Helping people understand how new technologies interact with the UK's data protection framework

Tech Horizons Report

December 2022

ico.
Information Commissioner's Office

Contents

Executive summary

From wellness apps to smart offices, data-driven technologies are creating huge changes to how we live and work.

As the UK's data protection regulator, the Information Commissioner's Office (ICO) seeks to foster trust in how organisations process personal information. We want to empower people to safely share their information and use the products and services that will drive our economy and our society.

In our ICO25 strategy, we committed to set out our views on emerging technologies to reduce burdens on businesses, support innovation and prevent harms. Our first annual Technology horizons report examines the implications of some of the most significant technological developments for privacy in the next two to five years.

Our analysis focuses on the following four technologies:

- **consumer healthtech**: wearable devices and software applications that help people assess their health and wellbeing;
- **next-generation Internet of Things (IoT)**: physical objects that connect and share information, with the ability to sense, respond to or interact with the external environment;
- **immersive technology**: augmented and virtual reality hardware that creates immersive software experiences for users; and
- **decentralised finance**: software that employs blockchain technology to support peer-to-peer financial transactions.

We will examine a fifth area, **neurotechnology**, in a separate foresight report that we will publish in spring 2023.

These technologies advertise significant opportunities to make our lives easier, safer, more comfortable, efficient and fun. They also present a range of risks that may harm people's privacy and their trust in these technologies, if they are not addressed as the technology develops.

We observed a common set of challenges:

1

A growing set of technologies are collecting personal information in ways that may not be transparent to people and they may not have meaningful control over. This is particularly true when information is captured about third parties other than the intended user, for example by augmented reality devices or smart home systems.

2

The complexity of some data ecosystems may make it difficult for people to understand how organisations are processing their information and hold them to account. Attention is needed to ensure that people can exercise their information rights in decentralised finance systems and complex IoT networks.

3

Some technologies are collecting more information than they may need for their primary purpose. For example, they may track people across consumer healthtech or virtual reality devices in ways that may not be transparent or necessary for their primary purpose.

4

Many technologies are collecting information about sensitive personal characteristics, that may require additional safeguards. Organisations need to understand when this information is classified as special category data (eg biometric or health data) and put appropriate measures in place.

Further challenges include the accuracy of inferences made by some devices and the security of information processed by others. Some challenges were specific to certain technologies, for example the difficulty in exercising the rights to rectification and erasure if information is held on a blockchain.

Some organisations are exploring new and innovative ways to engineer privacy into the design of these technologies. For example, some manufacturers are embedding redaction technology into extended reality devices to minimise unintended processing of information about bystanders.

Other organisations are failing to imagine how privacy could be engineered into their ideas. We will not allow businesses that are doing the right thing to be outcompeted by businesses that fail to comply with data protection law.

To prepare for our role in regulating these technologies, we will:

- work with the public about the benefits and risks of these emerging technologies and how we will approach them as a regulator;
- invite organisations to work with our Regulatory Sandbox to engineer data protection into these technologies;
- develop guidance for organisations where needed, starting with guidance on data protection and IoT; and
- proactively monitor market developments so we can take action on uses that cause concern.

Introduction

As the UK's independent data protection regulator, the ICO's work spans the whole of the economy. We regulate organisations from finance, commerce and communications, to employment, education and the public sector. Collecting and processing personal information in each of these areas is driving huge changes to how we live and work.

As new technologies emerge, we aim to help innovators develop safe and trustworthy new products and services. We also work to protect consumers from any harm arising from organisations processing their personal data. These are complementary aims. Clear regulatory frameworks allow organisations to innovate with certainty. Consumers are also more likely to adopt technologies that handle their information securely and appropriately.

To support this, our report examines the implications of some of the most significant technological developments for privacy in the next two to five years. We have conducted wide-ranging research, dozens of interviews with world-leading academic and industry experts, workshops and scenario creation. It indicates our emerging thinking and our ambition to help innovators consider privacy at an early stage of design.

The report focuses on the following four technologies that we believe have the most tangible, and often novel, implications for data protection:

- consumer healthtech;
- next-generation Internet of Things (IoT);
- immersive technology; and
- decentralised finance.

A fifth area, neurotechnology, will be the focus of a deep dive report will be published in spring 2023, similar to our recent deep dive reports on biometric technologies.

Technology selection

This report examines technologies that present novel and significant implications for privacy in the next two to five years. We have chosen these developments because they will soon affect how people consume and use technology, but organisations still have scope to integrate privacy considerations before deploying them on a wide scale.

We focus on applied technologies. These are innovations with discrete effects across the economy in a range of uses. It is these technologies that have more certain implications for privacy in the near future. We have not examined the technologies that they are built on: foundational (eg artificial intelligence) or connectivity technologies (eg fibre or 5G). We have also excluded 'use-case technologies'; single application technologies with a narrow use-case (eg smart glasses), that are not found at scale across the sectors that we regulate.

Figure 1:
Foundational technology

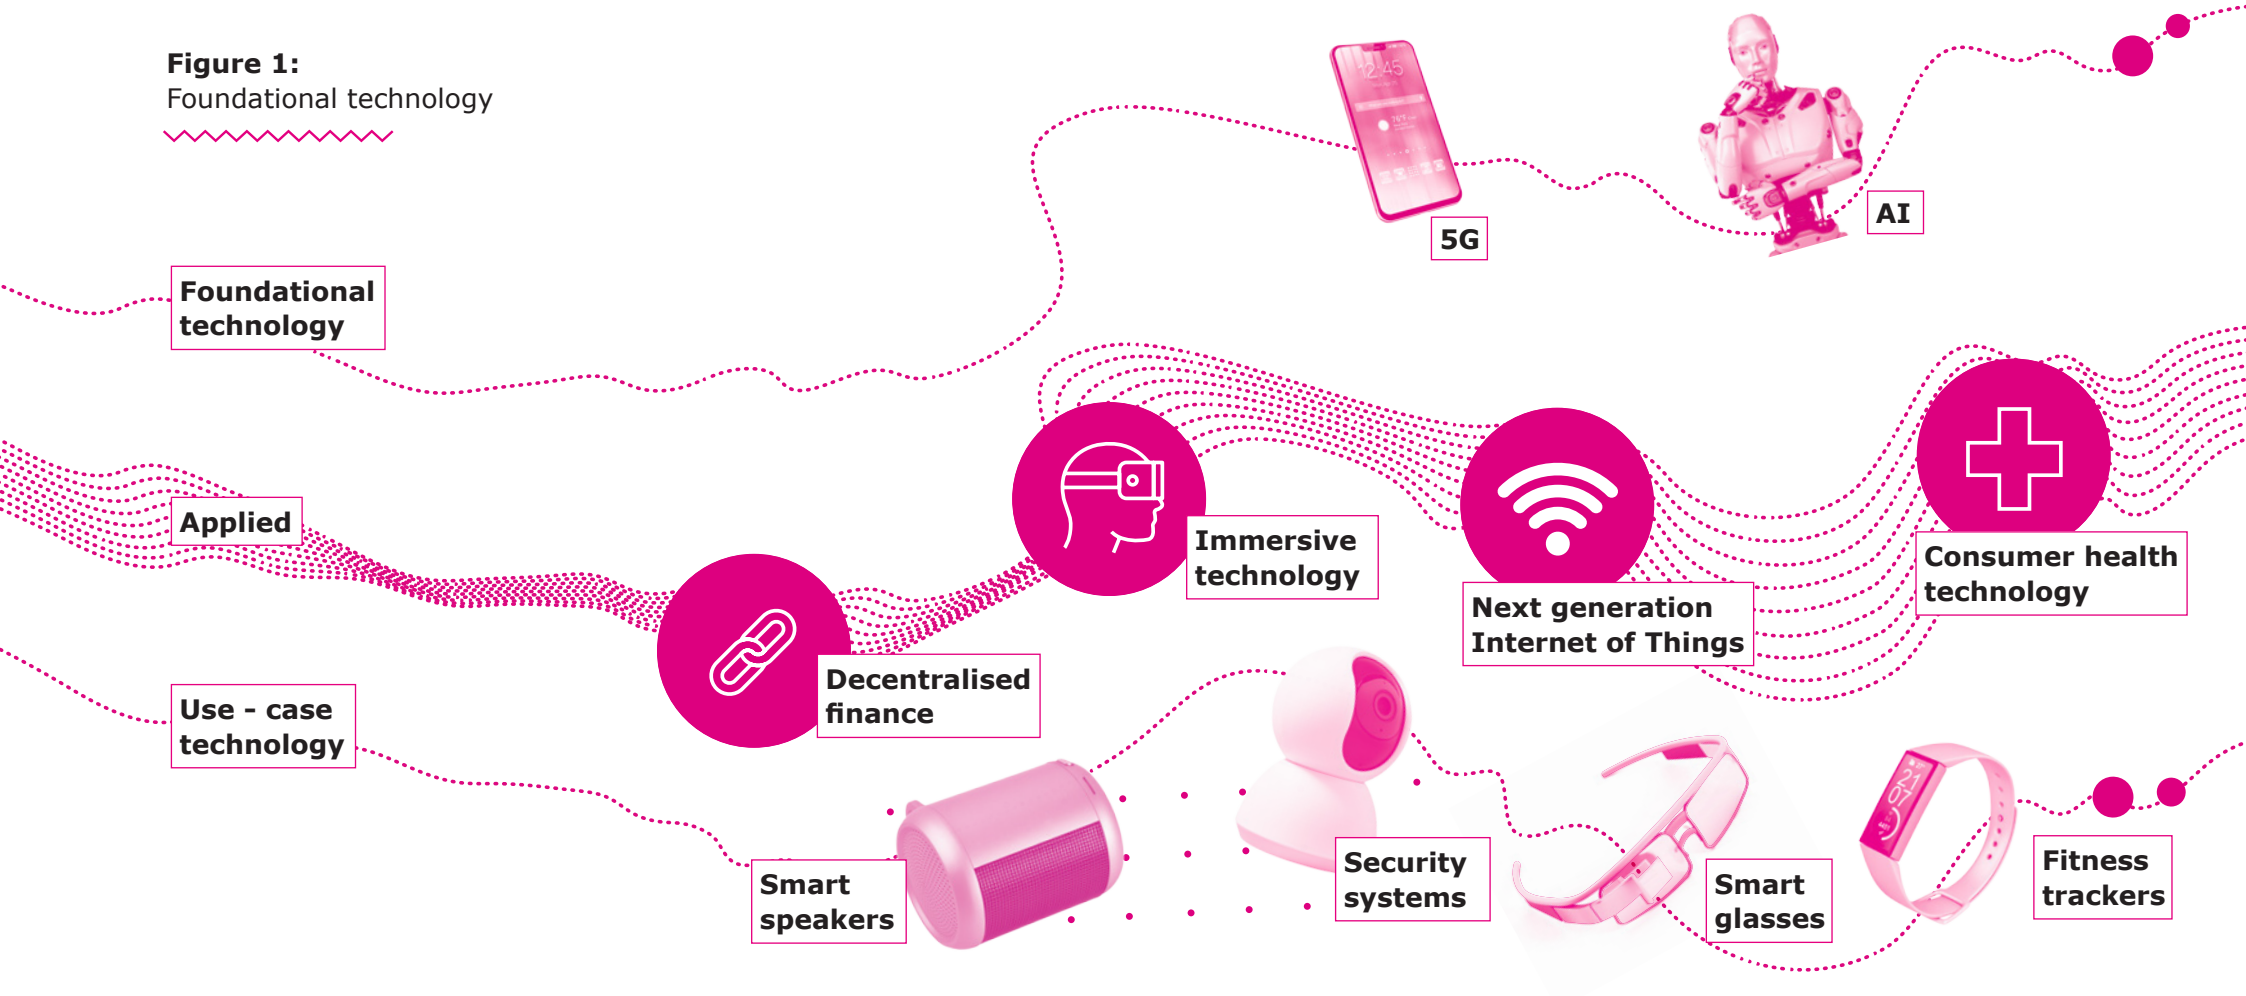

We applied **foresight methodology** against a list of 65 emerging technologies, scoring each against a prioritisation matrix that ranked the probability, scale, and associated harms and benefits in relation to privacy law. In order to identify drivers and relevance of technologies we considered factors including the:

- number of people likely to be impacted;
- sensitivity of the data being processed;
- magnitude of risks and harms to both people and society;
- magnitude of benefit to both people and society;
- likelihood of the issue to impact on vulnerable groups (such as elderly people or minority populations);
- level of current technical and commercial maturity in the UK; and
- relevant regulatory, commercial and investment factors within the UK.

This scoring process generated a shortlist of 11 technologies most likely to impact privacy in the near future.

The most significant of these, neurotechnology, is the focus of a separate deep dive report to be published in February 2023, and is not covered in this paper. This report focuses on the following four technologies deemed to have the most tangible, and often novel, implications for data protection:

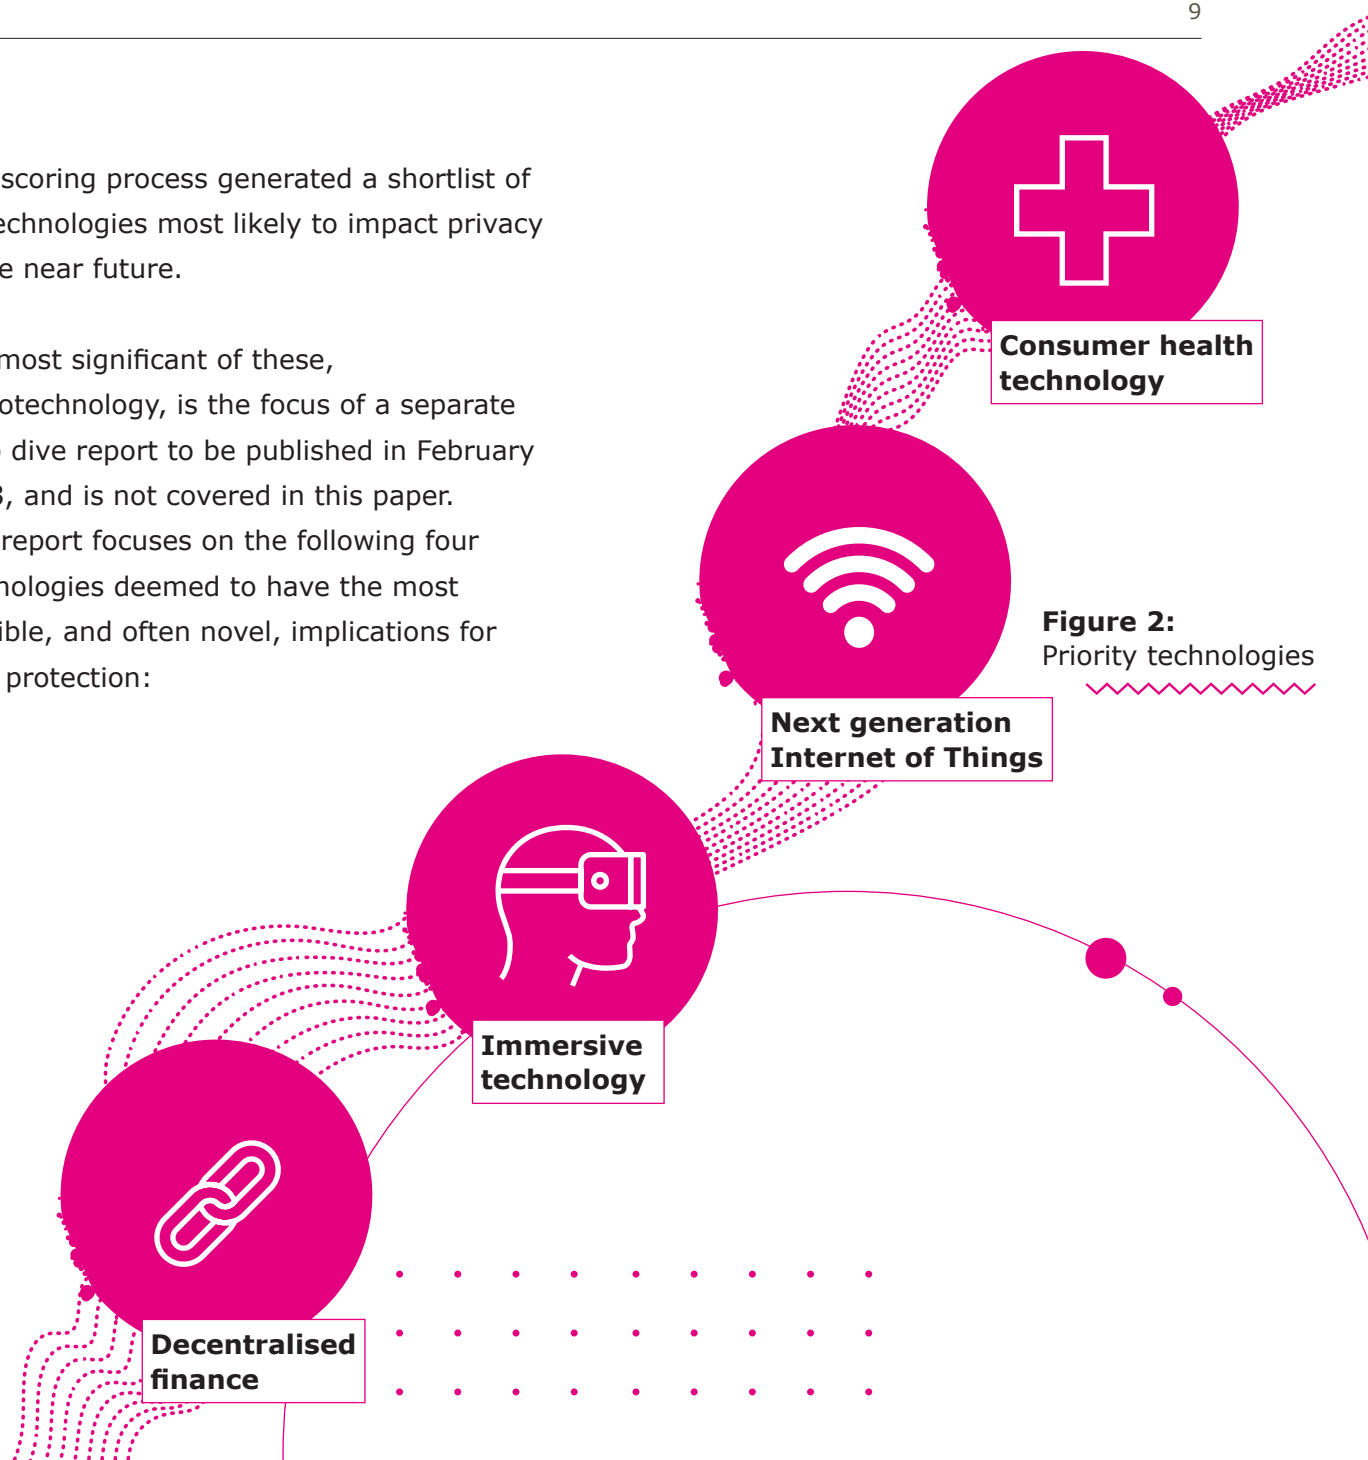

Figure 2:
Priority technologies

In selecting and assessing the four technologies discussed in this report, we conducted in-depth interviews with academic and industry experts, and held a series of workshops to ascertain how these emerging technologies will develop in the near future.

There are six technologies we did not analyse further. Although still important for privacy we assessed these to have a less **immediate** impact. Our emerging technology programme will continue to assess these technologies (and others) in future horizon-scans:[1]

[1] This does not mean the ICO does not have active programmes assessing the impact of these technologies, rather the resources of the emerging technology team was prioritised to investigate the privacy dimensions of the final 5 technologies selected.

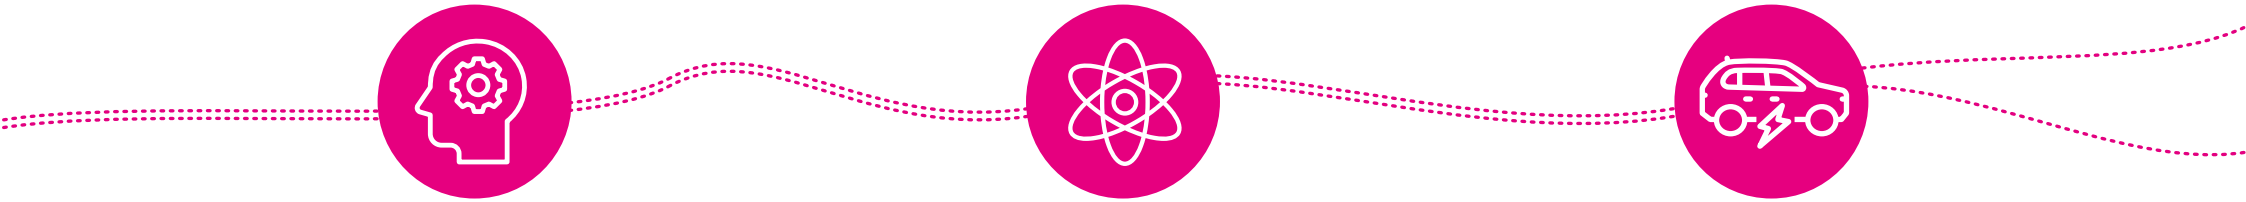

Behavioural analytics, machine-learning driven processes like sentiment analysis or emotional AI that can detect subject moods, intent or feelings. We consider this technology to be a high-risk development due to its capacity to reveal sensitive data through subconscious behaviours and responses, interpreted by highly contested forms of analysis. The ICO has cautioned organisations to assess the risks of using these technologies.

Quantum computing, roughly defined as systems that harness the laws of quantum mechanics to solve problems more effectively than classical computers. While we are still considering associated data protection issues, the possibility of quantum computers to break conventional encryption standards with relative ease is of particular interest.[2]

Digitised transport, captures connected and autonomous vehicles (AV) that use sensors and smart city infrastructure to navigate or augment passenger experiences. This issue, particularly AVs, remains an area of ongoing interest to us. For example, we are interested in the capabilities for harvesting driver and bystander data to enable vehicle operation, and the transfer of this data to third-parties to determine accident liability.

[2] The ICO acknowledges important work being undertaken on 'post-quantum encryption' to ensure cryptography can continue to secure sensitive information as quantum computing technology matures.

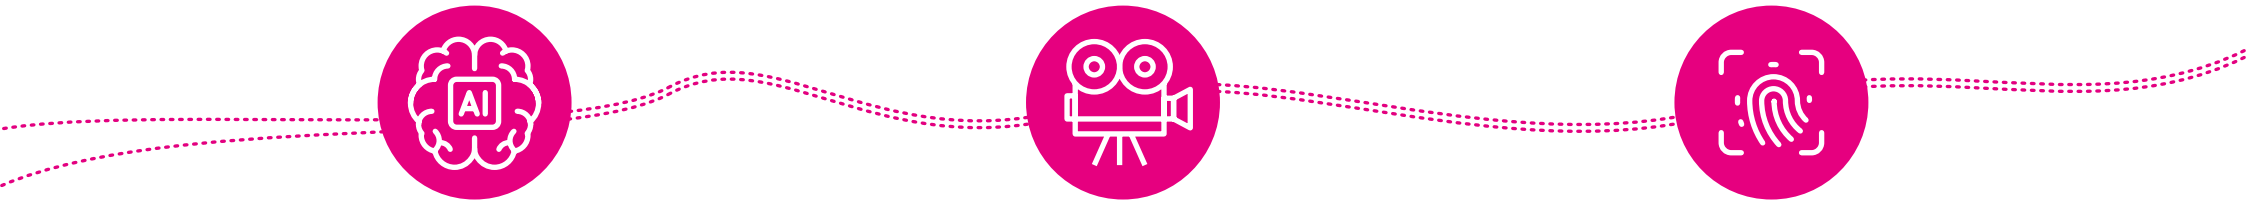

Generative AI includes deep-learning systems that create unique content, encompassing popular large language models such as DALL-E[3] and GPT-3.[4] It also covers generative adversarial networks that test AI models against each other. Although this technology is currently deployed, many of the privacy issues presented are similar to those of other AI systems. We note the ability of these models to 'scrape' information from the public internet; a means of rapidly gathering public information for commercial or private use in ways that may not have been originally intended.

Synthetic media, are often a product of the generative adversarial networks mentioned above. These are videos or images where a person has been digitally altered to create 'fake' events. While these can be benign, there is a growing practice of misappropriating someone's identify for malicious purposes, such as misinformation. Such 'deepfakes' raise novel challenges for both privacy and online safety. The Government is considering the risks posed by synthetic media in the drafting of upcoming legislation.

Digital ID covers digital representations of a person to prove authenticity and access services. We are working closely with the Government to further their commitment to using secure and trusted digital ID products in the UK. This includes helping the Government take a data protection by design and default approach to its trusted digital ID system.

The following chapters detail what types of personal information these technologies collect and how they process it, including any possible data protection implications. We undertook **scenario planning** to explore briefly how each area may develop.

We have included these scenarios to help you consider the privacy implications of our data-driven future. You should **not** interpret these scenarios as confirmation that we believe this may occur, or that this type of processing is either desirable or legally compliant.

[3] DALL·E: Creating Images from Text (openai.com)

[4] [2005.14165] Language Models are Few-Shot Learners (arxiv.org)

Consumer healthtech

Whether it's tracking exercise routines or analysing sleep and diet patterns, many people have embraced consumer healthtech in a bid to enhance their wellbeing.

As the technology matures, we expect its capacity to collect and analyse personal information to expand considerably. Already, applications detect and analyse sensitive biological data such as heart rate and blood oxygen level. Future devices may be able to gather new, arguably more revealing information like uric acid levels or biomarkers within tears. Collecting, transferring and interpreting this information between people, organisations. That might be manageable, and medical professionals may deliver positive benefits for personal health and wellbeing. However, poor data management may see information about a person's physical and mental health:

- leaked;
- processed for unsuitable purposes; or
- used to make inappropriate decisions about personal health.

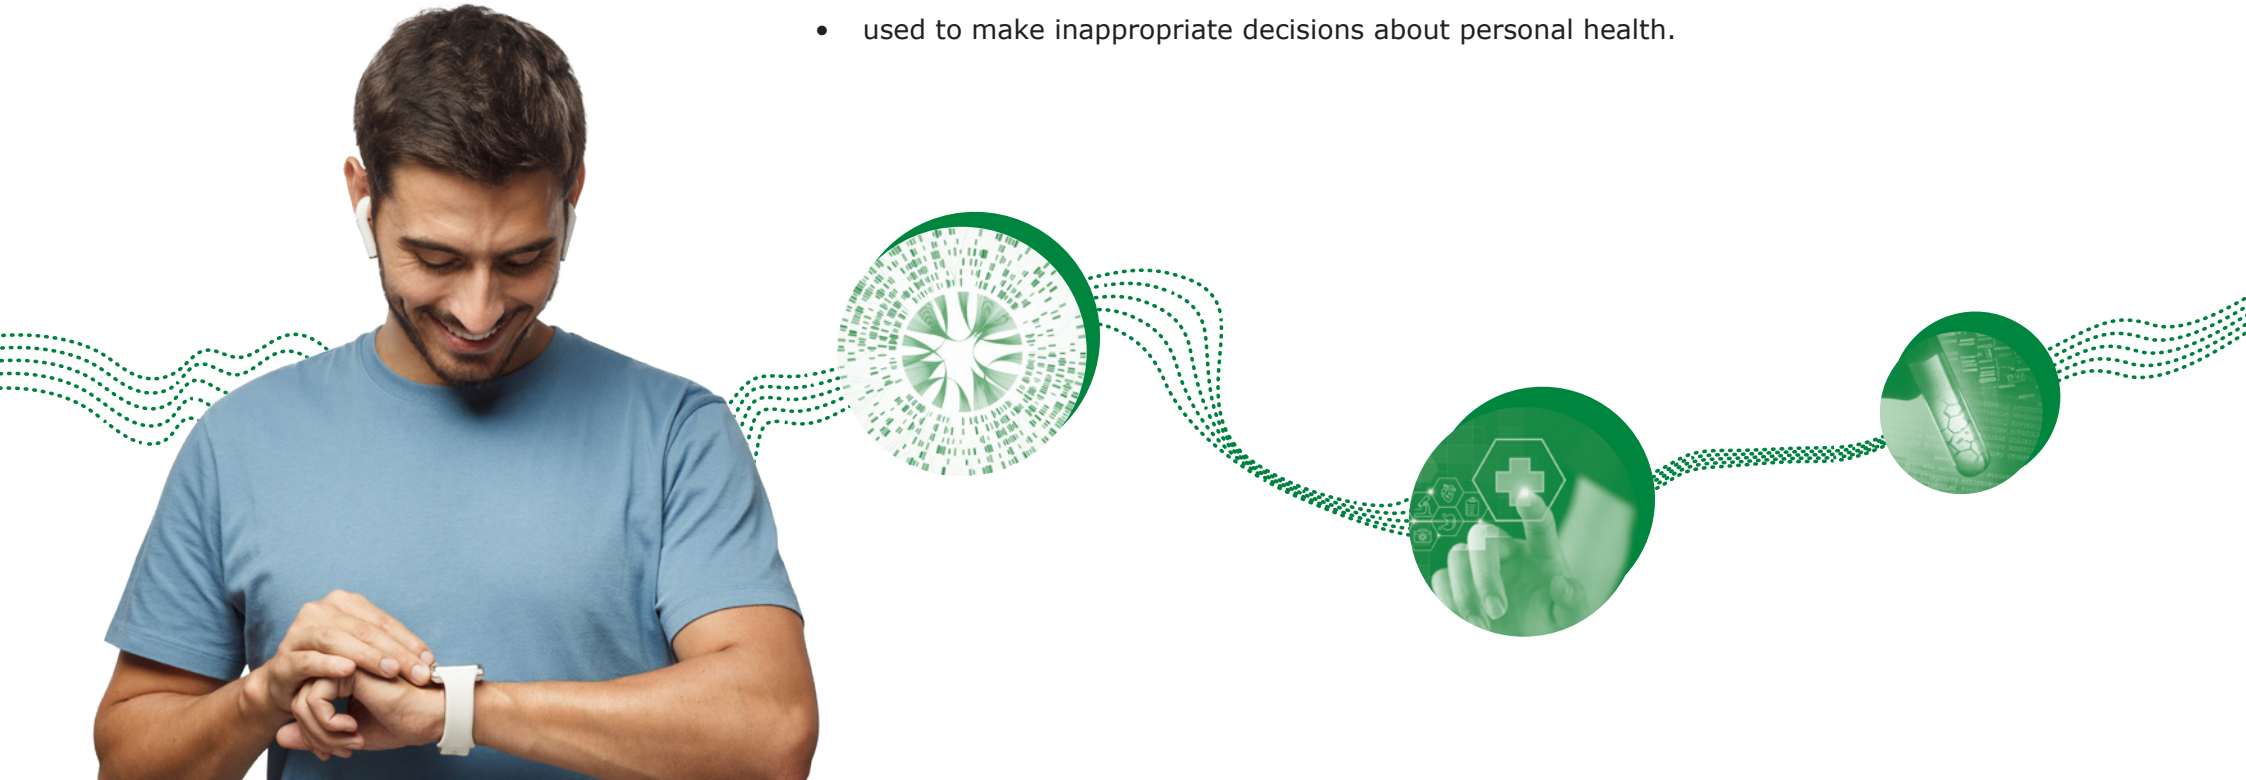

What is consumer healthtech?

Consumer healthtech refers to widely available consumer devices or applications which seek to assist a user with their health and wellbeing. This includes:

- wearables (eg fitness trackers, smartwatches and smart rings) that monitor activity levels and other health metrics such as heartrate and sleep quality;
- smart fabrics, which are textiles with electronic components that generate and record personal information from the wearer; and
- apps (commonly known as 'mHealth apps') that perform a wide range of health and wellness related tasks (eg AI-powered cognitive behavioural therapy, fertility, cough and sleep tracking).

This technology focuses on improving a consumer's wellbeing rather than medical devices or digital therapeutics (Dtx) specifically designed to provide evidence-based treatment by the medical sector.[5] However, we expect this to blur as consumer healthtech capabilities improve. For example, the NHS is already using some mHealth apps for treatment.[6]

Collecting and processing personal information

In general, consumer healthtech involves the tracking and analysis of a user's personal information. This includes sensitive biological and, in some cases, biometric data,[7] as well as a person's identifying information.[8]

The range of metrics the technology can currently collect is wide-ranging (all potentially personal information). Wearables such as smartwatches allow devices to scan for heartbeat and blood oxygen level with recent advances allowing for scanning atrial fibrillation. For example, the health and fitness capabilities of an iPhone and Apple Watch are currently reported to capture more than 150 types of health data. The integration of powerful new sensors, such as infrared spectrophotometers, may soon allow further measurements. These could include blood oxygen levels, blood pressure, core body temperature, hydration, blood alcohol and glucose levels.[9]

[5] In the UK context DTx may have been approved and classified by a medical agency (such as the MHRA). The report excludes consideration of DTx on the basis that such products are likely to have faced greater scrutiny and safeguards in the process of being designated a medical device. For example, a digital health provider which hopes to supply to the NHS will need to complete the Digital Technology Criteria (DTAC). As part of the DTAC, prospective supplies are assessed on their compliance with data protection legislation.

[6] Is it a wellness app or medical device? A critical boundary issue in the smart wearables sector (taylorwessing.com)

[7] The UK GDPR currently defines biometric data as personal information resulting from specific technical processing relating to the physical, physiological or behavioural characteristics of a person and which allows the person in question to be identified. One example is fingerprinting.

[8] ie name, address and some financial information associated with registering a mHealth app account or accessing the software layer of a device

[9] How Apple is empowering people with their health information - Apple, VitalSpex™ Biosensing Platform | Rockley Photonics

The steady introduction of smart fabrics into the range of consumer health devices will facilitate the collection of wider categories of personal information. Smart fabrics available to consumers can track hydration,[10] with future iterations likely be able to perform gait analysis[11] and track user distress.[12]

Processing this type of personal information by these technologies can occur both locally and on the cloud. It may be initially processed locally, at the device level, with the information later moving to the cloud when the device syncs with an accompanying app.

Underpinning many of these developments are advances in artificial intelligence (AI) and algorithmic processing. For example, AI can provide automated insights and recommendations about fitness based on user activity recorded by a wearable. It powers chatbots that offer support to people with mental health conditions and helps these bots to communicate in a more natural way.

[10] Gatorade launches new sweat patch for personalised hydration strategy | road.cc

[11] Smart textiles sense how their users are moving | MIT News | Massachusetts Institute of Technology

[12] Innovative 'smart socks' could help millions living with dementia | Research, Business and Innovation blog (uwe.ac.uk)

Consumer healthtech futures

Automated therapy

Many assert that the increased availability and sophistication of consumer healthtech and associated mHealth apps is likely to create more healthcare opportunities in the future. Although these devices are not specifically designed to replace medical treatment, people may turn to them to seek more immediate health outcomes, such as the provision of therapy. Equally, health authorities may seize on the capabilities they offer to promote healthy lifestyles.[13]

The steady advancement in these applications, particularly in mental health contexts, may prompt people to download and use them if they find speaking to a health professional difficult. It is claimed that training on successive sets of user data can significantly improve the natural language processing abilities of these apps. It is also claimed that people may be able to have substantive conversations with AI-powered therapy bots which have the capacity to instruct in basic cognitive behavioural therapy methods.

To provide tailored advice, we understand that apps like this will need to record and analyse discussions on the platform that reveal intimate information about a person's life, experiences and mental state. While many of these apps are free, their business models are supported by advertising and include subscription options for more personalised support. Further, apps may need to send usage data to the providers of third-party development kits in order to function as a person uses the app.

A future where these applications are widely used raises several important questions for data protection:

- Will the sensitive data provided be securely processed?
- If an app is free, is the person confident that the app is not monetising the sensitive information they provide or sharing it with third parties?
- Does the privacy policy adequately explain the way in which personal information is being processed?
- Are the insights derived from their personal information accurate and in keeping with professional best practice?

[13] This is an opportunity which has already been taken by the UK government which has launched a pilot scheme which linked healthy habits (as recorded by an app or a wearable) to rewards (such as cinema discounts) and Public Health England which has produced popular apps which encourage and track user activity

New wearables, new processing

One key driver of the uptake of consumer healthtech is the notion of 'the quantified self'. This term describes people who track various personal metrics to optimise their lifestyle, such as activity levels, nutrition, and other metrics. This behaviour both enables and creates demand for newer, more powerful devices and apps that are able to deliver ever more detailed behavioural insights.

We may see next-generation smartwatches that are able to track a range of other metrics such as blood sugar and alcohol or hydration.[14] These insights could lead people to adjust their diets and adopt healthier eating habits. However, it is possible that without a proper understanding of what the information is revealing, people may misinterpret the data they receive. For example:

- Non-diabetics that have never measured their blood sugar levels before may over-correct their diets or raise unnecessary alarms to NHS professionals as a result of readings generated by these apps;
- People may decide to drive home under the influence of alcohol after failing to properly interpret their blood alcohol levels.

In any event, the accompanying app will allow people to track these metrics and show their doctor, if necessary.

[14] Wearable devices measure a growing array of health indicators | The Economist

Data protection and privacy implications

Issue 1: Some consumer healthtech devices will generate special category data, requiring additional safeguards

We anticipate that increasingly sophisticated consumer tech is likely to process more personal information about fitness, wellbeing and health. It will therefore become increasingly important for organisations to understand when this information meets the definition of health data set out in data protection law and that requires additional safeguards.

Data protection legislation describes 'data concerning health' as:

> "personal information relating to someone's physical or mental health, including the provision of healthcare services, which reveals information about their health status."[15]

This data is classed under the UK GDPR as special category data. The legislation requires that:

> "this type of data needs to be treated with greater care because collecting and using it is more likely to interfere with these fundamental rights or open someone up to discrimination."[16]

A lawful reason is required whenever organisations are processing personal information. Alongside this, a controller needs to satisfy a separate condition for processing under Article 9 when processing special category data. Our guidance states that health data "includes any related data which reveals anything about the state of someone's health" and that this can include data from wearable devices.[17]

There are circumstances in which the source of personal information may be relevant to the determination as to whether personal information is health data. For example, a recorded assessment made by a doctor that someone needs to undertake more physical exercise is very likely to be considered health data. However, a similar inference made by a fitness wearable, and an accompanying notification advising the user to walk more (a common wearable prompt) may not necessarily always be considered health data – this will depend on the context. Another factor determining special category status is the purpose for which it is being collected. Organisations will therefore need to be aware of:

- what constitutes health data; and
- the circumstances in which personal information is likely to be health data.

This will help to ensure that they can satisfy an appropriate separate condition for processing, and that sensitive data is being handled appropriately.

[15] Health data | ICO

[16] Special category data | ICO

[17] What is special category data? | ICO

Issue 2: Action is needed to ensure that users have meaningful transparency and control of processing of their personal information

Article 5 of the UK GDPR states that personal information should be "processed lawfully, fairly and in a transparent manner". This means organisations must be clear, open and honest with people from the start about who they are, and how and why they use their personal information.

Several studies have now expressed concerns about how transparently mHealth apps have processed personal information. Although they considered these apps to have higher standards of privacy than non-health apps, they found that they continue to maintain persistent device identifiers and user contact information.[18] These allow for tracking people over time and over different services. Third parties also receive specific identifiers, such as advertising IDs.[19] Their findings also included a lack of compliance by organisations with their own privacy policies about third party sharing. Studies also found a considerable numbers of apps had provided no privacy policy whatsoever.[20]

One underlying issue is the use of software development kits (SDKs), often used by developers to add features to their own apps. For example, if developers wanted to increase traffic, shorten development time and reduce costs, they might use an SDK to allow people to log on through accounts on other platforms that people may already have.[21] This can provide third parties with insights into people's lives.[22] Our guidance states that organisations need to advise people:

- if they are going to transfer personal information to third parties;
- what the names or categories of recipients are; and
- the reasons for the transfer.[23]

However, major mobile operating systems allow third party code to have the same privileges as the app itself. The privacy settings the user applies to an app will therefore apply to the SDKs embedded within the app. As a result, people may have difficulty identifying whether the privacy permissions they grant will also apply to additional purposes such as profiling or advertising by third parties.[24]

[18] Mobile health and privacy: cross sectional study | The BMJ

[19] Mental Health Apps and User Privacy - Consumer Reports

[20] Assessment of the Data Sharing and Privacy Practices of Smartphone Apps for Depression and Smoking Cessation | Depressive Disorders | JAMA Network Open | JAMA Network; Mobile health and privacy: cross sectional study | The BMJ

[21] How SDKs, hidden trackers in your phone, work - Vox

[22] A Fait Accompli? An Empirical Study into the Absence of Consent to Third-Party Tracking in Android Apps, Proceedings of the Seventeenth Symposium on Usable Privacy and Security

[23] Right to be informed | ICO

[24] SDK Runtime | Android Developers

Issue 3: Some consumer healthtech devices present issues in terms of accuracy and bias

Under Article 5(1)(d) of the UK GDPR, organisations are required to take all reasonable steps to ensure the personal information they hold is not incorrect or misleading "as to any matter of fact".[25] This principle is relevant to consumer healthtech because of the concerns raised about the accuracy of the data produced by wearables. Some studies observed that consumer wearables can measure certain metrics such as steps, heartbeat and sleep duration accurately.[26] However others produced results that indicate potential accuracy issues for all users.[27] There are also concerns that the information recorded by wearables may be less accurate for people with darker skin.[28] These studies argue that the variation in accuracy is the result of insufficient demographic diversity in product test data.

Conclusion

Advances in consumer healthtech have the potential to make a considerable impact on consumer wellbeing. They would allow people to track a wider range of metrics and could potentially promote healthy habits and levels of activity. However, these perceived benefits need to be considered in the context of increased processing of sometimes intimate data by consumer health devices. In order that organisations develop consumer healthtech in a privacy positive way, organisations should consider the following points:

- Privacy notices should provide clear, intelligible information on how and why they are processing someone's data and what inferences they are drawing.
- Clarity about whether they are processing special category data, taking into account the factors discussed above. They should consider using the DPIA process where appropriate.
- Algorithmic processing and use of AI in conjunction with healthtech should be accurate, fair and checked for risks of systemic bias (our guidance on AI and data protection can be found here).

[25] Principle (d): Accuracy | ICO

[26] Accuracy of Fitbit Devices: Systematic Review and Narrative Syntheses of Quantitative Data - PMC (nih.gov)

[27] Evaluating the Validity of Current Mainstream Wearable Devices in Fitness Tracking Under Various Physical Activities: Comparative Study - PMC (nih.gov), Accuracy of Fitbit Devices: Systematic Review and Narrative Syntheses of Quantitative Data - PMC (nih.gov)

[28] How Accurate Is Smartwatch Heart-Data? It Depends on Your Skin Tone - American College of Cardiology

Next generation IoT devices

The Internet of Things (IoT) describes the network of physical objects ('things') that can connect and share information with other things and systems over the internet. These 'things' can sense, respond to or interact with the external environment.[29]

IoT is not a new technology, but it is evolving. With the aid of localised processing (commonly known as 'edge computing'),[30] improved hardware, software and interoperability, the next generation of these environments will be able to respond to people's needs in real-time. Possible examples include managing and optimising the work environment or receiving personalised recommendations from appliances.

Developments in IoT and edge computing are moving us into a future of environments that can anticipate and respond to people's needs in real time. The Roadmap for IoT research, innovation and deployment in Europe highlights that the technologies are "narrowing the gap between the virtual and physical worlds".[31]

[29] See, eg, What Is the Internet of Things (IoT)? | Oracle India. The NGIoT Report: A Roadmap for IoT in Europe – Next Generation IoT provides a more detailed definition. There is further discussion in: Trusted-Internet-of-Things-at-home-and-in-the-workplace.pdf (tas.ac.uk).

[30] In this report, edge computing is used broadly to refer to "distributed computing in which processing and storage takes place at or near the edge" of a network – ie close to the source of data (eg on a device or network of devices), rather than in the cloud. For further detail, see: Roadmap for IoT in Europe – Next Generation IoT, citing ISO/IEC TR 23188:2020.

[31] NGIoT Report: A Roadmap for IoT in Europe – Next Generation IoT

What are IoT devices and smart spaces?

IoT devices vary in function and sophistication. They may contain sensors, actuators (a means of interacting with the environment, such as engaging a switch),[32] software and other technologies.[33] When used in the home they can include:

- smart kettles;
- thermostats;
- heating, ventilation and air conditioning systems (HVAC);
- locks;
- smart phones;
- smart video cameras;
- TVs; and
- smart speakers.

'Smart' spaces are IoT-equipped physical environments, such as a home or office. These spaces can deliver significant volumes of detailed information about the condition of the spaces and how they're used. They can also actively manage the space and ambient environment (eg temperature and humidity). Promoters argue that the increasing integration of IoT into these spaces can:

- increase comfort, convenience, safety and security;
- support independent living; and
- enhance employee productivity and wellbeing.

IoT devices already present data protection issues. Soon to be published research by the University of Cambridge, funded through the ICO grants programme, has observed that:

- IoT devices can capture more information than many companies routinely disclose to consumers. This includes account registration data (through linked apps), sensor data, information about user interactions with a device, technical information about the device, and inferences or user profiling;
- information collected by IoT devices may flow to a range of servers, third parties and across jurisdictions; and
- manufacturers of connected devices can have inconsistent and limited understanding of transparency obligations.

The growing deployment and increasing interconnectedness of IoT devices is likely to exacerbate these issues, as will the wide variety of devices in use, scale and complexity of the ecosystem.

[32] An actuator is a part of a device that enables it to physically interact with the environment (eg, turning off a light, or pulling down a blind, or turning on the heating).

[33] See footnote 29.

Collection and processing of personal information

Many IoT devices already process highly revealing personal information at a large scale.[34] It is anticipated that next generation IoT setups in homes and offices will be 'smarter', and devices will be more integrated within spaces and better at talking to each other. We anticipate these developments will be aided by:

- improved network infrastructure, such as the rollout of 5G;
- more advanced machine learning; and
- in the longer term, edge computing, that offers opportunities for more responsive IoT devices capable of faster, more efficient, and more powerful processing.

As a result, some devices will offer greater capabilities for personalisation. When operating together, we expect they will be able to provide increasingly detailed insights into peoples' behaviours, lifestyle, habits and attitudes.[35]

Device sensors collect specific types of information from people and the environment. Depending on the functionality of the sensor, information collected may include:

- voice commands, video images, user location and movement;
- home layout, room, appliance or energy use; and
- temperature, humidity and air quality.

Inferences about personal habits can also be drawn from how an IoT device operates. For example, the time an appliance switches on or off may suggest information about people's whereabouts.

It will be personal information if:

- observed or inferred data is linked to a person (eg through an account or other unique identifier); or
- the information collected otherwise allows for the identification of a person (eg voice recognition).

[34] ICO response to DCMS call for evidence on "Connected Tech: Smart or Sinister?"

[35] ICO response to DCMS call for evidence on "Connected Tech: Smart or Sinister?"

Currently, many IoT devices transfer collected information to company servers (eg in the cloud), where it can be analysed and interpreted using machine learning. It may be combined with information from other sources and devices. The analysis allows manufacturers to understand, anticipate and respond to people's preferences. An IoT device then receives instructions and will take various actions such as adjusting temperature, ventilation or turning off a light.

Like smart homes, smart offices use a combination of sensors, actuators and analytics to learn and adjust the workplace environment. Examples of currently commercially available technology include systems with the ability to:

- automatically adjust lighting and air quality;
- optimise heating systems (eg by adjusting output to the number of people in a room, or allowing employees to tailor environmental preferences via a mobile app); or
- to control energy usage by forecasting the weather and a building's energy consumption.

Not all data processed will be personal information, but it may be linked to personal information. A smart heating, ventilation and air conditioning system, for example, may also use geofencing and location information.[36] Or, apps could feasibly collect information on individual worker environmental preferences or usage of spaces.

IoT platforms can draw on real-time sensor data to visualise usage patterns of spaces. Data analysis tools can inform space design, allocate desks, or monitor numbers in a space to facilitate social distancing. Smart access can be enabled via technologies such as geofencing, automated entry or QR-enabled self-service check-in. Where the information collected relates to, or allows for direct or indirect identification of a person, it will be personal information.

[36] See, eg, What is HVAC, and how can making it SMART create a healthier workspace?

Connected futures

The smart home

Many aspects of the home may become part of an IoT ecosystem. Advances in consumer tech will enable sleep monitoring, automated environmental adjustments, devices that can control other devices or appliances capable of making personalised recommendations. Examples may include:

- smoke detectors that can switch off an oven;
- a smart speaker that turns off lights at a specific time of day, without being asked; or
- appliances that can make meal planning or water saving recommendations.

Open source interoperability software will enable many different devices, using different network protocols, to communicate and work together locally. These devices won't need to connect to the cloud except for software or security updates. Other setups may continue to use a mix of cloud and local processing. It is predicted that homes will increasingly be equipped with smart meters and energy smart appliances that respond automatically to price or other signals (also known as 'demand response'). These devices could feasibly integrate into wider smart home setups.

It is expected that voice assistants deployed on more sophisticated IoT devices, such as smart speakers, are likely to be the way people access the increasingly digitised home and can support multiple user profiles. These assistants will deploy increasingly sophisticated AI, using machine learning, prediction and personalisation. This will enable people to ask their device more advanced questions.

Some devices will increasingly be able to predict what a person wants, and take decisions or complete a wider range of tasks without being asked. This is likely to be based on significant volumes of personal information generated in the home about people's behaviour, history and preferences.

The capabilities offered by IoT may be used to help people live independently, or in assisted living situations. For instance, the future home of an elderly person or a person with additional needs could include a range of sensors and other smart devices to assist them with home living. Examples include:

- managing energy consumption by automating things like lights and heating; or
- installing safety features, such as making sure electric and gas appliances shut down automatically if the smoke alarm is triggered.

Issues of consent and awareness of the people involved, including the individual, visitors and professionals working in the space will remain extremely important. How organisations may use this information is another important issue. Other organisations may seek to sell pseudonymised data to third parties for advertising purposes. For example offering device upgrades, or offers to switch energy suppliers based on projected usage. The issues of transparency, fairness and purpose limitation will be important here and are considered further below.

The smart office

The future smart office is expected to incorporate more of the functionalities of existing devices for different purposes, and increased adoption of back-end AI to control buildings. Organisations may deploy the technology with the aim of enhancing employee experience and productivity, adapting spaces to hybrid working, and enhancing energy efficiency.

Networks of sensors linked to IoT platforms may support open plan working. This might be by indicating what workspaces are available at any given time or allowing app-based or automated desk booking. This same technology could track individual desk and room usage to provide aggregated insights into how employees use the space, and allocate space more efficiently in a hybrid working setup.

It is anticipated that firms and commercial properties may invest in IoT-based energy efficiency measures to reduce carbon footprint and save costs. As in the smart home, this could include systems that automate HVAC, air quality, lighting (eg intensity of light, position of blinds or window tints) and electrical charge points. Systems will employ AI to determine the optimum balance between employee preference, productivity and savings.

Another hypothetical system could combine IoT data from a range of sensors and employee use of IT. This could be used to generate emotional and environmental metrics and automatically optimise uses of a space.

Analytics may increasingly take place at the 'edge' of the network. This is in response to industry demands for faster, more efficient and responsive processing. Other setups may be a hybrid of cloud and edge. Such a transition may allow devices to apply privacy-enhancing techniques before transferring data to the cloud, building further on any existing data minimisation processes used. Conversely, the move has the potential to open up privacy and security risks. This could include opportunities for physical attacks on edge devices to steal data, and challenges in determining accountability.

Data protection and privacy implications

Issue 1: Widespread use of next generation IoT devices may increase cybersecurity risks

Cybersecurity is a key issue for both legacy and next generation IoT systems. Organisations have obligations under UK GDPR and DPA 2018, and the Network and Information Systems Regulations (NIS) may apply in some cases.[37] For example, to cloud computing services that underpin current IoT systems.

The increased uptake and variety of connected devices and integration of legacy and next generation devices in smart spaces is likely to increase the cybersecurity risk. In the home particularly, low consumer awareness is likely to compound the challenges. Potential cybersecurity challenges include:

- additional routes for cyber-attack due to the variety of devices in use;
- software update cycles that leave older devices (such as appliances) without security support;
- computing power limits of some IoT sensors, which make them harder to update;
- insecure default settings (such as default passwords);
- the risk that some device manufacturers could prioritise speed to market over security and data protection by design; and
- IoT system design that allows for malicious third parties to observe mid-stream processing and remote access to video or audio.

Increased use of machine learning and on-device (edge) processing in next generation IoT may bring some security benefits. However, these technologies may also open other routes for attack on devices. Integrating multiple IoT systems with different security standards is also likely to be complex.[38]

[37] Together with the NCSC, we have also produced joint NIS resources for controllers and processors: Security outcomes | ICO; GDPR security outcomes - NCSC.GOV.UK

[38] Key developments to watch in this space include the Product Security and Telecommunications Infrastructure Bill, which proposes requirements for default settings for devices sold in the UK. The Department for Digital, Culture, Media and Sport is also reviewing options to enhance app security: App security and privacy interventions - GOV.UK (www.gov.uk)

Issue 2: Action is needed to provide people with meaningful transparency and control of their personal information

A key ongoing challenge for IoT designers and developers is how to effectively communicate privacy information to people in the absence of a user interface.

Data flows between IoT devices and organisations are already complex, as manufacturers, software providers, and other third parties can be involved. Greater interoperability and a wider range of devices on the market could further complicate these data flows. It will be increasingly important to explain these data flows to people in ways that give them meaningful control over their information while avoiding overloading them with detail.

These challenges will only become more acute in smart homes. For example, real-time profiles from individual appliances or a wider home IoT network can also be used for personalised recommendations and marketing. The aim is to encourage certain behaviours, or to inform prices or incentives. Multiple information sources can also be fused together to generate more detailed insights.

This raises several questions. How would a kitchen integrating a smart fridge, smart oven, smart smoke detector from a variety of manufacturers meaningfully communicate privacy information to people? In next generation smart homes designed to aid independent living, how should privacy information and choices be communicated to people with different needs or vulnerabilities, to ensure they can exercise their information rights?

Landlords also need to be mindful of the transparency considerations for residents and visitors in a smart rental property. This issue is particularly relevant if the landlord has access to the information generated.

In each circumstance, it will be important to ensure people have sufficient options to effectively exercise their rights over how their information is collected, retained and shared with each of the companies involved. The ICO grants programme has funded research looking at how to help people understand external data flows in a smart home,[39] and the accuracy of privacy policies in third party apps running on smart home assistants.[40]

Existing options, such as:

- using lights to show that a device is switched 'on';
- displaying privacy information through a connected app; or
- voice-activated options to exercise data rights may still have a part to play.

However, organisations will need to find ways to mitigate the risk of information overload preventing people from understanding or choosing how their personal information is used.

In a smart office, the presence of sensors and other connected devices is unlikely to be immediately obvious without direct intervention by the employer. They will need to make staff and visitors aware of their use, especially if devices are integrated into the fabric of a building or other amenities.

This is particularly problematic for more invasive sensors, such as:

- sensors tracking movement in real time; or
- devices collecting information that could be used to make inferences about employee engagement (eg analysis of tone of voice).[41]

Potential interactions between IoT integrated into the building and employees' own devices and wearables could also generate additional information. Employees may not be aware of the collection and processing of this information. Likewise, employees may not be aware of the use of back end AI to process information gathered by the building, which may or may not include their personal information.

Use of office space is likely to be a condition of employment. Given such a power imbalance, the fairness, necessity and proportionality of processing will also be an important considerations in a smart office. Employers must also identify their lawful basis for processing.

[39] Research on privacy by design in smart homes (University of Oxford)

[40] Evaluating third-party smart home assistant developers (Kings College London)

[41] Our recent biometrics insights and foresight reports discuss the issues associated with emotion analysis technologies in further detail, such as voice analysis: Biometrics technologies | ICO

Issue 3: Concerns exist about excessive collection or repurposing of personal information

Without adequate safeguards or transparency, there is the potential for smart space information to be used for secondary purposes. For example, employee performance reviews, tenant surveillance, property valuation, or marketing. Organisations must be clear and specific in setting out the purpose that they intend to gather and process personal information for. They must ensure that any further use of such information is compatible with the original purpose for which it was collected.[42]

There is also potential for excessive information collection, as more devices are deployed and connected together, or multiple sensors are deployed in one device. Determining the minimum amount of information needed to achieve the purpose for processing is likely to become increasingly important. This is particularly relevant as IoT systems' real time analytics capabilities increase.

Conversely, privacy outcomes may improve as device sophistication and processing power increase, combined with a move towards edge (localised) processing. For example, some devices (such as smart speakers or smart home hubs) could increasingly apply certain privacy enhancing techniques (PETs) to personal information before it is transferred. This could include data aggregation and encryption techniques, that make it far harder to identify individuals. However, the full feasibility and implications of different techniques for smart homes and offices are unclear. Moreover, not all PETs will be useful for every type of IoT setup.

[42] The ICO will be releasing guidance on employee monitoring that will discuss related issues.

Conclusion

Many claim that the next generation of IoT devices holds tremendous promise for reducing carbon footprints, assisting vulnerable people and improving the productivity and efficiency of both our domestic and professional lives. However, to maintain trust in these innovations, developers and regulators need to be aware of the significant challenges these technologies represent for data protection compliance.

Organisations should consider additional steps to take now to implement privacy-positive innovation, including:

- prioritise efforts to enhance security of IoT devices, for example by implementing default security standards outlined in the upcoming Product Security and Telecommunications Infrastructure Bill,[43] and considering the European Telecommunications Standards Institute's IoT security standard;[44]
- ensure high standards of privacy by default, with user-centred design of connected devices and, in the home, age-appropriate approaches;[45]
- continue to explore approaches to transparency and data minimisation in smart spaces, including homes and offices;
- remain alert to the potential benefits that edge computing in smart spaces offers, and be aware of the unique privacy and security challenges; and
- further explore the potential of PETs in the context of connected devices.[46]

The ICO intends to:

- develop guidance on the data protection aspects of IoT devices;
- continue our support for default security standards, age appropriate design and privacy by design and default in IoT, including fostering innovative approaches; and
- continue our engagement with stakeholders in the property sector on their privacy readiness for future IoT deployments.

[43] Product Security and Telecommunications Infrastructure Bill - Parliamentary Bills - UK Parliament. The Bill builds upon the UK's Code of Practice for consumer IoT security - GOV.UK (www.gov.uk)

[44] EN 303 645 - V2.1.1 - CYBER; Cyber Security for Consumer Internet of Things: Baseline Requirements (etsi.org)

[45] 14. Connected toys and devices | ICO

[46] ICO draft guidance on anonymisation, pseudonymisation and privacy enhancing technologies guidance: Chapter 5 ; ICO call for views

Immersive technology

After years of growth and investment, immersive technologies are beginning to play a role in how we work, learn and play. The newest smartphones are ready to use augmented reality applications and individuals are exploring ‘virtual arcades’, booking tickets to immersive entertainment experiences and purchasing their own virtual reality headsets.

Experts have predicted that immersive technologies are likely to play a role in a more interconnected future. This would mean the internet can be accessed not only through two dimensional (2D) rectangular screens, but as a more lived experience that people step into and spend time within.

The simulation of direct experience is a fundamental concept underpinning augmented reality, virtual reality and other immersive technologies. Virtual learning environments or training facilities can lower barriers to education whilst improving its quality and we are aware this technology is already being deployed in the workforce to overcome geographic limitations to collaboration.

Despite these benefits, the arrival of immersive technologies into our lives, homes and public spaces raises legitimate concerns about the implications for privacy if these devices become as universal as mobile phones are today.

What are immersive technologies?

Immersive technologies encompass a broad range of applications, labelled as 'extended reality' (XR). However, this analysis focuses on the nearer term and more common applications of augmented and virtual reality.

Augmented reality (AR) places a digital layer on top of the physical world. It is the real-time use of information in the form of text, graphics, audio and other virtual enhancements integrated with real-world objects.[47] AR technology can and frequently is employed by 2D devices such as smartphones and tablets today.

AR technology is now used in many mainstream smartphone apps, such as Instagram and Snapchat filters. The trend of integrating AR into potentially discreet wearable devices such as smart glasses and head mounted displays is raising novel privacy concerns.

Virtual reality (VR) generates a three-dimensional (3D) environment that surrounds a person and can respond to their prompts and movements.[48] Typically, VR achieves this via an immersive head-mounted display. This is supported by handheld controls that provide body tracking, and may include a degree of haptic feedback (where a tactile response is received by a user, such as force or vibration).

Since the 1980s, VR capabilities have evolved dramatically from basic stationary locations. They are now standalone wireless units with sophisticated visuals, tracking, sensors and cameras designed to enhance comfort and immersivity. Concerns have been raised that these improvements to the user experience are supported by an increase in processing sensitive personal information.

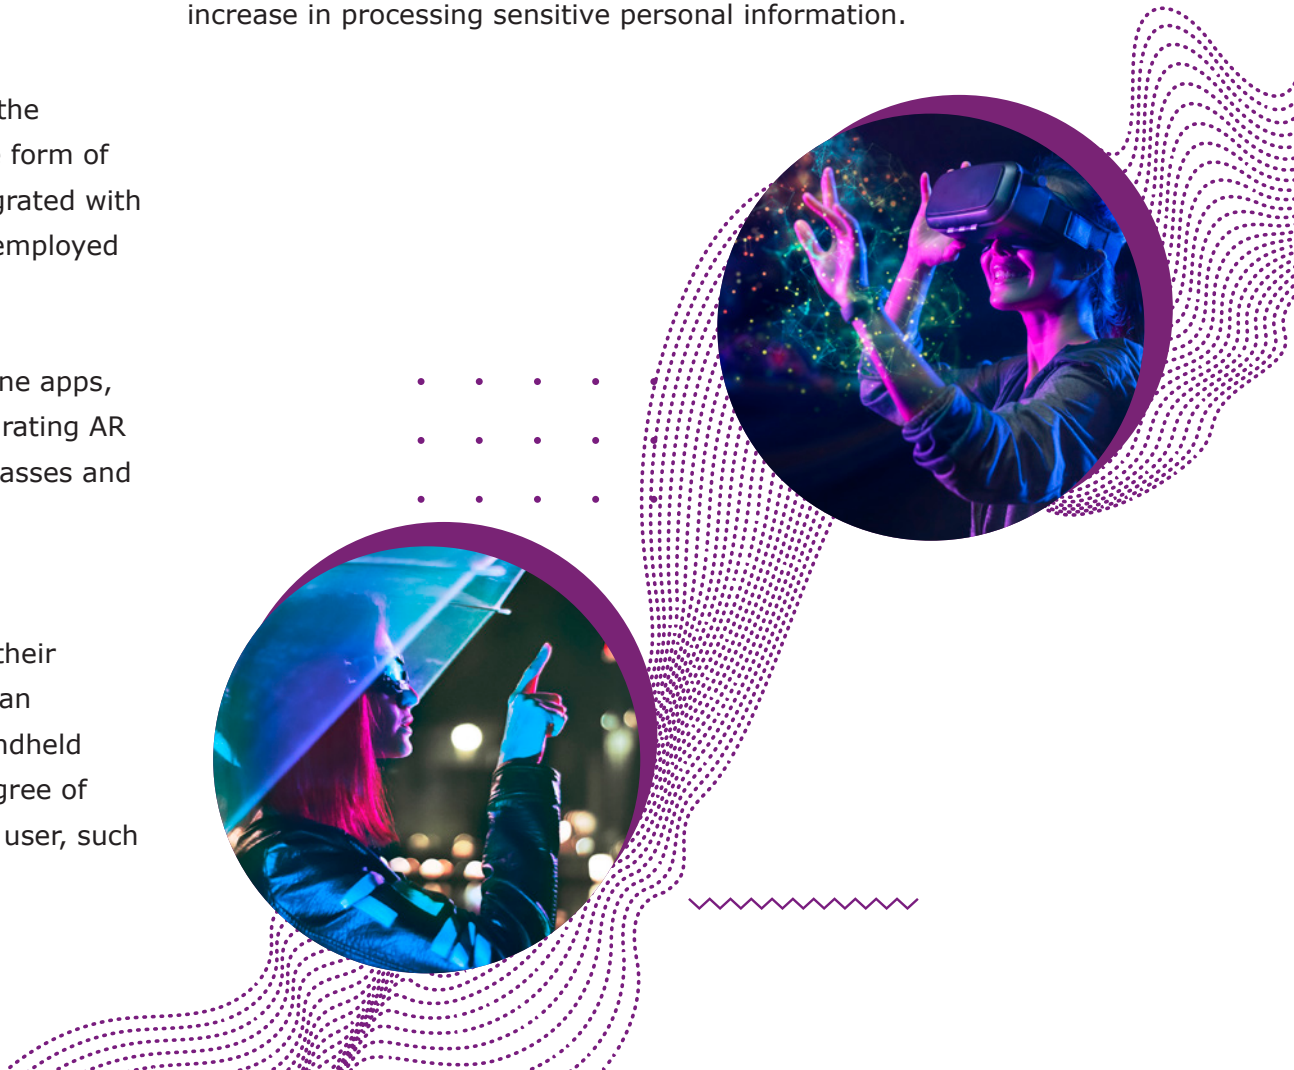

[47] Definition of Augmented Reality (AR) - Gartner Information Technology Glossary

[48] Definition of Virtual Reality (VR) - Gartner Information Technology Glossary

Collection and processing of personal information

Collecting and processing personal information in immersive technology starts with the AR or VR devices themselves and the functionality they provide.

User characteristics and biological information: Devices collect a variety of information on people, ranging from how they physically move throughout the space they are in, to their arm distance or wingspan, relying on a variety of sensors within a headset or handheld controllers. Higher specification models already include user facing cameras to monitor facial movements, and heart rate monitoring.

On some devices, sensors collect a variety of biological data relative to the optical movements of the person. This includes eye movement or gaze tracking and pupil size, with some devices having the capability to identify people by iris scans.

Audio information: Most devices also feature microphones within the headset. Microphones can capture a person's voice for interactivity, social communication and potentially to enable any noise cancelling features.

Spatial and location information: Cameras, gyroscopes and depth sensors are used to collect information on the environment surrounding the device including assessing the room or 'zone' of operation. This can include details of room dimensions and furniture or objects within the space.

Devices can collect approximate location information via the device's IP address. More exact geolocation information may be available through Wi-Fi and Bluetooth functionality and connected applications.

User interaction information: Software applications in phones and tablets can track people's digital interactions such as app purchases and log in or usage times. The same telemetry information can be generated by software applications within immersive technology devices. This includes any information a person proactively inputs into an application or device such as a mindfulness application that asks someone their mood to determine the most appropriate exercise for them.

Immersive futures

Entertainment and wellness

It is expected that the home entertainment and media sector will continue to be the primary area where consumers adopt immersive technologies.[49] In the coming years many expect traditional entertainment such as TV and video to blur into gaming, virtual worlds and immersive social spaces. VR presents a new and exciting opportunity to engage in entertainment through 'best seat in the house' formats. This means people are placed at the centre of an event or performance, such as a play or boxing match.[50]

Participation in wellness activities is also expected to be altered by immersive technologies in the home, with a possible shift towards virtual fitness classes, mindfulness and meditation exercises. Of particular interest is how these programmes may collect and use personal information either from the AR or VR device, or in combination with existing consumer healthtech wearables. Combining data points from multiple devices may enable developers to tailor content to people's needs. However, it may also build a detailed and accurate picture of a person's wellbeing, fitness and overall health.

VR apps are expected to have the capacity to capture increasingly sophisticated records of people's attention. Software may analyse a person's eye movement, pupil size, heart rate, voice and hand gestures as well as noting their choices in a simulated environment to understand what captures their interest.

As this industry matures, organisations may combine information generated within a virtual setting with datapoints generated in other online sources. This could include information from social media, search engines or online stores. Combining these datasets could create a revealing picture about a person's activity across a range of domains. This may allow detailed inferences to be drawn about preferences and behaviour, for example to attempt to increase the effectiveness of targeted advertising.

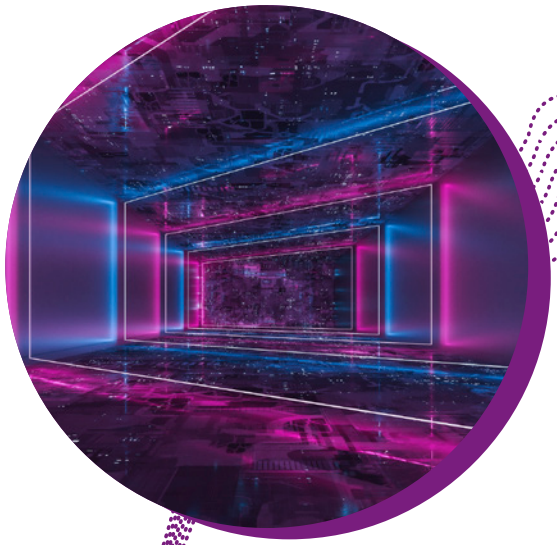

[49] YouGov 2022: 2022_08_YouGov_UK_US_Metaverse_Report.pdf p32

[50] Limina Immersive 2018 Immersive_Content_Formats_for_Future_Audiences.pdf (immerseuk.org) p69

Workplace

Initially, commercial uses of immersive technologies were largely restricted to sectors such as engineering, architecture, and construction. They used 3D modelling to support project planning for physical objects such as buildings, vehicles and turbines. These sectors remain prominent users of the tech.

However, it is expected that the future will see AR or VR devices integrate with drones and digital twin technology to recreate high-risk environments for vocational training purposes. This presents a benefit to the employer by reducing cost, risk and providing a higher quality of training. However, it comes at the expense of processing greater volumes of an employee's personal information, and types of personal information not typically processed for employment purposes previously. Employers could in future use this information to assess fear in virtual high-risk environments. For example, determining engineer comfort working at simulated heights or in high-pressure environments as part of a health and safety or risk management programme.

The growing adoption of immersive technologies into more white collar settings is expected by some to profoundly affect how and where we work in the future. To replicate in-office engagement in remote work settings, employers increasingly rely on video conferencing technology. Yet the experience is not seamless. Attendees can be distracted, body language and visual cues are difficult to pick up on, and 'zoom fatigue' remains an issue across the workforce. Immersive technologies represent a future of work that seeks to solve these engagement challenges. These technologies combine the benefits of remote employment with the presence of live, in-person interaction. Teams can:

- congregate around a virtual table;
- hold ideation sessions with whiteboards;
- gather around a 3D object; and
- work collaboratively without barriers.

While this vision of an optimised workforce promises benefits to industry, there are concerns that it may result in a new and more intrusive form of employee monitoring. For example, while it may be technically possible for devices to assess attention span or stress levels and then report those readings to a central employer, such practices would raise significant privacy concerns.

Data protection and privacy implications

Issue 1: Many immersive technologies will collect information about sensitive human characteristics, requiring additional safeguards

Technologies such as eye movement tracking, iris scans and facial movement monitoring embedded within current hardware devices collect information that may be biometric. Biometric data has special category data protection under the UK GDPR if it is used for the purpose of uniquely identifying individuals. Even if the information collected doesn't meet the threshold for special category biometric data under the UK GDPR, the data collected may be special category and when considered in the aggregate, can reveal a vast amount of information about a person.[51]

Many see this as an opportunity for the development of ever more sophisticated user profiles which combine a person's:

- eye movement information (tracking where and how long a person may observe an advertisement); and
- facial movement responses and biological markers (such as spiked heartrates).

Combining this information would allow for detailed inferences to be made about a person's likes, dislikes and desires. The concept of 'biometric psychography' details this privacy risk, when a person's biometric data is combined with targeted advertising.[52]

Future advancements are expected to include the integration of brain computer interface technology, which allows for a direct communication link between the brain and a computer, to be embedded within AR or VR wearables.[53] This could open up the opportunity to collect a new and sensitive form of neural data on an unprecedented scale. Previous applications involving the collection or processing of neural data have largely been confined to the healthcare industry. However, the potential expansion of processing information in this way and its use for the purposes of mainstream marketing could have profound implications for privacy. Should this personal information be combined and processed for the purpose of emotional analysis, previously raised concerns by the ICO including accuracy, fairness and transparency should and must be considered.[54]

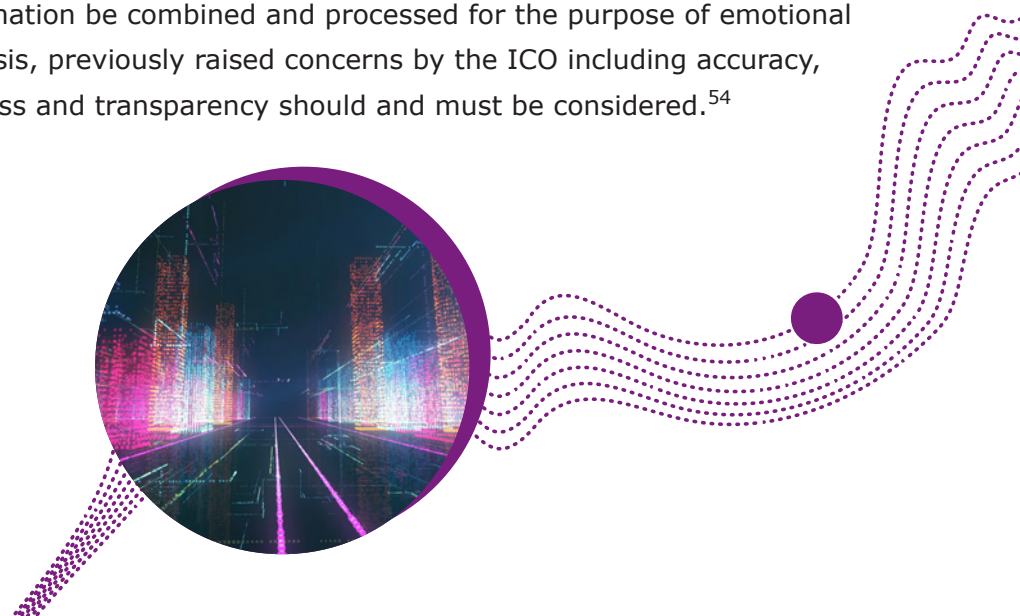

51 Link to ICO guidance

52 "Watching Androids Dream of Electric Sheep: Immersive Technology, Biome" by Brittan Heller (vanderbilt.edu)

53 Frontiers | Progress in Brain Computer Interface: Challenges and Opportunities (frontiersin.org)

54 'Immature biometric technologies could be discriminating against people' says ICO in warning to organisations | ICO

Issue 2: Immersive technologies collect large volumes of personal information, prompting questions about data minimisation

Immersive technology devices require a vast amount of information to generate interactive, simulated and augmented environments. The immersion provided relies on the continuous collection and processing of personal information, which raises a significant privacy concern (similar in many ways to the privacy concerns which have been raised in relation to the use of devices within the Internet of Things ecosystem). How might the principle of data minimisation be applied to technology that is by some measures 'always on' and therefore arguably always collecting and processing personal information?

Devices have many sensors and collect information very frequently (often up to 60 times per second). Immersive technology therefore has the potential to gather information about how a person moves, speaks and looks, on an unprecedented scale. Organisations should:

- consider the need to minimise information collection to what is required for the purpose of the processing; and
- take steps to ensure that they retain personal information for no longer than is necessary.

Issue 3: Consideration needs to be given to how to provide transparency for the designated user

As with all emerging technologies, people need to first be educated on how the device collects personal information to fully understand the privacy implications and what they are agreeing to. The use of lengthy 2D privacy policies in a 3D environment may not be the most engaging or effective way to provide this information.

Further, as these devices may be shared across multiple users, it is important to ensure privacy information and settings are appropriately targeted to all users. For example, children are a key audience for this technology. For products and services which come within the scope of the Children's code, developers will need to tailor content and parental controls to ensure children can engage with it safely.

Importantly, the code provides for in-scope services to have 'high privacy' settings by default (unless there is a compelling reason not to). All organisations developing immersive tools for children need to keep this requirement at the front of their minds. Design features and software prompts should also encourage informed consent.

Issue 4: Concerns exist about lack of transparency for third parties whose personal information may be collected

When used in public spaces, immersive technology devices may both deliberately and inadvertently monitor people who are in close proximity to the device. As devices become more discreet (eg AR head mountable units like 'smart glasses') some commentators anticipate that they will have the potential to become capable of an unprecedented level of covert surveillance.

While covert surveillance can currently be conducted using body-worn video cameras (BWV) and mobile phone applications, BWV processing is typically carried out by an organisation such as law enforcement or private security. Such organisations will be operating under data protection legislation that obliges them to demonstrate how the processing is necessary and proportionate, and how it meets the principle of transparency. Our guidance requires such devices to provide sufficient private information to individuals before using BWV such as clear signage, verbal announcements, lights or indicators on the device and have readily available privacy policies.[55]

The use of AR headsets by someone acting in a private capacity may fall outside the scope of the existing regulatory framework and consideration will need to be given to how individuals' information rights will be protected when these technologies are operated in public. Novel solutions may include:

- location-based notifications prompting users to minimise or cease recording in sensitive locations such as hospitals; or
- embedding redaction technology within devices that will blur or mask parts of the footage by default.

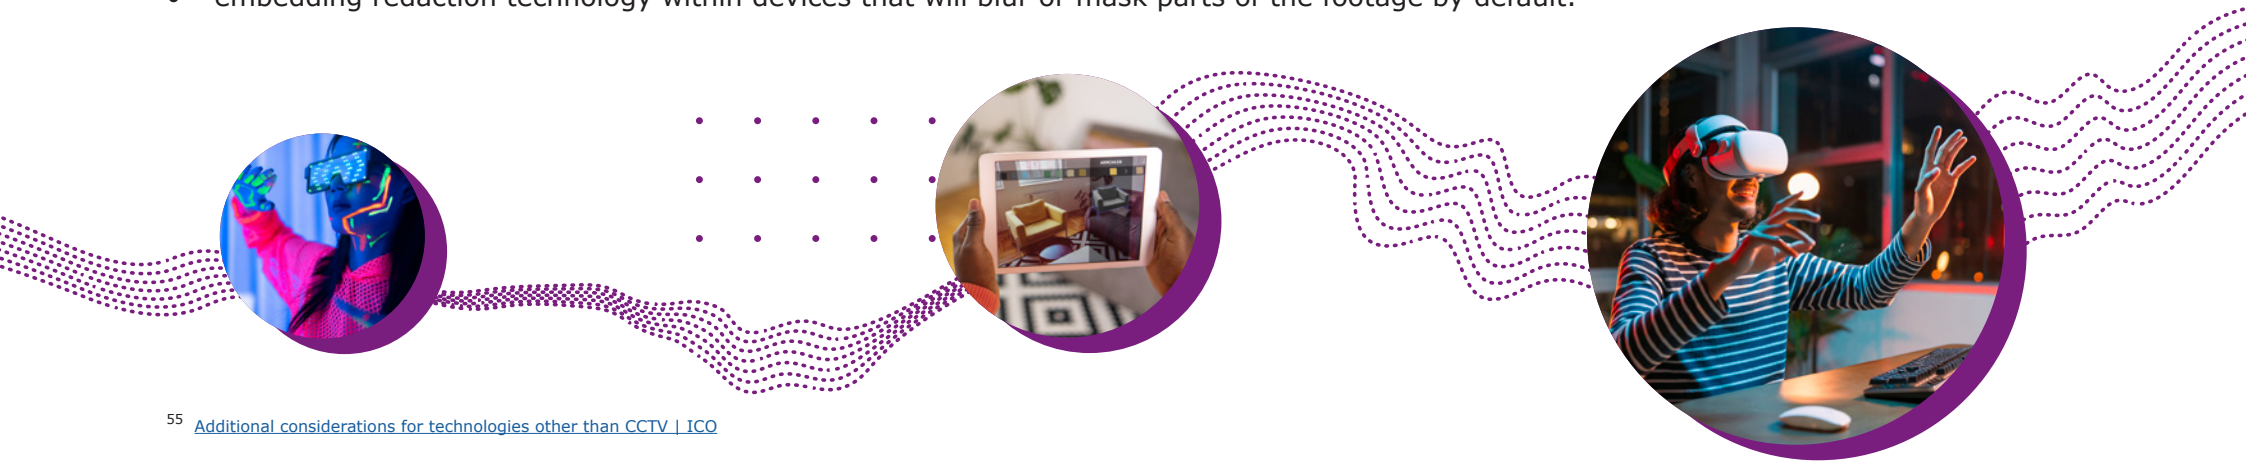

[55] Additional considerations for technologies other than CCTV | ICO

Conclusion

As public adoption of VR and AR grows, new privacy concerns may emerge. These include the prospect of targeted profiling in an environment where details of eye gaze and heart rate can be collected to make inferences about interest levels and, more generally, the collection and processing of significant volumes of potentially special category data.

Embedding privacy by design in infrastructure, policies and standards early on will be critical to ensure people's rights to privacy are safeguarded in the next generation immersive internet.

AR and VR devices are likely to process personal information of those who choose not to use these devices. Those that do engage with them are likely to be volunteering significant quantities of special category personal information to multiple parties. While the privacy implications are not always novel, they will exacerbate increasing privacy risks, if they are not designed and implemented in a privacy-positive manner. Businesses developing these technologies should explore technical and policy solutions at the same time. This will maximise the opportunities that AR and VR technology present, and manage the privacy challenges and risks they pose.

Decentralised finance

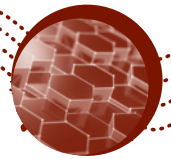

Decentralised finance (DeFi) refers to financial systems that remove centralised intermediaries from transactions and financial products and services. It enables continuous and independent access to financial services such as lending, borrowing and trading, without processing via banks, traditional exchanges or hosting providers. In these systems, control is both decentralised and distributed, and personal information is often permanently embedded in public transaction records. These features raise important challenges for transparency and the exercise of information rights.

At present, DeFi only represents a tiny share of total financial services markets (less than half a percent), and remains volatile. However, the potential for growth and the novel architecture underpinning DeFi services may create significant implications for people's privacy, as well as ownership and control of personal information in the near future.

What is DeFi?

Traditional centralised finance involves third parties moving money between consumers and merchants, primarily for a fee. These third parties, often banks, verify and approve the transaction, providing a layer of trust in the system. In contrast, DeFi employs software in the form of distributed ledger technology (DLT) to establish trust. This enables peer-to-peer transactions without a financial institution intermediary. Blockchain, the most widely-known application of DLTs, uses multiple interconnected database copies (or 'instances') to process transactions. These databases might be in a number of different physical locations, but are linked across a network, appearing to be a single database.

Information processed by these distributed databases is recorded in real-time within 'blocks' and verified by consensus. This consensus is usually reached by multiple people 'mining' (solving) cryptographic problems to achieve an agreed result (called 'proof of work'), or put into the hands of those who have the most assets stored on, and therefore most investment in the security of the chain (called 'proof of stake'). If the transaction is validated, the block is written to the ledger. It is this scrutiny of financial transactions that reproduces central authorities' traditional role of verification. Once validated, the process creates another block that has information about the previous block within it. This forms an unalterable 'chain' of transaction information. In a public blockchain it is this transparent, permanent and unchanging log of transactions that creates trust between everyone using the service.

Some blockchains use 'smart contracts' which are software programs triggered when predetermined conditions are met (eg a payment is executed automatically upon delivery of goods). As long as all contractual conditions are met, smart contracts provide immediate certainty for everyone involved in a transaction, without using an intermediary such as a bank.

In addition, often the governance of DeFi services and applications are distributed and autonomous, as is the case in decentralised autonomous organisations (DAOs). A DAO is an entity with no central governing body, that uses the blockchain to distribute decision-making power across potentially thousands of 'token holders'. Token holders are people with ownership in the decentralised protocol.

The evolution of DeFi and decentralised architectures could lead towards what some call 'self-sovereign systems'. That is, people able to be their own bank, master their own information or manage their own identity. From a personal information standpoint, this would allow people to not just manage their own information but also choose how to monetise it. Such a future model could be empowering, but also puts near total responsibility for managing our information on ourselves.

Collection and processing of personal information

Information collected and processed within DeFi usually describes transactions of a financial nature. This typically involves transferring funds between multiple parties.

However, it is possible to exchange other types of information, by sending transactions through smart contracts. The information collected and automatically processed within a smart contract will be dictated by the terms of the contract itself. Self-executing smart contracts will typically record at least the transaction addresses of the two parties and the value of the transaction. Additional information may be processed if it is captured by the front end of the decentralised application or consumed via an API from a traditional data source. That information might include:

- information about the smart contract itself;
- the IP addresses of parties, account balances and transactions;
- ownership of tokens (including DAO membership) and
- encrypted messages or text.

The customisable nature of smart contracts therefore means that any information available to parties could be stored, if written into the contract.

Where this information relates to an identified or identifiable natural person, it will be personal data.

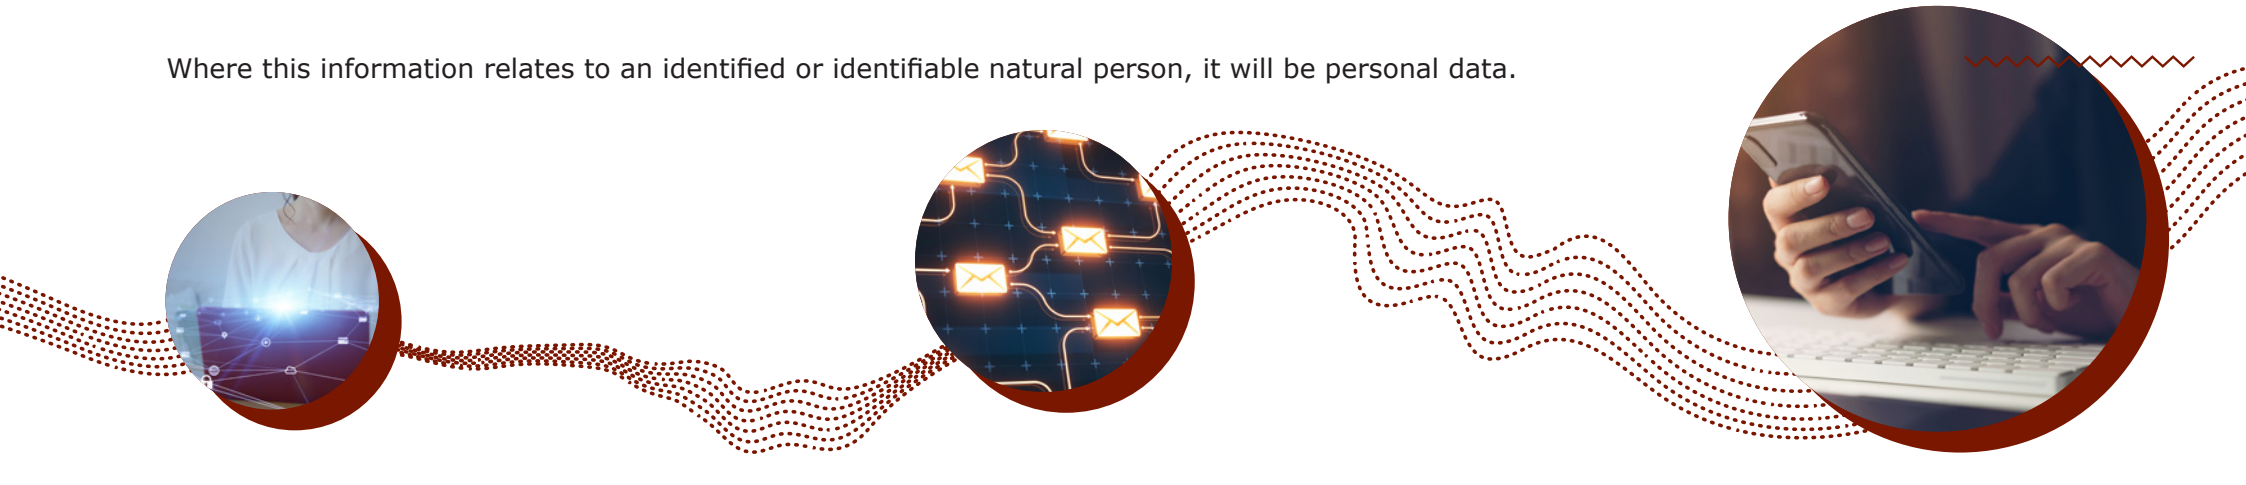

Decentralised futures

Connecting the dots

Organisations may use DeFi systems for a diverse range of transactions. The appeal includes:

- perceived benefits to security, accessibility and a lack of centralised control; and
- time and labour efficiencies from automatically executing smart contracts integrated with the blockchain.

For example, a trade union might employ DeFi to process its subscription fees. The union can create a transaction address for members to pay their dues, and publish this address on their public website for ease of access. This address is now intrinsically linked to the trade union, and it is for the purpose of membership fees. As a result, any other account transacting with it (particularly if there is also a known fee amount used), can be inferred to be that of a member.

Since anyone with access can inspect the chain and transactions, it is relatively simple for a bot or script to harvest information about the accounts transacting with the membership address.[56] This information can be combined with other sources to establish that someone is a union member. The information might be an external source, such as:

- the originating IP address of the transaction correlating with someone's home connection; or
- other transactions made by that wallet that could collectively be used to identify the person.

It might even be as simple as the person having their transaction address in their social media profile.

Even though the transaction is anonymous on one side, the fact that it is fully transparent and related to union membership has allowed special category data to be determined.

[56] This technique is more commonly used in DeFi to copy an experienced trader's transactions.

Personal, permanent and decentralised

In a possible future, DAO-managed DeFi apps could be used for everyday payments. These might include paying to stream a movie or buying a bus ticket home. However, the complexity of how the blockchain works may create a lack of understanding. This means some people may fail to realise that:

- these transactions are viewable by anyone with access to the chain;
- the transactions are stored permanently; and
- the more transactions they make, and the more information that's added and stored in the chain, the easier it may become to identify them.

People may become concerned about what this information might reveal, and seek to understand who has access to any personal information now stored on the blockchain the app runs on. In this case, they may have to identify who the accountable party is. Without a clear single responsible entity, people might be told that:

- no-one is accountable;
- it's the responsibility of all token holders collectively;
- it's actually the public blockchain the app is built on who is responsible (but as this is also a decentralised organisation, the question remains unanswered); and
- they themselves are responsible for their own information.

Alternatively, a DAO may nominate a single token holder as the point of contact for these matters. However, even if someone could identify an accountable party, it's unclear what might be done to erase any personal information. This is due to the permanent, unchanging nature of the blockchain.

Data protection and privacy implications

Controllership

UK data protection law specifies that the entity responsible for determining the purposes and means of processing personal information is the 'controller'. A controller processing personal data has obligations under UK GDPR and these include transparency, lawfulness, access, security and maintaining privacy standards. However, the decentralised nature of networks in the crypto-asset industry raise questions about who may be the controller (or joint controller).

This issue may be particularly acute in certain environments. For example, DAOs have no central governing body, and use blockchain to distribute decision-making power across token holders.

The operation of the DAO is defined within smart contracts. Token holders within the DAO vote on changes they perceive as being in the best interests of the whole organisation. If token holders are considering DAO policy about processing personal information, any participant could play a role in determining how and why it is processed.

The number of people voting in a DAO can be in the thousands or more. In effect, this could create a large set of joint controllers who share accountability for their processing of personal information. This can create several practical challenges:

- the person using the service the DAO provides might struggle to determine whom to contact to exercise their information rights;
- the DAO may find it difficult to ensure that all parties understand their regulatory obligations; and
- regulators may encounter challenges in engaging with or enforcing against non-compliant entities.

This forms part of a broader discussion about liability in decentralised structures.[57] The Law Commission is actively considering this question.[58]

Similar issues are encountered when considering who may attract obligations as a 'processor' under UK GDPR. A processor is the organisation acting on behalf of, and only on the instructions of, the relevant controller or joint controllers.

Additionally, there may be different data protection responsibilities for each service operator used in the DAO (ie, the entity responsible for the underlying public blockchain, smart contracts, or relevant cryptocurrency exchange). Depending on how the DAO is set up, the operators may be a controller, joint controller or processor.

[57] Calculations on identifying controllership may change when considering private, permissioned DLTs.

[58] Law Commission seeks views on decentralised autonomous organisations (DAOs) | Law Commission

International data transfers and applicable legislation

For many, an attractive feature of decentralised structures is its advertised inclusivity and global reach. These structures allow people to move crypto-assets wherever they want, including across jurisdictions. However, UK data protection legislation contains rules about transfers of personal data to receivers located outside the UK. People's rights in relation to their personal information must be protected or one of a limited number of exceptions must apply. Like any entity that spans jurisdictions, decentralised organisations will need to navigate compliance in all jurisdictions they operate in.

Where personal information is involved in transactions, those responsible for processing in decentralised systems will need to put in place mechanisms to ensure any international transfer of personal information comply with applicable data protection law.

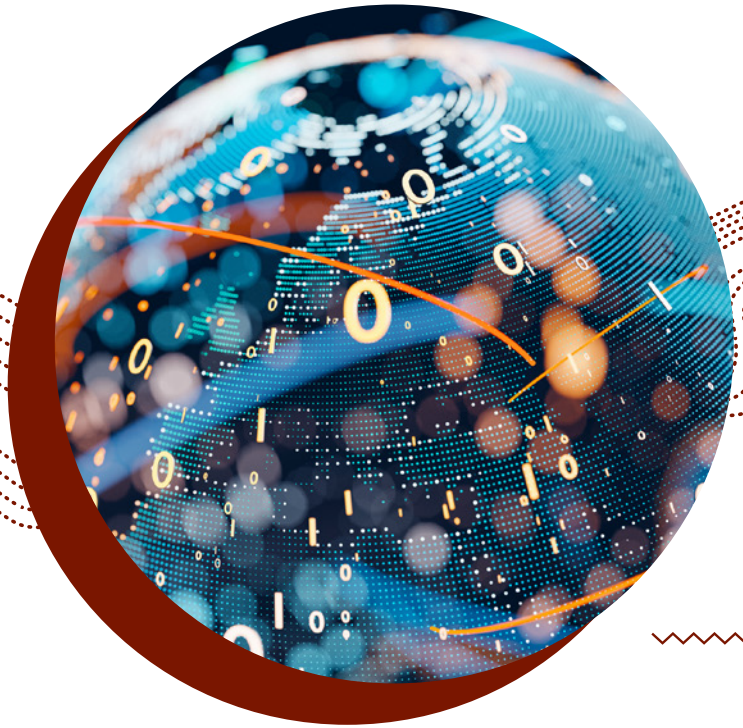

Transparency and identification of people

Anonymity of blockchain transactions is often held up as one of the technology's privacy strengths (although financial regulations often require large crypto exchanges to link an identity to an account). While blockchain transactions may not involve disclosing certain categories of personal information (eg an individual's name), that does not eliminate the possibility that people could be identified. This is because information recorded as part of blockchain transactions may be pseudonymised rather than strictly anonymous. Pseudonymised information does not directly identify people, but their identity can be determined by using additional information.[59]

Pseudonymised information is still personal information. There may be cases where it is reasonably likely that a link can be established between the information processed by a blockchain system and an identified or identifiable person. This information is likely to be personal information and falls under UK data protection law. It will still be personal information even if it does not enable the 'real world identity' of that person to be determined.

The risk of re-identification increases with the volume of information stored on the blockchain. As data points are progressively added through additional transactions, an increasingly detailed view of the wallet holder could be constructed. This might include their personal preferences, or behaviours and attitudes, for example. The risk is likely to increase as DeFi payments transition to paying for real-world goods and services. This increases further if a third party has additional information such as the time of a transaction, its value and the organisation a wallet holder transacted with. This is a risk that operators and policymakers need to be mindful of as associative analytical tools that could establish these links grow in sophistication.

Decentralised app developers must comply with their transparency obligations under data protection law and in so doing help people understand how their information is being processed. This will require explaining what information would be collected and processed during smart contract transactions, and that information added to a chain would be available for any third parties with access to use. Compliance with transparency obligations will allow people to make a more informed choice about the processing of their personal information.

[59] See ICO Draft Guidance on Anonymisation, Pseudonymisation and Privacy Enhancing Technologies, Chapter 3: Pseudonymisation.

Exercising information rights

The permanent, unchanging nature of storage on blockchains means that any information can be retained indefinitely. In cases where personal information is stored on the chain, that permanency raises potential concerns for the ability of people to exercise their information law rights.

The UK GDPR's principle of storage limitation says that personal information should not be retained for no longer than necessary to achieve the purposes of processing. This is a requirement that might be difficult to reconcile with indefinite retention on the blockchain. This same feature of blockchains may make it difficult for someone to exercise their right to erasure or rectification. For example, in cases where inaccurate information has made it into the ledger or they want to remove their personal information from the chain. As each block in a chain contains a cryptographic hash of the previous block, any intervention to change or remove the content of a given block would disrupt the chain and undermine the trust mechanism. It is likely therefore that storage limitation, rectification and erasure obligations may present particular challenges in current popular applications of the technology.[60]

Resilience and security

There have been several well publicised breaches, exploits, hacks and compromises on DeFi apps and infrastructures. In fact, the novel nature of decentralised finance has brought novel attacks. These include flash loan attacks, price manipulation and oracle attacks, and attacks on the process validating transactions. This documented timeline of DeFi hacks and compromises suggests that a total value of over $3bn has been stolen to date. Whilst the headline is the amount of currency that was stolen, compromise of DeFi organisations may mean that personal information held by those organisations is also exposed.

As a decentralised and distributed technology, DeFi can be more resilient and harder to attack at a macro level than some traditional applications with certain attack types. For example, a Distributed denial of service (DDoS) attack can be harder if there is no one main data processing point or there is a financial cost to execute transactions on a network. This is not necessarily the case at a micro level. If a person's security is insufficient, then their assets may be at risk. Even if accounts are secured through two-factor authentication, new attacks have been developed, such as sim-jacking.

[60] The ICO is aware that some technical approaches are being explored that may facilitate change or erasure without corrupting the chain.

Conclusion

The qualities of DLT that enable transparent, permissionless and permanent processing beyond centralised control structures also present clear challenges for privacy and compliance with data protection law. Obligations relating to data protection by design and default require organisations to implement privacy standards and information rights in the systems and services they develop. This applies to blockchain in the same way as to any other technology.

Encouragingly, there are already organisations developing privacy-positive capabilities to address data protection concerns. For example, services like privacy mixers might protect against re-identification in decentralised environments (although these present their own challenges around identification of controllership and the introduction of a new intermediary). Elsewhere, infrastructure organisations are researching implementation of privacy enhancing technologies. For example, 'zero knowledge proofs' ensure that personal information is kept off-chain.[61] By keeping the personal information off chain, but proving it exists, transactions can still take place without concerns around data protection compliance in a DLT-based decentralised governance structure. However, we are aware that some of these solutions are viewed by some as reducing the usefulness of the service. We therefore encourage organisations to continue to work to identify solutions that maximise privacy protections without compromising utility and effectiveness.

[61] See ICO draft anonymisation, pseudonymisation and privacy enhancing technologies guidance chapter 5

What’s next

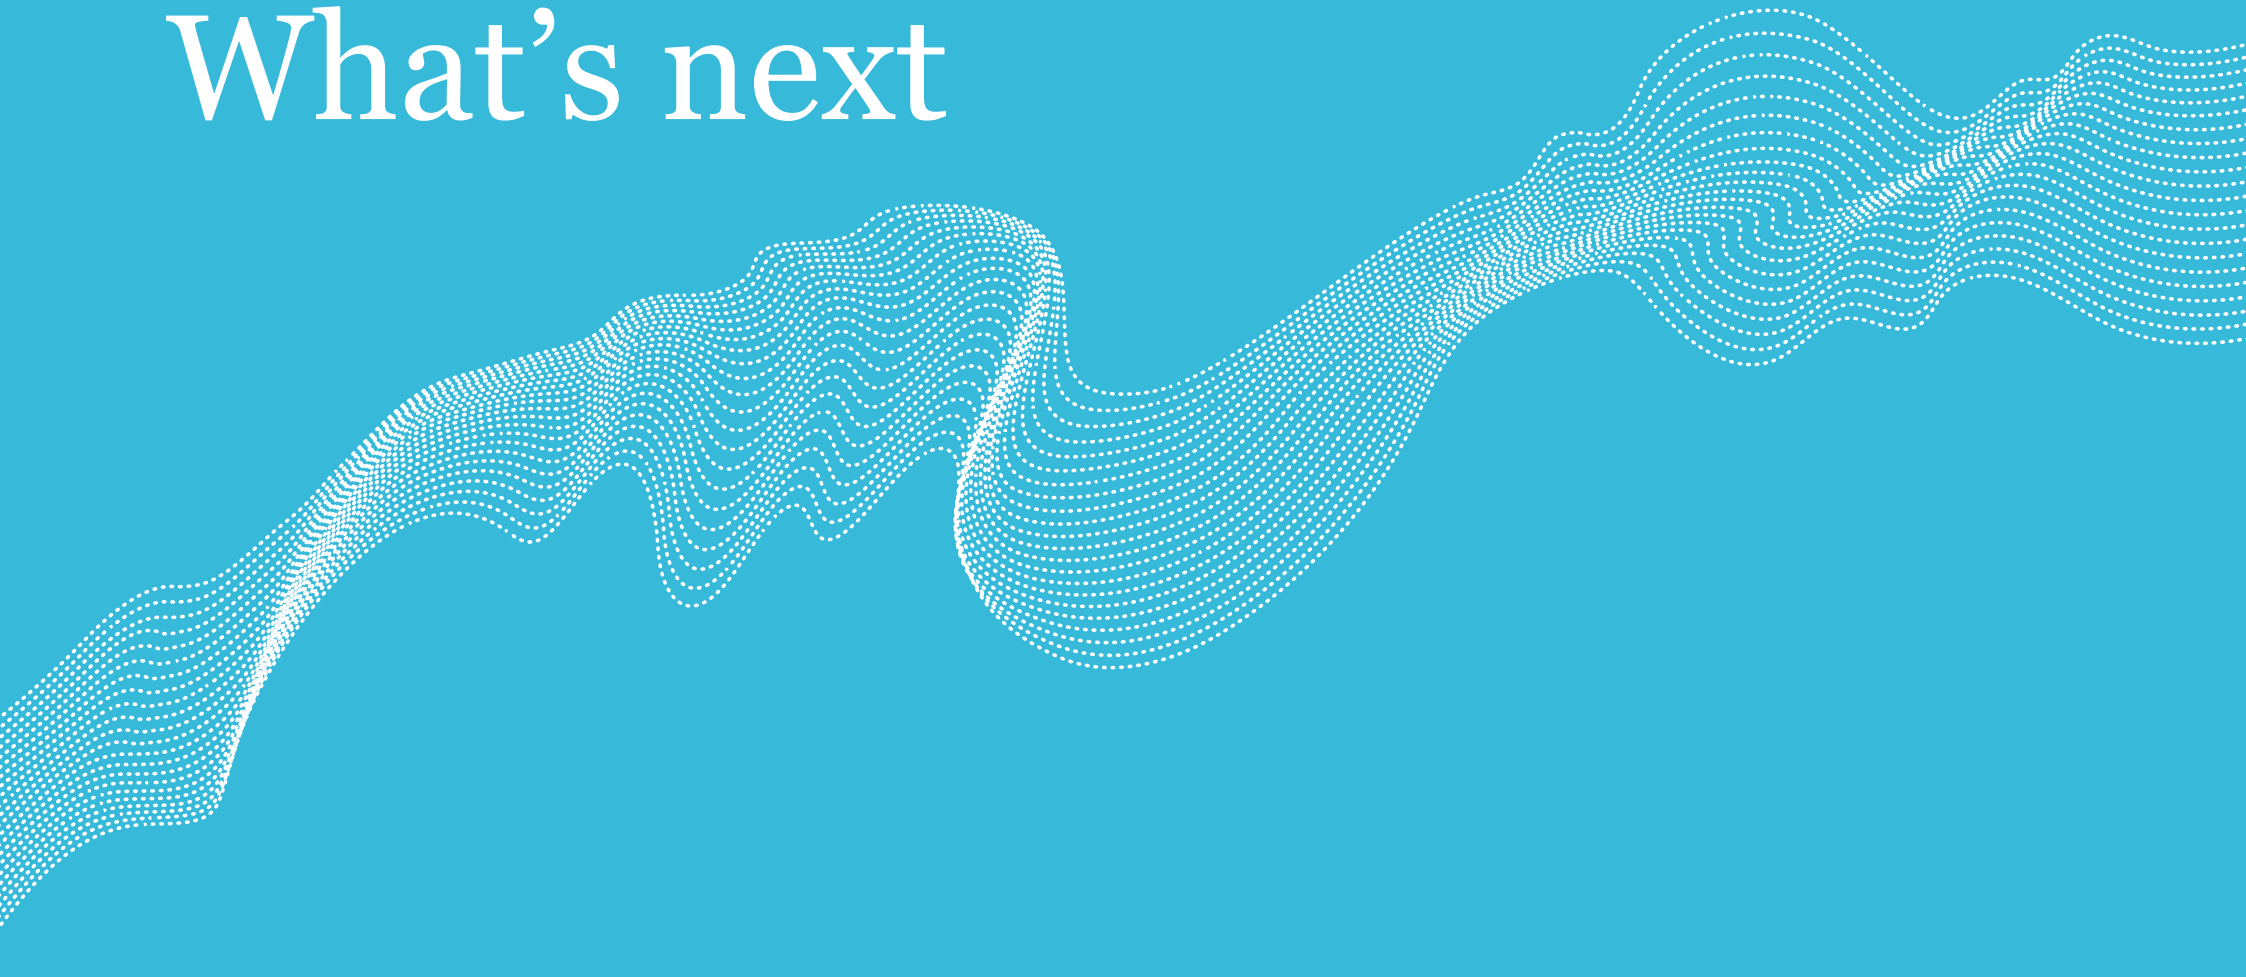

To prepare for our role in regulating these technologies, we will:

- work with the public about the benefits and risks of these emerging technologies and how we will approach them as a regulator;
- invite organisations to work with our Regulatory Sandbox to engineer data protection into these technologies;
- develop guidance for organisations where needed, starting with guidance on data protection and IoT; and
- proactively monitor market developments so we can take action on uses that cause concern.

We will assess the technologies covered in this report, and many others in our annual horizon scan. Our scan will ensure we can anticipate technology trends and any associated impacts on privacy.

Following the publication of our report into biometric futures in October 2022, we will publish our second emerging technology deep-dive on neurotechnology in the first half of 2023. It will consider:

- novel questions around processing subconscious neurodata; and
- applying the UK GDPR to brain-computer interfaces.

We are committed to building on the analysis in this report to help developers of emerging technologies understand the relevant data protection considerations. We will continue to work with stakeholders to explain the importance of privacy by design and how to use personal information compliantly.


ico.
Information Commissioner's Office