The Alan Turing Institute

Response to the Information Commissioner's Office consultation on its Age Appropriate Design Code

Response of The Alan Turing Institute to the Information Commissioner's Office consultation on its Age Appropriate Design Code

This document provides the response of The Alan Turing Institute to the Information Commissioner's Office (ICO) consultation on its draft code of practice for online services likely to be accessed by children. The Institute's response combines the perspectives of various Turing researchers and those in its wider university network. A list of researchers who contributed to this response can be found in the Appendix.

This response will be split into two sections. The first will make comments on the overall code including why it is necessary and what we believe is missing from it, while the second will address specific areas of the code.

Section 1

We congratulate the Information Commissioner's Office on the important work that has produced this code.

Overall, we believe the code is well developed and is not in need of significant revision. We welcome that it has a focus on children's rights, as they have their own priorities and complex challenges which are distinct from those of users more generally.

Why this code is necessary

Children in the UK are engaging with digital technologies at ever younger ages.[1] More and more young children are mastering the alphabet and basic numerical skills through mobile apps. While the effectiveness of learning through such technologies has not yet been fully assessed, there has been widespread concern around the potential impact and harms of these technologies for children. Research[2] has highlighted the vulnerability of young people online in relation to being harmed by content they may see whilst browsing, repercussions of the content they post themselves, and being harmed by the treatment of their personal data for the purposes of filtering and personalisation. These harms can take a wider variety of forms and can have long term consequences.

Areas that require further research include children's exposure to online promotions[3], data tracking and data surveillance[4], and persuasive design.[5] Although not directly applied to children, research undertaken by contributors to this response looked at over 1 million mobile apps from the Google Play

[1] Kidron, Beeban, et al. "Disrupted childhood: the cost of persuasive design." (June 2018). 5Rights, and 'Children and Parents: Media Use and Attitudes Report', Ofcom, 29 November 2018.

[2] Projects on Digital Wildfire: (Mis)information flows, propagation and responsible governance (ESRC), and Unbiased: emancipating users against algorithmic biases for a trusted digital economy (EPSRC).

[3] Zhao, Jun, et al. "I make up a silly name': Understanding Children's Perception of Privacy Risks Online." In CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK. ACM, New York, NY, USA

[4] Lupton, Deborah, and Ben Williamson. "The datafied child: The dataveillance of children and implications for their rights." New Media & Society 19.5 (2017): 780-794.

[5] Kidron, Beeban, et al. "Disrupted childhood: the cost of persuasive design." (June 2018). 5Rights.

Store, which showed that third-party data tracking is ubiquitous with 9 out of 10 mobile apps sending personal information to Google without the explicit consent or knowledge of users.[6] Indeed, explicit consent is an issue that becomes more salient when considering children, as it must be determined whether they are able to give such consent.

Other research into gaming further confirms that developers can use machine learning to infer a child's developmental stage or personal interests from data about children's interaction in games.[7] In-app or online promotions have also been identified as one of the primary means by which children discover new apps or online media resources.[8] A major concern is that children largely lack the knowledge or skills to recognise the implications of these promotions. As a result, they are more likely to opt for playing the promoted games, which can lead to negative experiences.

Parents who have been involved in research have been surprised at the extent of tracking that occurs[9], and expressed a strong desire for better control and more transparency in the tracking by providers from mobile devices.[10] They indicated that this is where regulatory frameworks should step in, and the code is an important part of that.

This code is therefore of great importance to the safeguarding of children operating in online environments and interacting with online services. There are nonetheless some considerations missing from the code that it ought to address.

Circumvention of regulations and technologies by children

The code does not address attempts by children to get around the proposed approaches for age-appropriate design. Children are increasingly tech-savvy, operating across multiple devices, identities and platforms. They may dislike parental and platform oversight, and perceive it as a limitation of their on-platform experience. Tactics to avoid any (perceived or actual) constraints include (i) using fake identities (with an over 18 date of birth presented) to access platforms, (ii) engaging in adversarial behaviour on-platform (e.g. manually setting data sharing settings to 'low privacy' and not informing their parents) or (iii) moving to niche unregulated platforms and the dark web. Older children (for example, those over the age of 11) are particularly at risk of these behaviours. There is a real risk of unintended negative consequences here if the new code standards are not implemented with input from children. Implementation must also be sequential so that the impact of the standards can be monitored.

---

[6] Binns, Reuben, et al. "Third party tracking in the mobile ecosystem." Proceedings of the 10th ACM Conference on Web Science. ACM, 2018. See also Binns et al. "Measuring third party tracker power across web and mobile". TOIT. 18 (4) p52.

[7] Newman, Joe, and Joseph Jerome. "Press Start to Track Privacy and the New Questions Posed by Modern Video Game Technology." AIPLA QJ 42

[8] Zhao et al. "What privacy concerns do parents have about children's mobile apps, and how can they stay SHARP?", KOALA Report 1. 2018

[9] Van Kleek, Max, et al. "Better the devil you know: Exposing the data sharing practices of smartphone apps." Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems. ACM, 2017.

[10] Zhao et al. "What privacy concerns do parents have about children's mobile apps, and how can they stay SHARP?", KOALA Report 1. 2018

Abuse of children's data by other users of online services

The code focuses on protecting child users from undue manipulation and exploitation from platforms and online service providers. However, it does not address the possible abuse and exploitation of children's data by other users of these services or platforms. Children would benefit from safeguards both to prevent platforms from manipulating users' data and invading privacy, and to prevent users of those platforms from exploiting each other's data. Both types of data manipulation should be covered by the code, because they both affect children's ability to make informed decisions and maintain their privacy, making them vulnerable to data abuse. Interactions between users occur within the platform environment, which makes this a platform-level problem.

Specifically, there is a considerable risk of children being doxxed, trolled or cyberbullied by both adults or children, and by people who are known or unknown to them. These activities often, but not always, involve data abuses such as (i) viewing or accessing users' data without their permission; (ii) sharing personal images and content without their permission, and (iii) encouraging users to click on malicious, illegal or harmful content (i.e. 'clickbait'). This can also transfer over into, and reinforce, offline forms of abuse. Platforms should provide children with support in tackling and preventing this sort of data abuse. At present, this is only limitedly addressed in standard 5 ('Uphold your own published terms, policies and community standards') and also standard 15, which discusses the need for 'Data protection impact assessments', which partly covers bullying (detailed in p. 82-88). Note that these issues affect all online users – but they are particularly worrying for children as they may not be well equipped to handle them and the negative effect on their emotional and mental wellbeing could be greater.

The other side of this is the issue of children receiving unwanted communication from other users of online services, for example explicit pictures. This issue of data or information transfer that flows not from children to others, but from others to children, does not appear to be covered in the code.

Digital literacy

The code's sixteen standards largely rely on those reading them to have a prerequisite understanding of ethics, privacy, data and security. These are complex concepts, which require younger users to be educated so that they can understand the guidelines created for their own safety. While there are multiple existing initiatives that offer such training, they will require funding for significant upscaling. For this, cross-sector engagement and cooperation is of fundamental importance and must continue to be supported.

Range of harms and long-term consequences

The code must recognize the different forms harms through the use of online services can take for young people, and the consequences they can have long term. It should take a wide-ranging view of practices that can void or mitigate harms when they occur. This should include the promotion of children's online resilience, in the sense that they can recover when they have encountered harms and develop strategies to prevent or deal with similar risks in the future.

Furthermore, a code of practice such as this should take a long-term view of harm and harm avoidance. The consequences of data tracking, personalization and other techniques are not necessarily seen or felt immediately, and can be cumulative. Steps to protect children therefore need to incorporate a long-term view and be put in place even when harm or the risk of harm cannot be discerned in the short-term.

Section 2

This section addresses various parts, but not all, of the code.

About this code

Under "Who is this code for?", use of the term 'Information Society Services (ISS)' may not be the most appropriate way to describe the audience of this document. It may be more appropriate to define the audience of this document more specifically as the providers of 'online products or services' including but not limited to:

- Apps
- Programs
- Websites
- Games
- Community environments
- Digital educational tools
- Connected toys or devices with or without a screen including personal trackers
- Smart home devices such as personal home assistants
- Any type of networked payment systems (e-wallets).[11]

Services covered by this code

As above, it may be beneficial to extend the services and products covered by this code. In its current iteration, the "At a Glance" section does not appear to include smart home devices, personal assistants, or networked services and products used in schools. The average knowledge of parents or guardians about personal data is relatively low. If systems are to be built that safeguard children's personal data (and their well-being), online services and products embedded in the functioning of physical places such as smart home devices and various networked products used at schools must be included.[12]

Additionally, "Information Society Services" is a vague term especially for the target audience of this code. Although various government and public organisations use "Information Society" to refer to contemporary societies, "Information Society Services likely to be accessed by children" can also include libraries or online databases. Therefore, to limit the possible variability in interpretation, the broadly defined products and services can be made more explicit. Established technology companies are unlikely to have any difficulty in interpreting this code. However, those with more limited experience in technology and its ethical implications designing services and products accessible by children for the first time may struggle to understand whether it applies to them.

1. Best interests of the child

Although the best interest of the child can be defined clearly, it is open to interpretation by the designers (and parents). There is also potential for this standard to be used to justify some potentially

[11] This includes ones used in schools across Britain for children to pay for their lunch with their fingerprint, or any other type of biometric data.

[12] Research carried out by Dr Didem Ozkul, a contributor to this submission, at schools in London including interviews with parents and teachers supported the conclusion here that the code should cover certain products and services used in homes and schools.

unwanted uses of personal data at the expense of the best interest of the child and in favour of various commercial interests. Further consultation could help explore how one defines the best interests of the child, along with a definition's potential limitations, and which steps should be followed to make sure that this standard is not exploited for commercial purposes.

3. Transparency

It is not necessarily clear how standards can be developed for age appropriate design, or how individual approaches should be evaluated, when developing content that will "attract and interest children". Multiple sources of subjectivity can be introduced when developing "diagrams, cartoons, graphics, video and audio content, and gamified or interactive content" to convey critical information. For example, reading diagrams or graphics may depend on children being able to interpret them, and also depends on the capacities for the creators to design that content appropriately. Each of these media are highly flexible and can take many forms which will not always be aligned with the goals of this code. For example, interactivity itself can mis-direct attention or disorientate users as much as it can aid understanding. This being the case, an increase in subjectivity in design relating to this code may introduce some challenges for developing and maintaining standards. Appropriate investment will be needed to evaluate and learn from the diversity of responses that could arise from this code.

User testing may play a key role in developing appropriate content and designs. That testing should be representative of the context of use if it is to support age-appropriate design, and also recognise that there may be substantial diversity in literacies within age bands and that multiple factors may impact on the interpretation of content. This testing should also measure the level of genuine engagement with this content and design. This is another area where academic involvement could be of assistance, especially in thinking about how user engagement with community guidelines can be reinforced, especially at the point of account creation.

Furthermore, one potential way to enhance the proposal for platforms to be "clear, open and honest with your users" is to provide an extension to this definition, adding a reference to a second 'responsible' individual who holds oversight or accountability over the child user. This second individual might not even be on the platform. Platform guardianship extends to those that the users are accountable to in the physical world. Those in a position of responsibility offline also need to understand the user/platform relationship in a transparent and clear way. This definition extension could be "clear, open and honest with your users and, when required, to those responsible for your users".

7. Data minimisation

The inclusion of data minimisation is welcome, however this is somewhat in conflict with the call for implementing more effective age verification. At present, robust age or person verification requires collecting more data, such as real-time photo verification, address, alternative platform use, or passport. More resources should be invested to understand how the requirement of data minimisation can be balanced against the need to collect certain types of personal data for user verification. Based on existing available technology, these ideas are mutually exclusive.

In addition, the data collected should be proportional to the service used. That is, mechanisms must be developed that notify users in a transparent way regarding the amount and type of data that is being collected for the particular service they would like to use. For example, if a user wants to simply stream music, the amount of data collected and processed will be different to that required if the user also

wants to find out about concerts happening in their area. An interface is required that makes it clear to children, as well as adults, what data is being collected depending on what they are trying to do with the service. This also makes it easier for providers to collect only the data required for a specific service.

9. Geolocation

The definition of geo-location data can be expanded to cover other means of obtaining the location of a user/device including the sensors of a device, such as near field communication and accelerometer, and other wireless connectivity such as cellular connection (e.g. GSM[13]) and Bluetooth. Additionally, the products and services described in this section may take into account connected toys and smart home personal assistants along with their sensors, which may also support geofencing.

Further, the code outlines that "Options which make a child's location visible to others should default back to off at the end of each session", which suggests it is focused on mobile apps. A 'session' can mean the use of an app or online service such as social media, and unless one quits an app or uses certain privacy-enhancing settings for browsers, the session may continue in the background. This becomes particularly problematic with connected toys, wearables and smart home hubs. Greater detail should be provided for how this is to be implemented for difference types of technologies.

10. Parental controls

An unintended consequence of parental controls is the knowledge by children that they can be tracked by their parents or guardians.[14] This can be psychologically harmful to the development of a child and their identity. This is not to say that children should not be made aware that they are being tracked by their parents through notifications. But the way the notifications are designed and the frequency with which children receive them should consider the age of the child, the child's personal development, and the rights of the child.

11. Profiling

The GDPR's transparency rights (Article 13–15) have been construed by some as a right to an 'algorithmic explanation', requiring that processors of data provide meaningful information about the logic of processing of certain data analysis of individuals (profiling) and automated decision-making based on these methods. In general, transparency is typically only required when decisions are 'significant' and 'solely automated', which are steep barriers in relation to children. In particular, 'significant' effects of personalisation and data analysis on individuals might be more about the effect they have on a person's environment over time, rather than a single incident that does not align with a person's expectations.[15] This would mean that technologies which might shape a child's environment over time, and significantly affect them that way, would be required to provide more granular information about the ways they function. This approach would be in line with other parts of the GDPR, especially Article 24, which states that the protection should become more stringent as the risks of data processing increase. All information should be presented in a form also amenable to oversight by

[13] Global System for Mobile Communications.

[14] This is based on ethnographic fieldworks with users of smartphones, connected toys and wearables.

[15] Sylvie Delacroix and Michael Veale (2019) Smart Technologies and Our Sense of Self: Going Beyond Epistemic Counter-Profiling. In: Life and Law in the Era of Data-Driven Agency (OUP 2019).

third parties, to avoid the 'transparency fallacy' risked by placing the burden solely on children to understand and enforce their rights.[16]

Further, the code appears to be targeted mostly at mobile and online services and products. If a code such as this one is designed only with a limited number or variety of technologies in mind, it will be more difficult to implement. For example, in some cases, profiling of children can start even before they are born through the use of maternity apps, and continues through the use of parental apps after birth.

13. Connected toys and devices

The standard on connected toys and devices would benefit from being divided into three separate categories:

1. Connected toys and wearables
2. Smart home devices and hubs
3. Services designed solely for use in education settings (for example, the use of children's fingerprints to pay for lunch at school)[17].

The consultation document also asks whether this standard requires a transition period of any longer than 3 months after the code comes in force. The code's implications may have an impact on the actual manufacturing of toys and devices as well as the software updates. It would therefore be reasonable to expect that it may take longer than 3 months to transition to implementation.

14. Online tools

The code would benefit from further clarity around data rights for children. Much of data processing that occurs in relation to children is undertaken on the basis of 'legitimate interests'. This is particularly the case with processing that would normally be carried out by contract, but cannot because children are not legally old enough to contract. Using 'legitimate interests' as the basis for processing children's data allows for the use of the right to object, which is not possible when processing under 'necessary for contract'. Therefore, providers should ensure that older children can object to the use of their data where they have this lawful basis to do so.

If providers' focus is on adult users who fall under a particular lawful ground for processing their data, they may not have considered that they need to provide an interface that allows children the right to object for certain processing operations. The right to object should be clear and actively provided to children at relevant points, and balancing tests should not require significant justification, which might quickly get legalistic and not be in line with the overarching approach of fairness. An electronic signal should be provided so that children and parents can automatically object to certain processing activities, for example by the use of plug-ins and browser signals.

The right to access is an important feature. Human–Computer Interaction (HCI) research can help build 'sensemaking' capacities into 'download my data' tools, noted as important to the successful

[16] Edwards L and Veale M (2017) Slave to the Algorithm? Why a "Right to an Explanation" is Probably Not the Remedy You Are Looking For. 16 Duke L Tech Rev 18.

[17] See https://www.theguardian.com/sustainable-business/2016/feb/19/surveillance-state-fingerprinting-pupils-safety-privacy-biometrics

application of the GDPR[18], to allow children to explore, erase, restrict and object to certain uses of data without having to navigate CSV files. APIs could potentially be used, so that third party platforms with the sole purpose of helping children use access, can help children understand and manage data across multiple devices and platforms.

Annex A: Age and developmental stages

Promoting age appropriate design in this field is extremely complex; there are nuances around determining and applying relevant age ranges for children, and it can be very difficult to identify and assess all potential risks and harms. Indeed, the developmental stages outlined here may not be appropriate in a digital, and therefore global, world. The terms used to define the age brackets and stages are rather loaded, and there is some research evidence to suggest that children in their early teens can be more sophisticated than adults in their understanding of the online world, and strategies for privacy.[19]

Further, the ethical imperative to protect children is weaker the closer the child reaches an adulthood threshold. That threshold is affected by other conventions, such as age thresholds for sexual consent, marriage and military service. In the UK, 16-year-olds are able to join the army with their parents' consent and are treated as honorary adults. Therefore, differences in stages of development must be accounted for and guidance provided for online service providers.

In early childhood, more granular categories may be required to reflect the different cognitive and physical developmental stages for children. The use of connected toys and wearables such as smart baby clothes, heartrate and sleep monitors require in particular the 0-5 category to be subdivided into 0-2 and 3-5 years.

[18] Veale M, Binns R, Van Kleek M (2018) Some HCI Priorities for GDPR-Compliant Machine Learning. CHI-GDPR 2018. https://ssrn.com/abstract=3143705

[19] See for example Boyd, D. (2014) *It's Complicated: The Social Lives of Networked Teens*. Yale University Press.

Authors

The following researchers contributed to this response:

Greg McInerny (Turing Fellow, University of Warwick)

Jon Crowcroft (Turing Fellow, University of Cambridge)

Michael Veale (Digital Charter Research Fellow, The Alan Turing Institute)

Alex Harris (Research Assistant, The Alan Turing Institute)

Bertie Vidgen (Research Associate, The Alan Turing Institute)

Cosmina Dorobantu (Deputy Director Public Policy Programme, The Alan Turing Institute)

Didem Ozkul (University College London)

Jun Zhao (University of Oxford)

Helena Webb (University of Oxford)

Tom Sorrell (University of Warwick)

turing.ac.uk
@turinginst