THE ROYAL SOCIETY

BBC

Generative AI, content provenance and a public service internet

Summary note of a workshop held on 14 – 15 September 2022

Background

This note provides a summary of workshop discussions exploring the potential of digital content provenance and a 'public service internet'. The workshop was jointly hosted by the Royal Society and the British Broadcasting Corporation (BBC) on 14 and 15 September 2022 at Carlton House Terrace and Broadcasting House.

The workshop was convened following the Society's report *The online information environment: Understanding how the internet shapes people's engagement with scientific information*, published in January 2022[1]. It was held two months prior to the release of OpenAI's ChatGPT (a chatbot powered by artificial intelligence) and several weeks after the release of Stability AI's Stable Diffusion (a deep learning text-to-image model). Since this workshop, the topic of artificial intelligence and large language models (such as OpenAI's generative pre-trained transformers, or GPT) has attracted significant interest in public discourse. High-profile debates within governments and industry have centred on how to regulate these technologies in order to prevent major societal harms and misinformation.

The concept of 'provenance enhancing technology' (renamed here as 'digital content provenance') was highlighted in *The online information environment* report as a solution which will become increasingly important as misinformation content grows more sophisticated[2]. In addition, the report identified 'new internet protocols' as a future trend that could lead to different versions of the internet, impacting the ability of the state to maintain the health of the online information environment[3].

This note summarises the workshop discussions, highlights key themes that arose and presents suggestions for action and further research. References are included in order to provide illustration of points raised in the workshop. This note is not intended as a verbatim record of discussions and does not necessarily represent the views or positions of any participants or organisations who took part. It was drafted by staff at the Royal Society and the BBC by considering comments, feedback, and references submitted by workshop participants.

The Royal Society

The Royal Society is a self-governing Fellowship of many of the world's most distinguished scientists drawn from all areas of science, engineering, and medicine. The Society's fundamental purpose, as it has been since its foundation in 1660, is to recognise, promote, and support excellence in science and to encourage the development and use of science for the benefit of humanity.

The Society's strategic priorities emphasise its commitment to the highest quality science, to curiosity-driven research, and to the development and use of science for the benefit of society. These priorities are:

- The Fellowship, Foreign Membership and beyond
- Influencing
- Research system and culture
- Science and society
- Corporate and governance.

1. The Royal Society. 2022 The online information environment. See: https://royalsociety.org/topics-policy/projects/online-information-environment/ (accessed 10 July 2023).
2. ibid.
3. ibid.

The BBC

The BBC is the world's leading public service broadcaster. Founded in 1922, its mission is to act in the public interest, serving all audiences through the provision of impartial, high-quality and distinctive output and services which inform, educate, and entertain. The BBC is established under a Royal Charter which sets out its public purposes.

These purposes are:

- To provide impartial news and information to help people understand and engage with the world around them
- To support learning for people of all ages
- To show the most creative, highest quality and distinctive output and services
- To reflect, represent and serve the diverse communities of all of the United Kingdom's nations and regions and, in doing so, support the creative economy across the United Kingdom
- To reflect the United Kingdom, its culture and values to the world.

Summary of key takeaways

- Digital content provenance is an imperfect and limited – yet still critically important – solution to the challenge of AI-generated misinformation.
- A provenance-establishing system that can account for the international and culturally diverse nature of misinformation is essential for its efficacy.
- Digital content provenance tools present significant technical and ethical challenges, including risks related to privacy, security and literacy.
- Understanding how best to embed ideas such as digital content provenance into counter-misinformation strategies may require revisiting the rules which dictate how information is transmitted over the internet.
- A 'public service internet' presents an interesting and new angle through which public service objectives can shape the information environment; however, the end state of such a system requires greater clarity and should include a wide range of voices, including historically excluded groups.

Digital content provenance: what is it, and what does it set out to achieve?

Digital content provenance refers to the use of metadata to present information to an end user about the origins of, and alterations to, a piece of digital content. This information can include: the date and time the content was created; the device and location it originated from; the IP address of the uploader; and a log of any edits made to the content. If used well, this technology could help platforms flag manipulated content, support fact-checkers in verifying contentious claims and provide context to those accessing media on their devices. It was highlighted in the Royal Society's 2022 report, *The online information environment* (pp. 8), as an innovation that 'shows promise' and will become 'increasingly important as misinformation content grows more sophisticated'[4].

Demonstrating the challenge that content provenance aims to address, the workshop began with a presentation by Andrew Lewis, University of Oxford, who designed and ran a deepfake detection survey for *The online information environment*[5]. The survey used two experiments to assess participants' ability to detect a high quality deepfake video of the Hollywood actor Tom Cruise from a set of genuine videos. The results showed that participants were no more likely to notice something out of the ordinary when they viewed a deepfake video than when they viewed authentic, unaltered videos. They also showed that most participants could not identify the deepfake, even when primed with a disclaimer that of the videos was altered.

The survey highlighted the risk that people may not be able to distinguish high-quality deepfakes from genuine content which, as generative artificial intelligence advances, could become a prevalent challenge in the future. By providing information on the authenticity of a piece of content, digital content provenance has been proposed as a solution to mitigating this risk.

4. Op. Cit. 1.
5. Lewis A, Vu P, Duch R, Chowdhury A. 2021, Do content warnings help people spot a deepfake? Evidence from two experiments. The Royal Society.

Workshop participants heard two presentations on digital content provenance: the first was from Andy Parsons, Senior Director of the Content Authenticity Initiative, Adobe), and the second from Laura Ellis, Head of Technology Forecasting, BBC. Mr Parsons outlined the challenges with online image content related to transparency, attribution, and trust, referencing the rise of photorealistic, image generation tools such as DALL-E, Midjourney and Stable Diffusion. The realism of these outputs means that trust in genuine images could be undermined, and that there may now be a need for an image verification system to support users in distinguishing between real and artificial content.

Adobe's Content Authenticity Initiative (CAI) seeks to develop such an image verification system[6]. The CAI is a coalition of media outlets, technology companies, non-governmental organisations, and academics working to 'promote adoption of an open industry standard for content authenticity and provenance.' Members of the CAI include the Associated Press, the BBC, ARM, Getty Images, Microsoft, the New York Times, Nikon and Qualcomm[7].

As Adobe's products (which include Photoshop and Premiere Pro) are used by creative professionals across the world, they are well positioned to experiment with and embed provenance tools in their own software. An example presented was the use of 'Content Credentials' in Photoshop. Adobe's Content Credentials enable the user to attach 'tamper-evident attribution and history data' to a piece of content as it is exported. This is known as a 'manifest'. The manifest contains assertions about the content, can be associated with the creator's identity and is cryptographically signed.

Provenance on a single company's software alone is unlikely to be enough to build a system of image verification. For this reason, the CAI has brought together media outlets, camera developers, technology companies and others to work together on this challenge. In the CAI's view, the broad adoption of provenance standards will require the participation of camera developers, smartphone producers, social media platforms, publishers, and operating systems.

Ms Ellis presented on Project Origin. The Project is an alliance of the BBC, CBC/Radio-Canada, Microsoft, and the New York Times to create 'a process where the provenance and technical integrity of content can be confirmed' with an intent of creating 'common open industry standards'[8]. Similar to Adobe's initiative, Project Origin presents a potential solution to the challenge of disinformation by helping internet users identify genuine content. To illustrate the challenge, an example of an inauthentic BBC News explainer video was presented. The video, formatted in the same style as BBC News' social media content with near-identical graphics and fonts, reported a fabricated news story: that Poland was poised to invade Ukraine. Due to the similarity in presentation style and branding, this type of content is a misinformation risk for casual viewers.

It is the BBC's view that establishing media provenance is one mechanism for reducing the risk of visual misinformation content. Their proposed process involves an 'asset' (defined as metadata plus content), which is cryptographically signed by the publisher. This manifest could be one of many attached to a single piece of media, representing different steps in the production of an asset. The similarity of this work and that of the CAI led to a joint formation of the Coalition for Content Provenance and Authenticity (C2PA)[9].

The C2PA describe the process of cryptographic signing as 'hard binding', in which the provenance data and the asset are two parts of a unique puzzle: any alteration to either part alters the algorithm. This is achieved using a list of hash algorithms (SHA2-256, SHA2-384 and SHA2-512) as set out in the C2PA's proposed specifications[10]. If done effectively, this process would prevent malicious actors from tampering with a content's manifest and enable a system for verification.

6. How it works. Content Authenticity Initiative. See: https://contentauthenticity.org/how-it-works (accessed 10 July 2023).
7. Members. Content Authenticity Initiative. See: https://contentauthenticity.org/our-members (accessed 10 July 2023).
8. Overview. Project Origin. See: https://www.originproject.info/about (accessed 10 July 2023).
9. About. Coalition for Content Provenance and Authenticity. See: https://c2pa.org/about/ (accessed 10 July 2023).

The use of cryptographic hashing for provenance was explained further in a presentation by Dr Charles Bennett ForMemRS, Fellow of the IBM Research Division. Using the example of his work on 'time-bracketed authentication' at IBM following the OJ Simpson trial in the 1990s, Dr Bennett outlined the importance of proving that an event happened at a specific time (accurate to within a few seconds) and how cryptographic hashing of content can help. In the OJ Simpson trial (a high-profile case in which a prominent former US National Football League player was accused of murder), members of the jury said they relied mostly on audio and video evidence[11]. Dr Bennett described how, at the time, accusations that the video evidence had been falsified led him to explore the idea of time-bracketed authentication.

Time-bracketed authentication works by importing a stream of unpredictable information from a trusted public source and using it to influence the scene being recorded[12]. It combines two ideas: the 'kidnapper's trick' of photographing a hostage holding today's newspaper to prove they are alive and the 'inventor's practice' of publishing an invention to prove the invention originated at a certain date[13]. In Dr Bennett's proposals, this involves using random laser movements on a video subject with audio-visual data being hashed at frequent intervals and exported to a trusted repository. In modern times, he suggested this could instead be done using distributed ledger technologies, such as blockchain.

In a similar approach to that taken by the CAI and Project Origin, Dr Bennett argued that a successful system for content provenance would require the cooperation of social media platforms, scalable forms of timestamping and tamper-resistant mechanisms for authenticating location.

Technical, ethical and adoption challenges of digital content provenance

Following the opening provocations, workshop participants joined group discussions on the technical, ethical and adoption challenges associated with digital content provenance.

Technical challenges

Significant technical challenges were raised, ranging from the definition of content manipulation to third party verification of provenance labels.

The manipulation of digital content can automatically occur on smartphones at the moment an image is captured (for example, to remove blurred elements or improve lighting) or later in minor aesthetical edits by the user (such as with filters and cropping). These minor manipulations are unlikely to be of interest when it comes to identifying content which has been purposefully edited as part of a disinformation campaign. However, in the context of shallowfakes[14] (crudely manipulated content), understanding whether context has been cropped or edited out of an image is likely to be of interest. This presents an element of subjectivity into what is included in a provenance label—decisions around which may undermine the efficacy of these labels as a solution to misinformation.

If the criteria for manipulation are too broad and content is flagged as manipulated too often, there may be a risk that the public will develop 'manipulation fatigue', akin to people's reactions to cookie banners.

It was raised that some edits (including image cropping and blurring) may be undertaken for privacy reasons and that actions which risk reversing these should be approached with caution. Furthermore, bad actors may decide to game automated detection systems by making minor edits to genuine content in order for them to be flagged as potentially misleading. Another potential method for gaming the system, raised by participants, would be to take a picture of a picture. For these reasons, the C2PA does not claim to verify the truthfulness of an image, but to provide information on its origins.

10. Hard bindings. C2PA Specifications. See: https://c2pa.org/specifications/specifications/1.1/specs/C2PA_Specification.html#_hard_bindings (accessed 10 July 2023).
11. CNN. Simpson jury: We relied on tapes, not witnesses. See: http://edition.cnn.com/2008/CRIME/10/05/simpson.juror/ (accessed 10 July 2023).
12. Bennett, C. 2003 Improvements to Time-Bracketed Authentication. https://arxiv.org/ftp/cs/papers/0308/0308026.pdf (accessed 10 July 2023).
13. *Ibid*.

The ubiquity of digital content presents another technical challenge. Media outlets do not always produce their own content, but may purchase or reuse content from others (such as freelance journalists or eyewitnesses). There is no guarantee that these external providers will have manifests for their content. To do so would require all image-capturing devices globally to have embedded provenance capabilities or, more realistically, for manual verification to be undertaken by media outlets. Given the speed at which misinformation content can spread online, this limitation may also undermine the efficacy of provenance labels as a solution.

The scale of content published online led some to state that social media platforms should ultimately be responsible for provenance labels in the long term and be incentivised to do so. However, this was met with the same challenges of subjectivity and lack of provenance capability in some users' devices. Despite this, manifests created by social media platforms could have utility when it comes to tracking the origins or dissemination of misinformation content online.

The example of YouTube's work on copyrighted content (which involves detecting and blocking copyrighted content uploaded by those without rights to do so) was raised as a concept similar to the provenance solution for misinformation. Although YouTube has had some success in doing this, they struggle to detect every infringement, despite having substantial resources as a major global platform. There were also questions around what should be done once manipulated content has been detected. It would be controversial and potentially undesirable to simply remove content in the way YouTube blocks copyrighted material.

Authentication of content and of provenance labels was the most significant technical challenge raised by participants. The World Wide Web Consortium's 'Verifiable Credentials' specification was highlighted as a possible model which could be replicated for content provenance. However, this specification is designed for verifying clear credentials (such as driver's licenses and passports) and relies heavily on there being an issuer of credentials[15]. The use of blockchain was also put forward as a potential solution, making use of a peer-to-peer network to verify the originality of an image, rather than relying on a single authority[16].

On Adobe's services, the company signs on behalf of the content creator and can act as an identity authority for its users with Adobe IDs effectively stamped onto outputs. Participants debated whether this could be replicated in wider society, given the countless image editing applications in existence. It may be feasible on a local level (for example, allowing the BBC to verify the identity of its own journalists), but on a national or international level this may require numerous identity authorities, presenting challenges to practicality and privacy. In addition, there is likely to be a need for third parties to verify that credentials have not been manipulated. Participants questioned who these third parties might be and how this could work.

Ethical challenges

Significant ethical challenges arise from the use of content provenance tools; these include challenges of interpretation when determining 'truth', people's rights related to freedom of speech and the implications of gatekeeping content for an open internet. Key challenges that emerged during the discussions include: the risk that provenance labels will be interpreted as markers of truth; that well-resourced media outlets will form 'knowledge cartels'; and that provenance labels will foster conspiracy theories.

Participants discussed the meaning of truth in the context of combating misinformation. While the term 'fact' was being used as an objective item of information, it is often the case that, in a rapidly changing environment, online content represents an opinion rather than a truth. In a news media context, it can be sufficient to simply know where information is coming from if the truth itself cannot be verified. Efforts to verify the origins of content can therefore be a helpful heuristic for journalists seeking to validate sources and filter out misinformation[17]. However, if a provenance label is represented publicly as 'verified' or similar, there is a risk that people will understand this content to be trustworthy, even if the label is in reference to origin rather than truthfulness.

14. The Royal Society. 2022 The online information environment. See: https://royalsociety.org/topics-policy/projects/online-information-environment/ (accessed 10 July 2023).
15. World Wide Web Consortium. 2022 Verifiable Credentials Data Model v1.1. See: https://www.w3.org/TR/vc-data-model/ (accessed 3 April 2023).
16. OpenOrigins. 2022 How Digital Provenance Can Combat Disinformation. See: https://www.openorigins.com/post/how-digital-provenance-can-combat-disinformation (accessed 3 April 2023).

Should provenance labels become a heuristic for truth, there is an ensuing risk that truthful content without a provenance label attached (for example, content was produced on an old device) would be considered untrustworthy. This feature could also be manipulated by bad actors, a phenomenon known as 'the liar's dividend', in which genuine content is dismissed as disinformation by those with an interest in discrediting them[18].

Questions were raised about the audience for provenance labels. If the audience includes people who harbour mistrust for mainstream news organisations, provenance labels may have limited or no impact. This could also deepen distrust as provenance labels could be perceived to be an authoritarian tool of 'the establishment'. If the audience includes those who do trust mainstream news organisations, the effect is likely to be marginal. However, building on Laura Ellis's presentation, participants discussed the cost of inaction. In the case of news media organisations, there are already attempts to mimic legitimate news (by using logos, fonts and other stylistic features) to spread disinformation. Examples highlighted people contacting news organisations and falsely accusing them of doctoring photos.

A system of using recognised news outlets as the basis for provenance labels presents another challenge: such a system would be akin to nutrition labels for trust and decisions over which media outlets should be included, which would likely attract controversy. Employees at trusted media outlets may also face higher pressure to produce accurate content as any mistakes may undermine trust—not just in a single media outlet, but throughout the system of provenance labels, affecting other media outlets as well.

Related to this was the risk that such a system might create a 'cartel' of trusted information sources with some news media outlets considered 'elite' compared to other outlets. These other outlets may still provide high quality journalistic content but have fewer resources and a lower profile. Local news outlets were identified as being particularly at risk. This risk could be exacerbated if provenance capability is used as a mechanism for filtering out content, or to influence ranking results in search engines.

Adoption challenges

Participants were asked to consider the adoption barriers of provenance technologies (aside from the issues of technical feasibility and ethical challenges.) The discussion centred around the themes of literacy, trust, and global content.

Should provenance labels become widely used, it will be essential to ensure the general population understand what they symbolise. This would require understanding: that provenance can assure trust and not truth; how labels cannot be falsified; and what an authentic provenance label looks like. This could form part of a wider literacy programme on misinformation or be led by platforms and media organisations. To aid people's understanding, and in light of people's experience with cookie banners, it was considered important that provenance labels should provide simple information and be user-friendly.

Finally, the challenge of global scale was raised. Misinformation content appears in all languages, cultures, and political contexts. Developing a provenance-establishing system which can account for this will be essential for its efficacy and will require global input.

17. A useful example can be found in Wikipedia's methodology to assess the reliability of various types of sources. See: Baigutanova A, Myung J, Saez-Trumper D, Chou A-J, Redi M, Jung C, *et al*. Longitudinal assessment of Reference Quality on Wikipedia [Internet]. Proceedings of the ACM Web Conference 2023; 2023. Available from: https://arxiv.org/abs/2303.05227. (accessed 10 July 2023).

A public service internet: what is it, and what could it look like?

The second day of the workshop included a discussion on an emerging model for a public service internet. It was based on a programme of work within BBC Research & Development investigating how the BBC can best deliver its public purposes in the internet age. As a trusted content provider in the online space, there is a strong argument that the BBC should actively work to uphold democratic values online and provide a trusted space underlined by principles of universal access, cutting through misinformation and abuse with an impartial voice.

The idea of a public service internet was first raised in 1998 at a Nordic Council conference of Public Service Broadcasters; this was an attempt to explore how the internet could be used to enhance public service television. Since then, the internet has hugely transformed the modern world, but was not developed with the requirements of public service organisations like the BBC in mind. As a result, there are aspects of today's network that make pose challenges to the delivery of public service outcomes over commercial outcomes: advertising is easy, while creating safe online spaces for debate remains difficult.

There has been broad and consistent BBC engagement with the internet since 1989, which has largely been defined by the desire to translate public service values into online behaviour. However, while the BBC has been an active user of the network, it has not had a significant influence on the internet's wider development. Over the last twenty years the nature of online space has changed dramatically and is currently monopolised by a small number of multinational, profit-oriented tech giants who may not prioritise user wellbeing.

The internet has also become a strong driver of behaviour and discourse offline, playing an integral role in democratic movements such as elements of the Arab Spring in 2010 and the ongoing Belarussian uprising. Governments have responded to this by trying to limit network access and imposing internet shutdowns. In 2016, the United Nations declared internet access to be a human right: an acknowledgement of how integral the internet has become in everyday life.

Since its first steps online in the late 1990s, the BBC's role has generally involved creating inspirational, high quality online services such as News Online and iPlayer, which have brought people online and set high standards for others to meet. It has also shaped underlying network standards in order to support the efficient delivery of content and services but has not been heavily engaged in other areas of the network's development or governance.

The BBC is now working with partners to explore the ways in which today's internet can be reimagined, changed or perhaps reinvented to support a digital ecosystem based on trust, accessibility, accountability and human values. Their goal is to identify potential changes to the way the internet operates and is managed, as well as to the applications and services it supports. They are collaborating with others to make interventions that drive these changes forward and create an online environment that aligns with the BBC's public service values.

There are currently many deterrents to the use of online services: people are concerned about spam, phishing, fraud and the impact on young or vulnerable users. Additionally, the relationship between providers and audiences is undermined by intrusive surveillance. Joining an online community entails the potential of being a victim of negative behaviour, and existing mechanisms to manage abuse, trolling, and hate speech are considered to be ineffective. Lastly, misinformation and disinformation are significant issues that need to be addressed on any public service network.

The Public Service Stack

In order to understand the role of provenance signals in countering misinformation, the BBC has outlined an architecture for a public service internet at all levels, from basic network connectivity to core protocols, data, applications and design, to governance and regulatory frameworks. They call this the Public Service Stack.

The Public Service Stack is similar to the Public Stack[19], but is more focused on the needs of public service media organisations. It also incorporates some of the thinking that informs the Public Media Stack[20], which is more concerned with classifying applications and services.

The Public Service Stack is made up of three layers::

Standards and protocols

This refers to the collection of network technologies that underpin the internet, the agreements around functionality and interfaces that create a functioning network of networks. The basic standards are over fifty years old, and while they have evolved, they have not been required to deliver public service outcomes.

Tools and applications

Tools and applications rely on the standards and protocols to operate. Every website, online service and network application makes use of the core network. While the public service internet does not necessitate direct development of these tools by the BBC, there is opportunity for the BBC to create public value by contributing to the ecosystem at this level and developing more of its own apps and services.

Governance, regulation, and social impact

Regulation—that is, rulemaking, rule monitoring, and rule enforcement, and governance: the system that provides a framework for managing an activity or organisation—are vital aspects of the public service internet. There are a wide variety of global attitudes towards managing the online space and how public service broadcasters already engage with this. This does not happen in a vacuum; political shifts, technology innovations and changing audience behaviours and interactions with the internet must all be considered in defining public value.

Provenance and the public service internet

As has been noted, misinformation and disinformation abound on the internet, a problem that is fundamentally social as well as technical. There are, however, ongoing efforts to combat these issues through technical means. The second-day discussions thus focused on how the work described on day one could be understood in the wider context of the public service internet.

In discussing the ways provenance-establishing technologies fit in to the public service internet model, contributions ranged more widely across related topics. There was particular concern around ensuring that any initiatives were fully inclusive, engaging a wide range of voices, including network users, younger people, and traditionally excluded groups. Some expressed concern that the thinking to date largely involved northern hemisphere organisations, and that those who had no or limited network access themselves were not being invited to shape the future internet.

There was also a conversation about the intended outcome, and a call to be much clearer about the end-state rather than talk in generalities about a ‘better’ network. What would it be like to engage with a public service internet? What types of applications and services might be possible, or not allowed? What would the benefits be for users of different types? At present, these objectives appear very aspirational but not very clear; it is difficult to identify where and how provenance technologies might have an impact or be best deployed.

One contributor noted that there were parallels between the public service internet and the BBC Micro (and micro:bit)[21]: the BBC was attempting to create a wider public good, using the tools at its disposal, but it could not, and should not, act alone. Calls for more detail would be better addressed by a wider (and more diverse) group rather than the BBC alone. Furthermore, can a notion of ‘ethics’ that works for everyone be agreed? What are the best ways to evolve this with an eye on citizen engagement?

18. Lewis *et al*. 2022 Do content warnings help people spot a deepfake? Evidence from two experiments. See: https://royalsociety.org/-/media/policy/projects/online-information-environment/do-content-warnings-help-people-spot-a-deepfake.pdf (accessed 10 July 2023).
19. Public Stack: Towards open, democratic and sustainable digital public spaces. See: https://publicstack.net/ (accessed 10 July 2023).

The Public Service Stack was generally well received, but some expressed the need for a more detailed model that clarifies: how overarching issues like privacy would be addressed across all levels; where the tensions between different outcomes might occur (such as privacy versus stopping criminal activity); and how they might be resolved. Persistent online identify is another key area to be addressed. This would serve not only media provenance work but also a wide array of use cases.

It is widely acknowledged that society is entering a period of regulation, with a key concern being the need to ensure a diversity of voices, approaches and expertise as new internet guardrails are built. How can the digitally-excluded and those who have been deeply immersed in the internet and its public service potential for many years be brought in? Where are the touch points that could help create a digital public good?

Lastly, some raised the issue of software being used to deliver regulatory work; a software's capabilities are limited to the capabilities of those who build it and their briefs. Also discussed was the 'transfer issue' of failures between ideation/design and delivery. A multi-functional team is needed to see such an idea through, from impact assessment to successful operation.

Annex

The Royal Society and the BBC would like to thank the following workshop participants who have contributed to the development of this note.

Workshop participants:	
Mansoor Ahmed-Rengers, University of Cambridge	Antonia Kerle, BBC
Henry Ajder	Alex Krasodomski-Jones, Chatham House
Rosie Akeroyd, DCMS Sub-Committee	Claire Levens, Ofcom
Victoria Baines, Demos	Andrew Lewis, University of Oxford
Charles H Bennett ForMemRS, IBM	Sonia Livingstone FBA, London School of Economics
Georgina Born OBE FBA, University College London	Bruce McCormack, Neural Transform
Ruth Brentnall, Reuters	Anna McGovern, BBC
Kate Brightwell, Adobe	Carl Miller, Demos
John Collomosse, Adobe	Steven Murdoch, University College London
Kevin Coombs, Reuters	Jade Nester, TikTok
Aidan Corley, Google	Judy Parnall, BBC
Diane Coyle, University of Cambridge	Andy Parsons, Adobe
Lucy Crompton-Reid, Wikimedia	Stephen Pattison, ARM
Jon Crowcroft FRS, University of Cambridge	Ziski Putz, Wikimedia
Andrew Dudfield, Full Fact	Brendan Quinn, International Press Telecommunications
Lilian Edwards, University of Newcastle	Yvonne Rogers FRS, University College London
Florence Enock, Alan Turing Institute	Nina Schick
David Frank, Microsoft	Nikki Soo, TikTok
Charlie Halford, BBC	Bertie Vidgen, Rewire
Ian Horrocks FRS, University of Oxford	Patrick Vu, Brown University
Pica Johansson, Alan Turing Institute	Adam Wright, British Academy
Ellen Judson, Demos	Erika Young, Reuters

Staff
Areeq Chowdhury, Head of Policy, Data, the Royal Society
Denisse Albornoz, Policy Adviser, Data, the Royal Society
Laura Ellis, Head of Technology Forecasting, BBC Research & Development
Bill Thompson, BBC Research & Development