

Tech Horizons Report 2024

How new technologies interact with
the UK's data protection framework



ico.

Information Commissioner's Office

Executive summary	3
Introduction	6
Genomics	9
Immersive virtual worlds	14
Neurotechnologies	19
Quantum computing	24
Commercial use of drones	30
Personalised AI	34
Next-generation search	39
Central bank digital currencies (CBDCs)	44

Executive summary

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 harness the benefits innovation has to offer. By identifying the privacy and data protection implications of emerging technologies before they are widely used, we are better placed to proactively set out our regulatory responses and enable innovators to consider these challenges in the design-phase of development.

In our first annual Tech horizons report, published in 2022, we highlighted the implications of four of the most significant technological developments for privacy in the next two to five years.

In light of continuing rapid technological advancement, this second edition considers a further eight technologies we believe may have a particularly significant impact on our societies, economies and information rights in the next two to seven years.



Eight priority technologies

- **Genomics:** The sequencing of the human genome to improve understanding of a broad range of traits, mostly used in healthcare. There is further potential for these insights to be used in fields such as employment, sports and education.

- **Immersive virtual worlds:** Highly immersive virtual environments, also known as the metaverse, in which users can interact with each other and make use of digital services, such as e-commerce and gaming.
- **Neurotechnologies:** Consumer, enterprise and healthcare devices and procedures, both invasive and non-invasive, that directly record and process neuro-data to gather information, control interfaces or devices, or modulate neural activity.
- **Quantum computing:** By taking advantage of phenomena at the atomic scale, quantum computing may in future be able to resolve highly complex computational problems that current computers cannot, but may also present serious risks to existing encryption.
- **Commercial use of drones:** Unmanned aerial vehicles (UAVs) used in commercial settings, for example in delivery and e-commerce, monitoring, and crowd control.
- **Personalised AI:** The customisation of large language models, based on individual users' search patterns and personal preferences and characteristics, to create more tailored user experiences and better-targeted outputs.
- **Next-generation search:** Next generation search engines incorporating new technologies, such as embedded AI capabilities, as well as voice-based, image-based and ambient elements.
- **Central bank digital currencies (CBDCs):** A new form of central bank-issued digital money, which would complement physical cash and other payment mechanisms in facilitating everyday payment needs.

Each of the eight technologies discussed in this report presents their own unique opportunities and challenges. They could improve our health, our social interactions, our environment and how we do business. However, we have also identified a number of overarching trends and dynamics that will need to be addressed as the technologies develop:

- **New innovations across fields like neurotechnology, immersive technologies and genomics allow organisations to collect novel types of intimate information about people.** As innovators seek to harness the significant benefits these new technologies may offer, it is important they build in appropriate safeguards from the outset.
- **Many new technologies collect and process increasingly large amounts of personal information to better personalise experiences.** The scale and complexity of this processing makes it increasingly difficult for people to understand how organisations are using their information and what the combining of different pieces of information may reveal about them.
- **We have seen an accelerated pace of innovation in the past 12 months, especially in the field of artificial intelligence.** The increased use of AI in other technology domains, from neurotechnology to search, has resulted in a rapid increase in new uses and developments that we, as a regulator, need to keep up to date with.
- **Multiple organisations are involved in processing people's information, leading to a lack of clarity.** This makes it harder for people to understand who is processing their information, and how they can exercise their information rights. Being transparent will be critical to maintaining people's trust.
- **The misuse of several of the technologies considered in this report could particularly impact people who already need extra support to protect themselves.** Organisations should ensure they put appropriate safeguards in place. New solutions should also not introduce new types of harm or discrimination.

Many organisations continue to build data protection by design into their innovations. Others may need

guidance, or to work with the regulator, to consider how they can engineer privacy into their ideas.

We will address issues proactively as these technologies mature and our role in regulating them develops.

We will:

- involve the public in decisions about how we, as a regulator, address the risks and benefits these new innovations may present;
- continue to invite innovators to work with our [Regulatory Sandbox](#) to engineer data protection into these technologies from the outset, focussing on the most innovative propositions;
- share further insights on priority technologies by developing foresight reports, starting with quantum and genomics in 2024;
- continue to work closely with our partners at the [Digital Regulation Cooperation Forum](#) (DRCF) to act on shared regulatory challenges and opportunities; and
- continuously scan the horizon for further developments that require our immediate attention and possible intervention.

Introduction

Much has happened since we released our first Tech horizons report in December of 2022. 12 months ago, few could have foreseen the sheer speed and breadth of developments that would take the artificial intelligence space by storm, nor the spillover effects these advancements would have on other areas of innovation. Technologies do not develop in isolation, but are shaped by constantly evolving economic, political, technological and societal forces. Responding to the speed and complexity of this constant change in the technology realm is a well-known challenge for regulators, policymakers, and indeed, innovators too.

By identifying the privacy and data protection implications of new technologies before they are widely used, we are better placed to proactively set out our regulatory responses and enable innovators to consider these challenges already in the design-phase of development.

The first edition of the Tech horizons report covered four emerging technologies in depth: the Internet of Things, immersive technologies, health-tech applications and decentralised finance (DeFi). We will consider eight further technologies in this second edition that we believe may see significant adoption in the next two to seven years:

- **Genomics:** The sequencing of the human genome to improve understanding of a broad range of traits, mostly used in healthcare. There is further potential for these insights to be used in fields such as employment, sports and education.
- **Immersive virtual worlds:** Highly immersive virtual environments, also known as the metaverse, in which users can interact with each other and make use of digital services, such as e-commerce and gaming.
- **Neurotechnologies:** Consumer, enterprise and healthcare devices and procedures, both invasive and non-invasive, that directly record and process neuro-data to gather information, control interfaces or devices, or modulate neural activity.
- **Quantum computing:** By taking advantage of phenomena at the atomic scale, quantum computing may in future be able to resolve highly complex computational problems that current computers cannot, but may also present serious risks to existing encryption.
- **Commercial use of drones:** Unmanned aerial vehicles (UAVs) used in commercial settings, for example in delivery and e-commerce, monitoring, and crowd control.
- **Personalised AI:** The customisation of large language models, based on individual users' search patterns and personal preferences and characteristics, to create more tailored user experiences and better-targeted outputs.
- **Next-generation search:** Next generation search engines incorporating new technologies, such as embedded AI capabilities, as well as voice-based, image-based and ambient elements.
- **Central bank digital currencies (CBDCs):** A new form of central bank-issued digital money, which would complement physical cash and other payment mechanisms in facilitating everyday payment needs.

Each chapter briefly:

- explains the technology behind each topic as well as its emerging use cases;
- sets out the current state of play surrounding its adoption and development; and

- discusses possible future privacy and data protection risks these developments may present.

These chapters represent our early view on often highly uncertain, evolving technology areas. You should not consider the data protection and privacy issues we've explored, and the recommendations we've set out in this report, as formal guidance. They do not necessarily reflect our current or future policy positions.

Our work on emerging technology

Our emerging technology team was established in late 2021, with the goal of bolstering our capacity to proactively respond to emerging and future innovation. This is our second edition of our annual flagship Tech horizons report. We have also released dedicated foresight reports on **biometrics** and **neurotechnologies**, which delve deeper into the future regulatory considerations surrounding these two key technologies. We also work with our partners in the Digital Regulation Cooperation Forum (DRCF) to prepare reports on **web 3.0**, **quantum** and **immersive technologies**.

We take our work on emerging technology and foresight seriously. By mapping out of possible future developments, we can better set our strategic priorities today. Insights derived from publications such as this also allow us to support innovators to address emerging data protection risks and opportunities early. All four technologies we identified in the first edition of the Tech horizons report are now the focus of a call for applications into our Regulatory Sandbox. The innovation landscape is vast; it can therefore be challenging to separate the signal from the noise. By practising foresight, we are better able to determine which areas of technological innovation require our attention. Equally, we can decide which ones may not yet have reached sufficient levels of maturity to warrant our immediate action.

Technology selection

The selection and evaluation process behind the eight technologies we've explored in this report was guided by a robust, year-long horizon scanning and foresight cycle. This cycle consisted of four consecutive phases:

- **Phase 1: Horizon scanning**

This initial horizon scanning phase primarily focused on identifying a longlist of emerging technologies for us to consider for inclusion in the report. This process relied on a so-called "scan-of-scans" approach, in which we collated and contrasted technologies and emerging trends explored in recent foresight and technology publications. Our scan considered a large number of high-quality sources from a wide range of countries and stakeholder communities to ensure a diversity of perspectives (including civil society, academia, the private sector, government, think tanks, media, consultancies and peer regulators). This exercise identified more than 90 possible technologies.

- **Phase 2: Prioritisation**

In the second phase, we narrowed down this longlist of technologies to the final eight priority topics. This shortlisting exercise involved a rigorous evaluation and stress-testing process, by which we scored potential technologies on the basis of several qualitative and quantitative indicators.

These indicators aimed to surface:

- The magnitude and novelty of the possible privacy risks associated with an emerging technology (with an emphasis on harms disproportionately affecting those groups most at risk of harm, and the potential processing of special category information).

- The expected maturity and market penetration of an emerging technology over the next two to seven years.
- The degree and pace of innovation driving an emerging technology’s development (with a preference for technologies that have undergone significant change over the past 18 months across a range of sectors).

- **Phase 3: Evaluation**

During the third phase, we explored each of the eight priority topics in more depth, and stress-tested conclusions about possible future privacy and data protection considerations with external experts.

- **Phase 4: Scenario-building**

The fourth and final phase explored the possible trajectories our priority technologies may follow in the years ahead. We developed scenarios which aimed to shed further light on the privacy and data protection implications these different possible futures may bring. We carried out different scenario-building exercises to not just map the most direct emerging impacts and use cases of a technology, but also the more speculative second-order impacts they may generate (“don’t predict the car, but the traffic jam”). Each of the eight chapters includes an example box describing what people’s day-to-day interactions with a new technology might look like in the future.

Genomics

Introduction

In recent decades, advances in genomics (the study of the genome) have greatly increased understanding of the entire DNA sequence, as well as the interactions between genes. The genome includes genes, which only make up 1-5% of a human's genome, and all of a person's other DNA¹. The field involves a wide range of different technologies and techniques that have continued to evolve rapidly since scientists first sequenced² a human genome in 2003.

These technologies pose a range of data protection issues. In addition to genomic sequencing, one technique of particular interest in this chapter is polygenic risk scoring. This looks at the potential impact of many genomic markers to estimate "an individual's genetic risk for some trait" or disease.³

In healthcare, advances in genomics are moving us towards a possible future of predictive and personalised treatments for complex diseases. These are based on a person's genome, and clinical use of disease risk scores drawing on genomic insights and other information.

Outside of healthcare, future applications of more sophisticated genomic insights remain speculative. However, they could extend beyond existing applications for current, more limited tests (such as many popular direct-to-consumer genetic tests which only analyse specific sections of DNA)⁴. It is still highly uncertain what depth and quality of insight future analysis of genomic information may reveal. Solutions that consider health and behaviour are however already under development. This brings to the surface growing concerns about issues around accuracy and fairness, as well as the privacy impacts for newborns and relatives such insights may present.



About genomics

As costs fall and sequencing technologies continue to improve, it is now possible to sequence and analyse the whole human genome of large numbers of people. Larger and higher quality datasets are available for analysis using increasingly advanced machine learning and deep learning techniques. In future, genomic insights may also be combined with further medical and environmental information to provide a more complete picture of a person.

As a result, wider understanding of how a person's genome can influence polygenic traits (traits that are influenced by multiple genes) is improving fast. Polygenic traits include physical characteristics such as height or risk of diseases such as cancer, but also behavioural characteristics and skills. For example, recent studies have aimed to quantify the extent to which certain behavioural traits are influenced by the genome. This includes susceptibility to certain substance abuse disorders (42-79%), conscientiousness (44%) and extraversion (53%).⁵ In 2022, Nature published the largest study ever conducted on genetic associations with "educational attainment" using a direct to consumer (DTC) genetic dataset of three million records.⁶ These studies are drawing complex, uncertain and contested inferences through the use of highly sensitive information. Making inferences about people from genomic analysis heightens the risk of inappropriate bias and discrimination.⁷

State of development

Most genomics research, investment, and market growth focuses on further developing underlying technologies and healthcare uses (primarily in diagnostics, drug development and precision medicine). Genomics in medicine is "already saving lives",⁸ for example by enabling diagnosis and treatment of rare conditions.

Some predict AI-powered genomic drug development and personalised medicine will play an increasingly important role in healthcare in five to 10 years.⁹ The UK has several major genomics projects and is aiming to become the first nation to "offer whole genome sequencing as part of routine care".¹⁰ The US and China are among other countries with significant genomic industries.

Outside of healthcare, direct-to-consumer (DTC) genetic testing is a well-established market. Providers offer tests for purposes such as:

- tracing ancestry;
- "polygenic scoring" for insights into wellness (eg, genetic markers of how much alcohol you have consumed);
- predisposition to disease; or
- skills such as recognising musical pitch.

Some third-party services also offer personalised health or fitness recommendations based on DTC tests.

Currently, investment in non-healthcare applications of genomics is still very limited.¹¹ However, as insights expand and technological underpinnings continue to improve, we may see novel use cases for health or behavioural polygenic risk scores in non-healthcare applications emerge in future. The Government Office for Science has flagged potential applications, such as interventions in early-childhood education, sports nutrition, and health or personality screening in employment. Such use cases are not yet scientifically validated and should be treated with caution. Timescales for potential developments also remain unclear. There are however examples of novel patents and tests, such as an overseas law enforcement trial using

genomics to predict the physical features and gender of suspects.¹²

There could also be future uses for genomic insights in insurance to determine the risk of disease. In the UK, a voluntary agreement between insurers and government limits the use of genetic information to very specific circumstances. However, these limitations can be revised in response to market changes and technological developments,¹³ so it is possible insurers may seek more access to genomic data in future.

Fictional future scenario

Nadia, 10, is a gifted, prospective professional athlete. Her genome was sequenced at birth to inform her future healthcare. Her parents share this information with a firm offering polygenic risk scoring for elite training and nutrition management. This then provides a tailored plan linked to a third-party app. The firm notes that the scores are an estimate, but don't provide much information about its scoring model. Nadia's parents share the plan and scores with her coach. They don't know that diversity gaps in the training information means the model significantly overestimates the significance of certain traits. This affects the validity of Nadia's training plan.

Under the contract, Nadia's parents also agree to the use of her anonymised genomic information for research purposes. Five years later, the firm sends Nadia's parents new insights. The scores predict Nadia is a risk taker, with an elevated future risk of depression, sleep disorders and several other traits. This worries them all. Could these new scores affect her chances of being picked for an elite training programme, if they ask for health information? How could the firm know so much more about her, if her personal information was anonymised?

Nadia's older brother Omar also isn't pleased the firm has shared Nadia's genomic information in this way – this information also reveals a lot about him. He wants the firm to delete the information from everywhere other than Nadia's medical record. Can he do this?

Data protection and privacy implications

- **Accuracy and fairness:** for some conditions, current genomic sequencing techniques can establish a clear genetic link. However, the predictive power of polygenic scoring for many other medical conditions remains limited. For example, this is due to the sheer number and complexity of markers that may be associated with a trait, and limited diversity in datasets.¹⁴


Given the current limitations of the science, the predictive capabilities of polygenic risk scoring for behavioural traits are even more contested.¹⁵ It is highly uncertain whether significantly more accurate scores for particular traits will ever become available. As we have seen in our work on biometrics and neurotech, there is a risk that we may see inaccurate (or low accuracy) applications used in future, particularly outside of healthcare. Like outputs of many other AI systems, a polygenic risk score is "intended to represent a statistically informed guess".¹⁶ As models improve, there is a risk that organisations (or people) over rely on their predictive power due to a lack of understanding or a failure to make limitations and biases clear. There is also a risk that new use cases could exacerbate power

imbalances or even lead to genetic discrimination.¹⁷

Should an organisation seek to process genomic information, or use polygenic risk scores, [they must ensure the intended processing is fair](#). That means people's personal information is being used in ways they would reasonably expect and will not have unjustified adverse impacts on them. Organisations must also ensure that any scoring is [sufficiently accurate for the purpose](#) and they explain any limitations.

- **Anonymisation and security:** Genomic data is special category information. It is difficult to effectively anonymise, given the uniqueness of the information.¹⁸ In addition, genomes may need to be stored for a long time for research purposes to obtain insights. Similar to [biometric information](#), our genomes stay with us for life, which means that the impacts of a data breach can be serious. It is hard to predict what future insights the genome may identify. Given the sensitivity of the information involved, strong security measures and controls over the use of the personal information are particularly important.
- **Third party privacy:** The disclosure and analysis of genomic information (including genomes of deceased people) also affects related family members. If uses expand significantly beyond medical settings, it could become even harder for family members to maintain control over their information. This is a complex issue, particularly when considering questions such as lawful bases for processing, and processing subject access requests for genomic information that could identify others.¹⁹

Recommendations and next steps

- We will continue to engage with and monitor this complex and rapidly evolving space. As a first step, we are developing a further in-depth tech futures report on genomics, to be published in 2024.
- We will also explore the regulatory landscape to build our knowledge and identify areas of critical intersection and collaboration.
- We will continue to remain alert to examples of misuse of genetic information, or genetic discrimination arising from current and novel uses of genomic information. Should use cases evolve further, particularly beyond medical settings, we will monitor the risk of misinterpretation or over-reliance on the predictive capabilities of polygenic risk scores.
- Employers considering genetic testing should refer to our updated employment guidance [What if we use genetic testing?](#) and insurers should refer to the [Code on genetic testing and insurance](#) .


Further reading


[Government Office for Science report on "Genomics Beyond Health" \(2022\)](#) 

[UK Government policy paper on "Genome UK: the future of healthcare" \(2020\)](#) 

[UK Government policy paper on "Genome UK: 2022 to 2025 implementation plan for England"](#) 

[Ada Lovelace Institute report on the use of AI in Genomics "DNA.I." \(2023\)](#) 

[Centre for Educational Neuroscience blog on "can polygenic scores predict educational outcomes?" \(2023\)](#) 

[Research paper on Predicting Physical Appearance from DNA Data—Towards Genomic Solutions \(2022\)](#) 

- ¹ [Genomics England webpage on Understanding Genomics](#) Genomics is distinct from genetics, which only looks at genes.
- ² Before DNA can be analysed, it needs to be sequenced – converted into the basic building blocks of DNA, represented by a series of four letters that can be read by a computer.
- ³ [Nuffield Department of Population Health article from the Frontiers journal on “Calculating Polygenic Risk Scores \(PRS\) in UK Biobank: A practical guide for epidemiologists” \(2022\)](#)
- ⁴ [Government Office for Science report on “Genomics Beyond Health” \(2022\)](#)
- ⁵ [Government Office for Science report on “Genomics Beyond Health” \(2022\)](#)
- ⁶ [Nature Genetics article on “Polygenic prediction of educational attainment within and between families from genome-wide association analyses in 3 million individuals” \(2022\)](#)
- ⁷ See, eg [Centre for Educational Neuroscience blog on “can polygenic scores predict educational outcomes?” \(2023\)](#); [Ada Lovelace Institute report on the use of AI in Genomics “DNA.I.” \(2023\)](#)
- ⁸ Gartner article on the Hype Cycle for Life Science Discovery Research, 2023
- ⁹ Gartner article on the Healthcare and Life Science CIO’s Genomics Series: Part 1 – Understanding the Business Value of Omics Data; [Ada Lovelace Institute report on the use of AI in Genomics “DNA.I.” \(2023\)](#)
- ¹⁰ [NHS England webpage about the NHS Genomic Medicine Service](#)
- ¹¹ [Ada Lovelace Institute report on the use of AI in Genomics “DNA.I.” \(2023\)](#)
- ¹² [Australian Federal Police \(AFP\) media release on advanced technology which allows APF to predict criminal profiles from DNA \(2021\)](#)
- ¹³ [UK Government corporate report on the Code on Genetic Testing and Insurance: 3-year review 2022](#)
- ¹⁴ [Genome Medicine article on “Polygenic risk scores: from research tools to clinical instruments” \(2020\)](#)
- ¹⁵ [Government Office for Science report on “Genomics Beyond Health” \(2022\)](#); [Centre for Educational Neuroscience blog on “can polygenic scores predict educational outcomes?” \(2023\)](#); [Ada Lovelace Institute report on the use of AI in Genomics “DNA.I.” \(2023\)](#)
- ¹⁶ [ICO guidance on AI and data protection: What do we need to know about accuracy and statistical accuracy?](#)
- ¹⁷ To date, no evidence of such discrimination has been reported in the UK. There have been limited cases overseas, for example where individuals have been denied insurance: [European Journal of Human Genetics article on “Genetic discrimination still casts a large shadow in 2022” \(2022\)](#)
- ¹⁸ [Ada Lovelace Institute report on the use of AI in Genomics “DNA.I.” \(2023\)](#)
- ¹⁹ [PHG Foundation report on The GDPR and Genomic Data \(2020\)](#)

Immersive virtual worlds

Introduction

The 'metaverse', next generation virtual worlds, and the '3D immersive internet' are all terms used to describe a vision for a future internet in which the physical and virtual worlds become increasingly blended. Though there are many different conceptions of the so-called metaverse, most immersive solutions being developed today focus on building virtual environments in which users interact in real-time.

Users are able to visit virtual spaces across the globe, engage with virtual objects and artefacts and explore new environments from their homes, offices or classrooms.

If implemented at scale, this technology has the potential to profoundly impact how people learn, play and work. But whilst these virtual worlds can bring many benefits, from allowing users to overcome language barriers to experiencing ideas and new worlds in an interactive and collaborative format, they may also require collecting a significant amount of personal information merely to function.



About immersive virtual worlds (“the metaverse”)

There is a current lack of industry consensus on the terminology and definitions, as well as the technology stack that might underpin future immersive virtual worlds (IVWs). We expect these solutions to be immersive, 3D-rendered digital environments, where users will interact through avatars. Although the details and use cases remain relatively conceptual at this point. The avatars will serve as digital representations of people, which may range from being entirely lifelike or cartoon in style, and will be able to move between different digital worlds and environments. For example, to visit different shops, games, or virtual representations of famous landmarks and interact with other similar avatars.

Beyond the 3D immersive element, the realisation of IVW's varies. However the vision is broadly dependent on the advancement of many technologies which will "stack" to provide a seamless experience in the future. Some of these technologies exist today, for example virtual reality devices, along with some early use cases across home entertainment and gaming. However, we expect IVW's to grow in sophistication, scale and degree of interactivity in the years to come. The technologies involved in this process may include:

- immersive technologies and integrated virtual world platforms;
- artificial intelligence (AI) including machine learning, large language models and chatbots;
- sensors including motion sensors, eye-tracking and audio sensors, haptic sensors and neural-sensors; and
- distributed ledger technologies such as blockchain and non-fungible tokens (NFTs).

Users are likely be able to access virtual worlds through a variety of mediums, from existing devices including smart phones and tablets, to specialised extended reality hardware. Different types of extended reality solutions will require different devices:

- **Augmented reality (AR)** overlays computer-generated visuals onto a user's perception of the physical world. Future AR solutions are likely continue to work through camera filters on smart phones, glasses or head-mounted displays (ie headsets).
- **Mixed reality (MR)** combines elements of AR and VR. It involves computer-general visuals that interact with the physical world beyond a simple AR overlay. For example, a user throwing a virtual ball against a wall that exists in the physical world and having it bounce back to them. Similar to AR devices, users experience MR through handheld devices or headsets.
- **Virtual reality (VR)** is made up of 3D computer-generated environments that replace a user's field of vision with total immersion, typically through a headset. This is the deepest and most "immersive" reality that can be created.

State of development

Similarly to the lack of clarity on what the metaverse is, there is some scepticism around the speed at which these visions may be realised. The barriers to bringing the metaverse to life span across the social and technological. This includes the ability to scale the number of concurrent users in virtual worlds at any one time, to the relatively limited number of use cases for immersive technologies currently available. This leads to more limited uptake when compared to other devices. These are not challenges that are anticipated to be overcome in the short term, with many industry leaders estimating the full potential of immersive futures may not be realised for another 10 to 15 years.

Fictional future scenario

Aoife is new to the world of immersive environments and is setting up her profile for the first time. She is interested in managing her data herself, so instead of opting for the more common third-party intermediary to manage her data, she chooses to use a decentralised data store she controls herself.

This way, she can pick what information she wants to share with the virtual worlds she enters – beyond the minimum required information to continue interacting in the space. A single, interoperable token (agreed upon by all major platforms) is used to provide authentication across the different virtual worlds, devices, platforms and systems.

Aoife chooses to share a high volume of her information with the platforms, as she enjoys receiving personalised recommendations for content and goods. This includes physiological information collected from her VR headset, including optical information (such as how long she looks at content and eye dilation reaction) and cardiac responses to stimulation. Platforms use this information to deduce that she enjoys horror-related content. Her streaming service begins recommending a new series within this genre.

The next time Aoife goes shopping within her virtual high street, the mannequins at the front of the shop are wearing merchandise about the new TV series. Upon purchasing, Aoife is offered the opportunity to buy an NFT t-shirt for her avatar to wear, or a physical copy of the t-shirt which she can purchase in the virtual environment. The physical t-shirt would be shipped directly to her home. Interested in the physical version of the t-shirt, Aoife uses the AR capabilities on her phone to try on the t-shirt and select the best fit so she can minimise needing to return items.

Data protection and privacy implications

Immersive technologies pose a number of significant data protection and privacy challenges. Primarily around the required collection of increasingly high volumes of sensitive information on users and transparency concerns affecting both users and (unwitting) bystanders.

- **Processing large amounts of information about sensitive human characteristics:** Immersive technology hardware collects vast amounts of information on users, including eye movement tracking, iris scans and facial movement monitoring. This may be biometric under the UK GDPR and therefore subject to special category information protections. Even where the threshold for biometric data is not met, the data when considered as a whole can reveal a vast amount of information about a person. This could provide organisations with increasingly sophisticated user profiles and centralises a huge amount of information in one place increasing security risks. These risks were explored in detail in our [2022 Tech horizons report](#).
- **Children’s information:** Children are already active users of early immersive virtual environments for gaming, socialising and learning purposes. A recent decision by a leading hardware and platform provider to lower the minimum age for their VR headset use from 13 to 10 years old indicates they will likely continue to be a growing user base. Research conducted by Global Counsel showed 56 percent of the UK public strongly supported rules which would require technology companies to protect children by restricting their experience of the metaverse.²⁰ Organisations should proactively consider putting in place safeguards in their IVW solutions. In particular, their application of age assurance technologies and compliance with privacy and safety by design requirements. Our [Age appropriate design code](#), which provides guidance for online services likely to be accessed by children, is particularly relevant here.
- **Interoperability:** Interoperability between different IVW providers will enable users to have an increasingly “frictionless experience”. This means they can move between different environments seamlessly. This may allow users to easily enter different shops in a virtual shopping centre or visit a number of virtual tourist attractions one after the other without having to ‘sign in’ to each location. In

practice, this would mean there is no need for multiple log ins, usernames and passwords for different platforms. It may also provide a more dynamic user experience than being blocked from entering new virtual spaces until terms and conditions are agreed with each virtual world operator. Therefore, interoperable virtual worlds may empower users to be more in control of their personal information. However, they will also rely on establishing common technical standards, protocols and systems. Any standardisation of these will rely on platform and industry agreement and collaboration.

Interoperability also has the potential to open up digital markets and allow a level playing field in which users can exercise more control over their personal information. However, if badly executed, it can also present new privacy, safety and security risks. This is because there is uncertainty about how personal information may flow between solutions. This question only becomes more complex when considering different jurisdictions. Data minimisation in exchanging personal information should be a guiding principle in the design of seamless experiences.

Within this complex ecosystem, other key considerations will include establishing:

- who the data controller or joint data controllers are, and who the processors are;
 - which organisations hold copies of a user's personal information; and
 - how users know how to enact their individual rights, and with whom.
- **Digital identity** is also likely to be a core element of future IVWs. Developers of interoperable environments may look to digital identities to help provide a frictionless experience in moving between virtual worlds. Users may enter IVWs using portable, unique identities, which may be lifelike representations of themselves or may be intentionally not lifelike at all. Platforms will need to find a balance between accountability and privacy, as many users may seek to be anonymous while online for a variety of reasons. While this challenge is one that is faced by platform providers today, the potential to interact via visible avatars in immersive virtual worlds adds a new dimension to this conversation.

Recommendations and next steps

We will continue to monitor developments in the space of immersive virtual worlds and scrutinise the market as new products and services emerge and use cases become clearer.

Regulatory cooperation: We will continue to engage with other regulators working across digital services, jurisdictions and industries as policies and standards emerge. We have in December published an [Immersive Technologies Foresight Paper](#) with our counterparts within the DRCF.

Privacy by design: As metaverse technologies evolve in the years to come, new privacy concerns may emerge. It will be critical for organisations to embed privacy by design approaches into the infrastructure of the metaverse as it develops and industry standards are set. This will ensure solutions are built with high data protection standards and safeguards in mind.

Further reading

[Further explanation on virtual reality vs augmented, mixed and extended reality](#)

[ICO guidance on special category data](#)

[Ball, M. \(2022\). The metaverse the metaverse: And how it will revolutionize everything. WW Norton](#) 

[ICO guidance on controllers and processors](#)

²⁰ [Global Counsel report on "Regulating the metaverse"](#) 

Neurotechnologies

Introduction

Mind reading; it sounds like science fiction, and it is, but the current reality (and near future) is just as fascinating. Implanted and wearable devices that analyse and even directly alter our brain patterns already can and will have a huge impact on medical treatments, the ways we work, and enjoy ourselves and will continue to do so. Unsurprisingly, interest in using and regulating neurotechnologies has grown significantly in recent years. This is set against a background of growing financial investment in a market estimated to be worth some 14 billion pounds by 2026. However, much of the current focus is on the social and ethical implications of these technologies, rather than the privacy challenges they present.



About neurotechnologies

Neurotech encompasses a wide variety of approaches but can be broadly defined as “devices and procedures that are used to access, investigate, assess, manipulate, and emulate the structure and function of neural systems”. This includes:

- medical devices from prosthetics and diagnostic devices, such as cochlear implants and MRIs;
- cutting-edge implants and wearable devices that may be used in medical settings, such as treating neurodegenerative conditions; and
- increasingly for commercial uses, such as workplace monitoring.

The personal information derived from these technologies, so-called “neural information”, can be described as “information relating to the functioning or structure of the human brain of an identified or identifiable individual that includes unique information about their physiology, health, or mental states”.

The delivery and use of neurotechnologies can be diverse in terms of devices and the ways in which they gather information. Not all of these directly present privacy concerns; the following are some of the key technological approaches and concepts of most interest to us:

- **Implanted or semi-implanted devices** are predominantly medical devices embedded into the brain, which offer a high degree of accuracy and granularity of information.
- **Wearable devices** are medical and consumer devices worn on the body, which offer cheaper, lower risk access to neurotechnology.
- **Read-only devices** are devices which only gather neural information for analysis, although outputs may directly feed into other devices such as VR headsets.
- **Read and write devices** are devices which gather neural information but also modulate or stimulate brain patterns directly, potentially affecting behaviour and individual's responses.
- **Closed loop processing** is the automated processing of neural information that gathers and analyses personal information and feedbacks to either the person or device without meaningful human intervention. For example, a medical device that automatically stimulates a part of the brain to prevent a serious seizure.

State of development

The medical and healthcare sector remains at the forefront of neurotech development, with rapid innovations happening around devices used in treatment of epilepsy and spinal injuries, to name just a few. More advanced research on potential treatments for neurodegenerative conditions, such as Parkinson's disease, as well as mental health therapies is proving promising. This may result in significant future uptake, if successful.

Other sectors that are likely to see market developments in the near future include:

- wellbeing and sports wearables designed to track and enhance sleep and physical performance;
- workplace deployment to track employee safety around heavy machinery, monitor attention levels, or possibly even as part of recruitment processes; and
- neuro-entertainment to enhance e-sport performance or develop games directly controlled or responsive to the user's brain patterns.

These developments, as well the general uptake of emerging neurotechnologies, appear to be driven by a variety of factors:

- The increased affordability of sensors and a growing trend towards lightweight, efficient, portable and non-implanted devices.
- The increased sophistication of supporting technologies which enable rapid scaling, such as AI and machine learning, 5G and wireless connectivity and cloud storage.
- The global lack of neurotech specific regulation combined with the appetite for the 'datafication' of people may also drive commercial uptake.

In response to these developments, calls for specific neurorights are also emerging, although the need for these remains debated. Considerations of neurotechnologies have also directly impacted emerging legislation. This is most notable in Chile, where constitutional reform has embedded distinct limitations and obligations about the use of neurotechnologies.

Fictional future scenario

Alex is a football player in the middle of negotiating a dream transfer to a top Premier League team. Keen to get the very best analysis of their performance and fitness, Alex agrees to the collection and use of not just biological information like their heartrate, but also their neural information. Using wearable devices, the team can track their sleep patterns, ability to focus and response to physiotherapy. This is then fed into a program of training to enhance their performance.

However, as the contract discussions continue, it becomes clear that the club consider the neural information theirs. They see it as integral to a wider training program they're developing and could give them a competitive edge. Alex is deeply uncomfortable with this. Surely the neural information is personal information and they should have access and some control? They're uncertain and worried about the complexity of the technologies and information sharing. These concerns only grow when they learn the club has started working with a new insurance company, which plans to make use of neural information to rate and evaluate players. This doesn't sound like anything they agreed to. Yet there are rumours that clubs can use the information to pinpoint when a player is no longer at their best. How can they challenge these practices and the potentially huge impact they might have on their career?

Data protection and privacy implications

The data protection risks posed by collecting and using neural information include the following:

- **Processing of novel, highly sensitive information:** The potential for organisations to collect large scale, complex neural information sets about a person presents increased security and privacy risks. This may allow organisations to draw detailed inferences about highly sensitive information, such as someone's mental health or sexuality. A related challenge is understanding when neural information counts as special category information under the UK GDPR. There is no explicit definition of neural information under the legislation and unless used for identification or medical purposes, it is not considered special category information. Even when revealing sensitive information, such as workplace performance, neural information is likely to be considered personal information without the additional protections given to special category information. However, organisations should process it in line with the requirements of the UK GDPR. However, they should remain extremely careful in how (and why) they process such information.
- **Consent and transparency:** Neural information is subconsciously generated. People have no direct control over the specific information that is generated and shared through neurotech devices. This is likely to make the use of consent as a basis of processing challenging to achieve if people cannot be sure of what information they are being asked to provide and what it may reveal about them.

Neurotechnology's potential to not only observe and collect neural information, but to modulate brain patterns and alter behaviour may inhibit transparency. This may fundamentally hinder the pursuit of rights under existing privacy legislation given the potential to inhibit a person's ability to evaluate their own personal information. It may also raise far more fundamental questions about freedom of thought

and personal identity. Organisations should embed privacy by design at the beginning of any technological or service development.


- **Neurodiscrimination:** neural information might be taken as a new path to discriminate against minority and marginalised communities, if inaccuracy and systematic bias remain embedded in technologies. This breaches the UK GDPR requirement for personal information to be processed fairly. It may also lead to discrimination against new categories of people based on their brain patterns who aren't currently recognised under parallel legislation such as the UK Equalities Act.
- **Regulatory and technical clarity:** Further collaboration is key to understanding possible gaps in regulation or guidance and potential clashes of interest. This should be between privacy regulators and those regulators who oversee consumer protections, medical devices and other relevant areas. Developing technical clarity will ensure that regulators have functional knowledge of neurotechnologies. Agreed definitions to allow for effective regulation will also be important. A key area is how and when automated processing of non-medical neurodata may be appropriate and what meaningful intervention would look like in order to ensure that decisions are made fairly and transparently.

Recommendations and next steps

- We have already published a longer and more in-depth review of neurotechnologies and their privacy challenges; our [ICO Tech Futures: Neurotechnologies](#) report.
- We will continue to engage with key stakeholders across industry, regulation, academia and civil society. This will include inviting organisations to work with our Regulatory sandbox to engineer data protection into these technologies.
- Given the potential impact of neurotechnologies on the public, we will work with them to better understand their knowledge and concerns about neurotechnologies and privacy through public engagement activities.
- We will develop guidance in due course on data protection expectations for neurotechnology. This will help bring regulatory clarity and set clear expectations about the responsible and compliant use of neurodata.

Further reading

Regulatory Horizons Council's report on [Neurotechnology regulation](#) 


Royal Society's report on neurotechnologies in healthcare - [iHuman neural interfaces report](#) 

Future of Privacy Forum's [Privacy and the connected mind report](#) 

The Council of Europe's report on [Common human rights issues raised by applications of neurotechnologies in the biomedical fields](#) 

UNESCO's report on [Ethical Issues of Neurotechnology](#) 

Knowledge Transfer Network's (UKRI) report on [A transformative roadmap for neurotechnology in the UK](#) 


Organisation for Economic Cooperation and Development's (OECD) [Recommendation on responsible innovation in neurotechnology](#) 

Quantum computing

Introduction

Quantum computing, one of several quantum technologies, has generated a lot of attention and investment in recent years. It is a technically very complex, but potentially revolutionary technology. Quantum computers make use of particle behaviour at an atomic level to run computations. A fully functional future quantum computer could, in theory, solve certain problems exponentially faster than the computers we use today. If the existing technical and engineering hurdles can be resolved, quantum computers could unlock significant advances across a range of industries. For example, medicine, finance, physics, materials science and artificial intelligence.²¹

Quantum computing has well-documented impacts on existing encryption methods and future information security. We are interested in the ongoing efforts to address the risks to personal information that quantum computing may present. As research into potential real-world use cases accelerates, it is also important to consider if and how these new computers may process personal information in future.



About quantum computers

In simplified terms, classical computers (the computers we use today) process information represented as sequences of 1s and 0s (called digital 'bits'). Quantum computers are built very differently. Rather than relying on simple 1s and 0s, they use quantum bits called 'qubits'. Qubits can represent two states at the same time, meaning they can be in both a position of 0 and 1. Qubits can be linked in a way that enables them to represent even more states at the same time. The phenomena responsible for this are known as superposition and entanglement. Due to these properties, as you add qubits, the processing power of a quantum computer grows at an exponential rate. They also solve problems in a different way to classical

computers. This means future quantum computers may be able to solve some problems much faster than classical computers, including some that classical computers currently effectively cannot.²² These could include the following:

- **Factorising very large numbers**, which has implications for long term information security, as factorising underlies many existing encryption algorithms that protect digital information and communications.
- **Modelling highly complex systems**, which offers possibilities for advances in areas like drugs discovery or physics research.
- **Solving problems with lots of different variables**, which may be used to improve fraud detection or optimise financial portfolios.
- **Increasing search speed** within large and complex datasets.
- **Accelerating machine learning** for certain applications, such as personalised medicine or autonomous vehicles.

Because they work very differently from classical computers, quantum computers can be used for specific, but not all, computational problems. Organisations are therefore more likely to use them alongside existing computers, rather than to replace them.

State of development

Potential timelines for the development of quantum computers and specific use cases remain highly uncertain. Any advances generate a lot of media attention, such as breakthroughs in the number of qubits a computer has, or improvements in error correction. But these hide the technology's very early developmental stage. Prototype versions of quantum computers are available via the cloud for research purposes and industry testing. We are also seeing early tests of hybrid uses, where quantum computers are used for specific functions to support classical computing. However, any large-scale uses are still very far off.²³

It is unclear exactly how far off the "quantum advantage era" is. This would mean that larger quantum computers (10,000+ qubits), would be able to reliably outperform classical computers on real tasks. There are still significant technical and engineering challenges to further development. Some projections suggest we may reach it within five to 10 years,²⁴ while some commentators question whether we will ever be able to gain quantum advantage for certain applications.²⁵ If the challenges can be overcome, the era of "universal fault tolerant computers", that is, fully functional computers with widespread quantum applications could develop within 10 to 20 years, or longer.²⁶

But, many countries, including the UK, see developing their own capabilities as critical for their long term national strategic, scientific and technological competitiveness.²⁷ Globally, there has been more than USD\$35.5 billion of public and private investment in quantum technologies, including quantum computing.²⁸ It is therefore important to think proactively about the privacy implications of the technology.

Fictional future scenario

Rhian, Head of Cyber Security at a health research facility, is called before the management board. The board are deeply concerned by newspaper reports that suggest a massive improvement in the capabilities of quantum computing. They are especially worried about an “imminent existential threat” to information security. Rhian notes that the NCSC have not confirmed the reports. However, the Board insist she contact the ICO and review other industry-recognised guidance to understand what they should consider, before going ahead with one of the many new “quantum secure” services. They want to understand how to implement protections as quickly as they can.

Rhian recalls a management board meeting five years previously, where she had sought funding to prepare for the transition to postquantum cryptography. She had struggled to engage the Board, who were not keen to invest their limited budget in addressing a cyber security risk that might never materialise. The organisation knows they hold high-value and sensitive personal information, including large health datasets for research purposes.

At this Board meeting, Rhian stresses that some of the company’s datasets are fortunately already protected. Their cloud provider has automatically upgraded to a standardised postquantum cryptography option. Some information is protected by encryption that is not at significant risk from a quantum computer. However, as she previously flagged, a full transition could take years. Rushing the transition is just as risky as doing nothing at all. She’ll need time to work with the data protection officer (DPO) and wider organisation to identify all places where personal information, systems or hardware may be at risk, and take phased action to address the risks. She also doesn’t have the resources to find out whether all third-party processors the company works with have transitioned to postquantum cryptography. Rhian also highlights recent reports of organisations who implemented non-standardised solutions in a hurry, unwittingly putting customer information, and the organisation’s wider cyber security, at risk. She sits down with a sigh. Will the Board listen this time and approve the resources she needs?

Data protection and privacy implications

- **Threats to encryption:** without appropriate mitigation in place, quantum computing has the potential to undermine specific types of encryption that protect most of today’s digital communications and information in transit. This ranges from the financial system to biometric information to digital signatures (used to digitally prove identity). Some commentators also raise the risk of “harvest now, decrypt later” activities. This is where malicious actors may already be collecting high value information as it is sent today, in order to decrypt it once they gain access to a sufficiently powerful quantum computer.²⁹
- It is highly uncertain when such a quantum computer may materialise, with estimates ranging from five to 30 years.³⁰ It is also uncertain exactly what information may be most at risk from a quantum computer. However, very high value information that will remain relevant for a long time is considered more at risk than others.³¹ From a data protection perspective, this could include sensitive personal information (including special category information) that needs to be protected for an extended or indefinite period of time (10+ years). For example, health and genetic information or financial information processed by regulated industries.
- **Postquantum cryptography:** there is already a significant amount of ongoing work about different technical measures to address the possible security risks. For example, in the field of postquantum

cryptography. Multiple agencies have produced relevant guidance, including the [UK's National Cyber Security Centre \(NCSC\)](#).

Under the DPA 2018 and UK GDPR, organisations have an obligation to:

- ensure the confidentiality, security and integrity of the personal information they process; and
- take appropriate technical and organisational measures to protect this information, considering the state of the art.

The complexity of IT systems and cryptographic infrastructure differs between organisations. For some, the transition to postquantum cryptography may be as simple as an automatic software update.³² For others, the transition to postquantum cryptography is likely to be lengthy, complex and expensive.³³ Because of this, and the uncertainty of timelines surrounding quantum development, it is important for organisations to start considering their risk exposure in the immediate and near future. [Our encryption guidance](#) emphasises that organisations should be crypto-agile. This means that they keep their encryption use under regular review and ensure they remain aware of updates and vulnerabilities.

The NCSC recommends that organisations should not rush to transition to postquantum cryptography before final standards-compliant products are available. However, they should be ready to do so, once they are. We support the NCSC's recommendation that large organisations who manage their own cryptographic infrastructure start planning their future transition now, such as by identifying at-risk datasets and systems that rely on public key cryptography.

- **Processing personal information:** There is still a high degree of uncertainty about potential use cases for quantum computers. They are not "all purpose" computers and many early anticipated use cases are unlikely to involve processing personal information. Therefore, they would not fall within scope of data protection legislation. For example, using a quantum computer to solve a materials science or physics research problem. However, where an emerging use case does involve processing personal information, organisations must comply with their data protection obligations. Some potential use cases being explored, that could involve processing personal information, include:
 - near real-time fraud detection;
 - customer targeting and prediction;
 - natural language processing; and
 - genomic data analysis.

In some cases, training datasets may be anonymised or pseudonymised. This, and how insights are applied to people, will also change an organisation's [data protection obligations](#).

- **International transfers:** Quantum computers are expensive and technically complex. Therefore, in the short term, accessing them is likely to involve at least some access to overseas quantum computing capacity, as the UK continues to scale its domestic infrastructure.³⁴ Much like current cloud services and other third party information processing operations, if organisations transfer personal information, they will need to meet international transfer requirements.
- **Transparency, accuracy, storage and retention:** Qubits do not last very long because their state 'collapses' when observed or measured. They cannot be copied and are currently fragile and prone to inaccuracies. It is not yet clear whether qubit properties could make it harder to fulfil certain data protection obligations, if organisations start to use quantum computers or quantum or hybrid data centres for processing personal information in the future.

We want to consider if, and to what extent, it may be more difficult to respond to requests for personal information (SARs), retain or correct personal information, or account for inaccuracies in computing outputs. As noted in our earlier work on quantum as part of the DRCF, we are also interested in the potential implications for future explainability in quantum machine learning. It is early days, and many relevant use cases may only develop over the long term. However, as testing of academic and real-world use cases accelerate (including classical-quantum hybrid applications), we plan to explore these questions further in our foresight report on wider quantum technologies, which we will publish in 2024.

Recommendations and next steps


- We will consider our policy and regulatory positions further in our foresight report on the broader space of quantum technologies, which we are due to publish in 2024.
- We are committed to supporting organisations to understand current and future cyber risks so that they can appropriately protect personal information. This includes the risk arising from developments in quantum computing.
- We will continue to engage with all relevant stakeholders on the transition to postquantum cryptography. We will also engage with our DRCF counterparts to ensure our regulatory approaches are aligned.
- Quantum computing is at a very early stage of development, but existing regulation continues to apply. We are committed to supporting innovators testing quantum computing use cases to identify whether they may be processing personal information and embed privacy by design as early as possible. At this early stage, and as the UK's ecosystem continues to develop, we are exploring further opportunities to learn and share our insights with innovators, the UK's quantum hubs, the Regulatory Horizons Council, academia, and other regulators.


Further reading

[DRCF Insights Paper on Quantum Technologies](#) 


[National Quantum Computing Centre \(NQCC\) overview of Quantum Computing](#) 

[NQCC Technology Roadmap](#) 

[National Cyber Security Centre \(NCSC\) whitepaper on next steps in preparing for post-quantum cryptography \(2023\)](#) 

[NCSC whitepaper on preparing for Quantum-Safe Cryptography \(2020\)](#) 

[United States National Institute of Standards and Technology \(NIST\) news article on call for proposals for Post-Quantum Cryptography Standardization \(2016\)](#) 

[European Data Protection Supervisor's TechDispatch on Quantum Computing and Cryptography \(2020\)](#) 

[ICO guidance on data security](#)

[NCSC guidance on GDPR technical security outcomes](#) 

[World Economic Forum insight report on State of Quantum Computing: Building a Quantum Economy](#)

[\(2022\)](#)

[McKinsey & Co. Quantum Technology Monitor \(2023\)](#)

[AWS Quantum Technologies blog on noise in quantum computing](#)

²¹ [DRCF Quantum Technologies Insights Paper](#)

²² [DRCF Quantum Technologies Insights Paper](#); [National Quantum Computing Centre \(NQCC\) overview of quantum computing](#)

²³ Quantum Advantage Conference Panel; DRCF Quantum Symposium Panel.

²⁴ DRCF Quantum Symposium Panel.

²⁵ See, eg, [Nature article on "The AI-quantum computing mash-up: will it revolutionize science?" \(2024\)](#); [Communications of the ACM article on "Disentangling Hype from Practicality: On Realistically Achieving Quantum Advantage" \(2023\)](#)

²⁶ DRCF Quantum Symposium Panel.

²⁷ [National quantum strategy - GOV.UK \(www.gov.uk\)](#)

²⁸ [World Economic Forum insight report on State of Quantum Computing: Building a Quantum Economy \(2022\)](#)

³⁰ [Inside Quantum Technology News article on the "Quantum Cryptographic Threat Timeline"](#)

³¹ [National Cyber Security Centre \(NCSC\) whitepaper on next steps in preparing for post-quantum cryptography \(2023\)](#)

³² [NCSC whitepaper on next steps in preparing for post-quantum cryptography \(2023\)](#)

³³ [World Economic Forum whitepaper on transitioning to a Quantum Security Economy \(2022\)](#)

³⁴ [techUK report on Quantum commercialisation: Positioning the UK for success \(2022\)](#)

Commercial use of drones

Introduction

Drones, also known as unmanned aerial vehicles (UAVs) or remotely piloted aircraft systems (RPAS), are airborne vehicles which are remotely controlled by a human pilot on the ground. In the future, this process may be entirely automated. Drones have long been used in military settings but have seen increased commercial use since the 2010s with the arrival of small, consumer priced models used primarily for videography and photography. The technology has also seen rapid adoption across a wide range of new uses, including policing and security, search and rescue, and in the creative fields.

Today, more sectors are showing an interest in embedding drone technology, seeing them as tools that may help carry out tasks in a faster, safer and cheaper manner. In the future, as wider regulatory barriers to deployment are overcome and the broader potential of the technology is realised, drones could be rolled out more widely. This could significantly impact our cities and streets, and with it our concept of what constitutes public and private spaces. Increased privacy concerns are a natural byproduct of this increased proliferation.



About drones

Drones can come in sizes ranging from the very small (under 250 grams), with rotary wings allowing for vertical take offs and landings, to large vehicles that are similar in size, look and operation to small aircraft. Future commercial drone uses are likely to include industries such as health, construction and transportation and therefore they will become more present in people's daily lives. They are likely to be able to fly in higher airspaces, able to cover further distances for longer amounts of time, with the ability to carry cargo of varying weights.

State of development

While we have seen growing use in the UK, the scaling up of the commercial use of drones has so far been restricted by aviation regulatory requirements. These limit the widescale use of drones flying beyond the visual line of sight (BVLOS) of the operator, in addition to other regulations about environmental health and safety. The readiness and maturity of drone technology and operations has also needed to advance to prepare for meeting these new regulatory requirements.

The [government's drone ambition statement](#), released in 2022, publicly commits them to delivering an enabling regulatory framework. This will safely and efficiently support the drone industry as it develops and work with the Civil Aviation Authority (CAA) to ensure it has the capability and capacity to provide guidance in areas such as BVLOS. In recent months a number of high-profile pilot projects about BVLOS have begun across the UK. This includes the Royal Mail delivering post in the Scottish Highlands and RNLI & Royal Life Saving Society UK's pilot on Cornish beaches to test the use of a new emergency response drone pilot rescue service.

There are a huge range of potential use cases for commercial drones, from moving goods and people to air advertisements and flying QR codes, to more timely and accurate weather forecasting. However, the future trajectory of drone use will heavily rely on the advancements in the technology and public perception in years to come.

At present, beyond the use of drones for search and rescue and emergency services, industry research shows high levels of support for drone use in infrastructure, agriculture and for the tracking of criminals.

In contrast, there appears to be limited public support for Advanced Air Mobility (AAM) programmes, such as flying taxi drones. Notably, there also appears to be a discrepancy between consumer and industry interest in drones for delivering packages, with industry viewing this more positively than consumers.

Despite public wariness, drone technology is beginning to be used for this purpose internationally, including for expedited deliveries and food deliveries. It remains unclear if drone deliveries would be economically viable for routine business-to-consumer deliveries in the near future. However, they have been successfully trialled in the UK for business deliveries, including during the Covid-19 pandemic to deliver vaccines and test kits to rural areas and shore-to-ship deliveries, as well as internationally for a broad range of uses.

As drones become more widespread and more present in our day-to-day lives, concerns around their privacy and data protection implications are likely to become more prominent.

Fictional future scenario

Priya is attending her first music festival with a group of friends and will be camping onsite for the weekend. The festival organiser has decided to use drones for crowd monitoring and health and safety during the event. They procure the services of a third-party drone operator which will provide real-time images and monitoring to help manage crowd patterns and improve safety. This also enables the organisers to make better-informed decisions about provisions and festival staff allocations. Furthermore, the drones are equipped with heat sensors that they use to monitor attendees who have

wandered into out-of-bounds areas on the site (especially helpful late at night).

The organisers are happy with the services provided, but not all attendees feel the same way. Priya sees one of the drones hovering low over the camping area of the festival site. This feels very intrusive, especially as she and her friends had deliberately pitched their tents closely together, to provide a more private area in the centre. She's not sure if a fellow attendee is using the drone to capture photo or video content, or if the drone is being used in an official capacity by the organisers. Priya uses the Remote ID app on her phone to identify the drone operator, but it's not clear that they are acting on behalf of the festival organisers as they aren't mentioned on the Remote ID app. This leaves Priya concerned about who has recorded her friend group, and what the footage may have captured before they realised they were being monitored.

Data protection and privacy implications

- **Surveillance:** One of the key concerns about drones is their discreet nature and potential use for surveillance. We have already addressed this [in our guidance on video surveillance](#) and we continue to monitor for advancements in this market area. Where the primary purpose of the drone is for surveillance, is therefore out of scope for this exploration
- **Inadvertent and mass data collection:** Drones are likely to collect personal information during their operations. This includes:
 - personal information programmed into the drone prior to operational flight;
 - audio, video or photos of people, captured either when on the ground or in flight; and
 - information that may point to an identifiable person, such as their number plates or address details.

Some of these activities may involve capturing personal information of employees working in the vicinity of the drone, people in public spaces, or people within the boundary of their private property.

Due to their bird's eye nature and wide aspect, drones also have the potential to collect large volumes of information, some of which may be personal information. This raises further concerns about transparency and facilitating people's rights, including the right to be informed if their personal information is being processed. Our guidance on video surveillance clarifies that organisations should consider switching on and off any recording system when appropriate, and unless necessary and proportionate any recording should not be continuous.

Where organisations cannot avoid collecting and processing personal information, they could consider obfuscation or blurring solutions as a proactive mitigation. Similar obscuring filters are already used by a number of street view and map applications to selectively blur out people or spaces and could be expanded to drone captured imagery and video.

- **Transparency requirements and privacy information:** Broader concerns remain for the use of drones by commercial organisations. These cover a variety of purposes, including aerial imaging, inspection, deliveries and other capabilities which may inadvertently collect personal information as part of their operations. Complying with transparency requirements and conveying appropriate privacy information to people at the right time may be challenging due to the mobile nature of the drones themselves.

- **Drone IDs:** One option may be linked to a remote ID programme for drones. This solution allows people on the ground to remotely identify airborne drones. This type of identification system would not be dissimilar to car number plates. However, the information would be wirelessly broadcast from the drone to someone on the ground seeking to identify the vehicle through a mobile application. This information could be received by people or authorised organisations, such as aviation safety bodies, law enforcement or national security agencies.

Countries such as the US, Japan and Switzerland have already implemented remote ID programmes. While their core aim is generally linked to safety and security, they may also be used to support compliance with data protection legislation. They do this by providing people with the ability to verify a drone's operator and potentially link them to the relevant organisation's privacy policy. Remote IDs could also include information about a drone's location and altitude, their take-off and intended landing location, and the type of sensors in operation on the drone.

Some jurisdictions have chosen to implement a "tiered access" remote ID system. This provides minimal information to the general public, but authorised organisations, such as law enforcement agencies and safety inspectors, are able to access more detailed information about a drone's specifics. The UK is currently exploring whether and how to implement a remote ID programme.³⁵


Recommendations and next steps

- It will be critical that regulators and industry engage to establish policies and standards about personal information collected by drones as they become more prominent in our everyday lives. This will provide commercial operators with regulatory certainty and increase consumer trust in commercial drone operations. We are committed to take part in these discussions and to continuously monitor developments in the rapidly evolving drone space
- Embedding privacy by design into the hardware and applications of drones will be critical in supporting organisations and drone operators to evidence their compliance with data protection legislation. We are committed to supporting innovators to proactively embed privacy-enhancing mechanisms into their solutions through [our array of innovation services](#).

Further reading

[ICO guidance on transparency](#)

[PwC report "Skies Without Limits" v2](#) 

³⁵ [Civil Aviation Authority \(CAA\) update on Remote ID \(RID\) system in the UK](#) 

Personalised AI

Introduction

In the past year, generative AI has captured the public imagination. Freely available models allow users to generate unique images as well as a wide variety of text-based and other outputs.

They have gained a considerable numbers of users in a short period of time, as well as a flurry of new investment. Generative AI models are increasingly offering users the ability to personalise outputs.

Generative AI refers to systems which use deep learning and other technologies to generate novel content, usually in response to a prompt provided by a user.³⁶ The systems often use enormous datasets (eg images or texts), some of which may have been scraped from the public web. They use these to learn about the way sentences and images have been constructed. This allows the system to 'understand' user enquiries and generate appropriate responses. Responses can subsequently be further improved through user feedback.³⁷

Although open questions remain about possible use cases and the potential for monetisation of these solutions³⁸, the generative AI sector as a whole is seeing considerable investment. Some see the technology as crucial to future economic growth with consultants estimating that AI could generate between \$2.6tn to \$4.4tn.³⁹



About personalised AI

Various iterations of personalised AI systems already exist and work in a variety of different ways. These systems can offer outputs designed to be useful for specific users, based on information related to the user. For example, a person's search history, preferences and other information. They may be underpinned by a

large language model and then fine-tuned to better suit a user or business.⁴⁰

Other systems allow users to create bespoke models and outputs principally grounded on the information provided by individual users.⁴¹ Organisations or individuals could use this to automate the creation of instant messages or emails.⁴²

Generative AI systems, which allow for a greater degree of personalisation, have recently received increased attention. Leading tech figures have argued that the use of these is likely to grow substantially in the coming years.⁴³

State of development

A variety of services which market themselves as personal AIs are already on the market and in use. Existing products offer users the ability to train a personalised generative AI to respond to emails automatically or semi-automatically. Such models can utilise a user's own writing style and vocabulary to automate email or instant message responses. These iterations sit alongside other offers of personalised AIs which offer users assistance with tasks, such as shopping and booking travel.⁴⁴

Over the next few years, the usage of such systems could increase substantially. In particular, it is likely that the abilities of these systems will further improve, allowing users to automate a wider range of personal tasks across a wide range of sectors.⁴⁵

Likely future use cases include AI systems performing a greater educational role. Online educational platforms are already offering such generative AI powered tutors and similar systems could see wider adoptions in classrooms at all educational levels.⁴⁶ These systems could offer users a tailored learning experience, one which is based on the personal strengths and weaknesses of an individual student rather than generalised approach. Leading technology figures are citing this possibility as a means of improving educational quality.⁴⁷

As generative AI systems become increasingly sophisticated, other areas of use may include personalised systems assisting users in creating digital memory banks to assist in battling memory loss.⁴⁸ AI systems could also act as a personal life coach, offering guidance on how to approach difficult life events, taking into account a user's unique circumstances.⁴⁹ Some companies even go as far as to claim they will soon be able to replicate a user's personality and allow users to speak with clones of themselves.⁵⁰ Slightly nearer term, we could see personal AI solutions act as sophisticated financial assistants. These could analyse a user's financial situation and subsequently provide personalised savings and investment advice.⁵¹ Some experts go as far as to argue that digital assistants could soon "potentially do almost anything on the internet".⁵²

Broadly, personalised generative AI systems could offer a wide range of new and creative use cases. However, wider use will have implications for data protection and privacy.

Fictional future scenario

As part of his secondary schooling, Kwame can interact with a virtual tutor which uses generative AI to provide a personalised tutoring experience. This allows him to be taught at the most educationally

appropriate level for his skills and abilities, as the content it provides can be based on his previous experience and ability.

Kwame can therefore experience a more personalised experience than in a standard classroom. Kwame's parents sometimes express concerns about transparency and how the system makes decisions about his education. The wider availability of this technology increases the viability of homeschooling for his parents and many others. To offer this personalised experience, the virtual tutor will need to collect personal information about Kwame. This could include information about academic ability. Personalised tutoring systems which offer education in religious and social matters could process special category data about Kwame.

Data protection and privacy implications

The development of personalised AI has the potential to bring considerable benefits to users. It could assist in improving:

- workplace productivity;
- educational outcomes; and
- some of the features of existing generative AI.

Whilst there are potential benefits from personalised AI, its expansion presents privacy and data issues. Developers need to ensure that personalised AI is expanded in a privacy positive way.

- **Generative AI:** If reliant on a foundation model personalised AI systems, like generative AI more generally, will need to be trained on a large body of information in order to produce high-quality outputs.⁵³ In order to be adapted for specific purposes, models may also need to be 'fine-tuned'.⁵⁴ They will therefore share most of the data protection issues posed by generative AI more generally. We have already noted these issues and produced targeted communications to advise organisations about their obligations. Our recent blog describes eight privacy related issues which developers need to consider when [developing generative AI systems](#).

But the inclusion of more personalised information and approaches will bring additional issues that are specific to personalised AI systems, including the following:

- **More personal information:** Personalised AI solutions process a greater quantity of personal information than other generative AI tools. This information is usually provided by the user. It may be necessary to ensure that the outputs of the generative AI are useful for that user. For example, a system which offers users the ability to automate or semi-automate email responses may need to process information about a user's writing style in order to replicate it convincingly.

Similarly, educational systems are likely to need information about a user's educational attainment and progress. Depending on the educational content and the way in which the system is structured, this information could, for example, also be used to infer information about a person's religious or philosophical views. If this is the case, then personalised AI systems may be processing special category information, which requires extra protections because of the data's sensitivity. This could also apply to systems which process or infer personal information about students' special educational needs. Developers therefore need to ensure that they have appropriate conditions in place for processing such information and have a separate condition for processing under Article 9 UK GDPR.

- **Risk of model inversion attacks:** Generative AI models which offer higher degrees of personalisation may be at greater risk of model inversion attacks. This is because of the amount and nature of the personal information they process. These attacks try to extract the information used to train that model by exploiting its outputs.

These attacks could therefore risk leaking information the user has provided to personalise the service. This could include intimate details on users' personal lives, such as their finances or aspects of their identity. Organisations therefore need to implement appropriate technical and organisational measures in order to process data securely and the UK GDPR's requirements concerning encryption in order to mitigate this risk.

- **PETs:** Privacy enhancing technologies (PETs) could help to keep information secure and to implement privacy by design in these models and tools at the outset. For example, differential privacy could be used to obscure the fact that a particular person's data has been used to train a model. Another PET which could be relevant for personalised generative AI is federated learning. Federated learning may allow developers to minimise the amount of personal information they need to train the model, provide appropriate measures of security and reduce the impact of any potential data breaches. Here, developers face a trade-off, as the use of certain PETs, such as differential privacy and synthetic data, could reduce the effectiveness of the models themselves.⁵⁵ As the personalised generative AI space continues to grow, experimenting with and the developing of further safeguards and privacy enhancing techniques will be important.

Recommendations and next steps

We are considering further steps to take about personalised AI:

- We recommend that those developing and deploying personalised AI systems consult our AI guidance and our other publications which specifically concern generative AI.
- We will continue to monitor the rapid developments in AI generally and will respond accordingly. This includes reviewing and updating our current AI guidance. As the generative AI space continues to develop rapidly, we will continue to monitor new use cases as part of our wider portfolio of AI work.

Further reading

[ICO guidance on special category data](#)

[ICO glossary for AI and data protection](#)

[ICO guide to data security](#)


















[ICO guidance on how PETs can help with data protection compliance? \(ICO\)](#)

[ICO guidance on federated learning](#)

³⁶ [IBM Research Blog on Generative AI](#)

³⁷ [Simon Attard article about grounding Generative AI](#)

³⁸ [Gary Marcus article concerning economical potential of generative AI](#)

- 39 [Financial Times article about the sceptical case on generative AI](#) 
- 40 [Simon Pollington article about fine-tuning LLMs for Enterprise](#) 
- 41 [Personal.AI article entitled Differences Between Personal Language Models and Large Language Models](#) 
- 42 [Personal.AI article entitled Your True Personal AI](#) 
- 43 [CNBC article about Bill Gates' predictions for the educational possibilities of generative AI](#) 
- 44 [Homepage for Maya - Your AI travel assistant](#) 
- 45 [New York Times article entitled How 'A.I. Agents' That Roam the Internet Could One Day Replace Workers](#) 
- 46 [Cousera article about personalized and interactive online learning with generative AI, machine learning, and virtual reality](#) 
- 47 [CNBC article about Bill Gates' predictions for the educational possibilities of generative AI](#) 
- 48 [Personal AI article entitled What is Personal AI?](#) 
- 49 [The Guardian article about Google DeepMind testing 'personal life coach' AI tool](#) 
- 50 [VentureBeat article about AI clones](#) 
- 51 [MSN article about Morgan Stanley plan to launch AI chatbot to woo wealthy](#) 
- 52 [New York Times article entitled How 'A.I. Agents' That Roam the Internet Could One Day Replace Workers](#) 
- 53 [CMA Short Report about AI Foundation Models](#) 
- 54 [CMA Short Report about AI Foundation Models](#) 
- 55 [Leidos white paper about Privacy Enhancing Technologies](#) 

Next-generation search

Introduction

As the use of the internet has expanded, search engines have become a core feature of everyday life for most internet users. Indeed, for a long time, the traditional search engine was the main interface through which users accessed the internet and found new information online.

However, this may be changing as novel methods of search are gaining traction. These often move these activities not just away from the traditional search engine interface, but off our devices altogether.

Instead, search results could increasingly be impacted by the way in which users interact with their surroundings. They could generate recommendations from information drawn from the digital and physical world around us. The further embedding of emerging technologies, such as generative AI in particular into search interfaces, is similarly likely to transform how we find information online. These changes may have privacy and data protection implications. Given the centrality of search to our experience of the internet, any such impact has the potential to be significant.



About search

Traditional search engines grew in popularity as web 1.0 and web 2.0 gained momentum during the late 1990s and early 2000s. They responded to text-based queries, for example about sport scores, the latest news or local restaurants. They did so by providing ranked lists of websites which might provide an answer to those queries.

The later use of algorithms helped these engines to provide more relevant and useful results. Search engines traditionally monetise their products by embedding advertising into search results and ranking

them accordingly. The data processing⁵⁶ involved with embedding targeted advertising into search has seen some raise privacy concerns.⁵⁷ In recent years, both the ways users search and the methods of providing these results have undergone a significant evolution. These new developments aim to provide a more personalised experience to users, better tailored to their individual wants and needs, but are likely to involve the collection of even more personal information in the process.

State of development

A number of existing and emerging technologies may radically change how we use search. A combination of different methods of search, lots of new data points, sensors, ambient tech, and AI can make search a more personalised experience and allow these solutions to interact with the world around us.

One such development is the increasing use of voice-based search, particularly through voice-based assistants on smartphones or dedicated IoT devices. A quarter of UK citizens now own a smart speaker.⁵⁸ Their use as search tools may further increase as the ability of these systems to converse with users in a natural and fluent way improves. These improvements could include an increased ability of the voice assistant to remember conversations and the ability to analyse a user's emotional state and adjust its responses accordingly.⁵⁹ As deepfake technology becomes more widely used, users may be able to change the voice used by a voice assistant to that of a loved one.⁶⁰

Increased use of multi-modal search will similarly allow search to move away from purely text-based search. Current iterations include the ability to input an image of a particular type of food or service. The search engine then provides relevant local information to the user, such as local restaurant or service recommendations. Other multi-modal options allow the user to use their phone to capture images which will then generate information visible to the user on an augmented reality (AR) overlay.⁶¹ For example, users can already use this feature as an image recognition tool to identify the pair of shoes on a passerby or recognise a songbird in a tree. This method of search may become increasingly accessible if technologies such as smart glasses see wider adoption.⁶² As more users take advantage of immersive environments, searching via virtual reality (VR) technology may become increasingly common.

Queryless search is another innovative search method which could see development in the next few years. Queryless or "ambient" search refers to information presented in varying forms to the user without the need for a specific user input, such as a question or image. Information can therefore be presented to the user based on information gathered about them at other times or places. For example, options for restaurants which are presented on a smart home device screen during mealtimes. This could be based on a user's search history, combined with third-party and other personal information. The results could be personalised by information about a user's dietary requirements. These solutions may develop alongside the increased ability of ambient computing. This is when computers are embedded into our immediate environment, such as sensors.⁶³ This allows more information about our current, direct surroundings to feed into search algorithms, to further adapt what users are presented with.

Using generative AI as a search engine and embedded within established search engines is perhaps the most visible recent development. Instead of using a search engine to navigate to a website in the hope of accessing a piece of information, generative AI can provide the information directly to the user via a chat-based interface. As generative AI systems become more personalised, they may be able to refine their outputs based on information about the user.

The new methods of search described above are not expected to develop in isolation but will intersect with

each other. The search ecosystem could become more complicated and we may see new market entrants and users each using these tools in different ways. Therefore, there is an increased risk to transparency and people's ability to understand what information has been used and why.

Fictional future scenario

Miguel uses voice-based search. This means the search engine can use emotional analysis to assess his voice, provide insight into his stress levels and so prioritise restaurants likely to improve his mood. The search engine experience is more personalised and, in some ways, is better able to respond to his preferences than a desktop-based search experience in 2023.

However, the vast amounts of data collected about Miguel means that any data breach is likely to be riskier and could reveal a detailed picture about his life. Voice-based searches are likely to return fewer results, which could mean reduced user choice. The highly personal nature of the search means that the potential for serendipitous discovery is reduced.

Miguel knows that his watch collects lots of information about him but enjoys the highly personalised results the watch provides and the efficiency with which this allows him to choose a restaurant. But after several meals, he does sometimes wonder what other options he might be missing out on.

Data protection and privacy implications

These innovations are likely to have an impact on data protection and user privacy and will present novel issues for digital regulators.

- **Collecting vast quantities of personal information:** One central concern is about the quantity of information these new modes of search collect. For example, in voice-based search methods, such as smart speakers. The information collected helps this search method provide more personalised results to the user. This can compensate for the more limited number of results that voice-based search can offer in comparison with traditional search engines, which can provide at least a dozen per page.

It is possible that large amounts of personal information will also be required for the level of personalisation potentially offered by queryless search. The large amounts of information collected from a user and various sources, have the potential to make any data breach more harmful to the user. Organisations who develop these innovative methods of search need to be aware of their obligation to conform with the principle of data minimisation. This means that they need to ensure that the information they process is adequate, relevant and limited to what is necessary and that sensors, speakers and other means of collecting information only collect what they need to function.

- **Transparency:** Numerous data points, from a variety of sources, may be necessary to feed into the personalised results offered by ambient search solutions that interact with our environment. Users may therefore find it more difficult to exercise key data rights such as the right to be informed and the right of access. A related concern is the reported lack of transparency about sharing information gathered by voice-based search methods with third parties.⁶⁴ Organisations that provide new methods of search

need to process personal information in a transparent manner and ensure users are able to exercise their information rights.

- **Hallucination:** As noted above, there is a trend towards the increased use of large language models. These are used to search for information as well as embedding this type of AI into traditional search engines. The privacy implications of personalised generative AI are set out elsewhere in this report and more generally in our [previous publication on generative AI](#).

However, there are concerns about the accuracy of the results provided. It has been noted that such systems can “hallucinate”⁶⁵ and provide inaccurate information to users. While efforts are being made to make errors less likely, some have questioned whether such efforts can be entirely successful.⁶⁶ This is further complicated as a growing number of websites prohibit generative AI developers from using their websites to develop their models. Therefore, potentially this reduces the diversity and balance of the relevant training information (which can help to reduce the risk of hallucinations).⁶⁷

If such hallucinations contain personal information, this tendency could complicate an organisation’s accuracy obligation to “take all reasonable steps” to ensure the personal information they process is “not incorrect or misleading as to any matter of fact”. The future of search and generative AI are linked in other ways too. Recent reports have highlighted a tendency for generative AI content to feature prominently in search results gathered by traditional search engines⁶⁸, [raising concerns about false information spreading across the web](#).⁶⁹

- **Intersections with immersive technologies:** Multi-modal methods of search, noted above, offer users another way to search and gather information. However, these features may be subject to similar privacy issues, especially if used by next generation smart glasses and headsets. These affect the usage of AR which were noted in the [immersive tech chapter of our first Tech horizons report](#). Immersive technologies may therefore be used, inadvertently or otherwise, to collect information about people in close proximity to the device. This is likely to become more of an issue as relevant devices become more discreet. Other data protection issues include collecting considerable amounts of special category information and targeted profiling.

Current and future methods of search have the potential to offer users a wider variety of ways to seek information. They could offer the user more personalised results and an improved user experience. They could also have an impact on user privacy and will have data protection and privacy implications that organisations should consider from the outset.

Recommendations and next steps

- We are planning a foresight report on the future of search and discovery which we aim to publish in 2024.
- We will continue to monitor developments in the search space and seek to understand how the various technologies and elements that may drive change interconnect and intersect. We will continue to be look out for new privacy implications which may arise.
- We will also aim to work with other regulators, notably the DRCF, to bring more definition and regulatory clarity to the search space.

Further reading

[ICO guidance on the data minimisation principle](#)

[ICO guidance and resources on individual rights](#)

[ICO article on generative AI: eight questions that developers and users need to ask](#)

[ICO guidance on the accuracy principle](#)

[ICO Tech horizons report](#)

56 [Forbes article entitled How Much Does Google Really Know About You? A lot.](#)

57 [Reuters article entitled Google faces \\$5 billion lawsuit in U.S. for tracking 'private' internet use](#)

58 [YouGov study regarding Smart Speaker usage](#)

59 [Daily Upside article about Google smart speakers](#)

60 [NPR article about Alexa capabilities](#)

61 [Homepage for Google Lens](#)

62 [Meta page about new Ray Ban smart glasses](#)

63 [ZDNet article about Ambient Computing](#)

64 [Mozilla Foundation about Amazon Echo Dot](#)

65 [IBM article about AI hallucinations](#)

66 [Fortune article about AI hallucinations](#)

67 [IBM article about AI hallucinations](#)

68 [ITPro article about AI's impact on Google](#)

69 [Article from The Guardian entitled Does Australia exist? Well, that depends on which search engine you as](#)

Central bank digital currencies (CBDCs)

Introduction

A central bank digital currency (CBDC) is money that a country's central bank can issue in digital (or electronic) form, rather than as physical money, such as cash and coins. For example, in the UK this digital money (a "digital pound") would be issued by the Bank of England. It would hold the same value as physical money and could be used in similar ways as money stored in a bank account and be used for everyday payments.

Governments and central banks around the world are investigating how CBDCs can be introduced into existing monetary systems, to cater for changes in a payments landscape where physical money is used less often. In order to promote trust in their development and use, it is important that data protection is built into the development process of CBDCs from the outset. This chapter explores some of the emerging data protection considerations associated with the technology, internationally and in the UK.



About CBDCs

Whilst the concept of an electronic payment system backed by a central authority is common to CBDCs, there are different types of CBDC, underpinned by different deployments and technologies.

CBDCs are commonly described as "wholesale" and "retail". Wholesale CBDCs are not new and are granted by the central bank to financial institutions to, for example, settle high-value inter-bank transfers. Retail CBDCs are digital money issued by central banks for use by private sector businesses and individuals to make everyday payments.

Retail CBDCs are backed as an electronic form of a nation's currency by central banks. In this way they are

different to cryptocurrencies which are issued by private sector organisations. Cryptocurrencies are usually by nature decentralised. In contrast CBDCs have a centralised governance and decision-making architecture, managed by central banks and governments. Further, CBDCs are not necessarily based on distributed ledger technologies (DLTs) and are more stable and likely to retain value over time.

State of development – CBDCs around the world

To date, approximately 130 national governments (representing approximately 95% of world GDP and including all of the G7) have begun [looking into the application of CBDCs](#). More than 15 national schemes are in a pilot or later stage, including in France, Canada, India, Singapore, and China. The implementation of a “digital Euro” has entered a preparation stage, with the European Central Bank (ECB) currently working with European national central banks on the topic.

Globally there are different models and deployments of CBDCs being progressed. Arguments in favour of their adoption include:

- That CBDCs are more inclusive and will help the unbanked engage in financial transactions more effectively and safely.
- CBDCs have been suggested to improve the process of cross-border payments, especially where those payments involve exchanges between two respective domestic CBDCs.
- The increased ease of providing payments from governments to people has been cited. Practical implementations might include making rapid stimulus payments during financial or other crises, which is being considered by some countries (though not in the UK).

State of development – a CBDC in the UK

HM Treasury and the Bank of England have been assessing the case for retail CBDCs, in response to the [changing ways people and businesses use money in the UK](#). The following are potential benefits which are driving interest in the development of a digital pound.

- The introduction of a digital pound could act as an anchor for the wider monetary system by promoting trust and confidence in money and payments.
- A digital pound could provide a platform for private-sector financial innovation. Central banks could support new organisations offering CBDC-based financial products and services, in the same way that they support retail banks.
- A digital pound backed by the Bank of England would mitigate the risks of forms of other kinds of electronic money which are locked into a “walled garden” by a single provider, where users cannot transfer that money elsewhere.

CBDC policy development in the UK

[The February 2023 consultation on a Central Bank Digital Currency](#) released by the Bank of England and HM Treasury concludes that it is likely that a digital pound would be needed in the UK in future, however no decision has yet been made to launch one.

Exploration of how a digital pound might be designed is ongoing, however the Bank of England's [2023 digital pound technology working paper](#) discussed a possible model built around a secure centralised core ledger with access provided to third-party Payment Interface Providers, through which users engage with CBDC payments and services.

[In January 2024](#) [HM Treasury and the Bank of England published a response](#) to their 2023 consultation. This response states that while it is too early to decide whether to introduce a digital pound, further preparatory work will be carried out.

The response also sets out that **privacy would be a core design feature of any future digital pound**. It confirms that legislation would be introduced to Parliament, guaranteeing that neither the government nor the Bank of England would be able to access users' personal data, and that further technological options would be explored to prevent the Bank of England accessing any personal data through the CBDC's core infrastructure. The digital pound would be at least as privacy-preserving as current forms of digital money, such as money stored in a commercial bank account.

Data protection and privacy implications

The concept of the provision of digital money by central banks and governments has led to understandable questions about privacy, data protection and control of payment flows. This section discusses some of the general data protection and privacy considerations associated with CBDCs.

As is the case with how people use money today, providing a high standard of data protection is critical to building and maintaining public trust and engagement with CBDCs. If the public loses confidence in the security and confidentiality of their personal information in monetary systems, trust could be undermined.

As a general principle, those developing CBDCs need to consider if personal information might be processed within the deployment of these systems and ensure users can exercise their data protection rights, if such processing is unavoidable. They need to identify controllers and processors, so they are clear about their obligations and ensure that systems remain fair and transparent, and to ensure users trust in those systems.

- **Access to information:** as can occur in current digital payments systems, the information collected and processed for CBDCs to operate may be made available to a range of intermediaries (eg possible CBDC digital wallet providers, which is one mechanism through which CBDCs might be dispensed). Existing digital payment systems operate under established mechanisms to limit the information available to parties, even if a confirmation of identity is needed. An example of this is the [confirmation of payee scheme](#) which establishes that a payee's identity is correct without the need to share that same identity. The development of CBDCs can consider similar mechanisms to minimise the amount of personal information that is processed and to promote privacy.
- **Processing for law enforcement or anti-money laundering requirements:** for fraud and money laundering to be identified, or for law enforcement organisations to track illegal payments and the proceeds of crime, relevant transaction information must be monitored and analysed. As is the case for existing bank accounts, considering how access to transaction information is provided for these purposes without impacting the privacy of third parties is also a key technical consideration.

- Risk of re-identification: while the information about people’s CBDC use may be pseudonymised or encrypted, central banks will need to ensure that the information they process is not combined with other sources in a way that may lead to reidentification of users. Otherwise, there is a risk that sensitive information about a person and their spending habits could be revealed.
- Immutability of records: no decision has been made about using DLTs in a possible digital pound. However, in CBDCs which do involve the use of distributed ledger technology, trust in the system is supported by its permanent, unalterable nature. This has implications for users’ data protection rights around accuracy and rectification, as well as erasure, as discussed in the [decentralised finance chapter of last year’s Tech horizon’s report](#). If information needs to be corrected or removed to comply with these rights, this will consequently impact the trust in the chain and the provenance of the information held on it. Without a mechanism in place to resolve this, the tension between the rights of users and the trust in the chain may also have secondary effects, as users feel less able to raise concerns and exercise their rights.

Cross-border information sharing: Where cross-border payments are made by users today, transaction information could become available to actors in other jurisdictions which might have different data protection regimes. This is also the case where the payments are made between two different CBDCs, as information about a person’s user transactions could now be present in two different CBDC systems. If personal information is being transferred internationally, international transfer requirements would need to be met.

Recommendations and next steps

The ICO has and will continue to engage collaboratively with HM Treasury and the Bank of England about how data protection and privacy might be best preserved in a possible digital pound. We welcome the emphasis placed on data protection, and its recognition of the opportunities to preserve privacy and support of user control of personal data when using digital money.

- As the Bank of England and HM Treasury continue to explore the concept and design of a CBDC – and their legislative commitments to guaranteeing users’ privacy - we will continue to engage and provide a data protection perspective to this process.
- Where organisations introduce new or novel processing of personal information which is likely to result in a high risk to people, a data protection impact assessment (DPIA) is required. This is also true when major new projects are undertaken which require the processing of personal information. Even if there is no specific indication of likely high risk, it is good practice to complete a DPIA for any major new project involving the use of personal data. In designing and developing a potential digital pound, the Bank of England and HM Treasury should consider whether a DPIA is necessary.
- Data protection law will help ensure that CBDCs and associated applications are developed in a way that respects people’s rights and promotes trust in this technology. Following a data protection by design and default approach when developing the CBDC regulatory framework and technological infrastructure is key to achieving this outcome.

Further reading

[Response to the Bank of England and HM Treasury Consultation Paper](#) 

[Digital euro - The next step in the advancement of our currency \(europa.eu\) !\[\]\(082f818d99f166a3ba574d9284d73064_img.jpg\)](#)

[The IMF central bank digital currency handbook !\[\]\(d263118e0bfd47dc6bc704167d936b83_img.jpg\)](#)

[ICO guidance on DPIAs](#)

[ICO guidance on international transfers](#)