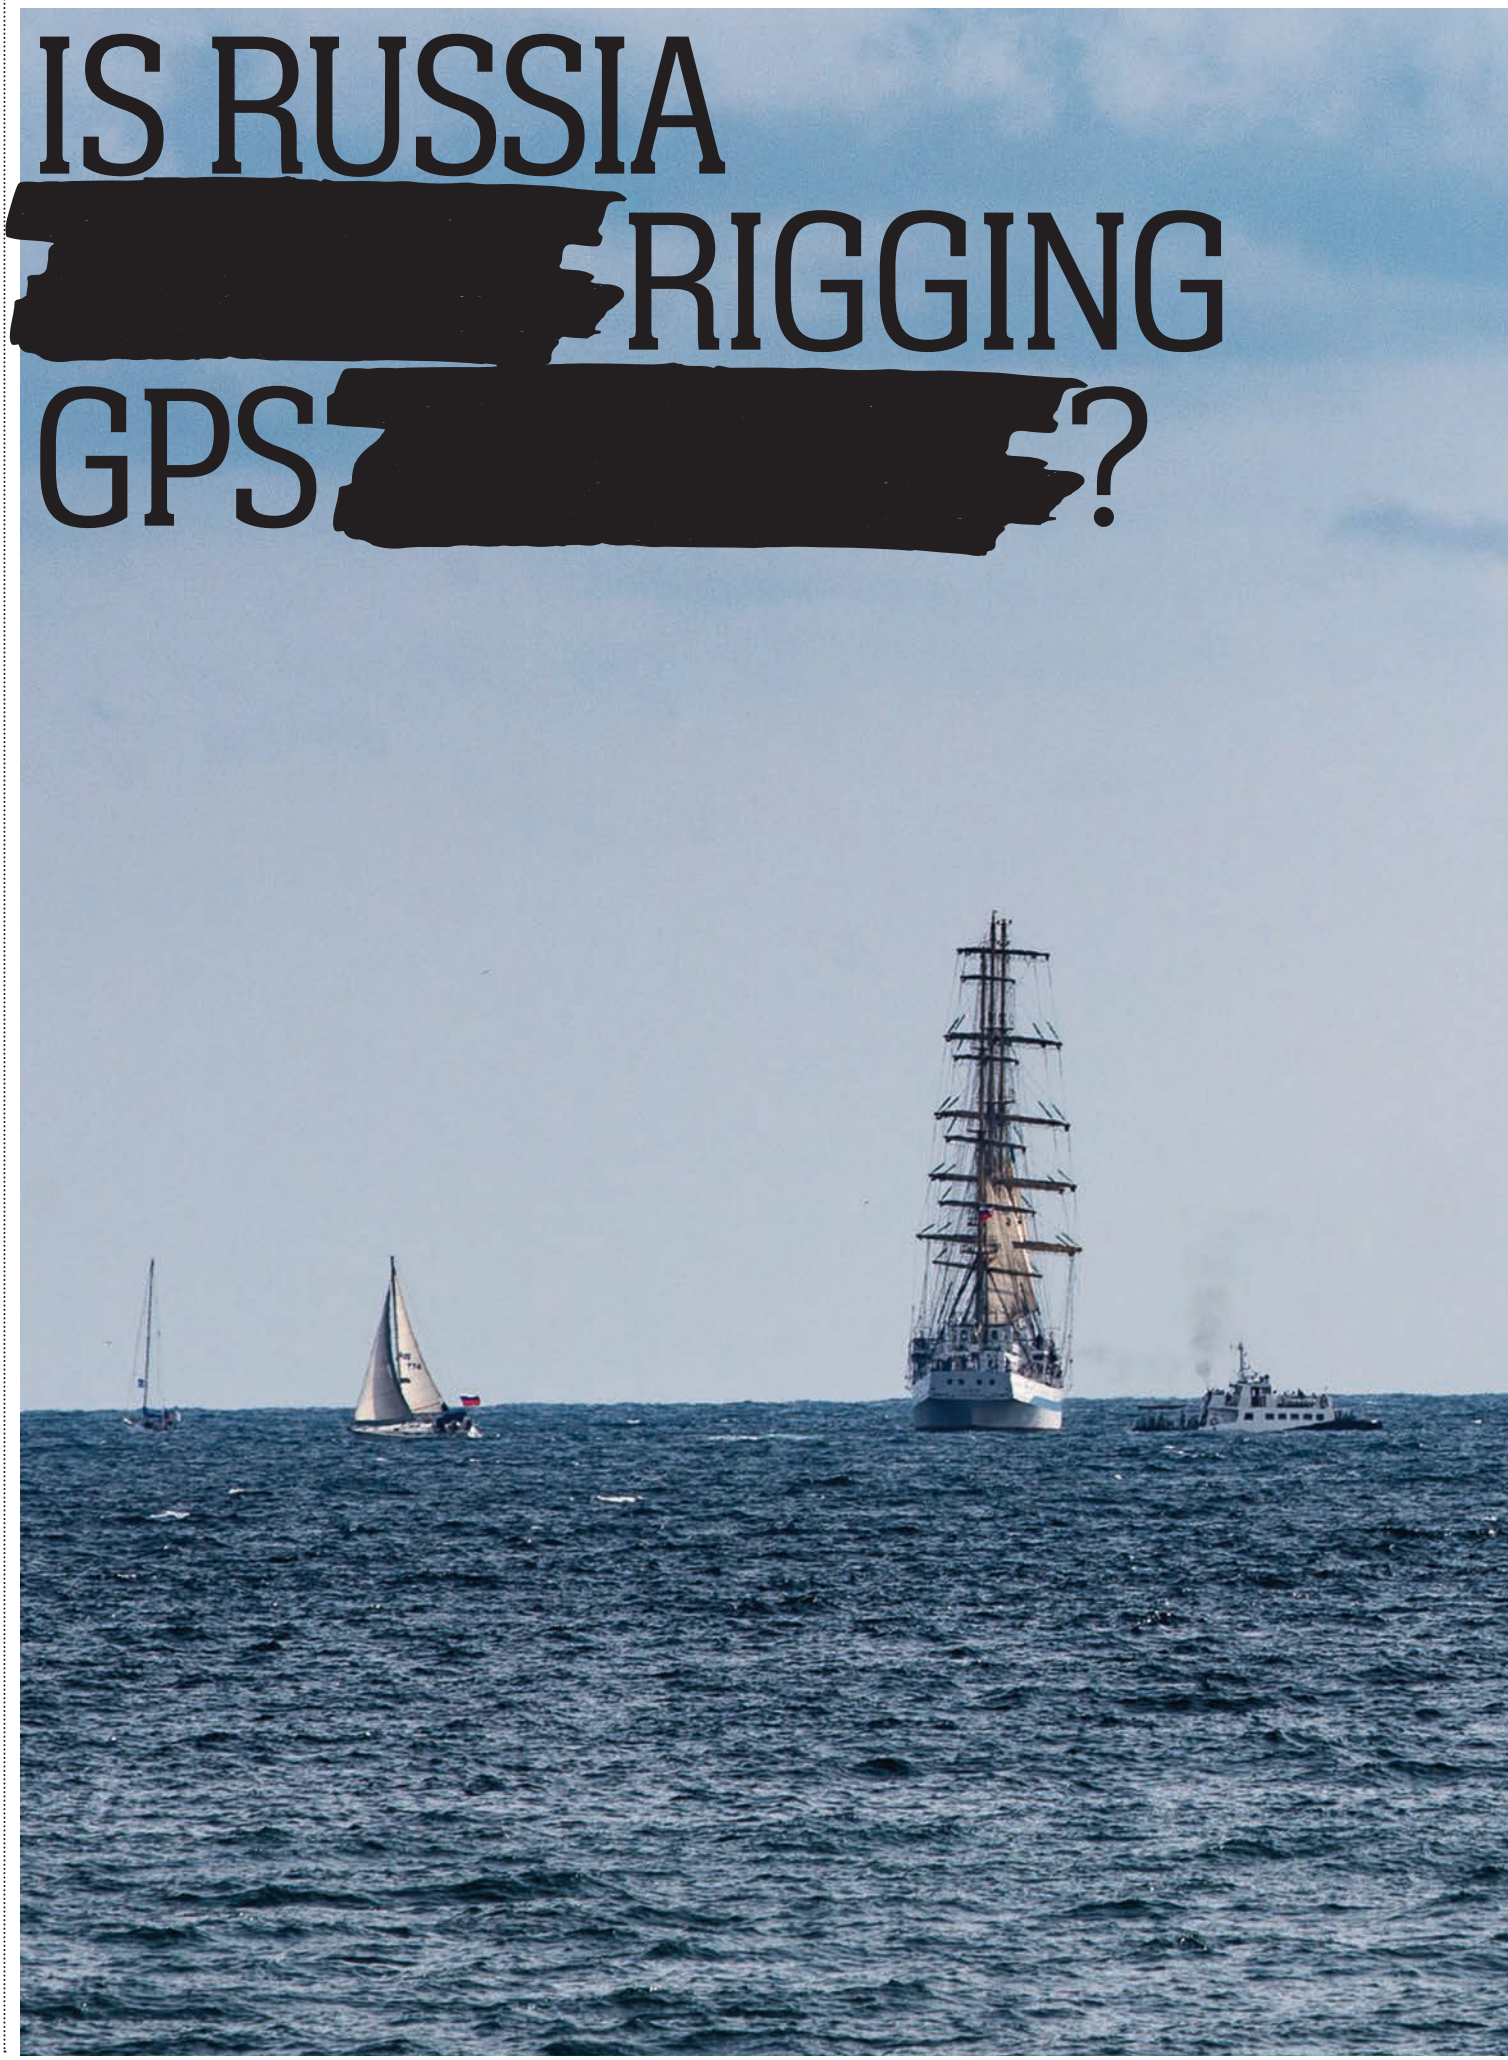

GETTY IMAGES, IMAGE SOURCE, REX FEATURES

Russia's Khersones training ship during the Novorossiysk stage of the 2016 SCF Black Sea Tall Ships Regatta

More than 20 years after the first Global Navigation Satellite System became fully operational, incidents in Russia seem to confirm that the most used network, the USA's GPS system, can be spoofed. Given the heavy reliance on GPS by transport, energy networks and even banks, many experts now believe we need systems to combat this interference.

By **Hilary Clarke**

SOMETHING CURIOUS happened at last year's Black Sea Tall Ships Regatta. As the international racing fleet sailed past the coast of Russia near Sochi, some of the ships began to notice that their GPS satellite navigation systems were starting to jump in strange ways. "They were giving spurious positions," explained race director Paul Bishop, "so positions were moving 50 miles in relatively quick succession, making the [data] coming through completely inaccurate."

Some of the ships' GPS units placed their location temporarily near Sochi airport, which lies a few kilometres inland.

"It was certainly very curious, but we haven't had a problem since so hopefully it was a one-off," said Bishop.

It wasn't. In June this year, the US Maritime Administration issued a safety alert showing that more than 20 ships had witnessed similar phenomena with their satellite navigation systems while sailing in the same waters.

A vessel reported to the US Coast Guard that its GPS equipment had intermittently been unable to obtain signals since nearing the coast of Novorossiysk, Russia. The ship's master then discovered his GPS put him in the wrong place – 25 miles inland, at Gelendzhik airport.

The Coast Guard tested GPS signals to make sure nothing untoward was happening; the ship's master was instructed to check the software updates on his equipment, which he did and confirmed it was working properly. However, he then startled the Coast Guard by telling him that more than 20 other ships in the area had the same problem, with their GPS sending them all to the airport.

A photo of the ship's GPS information screen convinced navigation experts to conclude this was a fairly clear case of 'spoofing'. "The nature of the incident is reported as GPS interference. Exercise caution when transiting this area," the US Maritime Administration said on its website.

The Tall Ships incident was not officially logged, but the one in June 2017 was evidence that Russia has the capability to cause GPS systems to give false readings.

It's not surprising, then, that there has been speculation as to possible GPS signal interference when the American naval ship USS John S McCain crashed with an oil tanker in the South China Seas in August. It was the second accident involving a large US naval ship in as many months in the >

TECHNOLOGY

THE JAMMING SPIRES OF TALLINN

The spire of St Olaf's Church in Tallinn, Estonia, used to conceal a powerful radio transmitter to block foreign TV stations

By Vitali Vitaliev

"Glushit' uzhe pozdno" – "Too late to jam" – is the ongoing call-sign of Radio Liberty's Russian Service, 'Radio Svoboda'. It is a Prague-based radio station financed by the US government, whose impartial and high-quality radio journalism played a crucial role in the collapse of the USSR.

We, the station's loyal Soviet listeners, knew only too well that to be able to tune to its heavily jammed – and therefore obviously trustworthy – broadcasts, we had to be fiddling with our short-wave radios at 4am, when the multiple Soviet jamming devices were somewhat less effective. Little did I know then that many years later I myself would be part of Radio Svoboda's broadcasts as the radio station's first Australian correspondent.

Much later, I discovered the technological reasons for that short-lasting early-morning relief in jamming. In trying to silence Radio Liberty, the BBC Russian Service, Deutsche Welle, Voice of America and other hostile 'anti-Soviet' broadcasts, the USSR used three main types of jamming: electronically generated noise signals, interference (whereby Soviet radio programmes were transmitted on the same frequencies) and the so-called speech-resembling signal. The last two methods would tend to 'calm down' a bit during the night, and millions of sleepless Soviet listeners were able to discern the crackling voices of their trusted radio journalists.

Our logic was simple: if the Kremlin was trying to silence all those distant foreign voices, that meant they were telling the truth.

It was only in late 1988, with the advent of glasnost, the Soviet Union stopped jamming Radio Liberty's and other foreign radio stations' Russian-language broadcasts.

Yet, jamming of western TV programmes, accessible in some Soviet border areas, continued until late 1991, right up to the USSR's final collapse. Tallinn, the capital of Estonia – formerly one of the 15 constituent republics of the Soviet Union – with its close proximity to Finland and its capital Helsinki, continued to remain the Soviets' main jamming hub.

No visitor to Tallinn who would bother to climb up the Toompea Hill for an impressive view of the Old Town and the bay could be spared the sight of two huge radio masts, dominating the cityscape. Everybody knew their main purpose was to send out waves of fuzz on the same frequencies as Finnish and Swedish TV broadcasts, to which many Tallinn residents had clandestinely tuned with the help of crude homemade receivers. The strength of jamming was such that it affected the quality of the official Soviet broadcasts from Moscow. I could testify to that myself after a short stay with distant relatives, who resided in Tallinn's main street – Parnu Maantee – in 1966. Due to constant thick fuzz on their TV screen, the normally clear territory of the USSR picture of the official daily news programme from Moscow appeared blurred and jumpy, which somehow made it much more trustworthy in our eyes.

My relatives said they were used to the constant interference, caused by two giant jamming towers less than a mile away from their block of flats. Later I learned that so strong was the jamming, it often affected the reception of TV programmes in Finland itself.

Used as they were to the constant presence of the two eyesores in the cityscape, very few residents of Tallinn knew that, in actual fact, there were not two towers, but three, with the third and the most powerful one hidden inside the spire of Tallinn's tallest and oldest church, St Olaf's (or *Oleviste Kirik* in Estonian). When the church was first built in the 12th century, it was officially the world's tallest structure. The Soviets took it over in 1944, showing little concern for Estonian history and religious values.

The jamming towers are now long gone, and one has a much better view of Tallinn's beautiful Old Town from the vantage point of the Toompea Hill. The spire of St Olaf's Church, the city's tallest structure, still looms above the capital of free independent Estonia. It is good to know that nothing sinister is hidden inside it any longer.

< same region, the first being the USS Fitzgerald. Seventeen US sailors were killed in the two incidents. The US military has said it is looking into whether the McCain's computer systems were compromised.

The USS McCain would have used the military GPS, which is separate from the civilian one and much more protected, but what if the oil tanker's GPS system had been hacked?

"The collision was with a commercial vessel that would have been a great deal more vulnerable," says Professor David Last, former president of the Royal Institute of Navigation, and one of the world's leading experts on the subject.

It is unlikely that in the two Black Sea disruptions, Russia was deliberately interfering with GPS. Russian sailors took part in the Tall Ships race and Russian President Vladimir Putin attended the award ceremony.

The fact that in both cases the GPS location showed up as an airport seems to point to something else – Russia has developed its cyber-warfare capacity to cause GPS satellite navigation systems to lie, perhaps to protect its airports from enemy drone and missile attacks.

Indeed, in 2016 Muscovites and foreign journalists began to notice that when they drove around the Kremlin their GPS devices showed their position as Vnukovo airport, around 30 miles away. Experts believe the phenomenon is caused by a Russian secret-services signal intended to disorient devices near the capitol building that navigate via GPS, most likely to protect the Kremlin.

System weakness

For a long time, western experts thought it would be too complex and expensive for a hacker to build a transmitter capable of interfering with GPS so it could emit false positioning, navigation and timing signals, and that such technology was the domain of thriller writers.

Indeed, the 1997 James Bond film 'Tomorrow Never Dies', features a pathological media mogul who hires a cyber terrorist to trick a British frigate into straying into the South China Sea in order to provoke a war between the UK and China.

GPS depends on a network of satellites orbiting the Earth at an altitude of 20,000km. At least four GPS satellites are 'visible' at any one time, wherever you are.

Each satellite transmits information about its position and the current time at regular intervals. These signals, travelling at the speed of light, are intercepted by a GPS receiver – such as a ship's navigation system or a mobile phone – which uses the differential time signals to calculate how far away at least three satellites are, and from this to pinpoint the location using a process of trilateration.

The GPS signal jammer works by sending out its own signal on the same frequency as the GPS unit, a noisy signal that prevents it from receiving or transmitting any useful information, either in bursts of sound or a continuous wave.

In 'Tomorrow Never Dies', 007 faces an evil media mogul trying to provoke a global war by tricking a British frigate to stray into the South China Sea

One of the system's drawbacks is the weakness of the signals. "GPS satellites are powered by solar panels as you would expect," says Dana Goward, president of the US Resilient Navigation and Timing Foundation. "Most satellites store the energy and pulse out transmissions. GPS satellites have to transmit all the time, so the net result is the cosmic hum going on all around space is louder than GPS signals," he explains.

By the time these signals reach the Earth, they are even fainter. It is then relatively simple to drown them out by emitting a loud noise on the same frequency.

GPS jammers can be bought on the internet and range in size from that of a small cigarette lighter to a suitcase, with a range of a couple of metres to a couple of hundred. You can buy combination jammers that also block mobile phones, all for a few hundred dollars. Customers for these devices are drivers seeking to shut off their employer's GPS tracker from their vehicle, car thieves and other criminals.

"If you set up a review by a motorway you will get up to 200 hits a month of passing vehicles carrying jammers," says Last, who is frequently called to be an expert witness in such criminal cases. "It is the same in Europe and the US."

Although it is not illegal to own such a device in the UK, it is illegal to sell one or to use one. The reason is that they rarely just jam the GPS device required, but can block signals to other GPS users around, endangering the public.

Insidious and dangerous

The fact that GPS jammers can be bought openly on the internet shows that most countries are prepared to live with the problem.

A ship's master or an aeroplane pilot would know immediately if the GPS signal was being jammed as the system would simply stop working or show a position way too far off to be credible, and alternative navigation methods would be used. GPS spoofing, however, is more insidious and far more dangerous.

"One of the dangers of spoofing, for example, is not just going off course but – for ships – crashing into water-borne obstacles," says Vlad Gostomelskly, managing consultant with Spirent Communications, which makes counter-jammers. Being tricked off course for just a few metres could cause a ship to run aground or puncture its hull, or run into the arms of waiting pirates.

The first person to show how to spoof GPS was Professor Todd Humphreys, an associate professor at The University of Texas's Aerospace Engineering and Engineering Mechanics department.

In 2012, with a group of brilliant students, he successfully hacked the GPS of a flying drone, steering it around a kilometre off course using a GPS spoofing device they made themselves that cost about $1,000.

Humphreys later successfully sent a luxury yacht off course in the waters of southern Italy, but the expertise to make >

President Vladimir Putin's presence at the Nadezhda Black Sea Tall Ships regatta is taken as evidence that Russia did not GPS-spoof this particular event

< the technology was so great that no one really expected it to be replicated.

Three years later at the 2015 DefCon hackers' conference in Las Vegas, Lin Huang and Qing Yan, two young Chinese electronics engineers working for Chinese internet search company Qihoo 360, showed they could spoof GPS using a software-designed radio costing just a few hundred dollars and find the software on the internet, the links to which can still be accessed today.

"Now we had a way accessible to the upper reaches of the hacker community at the very least," says Last.

Spoofing works like this. In any Global Navigation Satellite System (GNSS) signal, there is a peak inside that corresponds to the authentic positioning and timing signals with tracking points.

If you send a false signal, make a false peak, and then align the two, the tracking points can't tell the difference and get hijacked by the stronger counterfeit signal and can be sent off course without the captains, pilots, control towers or ships' masters even realising it.

"If the spoofer signals are stronger than the real signals at receiver end, the receiver locks on to your spoofer and not the satellite," Last explains.

Goward points out that the Kremlin announced in 2016 it had equipped over 250,000 cell towers with GPS phone-jamming devices as a defence against attack by US missiles, which usually rely on satellite navigation.

"Moscow has been very open about its ability to frustrate US and other forces by disrupting GPS," he says.

For more than a year now, Russia has been deploying its 1RL257 Krasukha-4 ground-based electronic warfare system on the battlefields of Syria.

When US fighter planes bomb, a list of GPS coordinates of 'red flag' targets such as hospitals and schools are checked against those of the operations target, such as an ammunition dump. When pilots are given new targets mid-flight, reconnaissance aircraft transmit the GPS coordinates. It is these signals that are disrupted by the Krashuka-4.

The Organisation for Security and Co-operation in Europe, an international conflict-monitoring group, has reported on several occasions that its drones watching the conflict in eastern Ukraine have been "subject to military-grade GPS jamming", forcing monitors to scrap missions observing the war below.

Cyber war

The result of Russia's advances in cyber-warfare and technologies for hacking GPS has given new impetus to develop a 'bullet-proof vest' for GPS navigations, says Goward, eliminating its potential to be a single point of failure for critical infrastructure.

As well as the prospect of some malevolent force sending ships crashing in the English Channel, according to a recent report by London Economics, the economic impact on the UK through loss of GNSS for five days would be £5.2bn through direct and indirect channels.

One solution is the eLoran terrestrial navigation system, which operates regionally within an 800-mile radius.

Developed from the Loran navigation systems that go back to the Second World War, using terrestrial transmitters and operating on a low-frequency radio band, eLoran can provide alternative position and timing signals for navigation. eLoran systems can now provide accuracies of between eight and ten metres. Its ±100 nanosecond accuracy at the receiver antenna underscores its unique timing ability, making it equal to GPS or certain atomic clocks.

"eLoran is the opposite in technology but is compatible in use to GPS. You can switch from one to the other seamlessly," says Last.

While eLoran doesn't provide altitude data and only works regionally within a range of 800 miles, its powerful low-frequency signals are far less susceptible to jamming or spoofing. The signal from eLoran beacons is 1.3 million times stronger than GPS.

The US dropped its old Loran systems in 1994 when GPS became fully operational, turning off the last transmitter in 2010, but there are now moves in Congress to again develop an eLoran back-up system, which already had support from the previous two US administrations, but has faced obstacles in Congress.

Russia is aiming to establish its own version of eLoran, called eChayka, and South Korea – which has long suffered from GPS jamming – is also developing its own system.

In Europe, a British-led initiative to establish an eLoran system that had already been tried and tested was pulled in 2016 after failing to garner interest from other European countries, which shut down their transmitters. The UK still has one transmitter in Cumbria but eLoran needed those in France, Norway and Denmark as well. They were all shut down.

One of the problems has been the EU's emphasis on developing its own GNSS, Galileo. "If you jam frequency bands, you jam all the GNSS systems. It is good to have extra ones, but the idea that Galileo can take over when GPS is jammed or spoofed is from cloud cuckoo land," says Last.

George Shaw, principal development engineer with the General Lighthouse Authority, which was behind the e-Loran project, says there are also moves to establish common standards with Russia.

Russia is keen to develop the Arctic Sea route for shipping, so compatible systems would make sense for trade.

Pressure is also coming from makers of driverless cars. Next year, Norway plans to launch the world's first ship with no crew.

All financial transactions rely on GPS at the moment for precision timing required by regulators. Substantial damage could be done in a covert cyber-war to the enemy's economy, without a shot being fired. *

IET Events
The Institution of Engineering and Technology

PETRAS iotUK

14 CPD HOURS

28 – 29 March 2018 | IET London: Savoy Place

CONFERENCE

Call for papers deadline: 10 November

Living in the Internet of Things

CYBERSECURITY OF THE IOT - A PETRAS, IOTUK AND IET EVENT

The IET, PETRAS & IoTUK are delighted to welcome you to the inaugural Living in the Internet of Things event - two days incorporating a call for papers conference, exhibition, workshops, seminars and demonstrations.

This call for papers event will address the cybersecurity of the Internet of Things exploring critical issues in privacy, ethics, trust, reliability, acceptability, and security through both social science and technical disciplines.

Benefits of submitting:

- Indexing on IET *Inspec* – which is accessed by 84% of universities worldwide!
- Published in the IET Digital Library
- Published in the Conference Proceedings
- Submitted for indexing on IEEE Xplore and EI Compendex

See the full course programme and book your place at www.**theiet**.org/cyberiot

Supporters:

#IETCyberIoT

The Institution of Engineering and Technology is registered as a Charity in England and Wales (No. 211014) and Scotland (No. SC038698).

IET Publishing
The Institution of Engineering and Technology

*"A **must have** for all radar researchers"*

Novel Radar Techniques and Applications

2 Volume Set

Editors: Richard Klemm, Ulrich Nickel, Christoph Gierull, Pierfrancesco Lombardo, Hugh Griffiths and Wolfgang Koch

Novel Radar Techniques and Applications presents the state-of-the-art in advanced radar, with emphasis on ongoing novel research and development and contributions from an international team of leading radar experts.

All IET members are entitled to a 35% discount off the retail price.*

Pre-order your copy by visiting

www.**theiet**.org/books-nov2volset

Publish date: October 2017

Volume 1:
Retail price: £130/$195
Hardback, 952pp • ISBN: 978-1-61353-225-6
Product Code: SBRA512A

Volume 2:
Retail price: £130/$195
Hardback, 552pp • ISBN: 978-1-61353-226-3
Product Code: SBRA512B

2-Volume set:
Retail price: £250/$395
Hardback, 1504pp • ISBN: 978-1-61353-229-4
Product Code: SBRA512X

* Please note, the promotional discount code set out above cannot be used in conjunction with any other discounts or promotions offered by the IET from time to time including IET member discounts. Any discount/promotion codes used will be void and the member discount will take precedence.

The Institution of Engineering and Technology (IET) is working to engineer a better world. We inspire, inform and influence the global engineering community, supporting technology innovation to meet the needs of society. The Institution of Engineering and Technology is registered as a Charity in England and Wales (No. 211014) and Scotland (No. SC038698).

CARELESS TALK COSTS [redacted] PRIVACY

A new generation of voice-activated digital assistants gives tech companies and hackers another way to collect sensitive personal information.

By **Martin Courtney**

VOICE-ACTIVATED digital assistants are a fast and easy way to look up information, initiate web-based communications and keep on top of busy schedules. But users may need to start censoring what they say, or face the very real prospect of a digital spy leaking more information than they care to divulge.

Samsung's newly acquired Viv artificial intelligence platform is about to join the growing number of voice-activated digital assistants available to consumers – these already include Apple Siri, Google Assistant, Microsoft Cortana and Amazon Alexa on their smartphones, dedicated Wi-Fi-enabled smart speakers like Amazon Echo and Google Home, and also Microsoft's Xbox One entertainment console. Facebook added similar capabilities to its Messenger platform in 2015 while Google embedded an intelligent agent to its new messaging app chatbot, Allo, this year.

Scale of adoption

Our usage of those digital assistants is steadily increasing. Google estimates that one-fifth of the searches on Android phones in the US are by voice. Microsoft said in May this year that 141 million users a month were using Cortana while Apple claims it has reached two billion Siri interactions a week, with an estimated 41.4 million monthly active users in January.

The vast majority of interactions with digital assistants involve smartphones, laptops, PCs or entertainment consoles. But data collected by independent research company Verto Analytics from over 20,000 US consumers between May 2016 and May 2017 highlighted the growing popularity of voice-activated digital assistants like Amazon Alexa and Google Home.

Strategy Analytics estimates that the total number of digital home assistant devices shipped by Google and Amazon will reach about three million in 2017, numbers very much in the 'early adopter' rather than mass take-up ball park.

Whatever the host device, the expected expansion of the digital assistant application and service ecosystem – whereby third-party providers integrate their own apps – will do much to drive their broader adoption. Ocado recently became the first UK supermarket to launch an app for Amazon Alexa, allowing customers to order their shopping using voice commands.

What are the threats?

Digital assistants use ML software to process natural language requests for information, goods and services, which offers considerable time saving for internet users >

AMAZON

Much is said in the privacy of the home or office, but what are digital assistants capable of sharing? And can they be hacked to share those secrets?

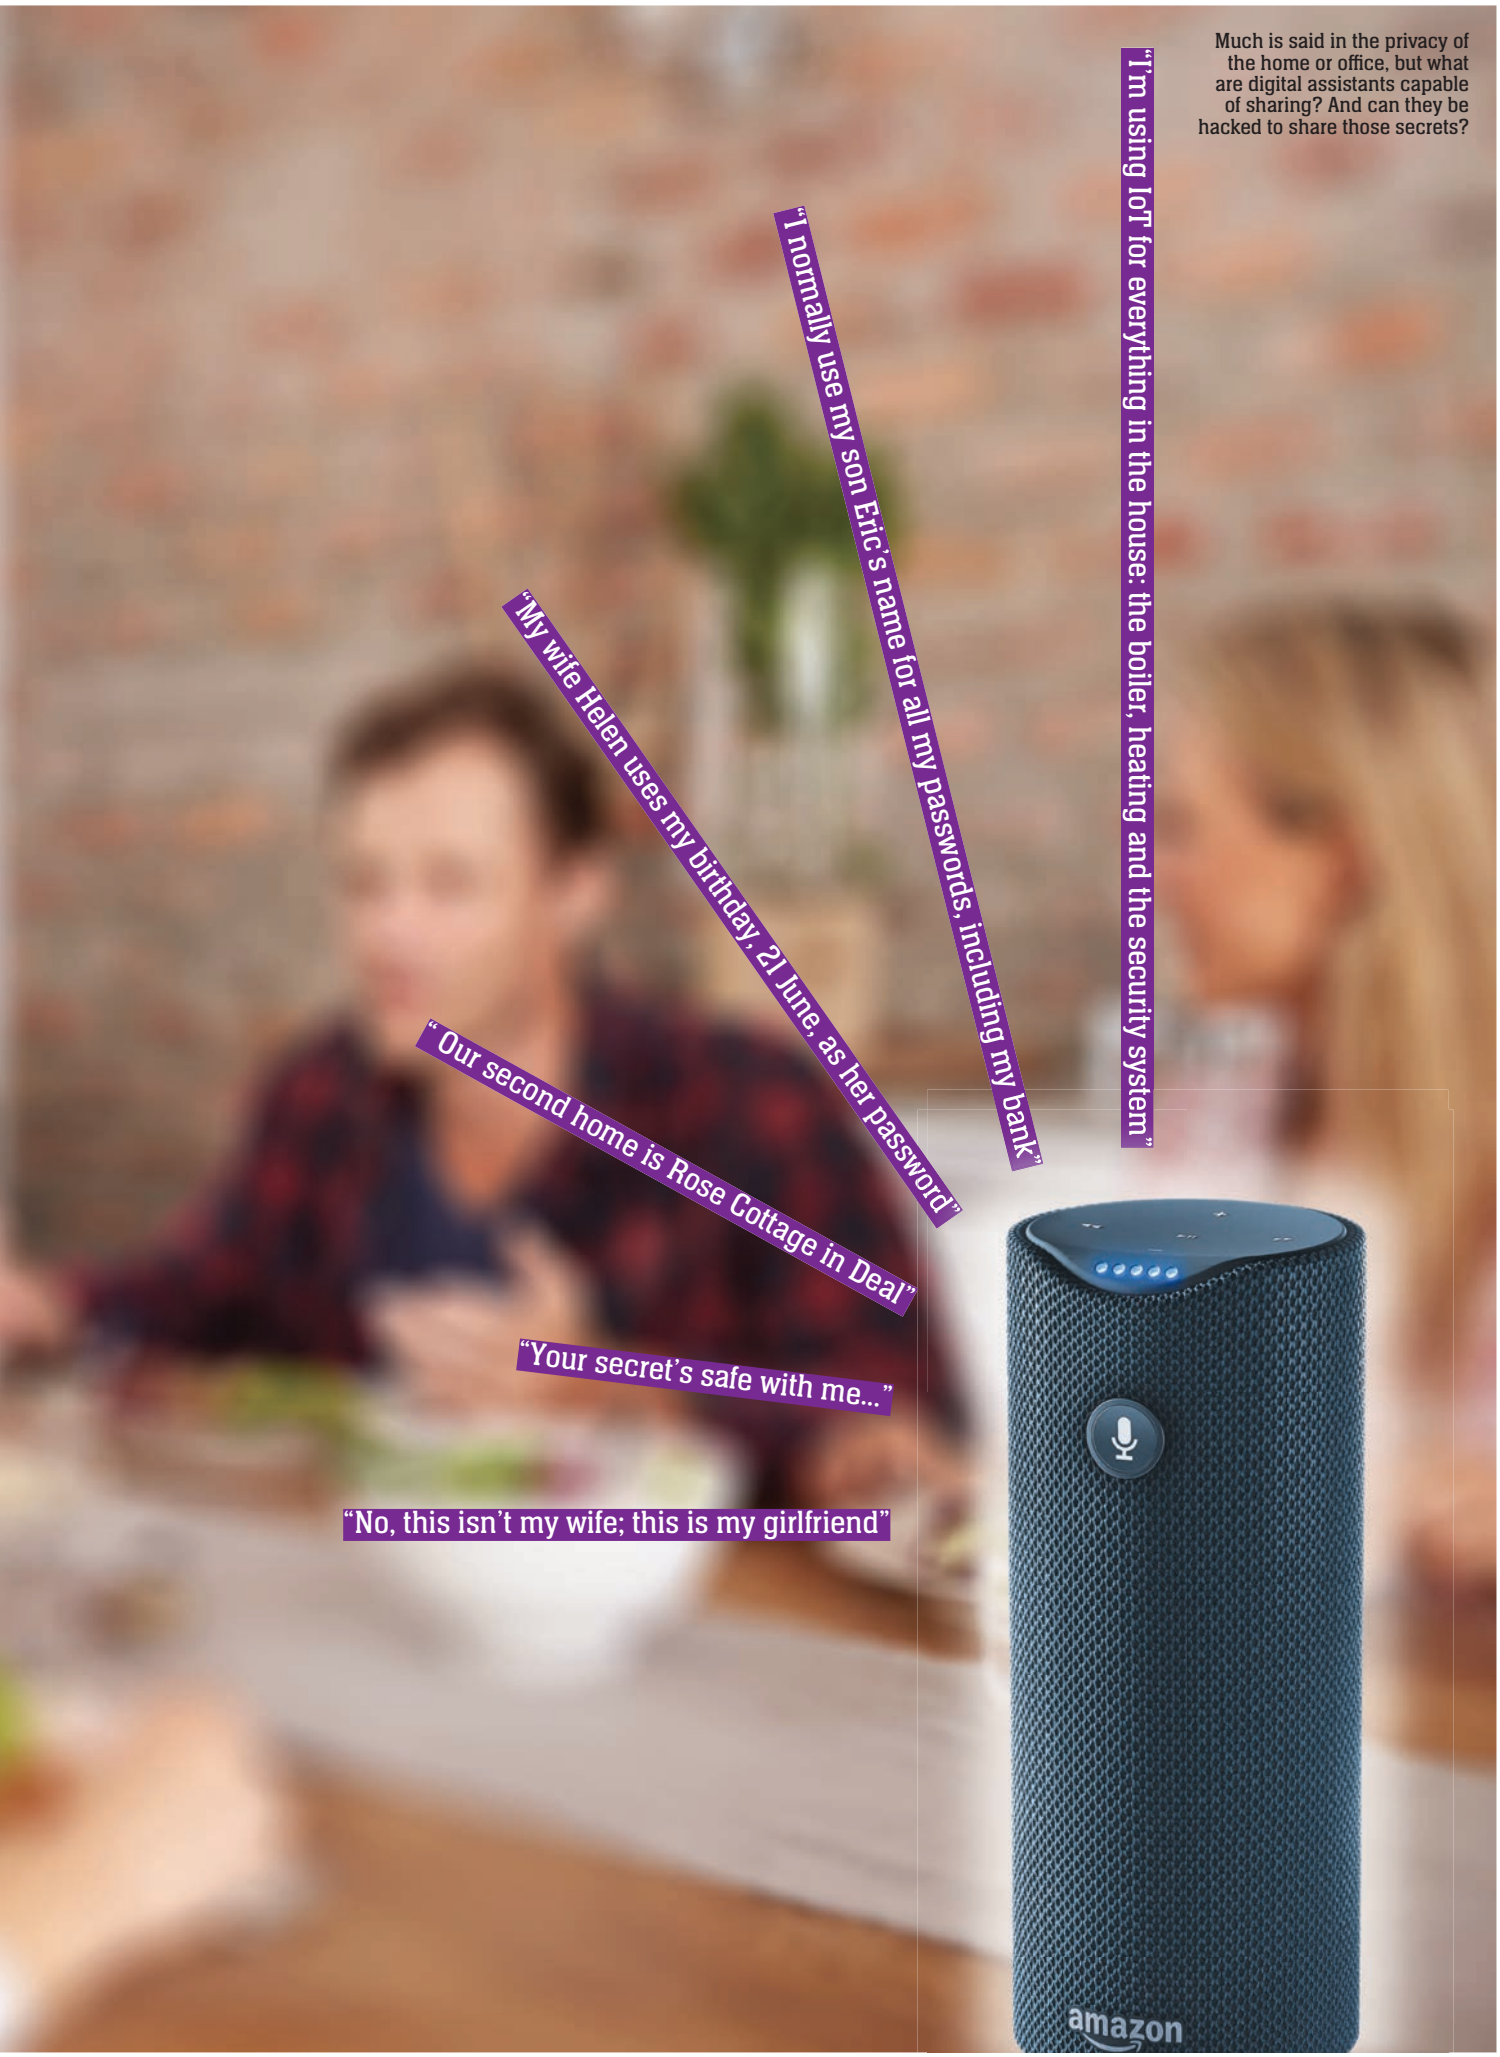

< who no longer need to type the same requests into browsers and search engines.

But that convenience may come at a price in the form of greater risks to data security and individual privacy, particularly when digital assistants are pre-configured to record all of our conversations and send transcripts of them to remote servers.

"The difference is that you are choosing what you want to put into that keyboard," said Simon Edwards, cyber security architect at Trend Micro. "The biggest problem with digital assistants like Alexa and Siri is that if you choose to let them do it they are listening all the time – do you really want Apple or Google to listen to all of your conversations?"

Verta's research also offers detailed insight into which apps US consumers access via their digital assistant, dominated by web browsers, maps, app stores and social media sites – all of which record personal information, preferences and location.

While most of the information shouted at the digital assistant might appear mundane – suggestions for restaurants, planned meetings, recommended temperature settings to optimise home energy consumption, for example – the data offers important clues as to the users whereabouts.

Access to messages, contacts and photos make it possible to snoop on other people's phones and gather information that can orchestrate other cyber attacks by providing criminals with locational information (so they know your house might be empty) and identity details which can be used (or sold) to form phishing attacks or fraud. Elsewhere IP addresses and unique device identifiers can be harnessed to launch distributed denial of service (DDoS) attacks that bring down networks and websites by flooding them with spurious data traffic.

The unwary could inadvertently disclose usernames, passwords, national security numbers and bank accounts. There is a small chance that information could be accessed by hackers, either during activation or transmission, but those odds increase significantly if it is then stored for long periods on cloud-hosted servers.

Another concern is the integration of digital assistants with IoT-connected home appliances – primarily lighting and heating controls, but also door locks and thermostats, which could be used to gain access or start a fire. Research company Gartner predicts that digital assistants will serve as the primary interface to IoT-enabled connected home services within 25 per cent of households in developed economies by 2019 – food for thought given their remote control capabilities.

Embedded protection not immune

Apple, Google, Microsoft *et al* point out that nothing is recorded until the digital assistant is activated using the 'hotword' (Hey Siri, Hey Cortana etc), and that if no vocal match is recognised, the feature will not activate. In some cases (Apple Siri, Microsoft Cortana) you can simply choose not to enable it in the first place.

Google Home – one of many wireless smart speakers on the market

The situation is very different for smart speakers and digital home assistants, where anybody can request information with no authentication. That means that if you have given Google Home access to your calendars, messages or other personal information, anyone can call them up. Even where the owner's voice is the authentication mechanism, anybody that can successfully mimic the voice gets full access.

> 'If people have concerns about the privacy issues of such services they should ensure they have thoroughly read and understood any privacy policy.'
> **ICO**

If a stranger is in your home or office you may be the victim of crime already. But the 'inside job' is not uncommon in cyber-security incidents, and theft of electronic data can prove a more valuable enterprise for a criminal than making off with the TV.

On any mobile device vulnerable to theft or accidental loss the situation is more perilous – one reason why Google has different policies for Google Home than for Assistant or Allo. Digital assistants on mobile devices and PCs are usually afforded some degree of protection by standard authentication processes, which involve account numbers, pin numbers, passwords and even biometric validation in the form of fingerprint readers and facial recognition.

Evidence suggests that it is easier to gain access to somebody else's data through their digital assistant than the internet companies would have us realise – just because the device is turned off, for instance, does not mean it is not recording. Cortana listens in and provides access to calendar, email, messages and other content even when the device is locked by default, although the feature can be disabled.

Security company Trend Micro found a passcode override for Apple Siri in 2015 when it discovered that certain questions would provide personal information even when the mobile device it was running on was password protected, including "what's my name", "text name/number message", "call name/number", "post Facebook status message", "first name", "what's my email address", and "show me date/timeframe schedule" etc.

More recently, researchers at China's Zhejiang University hacked digital assistants running on mobile devices and PCs using voice commands outside the range of human hearing. Dubbed DolphinAttack, the technique allowed them to translate normal voice commands into ultrasound which were then tested against over a dozen voice assistant systems including Siri, Alexa, Google Assistant and Cortana. The commands – which included the activation phrases and prompting a Macbook to open a malicious website containing malware – were universally obeyed, even though the humans could not hear the communication happening.

We do know that other smart devices with recording capabilities have been hacked in the past. Samsung, for example, was famously the victim of the Weeping Angel CIA hack that allowed the US government to compromise its F8000 smart TV and listen in to what its owners were saying.

While no demonstrable exploits for digital assistants have so far been discovered outside of the labs, that is not surprising at this stage. Hackers tend to maximise their chances of success by targeting operating systems and applications with the largest number of end users. Digital assistants are a long way off the cyber criminal's radar, but that is certain to change as their popularity spreads.

What data is being shared

We are all well accustomed to personal data being collected – it is nothing that we do not already type into browsers and search engines. But the risk of the digital assistant could be more sinister given that we may not always be aware our conversations are being recorded, processed, analysed and stored.

In a recent report, 'The New Privacy: It's All About Context', research firm Forrester notes that companies often collect far more data than they actually need or use, and were promoting a culture of 'collect if you can' among online and mobile applications irrespective of if they thought they could use it to good effect.

Apple has been careful to take a comparatively tough line when it comes to how much data Siri records or sends back to its servers. The information Siri transmits is anonymised using a random identifier rather than an Apple ID, email address or other personal data and deleted when the user turns Siri off. The company says that voice clips are only retained for the purposes of training or improving the accuracy of Siri's voice recognition engine. Nevertheless they are retained for a full two years, leading to concerns about what happens to them if Apple's hosting infrastructure is hacked.

Google Assistant, available in Android, Home and Allo versions, provides access to contacts, storage and calendar, name, search

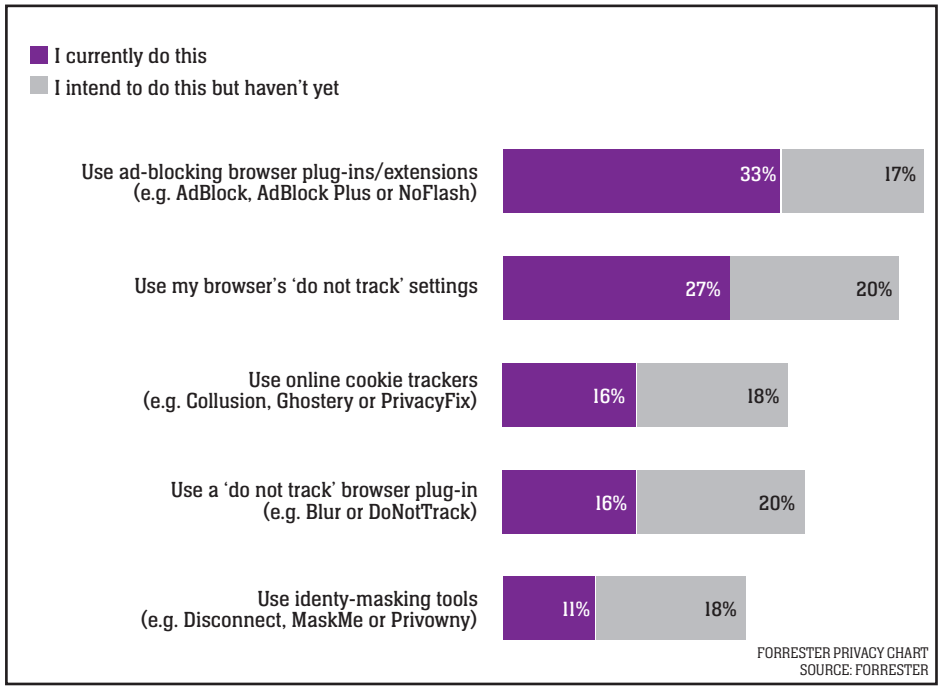

FORRESTER PRIVACY CHART
SOURCE: FORRESTER

history, voice and audio activity and other information on user accounts. Google is honest about using people's browsing history to deliver targeted advertising (which help keep its services free) and also admits to storing conversations on servers in its own data centres.

The company insists users can view and delete past interactions with the Google Assistant in My Activity, but the data is only permanently deleted from the user's Google Account – certain 'service-related' information concerning the use of Google products is kept, ostensibly to prevent spam and abuse and improve services. There is no time limit on how long the data is kept, only until the user chooses to delete it.

While Google does not sell personal information to anyone, it admits to sharing information with its affiliates and business partners about which the user requested information (e.g. a restaurant or airline).

Amazon's Internet Privacy Policy explains the type of information it collects through Alexa, its website and browser extension software, and what Amazon does with it. In some cases, that information may be personally identifiable, though Amazon insists it does not take active steps to determine the identity of any Alexa user.

That data can include a name, email address, country of origin, nickname, telephone number, website, company and title for example, as well as browsing history, information about your operating system, a unique identifier enabling Alexa to identify the user's device, and the date and time the information was logged. A history of the online advertisements displayed on the websites Alexa users visit is also retained, including text, source and URL, alongside the terms entered into search engines and their results.

Not only that, but Amazon can gather other information from third-party websites, such as social media sites, used to interact with Alexa, including much of the personal content posted there.

Regulatory clashes

Given the type and volume of personal data involved, it is no surprise that companies collecting it occasionally clash with national and regional privacy regulation.

Some experts have warned that recording the voices of children could contravene the Children's Online Privacy Protection Act (COPPA) in the US, for example, which was originally devised to protect young people from pervasive data collection. Video advertising from Amazon and Google frequently features images of children communicating with digital assistants, seemingly unaware or unconcerned that these interactions may contravene data protection laws.

COPPA precludes the storage of a child's personal information, including recordings of their voice, without the explicit consent of their parents. The forthcoming EU General Data Protection Regulation (GDPR), compliance with which becomes mandatory in May next year, provides strict guidelines on obtaining parental consent before storing and processing personal details (including audio recordings) of EU minors.

Any organisation offering voice-activated services to UK citizens needs to comply with the Data Protection Act 1998, says the UK information commissioner's office (ICO).

"This means that users need to be informed of how their voice recordings will be used, particularly where those uses might not be expected, such as disclosures to any third parties," said an ICO spokesperson. "Organisations that store voice recordings abroad will also need to ensure they have a proper lawful basis for doing so."

Recordings used as evidence

Consumers should also consider whether conversations recorded by digital assistants will ever be made available to police and other enforcement authorities on request to help with investigations or be used as evidence in criminal cases.

Google makes it clear that it shares information for legal reasons – where disclosure is necessary to meet applicable laws, regulations, legal processes or government requests, for example, but also to help detect, prevent or otherwise address fraud, security or technical issues, or to protect against harm to the rights, property, safety of Google, other Google users or the public as required or permitted by law.

Earlier this year, Amazon was forced to hand over data recorded by its Echo smart speaker to the Arkansas police for use as evidence in a murder trial. The company initially refused, arguing that such a move would constitute a violation of consumer rights and that the investigator's case did not merit sufficient cause. The dispute was ended prematurely when the defendant agreed that the audio files could be accessed.

Amazon still maintains that protection of free speech is enshrined in the First Amendment, and as such the police need to follow very specific legal procedures to gain access to any recording made by Alexa and prove that they have a compelling need for the information to be disclosed. Legislation will inevitably take time to catch up with new digital assistant technology, but in the meantime it is highly likely that similar cases will advance a definitive conclusion and that the feds will eventually get their way.

Privacy – an ongoing concern

In 2010, Facebook chief executive Mark Zuckerberg argued that consumers' social norms regarding privacy have "evolved over time", which justified ongoing changes in his company's approach to privacy. While privacy concerns are far from dead, it is true that people who have grown up sharing many details of their lives with social media sites appear more comfortable with IT companies' collection, storage and processing of their data than others.

Wading through the privacy policies for Siri, Cortana, Alexa etc is an onerous task, one which few people bother to undertake when setting up digital assistants, or devices with digital assistants enabled by default, for the first time.

But as the ICO advises, it is well worth the time to familiarise ourselves with the exact terms of those policies to set expectations and help us understand exactly what information is being collected and how it is being used.

"If people have concerns about the privacy issues of such services they should ensure they have thoroughly read and understood any privacy policy before purchasing any device that includes such capability, or before deciding to use any such service that may be available in a product they already have," it says.

Those that don't mind the idea of their home becoming the equivalent of a station waiting room or café, where they have no idea who is listening at the next table or what those eavesdroppers could or would do with any snippets of private information overheard, will no doubt carry on regardless.

Otherwise maybe it's time to either switch off the default 'always on' listening mode in the digital assistant or start to consciously censor what is said within its earshot.

"At the end of the day it is another incursion into people's privacy and the more we get used to having these things around, the more lax we will be," says Trend Micro's Edwards. •