SUPER AWESOME

SuperAwesome contribution to ICO Consultation: *Age Appropriate Design Code*

CONTENTS

1. AGE VERIFICATION

Principle: Service providers should apply best practices and best available technologies to verify whether their users are children, and implement appropriate protections.

To date the focus of regulators has been on improving protections for children when they are accessing child-directed online services. But before these rules and best practices can be effective, we need to address the elephant in the room—the fact that most kids' digital activity is on services that do not know whether their users are children, or believe that they are not children, or pretend that they are not children. Whilst the GDPR requires services to make a reasonable effort to verify the age of users, the truth is that current best practice in this regard is wholly inadequate.

It has been widely reported that age gates are frequently circumvented or ineffective[1], while most available methods of verifying a user's age[2] are based on increased data collection and hence incompatible with the GDPR's data minimisation principle.

In this section we set out how an age-appropriate design code can mitigate some of these issues, by guiding the industry toward best practices and pointing the way to technology solutions for age verification.

a. Age gates

The use of age gates in games and kids' services has proliferated in recent years, mainly because in the U.S., the Children's Online Privacy Protection Act (COPPA) allows services that consider themselves mixed-audience (as opposed to primarily child-directed) to segregate their audience into kids and adults by asking the users' age. Whilst well-intentioned, this approach has had three significant unhelpful effects that have made kids less secure online:

1. Having been exposed to dozens of age gates, most children have realised that '13' is a magic threshold that unlocks a grown-up experience. As a result they have learned to lie about their age in order to get access to those services, in particular to social media platforms. This has in many cases been exacerbated by parents who have become complicit in helping their kids set up profiles on 13+ platforms.[3]
2. The mixed-audience concept has allowed thousands of services (especially games) to avoid being categorised as child-directed, and hence to continue using data-driven monetisation strategies that end up profiling children on a vast scale.[4]

[1] More than 80% of children lie about their age to use sites like Facebook (*The Guardian*, 26 Jul 2013)

[2] Age verification – An age-old problem (*Parcelhero blog*, 1 Sep 2017)

[3] Why parents help their children lie to Facebook about age: Unintended consequences of the 'Children's Online Privacy Protection Act' (*FirstMonday*, 7 Nov 2011)

[4] This issue has been widely reported in 2018, starting with the publication of an academic research paper into the data collection practices of Android apps: Thousands of Android apps potentially violate child

3. The application of age gates or—in the case of social media such as YouTube, the use of Terms of Service limiting use to those 13 or over—has effectively allowed service providers to absolve their legal responsibility regarding children and reduced their incentives to improve the way they protect child users.

Our design guidance seeks to mitigate the impact of these effects, as a bridge to the industry developing improved technologies and methods for verifying the age of users.

First, the bias should always be toward protecting all users of a service that is likely to have many children using it. If the content, design and marketing of a service means it will be attractive to children, then it should not be permitted to try to segregate its audience using an age gate. This aligns with the GDPR's current definition of an "ISS offered to a child".

Second, if a service is primarily not child-directed, and is not actively marketed to children, then the use of an age gate can help isolate the minority user base that needs to be treated differently. In that case, however, it is critical to design the age gate to be 'neutral', which means it elicits the actual age of the child and avoids encouraging kids to lie about their age, specifically:

- The question is posed in a neutral manner, eg: "How old are you?" rather than "Are you over 16?"
- The age selector is designed in such a way that it does not default to an adult age, and does not encourage selection of an adult age.

Examples

In order of best to worst, see below examples of age gates currently being used by different apps and sites:

Best. The default setting is neutral (0 or --) and the question asked is neutral: 'What age are you?' or 'What is your age?'. The method to select your age is in a native scroll or a very clear and easy to use interface:

protection law (*The Guardian*, 16 Apr 2018); followed by a complaint filed in the U.S. under COPPA: How Game Apps That Captivate Kids Have Been Collecting Their Data (*New York Times*, 12 Sep 2018)

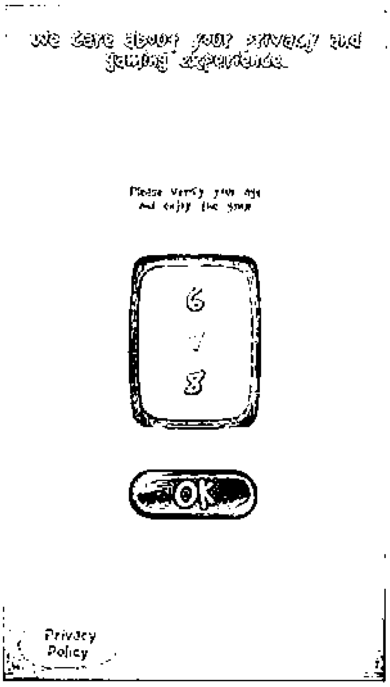

we care about your privacy and gaming experience

Please verify your age and enjoy the game

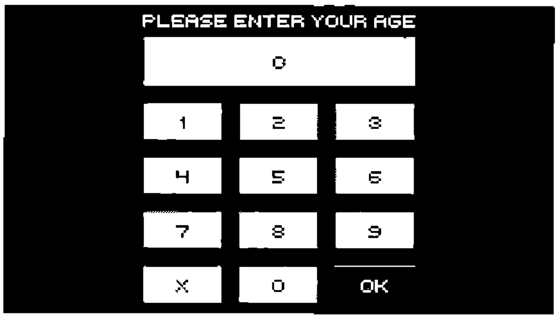

PLEASE ENTER YOUR AGE

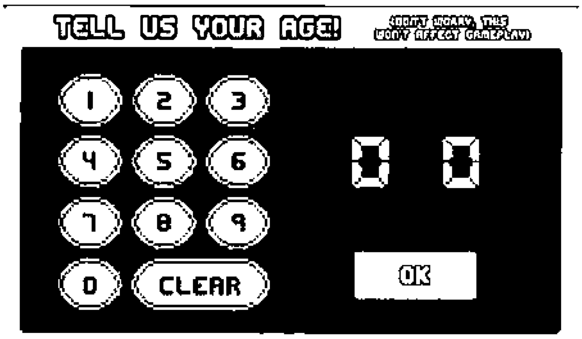

TELL US YOUR AGE! (DON'T WORRY, THIS WON'T AFFECT GAMEPLAY)

Good. The default setting of the age is set to a number under 13 and asks the question, 'How old are you?':

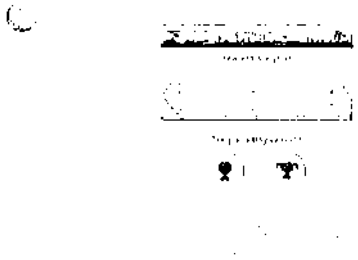

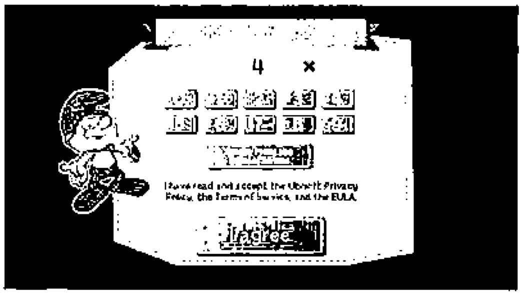

I have read and accept the Ubisoft Privacy Policy, the Terms of Service, and the EULA.

Good. The question asked is, 'When were you born?':

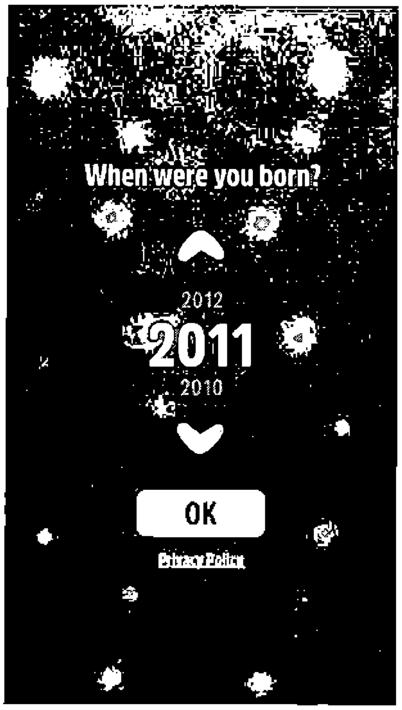

When were you born?

Bad. The question is good ('How old are you?'), but the default setting is set to over 13:

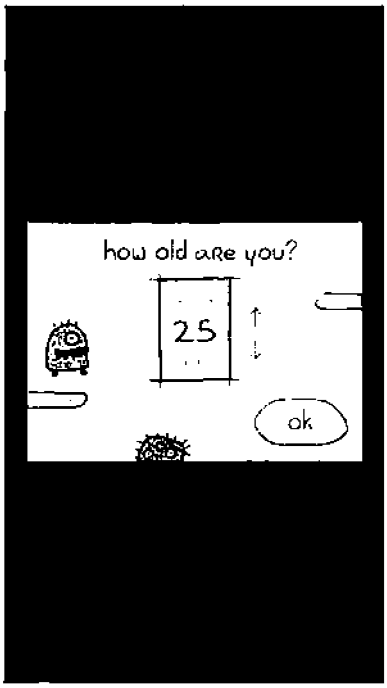

Bad. The child is asked whether they are over a the age threshold. Children quickly realise that unless they answer ‘true’, they will not be able to access the service, so they are highly incentivised to lie:

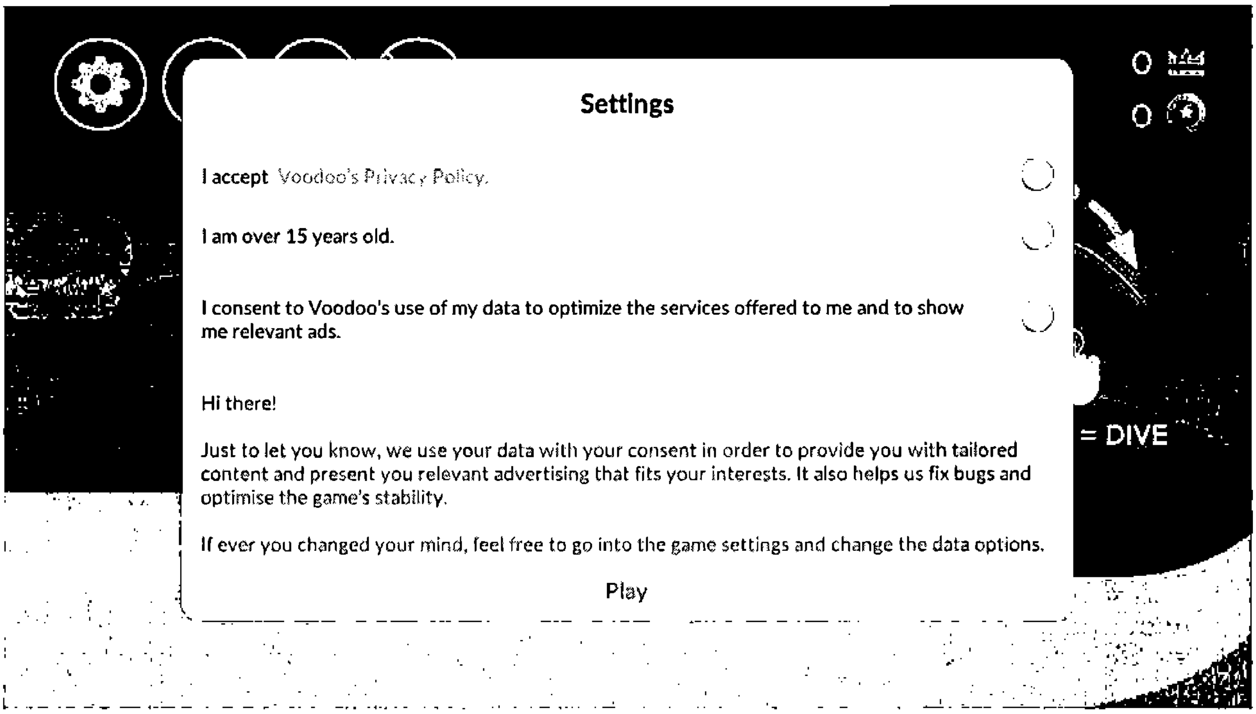

Settings

I accept Voodoo's Privacy Policy. ○

I am over 15 years old. ○

I consent to Voodoo's use of my data to optimize the services offered to me and to show me relevant ads. ○

Hi there!

Just to let you know, we use your data with your consent in order to provide you with tailored content and present you relevant advertising that fits your interests. It also helps us fix bugs and optimise the game's stability.

If ever you changed your mind, feel free to go into the game settings and change the data options.

Play

Bad. Similar example, from the recent update to WhatsApp in the EU, designed to bring them into compliance with GDPR:

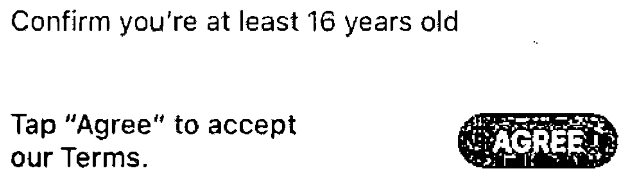

Confirm you're at least 16 years old

Tap "Agree" to accept our Terms.

b. Bumpers and parental gates

The flip side of age verification is how to segregate an adult section within a kids' service, or how to protect kids from leaving the service and following a potentially unsafe link.

The first mechanisms in widespread use arose following the 2014 settlement between Apple and the Federal Trade Commission (FTC) in relation to in-app purchases made by children on Apple devices that were logged into their parents' iTunes accounts.[5] Apple agreed to reimburse affected parents, and to implement more stringent methods to obtain consent for purchases. As a result, both Apple and Google amended their policies regarding children's apps.

In particular, Apple has since required all apps that are listed in the Kids category must "not include links out of the app, purchasing opportunities, or other distractions to kids unless reserved for a designated area behind a parental gate."[6]

In our experience, parental gates can be an effective way to ensure children stay within the safe environment of the experience designed specifically for them, especially when the age range of the audience is narrow enough for a particular age gate design to be effective. For example, a common approach is to pose a maths question that mostly kids over the relevant age can only answer:

[5] Apple Agrees to FTC Consent Decree Over In-App Purchases (Memo) (*Recode*, 15 Jan 2014)

[6] See section 1.3 Kids Category, App Store Review Guidelines (Apple)

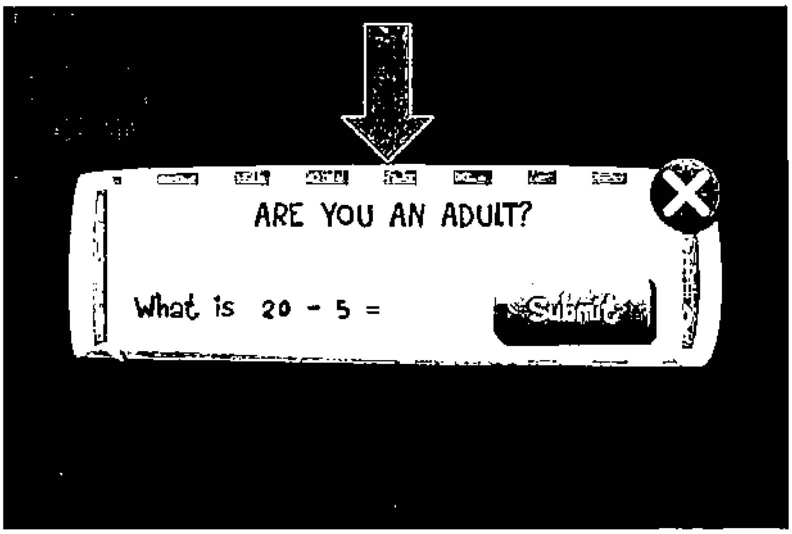

ARE YOU AN ADULT?

What is 20 - 5 =

Another approach commonly seen in pre-school apps or sites is to ask the user to complete a physical action that only those able to read can understand:

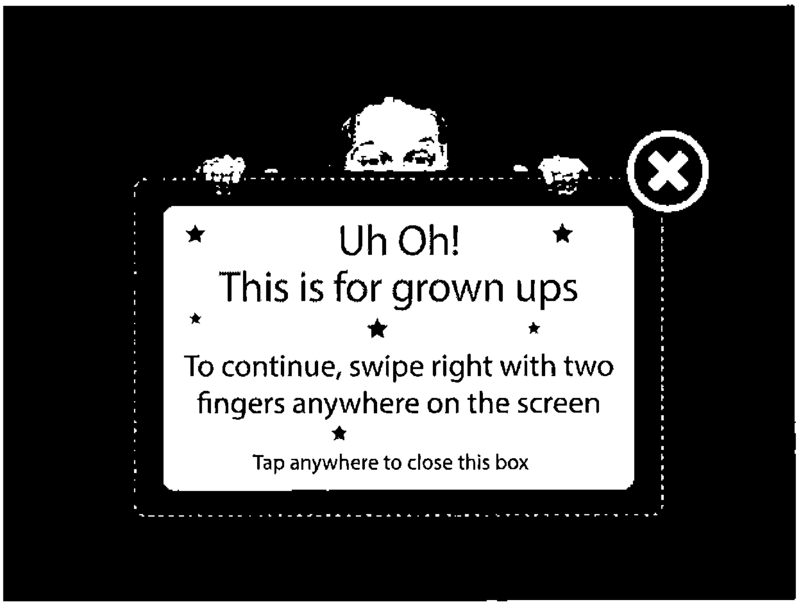

Uh Oh!
This is for grown ups

To continue, swipe right with two fingers anywhere on the screen

Tap anywhere to close this box

These methods work well for younger audiences,[7] but they tend to become increasingly ineffective as children reach the age of, say, 9 or 10 years.

Whereas children's apps on Apple devices are bound by this requirement, the universe of other kids' experiences—on the web and elsewhere—is far less regulated. For these and for services

[7] For more examples, please see How are kids' app developers communicating to parents? (*Moms with Apps*, 20 Aug 2013)

that appeal to older children, we always recommend content owners as well as advertisers to apply a warning 'bumper' as a matter of best practice.

This means that whenever a link leads out to a non-kids site—such as an ecommerce site, or a social media platform—a bumper window pops up for a few seconds with a clear notice that the child is leaving a safe zone and entering the broader web, for example:

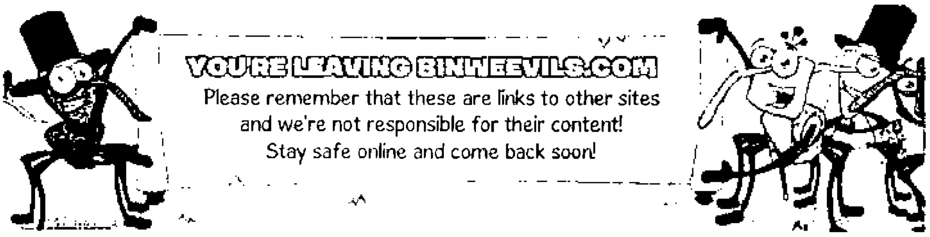

YOU'RE LEAVING BINWEEVILS.COM

Please remember that these are links to other sites
and we're not responsible for their content!
Stay safe online and come back soon!

In general we believe that the best combination of safety and self-responsibility comes from providing children and parents with the clearest possible guidance on the safety of their internet experiences and encouraging them to communicate about their web browsing activities.

c. *LOOK AHEAD:* The future of age verification technologies

Article 8 of the GDPR requires service providers to "make reasonable efforts to verify" a user is over the age of consent. The ICO's own consultation earlier this year on Children and GDPR highlighted the challenge of complying with this requirement given the state of current technologies, *in particular* without collecting yet more personal data potentially in conflict with the data minimisation principle.

We believe that, for the most common data processing activities, any age verification technique that requires the collection of more personal data—such as a national ID card, or a national insurance number—is overly intrusive and antithetical to consumer's increased desire to protect their privacy. In addition, these types of verification can only really be implemented to gate a specific service the user wants to access. But that use case represents a tiny fraction of the many opportunities to collect personal data from users as they browse the web. It is not feasible, for example, for a website to conduct this level of verification before deciding whether to serve an advertising impression.

We believe that the next generation of age verification techniques must be based on automated ways of detecting whether a user is likely a child or not, and providing that assessment, along with a confidence score, to the service provider so that they can implement appropriate policies.

Our research team at SuperAwesome is currently testing a number of approaches that could be promising solutions. As the largest kidtech platform in with the world, reaching some 500M kids per month around the world, we have been able to recognise patterns of behaviour that are most likely to be associated with a child, based on aggregated anonymous usage data. When a user arrives on a given website or app, our system can look at multiple signals (locally, on the device, without collecting the data)—whether the device has been seen before in our marketplace, how often, what is the current time of day, etc. Using these signals (and without relying on any personal data or specific browsing history), an algorithm can determine the probability of the device being, at that moment in time, operated by a child or an adult parent.

Such a dynamic score can then be passed to the publisher, who can verify whether it matches what the user provided via the age gate, or who can potentially avoid age-gating altogether and simply tailor the experience for the child to be safe. Due to the real-time nature of our analysis, such a signal can also be used by advertisers—for example those who specifically do not want to reach children—to stop the serving of (and data collection from) an impression if the 'child score' is positive. This would address one of the biggest challenges conscientious advertisers face when attempting to comply with the CAP Code to avoid targeting children.[8]

Whilst our work is still in its early testing stages, we believe this approach carries a lot of promise and we would welcome the opportunity to share our findings and proposed solutions with the ICO in due course.

2. PRIVACY CONTROLS & NOTICES

Principle: children and parents should fully understand and be able to control the data collection practices of digital services.

The GDPR's principles of transparency and informed consent reflect widespread recognition that privacy notices have become (or always were) ineffective. Those on the most widely-used social platforms can run to 5,000 words of densely-written legalese. It was common for adults to have no idea how their data might be used by different service providers. Most sites require that you accept terms and conditions before you can use them, and many users will consent without reading or understanding what they are consenting to[9].

The GDPR's transparency requirement means that any child-directed site must have privacy notices that are easily read and understood by children. This is challenging, in particular across different age categories, and so we set out below our recommendations regarding the content and language of child-friendly privacy notices; how to design and 'layer' notices; easy-to-use privacy controls; and mechanisms for ensuring appropriate parental involvement.

a. Content

[8] A good example is this recent ASA ruling in favour of Walker Snacks Ltd, which highlights the difficult and extensive effort HFSS brands have to make to try to avoid reaching kids: ASA Ruling on Walkers Snacks Ltd (22 Aug 2018)

[9] Click to agree with what? No one reads terms of service, studies confirm (*The Guardian*, 3 March 2017) You're not alone, no one reads terms of service agreements (*Business Insider UK*, 16 Nov 2017)

First, we must reduce the GDPR's notice requirements to the most critical elements that are relevant to the child user, and then attempt to translate that disclosure into language the child can understand. We suggest the following sections (with examples of language):

Exactly what the ISS's approach to data collection is

It's important for publishers to set out their data collection 'philosophy' in order to give context and comfort to the user.

For example: We'll never ask you for personal information, but our app needs to collect some data from the way you use it in order to work. We'll always tell you what we're collecting and why, and we'll do our best to keep your information safe. You can help by not sharing any personal information on the app!

Exactly what personal data is being collected

Within this section the types of data should be detailed, and explained in simple terms.

For example: We need to collect your email address and username to create your account, and information about your device so that we can make the app look great.

Why their personal data is being collected

The user should be able to identify the purpose of processing, whether it is required for the service to work, to improve features, or to deliver advertising, etc.

For example: We collect non-personal info to give you the best app ever, so it looks good, contains everything you love and we know how to help you with any bugs.

If and how their personal data may be shared with third parties

For example: If the police or government ask us to help stop or investigate a crime we may have to give them your username and internet address.

The rights of the user and how they can exercise them.

For example: You or your guardian can look at, change, correct or delete any information about you on the app. Just ask your parent or guardian to contact us.

b. Design of notices and 'layering'

It is important to consider how the information contained within the policy is presented.

Consent needs to be informed, that is users need to know what they are consenting to and why. A balance needs to be struck between giving sufficient information and not overpowering the user with a 'wall of words' which could have an adverse impact on readability, particularly for younger readers.

For this reason layered privacy policies should be encouraged, whereby key statements or information are offered in a concise manner with the option for the user to review fuller text if desired. Layering could also be supported through images, video or other graphics

Topline information, such as that suggested above, can be linked using "hover" functionality, or click-throughs, to a more comprehensive document.

The ICO has provided useful examples of layering, so we won't repeat them here, except to say that this approach works particularly well to address the challenge of communicating privacy notices to children.

Another good example is that applied by uSwitch:

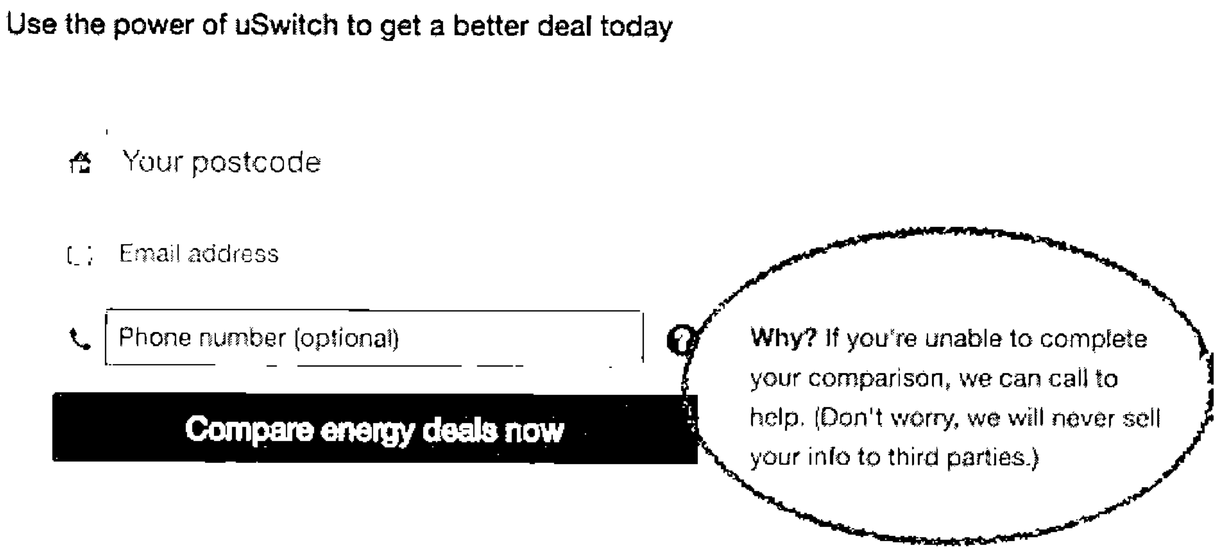

Use the power of uSwitch to get a better deal today

Why? If you're unable to complete your comparison, we can call to help. (Don't worry, we will never sell your info to third parties.)

In general, we are seeing more frequent attempts to innovate when it comes to bringing privacy to the attention of younger audiences, either by adding a separate child-friendly statement to a regular privacy policy (for example, Beano's Privacy Policy), or a specific child-facing information page such as the PopJam privacy policy re-written for kids. Other examples of child-friendly privacy policies are:

- TutoTOONS
- SuperCell

This remains an area that is challenging for publishers, in particular independent content owners who may have limited ability to work with legal teams to develop the best approach. This is a topic where we strongly encourage the ICO to provide as much guidance, specific frameworks and comprehensive examples as possible.

c. Privacy Controls

Regardless of whether an app or website is general audience or child specific, clear labelling of privacy controls should be encouraged.

Most apps and websites offer some degree of user control comprising one or more of:

- user profile visibility
- user blocking
- search history deletion
- notification management
- ad enablement
- cookie management

In terms of design, there is no current uniformity among toggles, controls or dashboards for users to exercise personalisation or privacy controls. This creates an opportunity for a code of practice to offer a common interface so that children and parents can easily identify when a control is offered, what choices are available and whether a given choice has been made. Given that a child may have several apps on their device(s) with essentially the same control features in different manifestations across those apps, identifying and selecting those controls may be overwhelming or confusing, particularly at scale with different toggles, buttons, colours or sliders.

For example, if a user has exercised a choice to have a feature active or inactive, that choice should be clearly identifiable—we recommend clearly labelled large toggle switches (see example below from the BBC).

Using large bold titles and short descriptions coupled with symbols such as ticks and positive colours in familiar colours, (e.g. green for on, greyed out, red for off) should help children easily understand whether they have a choice, the selection options they have and (at a glance) understand what choices they have made.

Each control should be clearly labelled. Plain, simple language such as 'on or off' should be used instead of adult-themed vocabulary such as 'enabled' and 'disabled'. The label should be placed in close proximity to the control. An example of a child-directed application (Roblox) where the language is quite adult-oriented and the toggle controls distant from the description shown below:

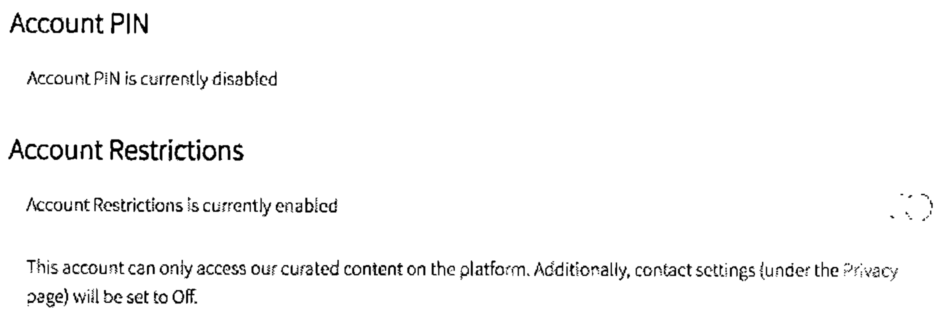

Account PIN

Account PIN is currently disabled

Account Restrictions

Account Restrictions is currently enabled

This account can only access our curated content on the platform. Additionally, contact settings (under the Privacy page) will be set to Off.

The target audience of the service and the sophistication of general users should also be taken into consideration when designing interfaces. Toggles may be supported by additional visuals—such as a padlock or graphic to show that a feature has been unlocked or locked. For example, Animal Jam's controls visually explain some of the settings. Although in this example, the descriptions supporting each toggle are high level and may not provide sufficient information for a user to make an informed decision:

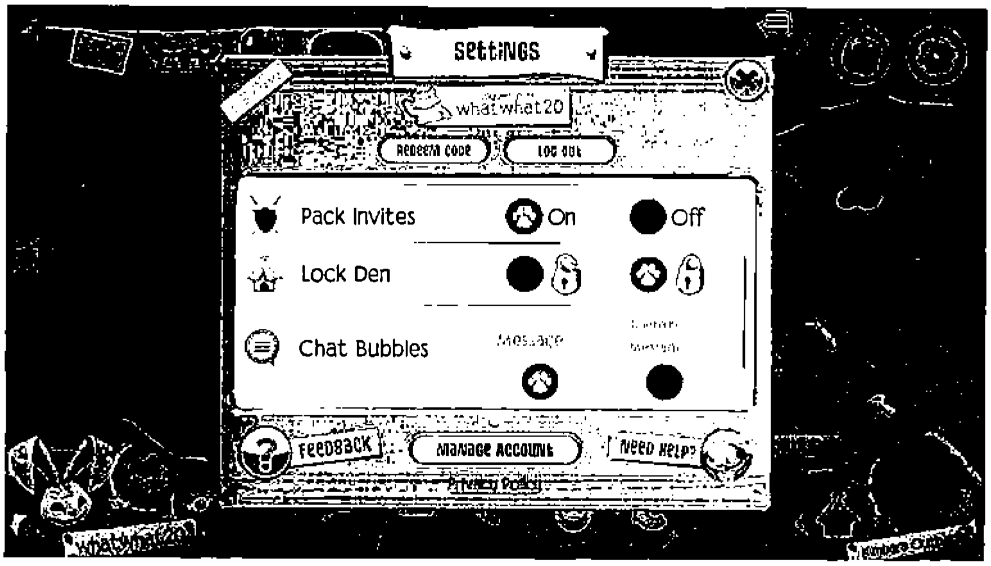

Even if there is no specific call to action for the user, the use of colour and graphics is a useful approach to express key messages, whether house rules or safety messages. For example see Club Penguin's approach to engaging users with regard to their house rules:

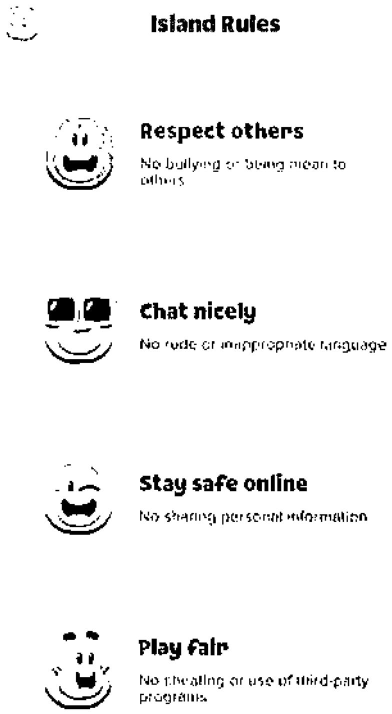

Island Rules

Respect others

No bullying or being mean to others

Chat nicely

No rude or inappropriate language

Stay safe online

No sharing personal information

Play fair

No cheating or use of third-party programs

Settings Options

There is a balance to be struck between offering granularity of choice and having a sufficiently clear dashboard to enable users to exercise their choices in a timely manner. There is no definitive approach whether a single page of controls (with simple descriptions), or several pages of more detailed controls is the right approach for younger audiences. For example, Animal Jam offers a clear but light set of controls:

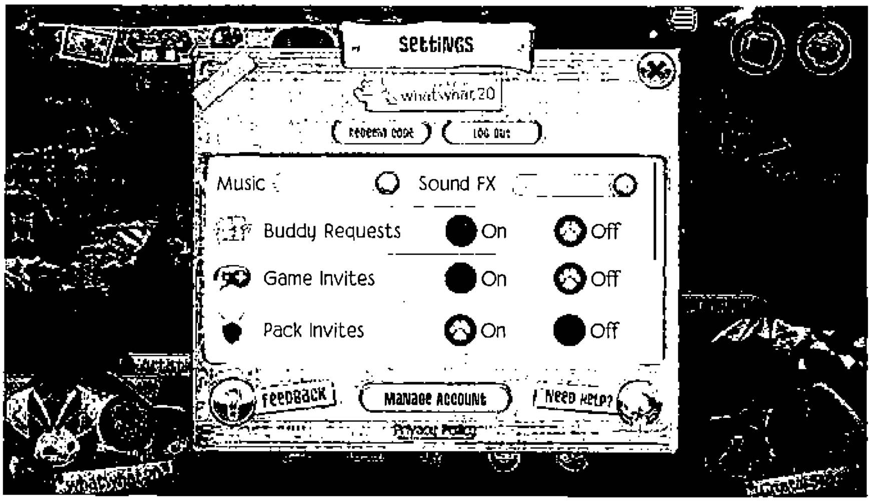

Music ○ Sound FX ○

Buddy Requests ● On ○ Off

Game Invites ● On ○ Off

Pack Invites ○ On ● Off

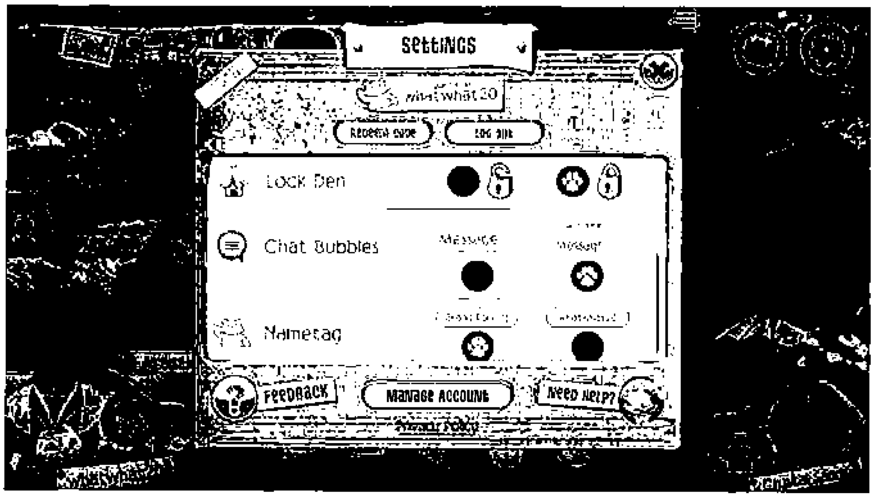

Lock Den ● ○

Chat Bubbles Message ● Message ○

Nametag ○ ●

As an example of apps designed for older audiences, Wooz World offers tick boxes and detailed descriptions where the target audience reading age lends itself to more detailed descriptions:

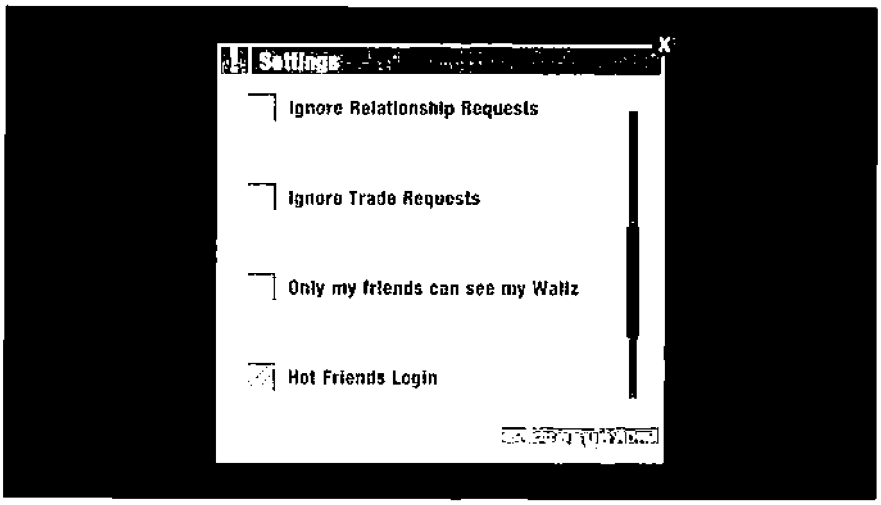

□ Ignore Relationship Requests

□ Ignore Trade Requests

□ Only my friends can see my Wallz

□ Hot Friends Login

Balancing the need to inform users of their choices with the age-based understanding of the audience is challenging, but we recommend the following best practice approach:

For all age ranges privacy controls should have clear bold titles and succinct descriptions. For pre-school users images may help illustrate their purpose. Toggles are a clear and familiar way to give control to all age groups. For 3-12 year olds they should be coupled with ticks and a positive green colour to clearly indicate that setting has been set to 'On'. Plain, simple language such as 'on or off' should be used beside the toggle. These recommendations should help kids easily understand what state the settings is in and make it obvious if they adjust a setting.

3. PARENTAL CONSENT

Principle: the parental consent requirement should be implemented in proportion to the risk of data processing, whilst not unduly reducing children's access to digital services.

a. Proportionality

According to the GDPR's Article 8, where the legal basis is consent and the data subject a child:

> *you must make reasonable efforts, taking into consideration available technology, to verify that the person providing parental consent does, in fact, hold parental responsibility for the child. A reasonable effort [...] might therefore entail simply asking for a declaration that the user is old enough to provide their own consent, or a declaration of parental consent and responsibility, via a tick box or email confirmation.*

The Guidelines on Consent under Regulation 2016/679 (wp259) also recommend a "proportionate approach" when (a) confirming that a data subject is over the age of digital consent; (b) seeking parental consent; and (c) establishing the parental authority of the consent provider. This is to ensure sufficient verification whilst minimising the collection of personal data.

The GDPR's risk-based approach to verification is welcome. It echoes the US experience with proportionality under COPPA. In order to facilitate the design of appropriate user consent flows, we recommend the ICO include in its guidance specific examples of how to match appropriate levels of verification to the actual risk of the data processing. This is particularly important in light of the GDPR's data minimisation requirement.

Based on our experience of working with children's online services (as well as building Verified Parental Consent (VPC) workflows and technology for COPPA compliance), we propose practical framework along the lines below (examples only, not exhaustive) as a design guideline:

Type of data being collected	**Sensitivity**	**Examples of sites or apps**	**Appropriate method to verify user is over age of consent**	**If not, parental consent required?** **Appropriate verification method**
Sensitive Personal Information (health, ethnicity, tied to a name or ID number, etc)	Very high	Ancestry or healthcare service that stores user profiles with identity information and demographic/ ethnic/health data.	Neutral age gate, plus Database check against national registry, or Copy of photo ID submitted	**Identity-Verified Parental Consent (w/ database)** 1. Parent provides consent 2. Statement by parent that he is the holder of parental responsibility; 3. Parent identity checked against national ID database, or by submitting copy of photo ID
Identifiable personal information, eg full name, address, national ID number; image/video uploads; free text content. Combination of online identifiers and profile information that can be used to identify a natural person	High	Social media app that allows use of real names, connections with strangers, free-text chat rooms Virtual assistant that records voice & stores it in cloud, builds useage profiles.	Double confirmation, eg Neutral age gate, plus Reconfirmation of birthyear; or, Two-factor confirmation, eg Neutral age gate plus Confirmation provided by email or text message	**Identity-Verified parental consent (no database)** 1. Parent provides consent 2. Statement by parent that he is the holder of parental responsibility; 3. Identity is confirmed by requesting credit card details and matching them against information provided (no transaction). Credit card information is then immediately deleted.
Technical online identifiers that cannot easily be resolved to a natural person, but are (a) shared with third parties, and/or (b) used for behavioural advertising & profiling, including geo-location	Medium	Websites that allow behavioral or profile-based advertising. Virtual world, or games app that includes username registration, leaderboards	Double confirmation, eg Neutral age gate, plus Reconfirmation of birthyear; or, Two-factor confirmation, eg Neutral age gate plus Confirmation provided by email or text message	**Verified Parental Consent** 1. Parent provides consent; 2. Statement by parent that he is the holder of parental responsibility.

Creation of a unique username (not PII)				
Enabling of notifications (eg, push) City-level geo-location information	Low	Apps that request permission to send push notifications; provide services based on city location (eg transport)	Confirmation that subject is over age of consent, via a simple, neutral age gate	**Direct Notice**. Opt-in, and direct notice sent to parent, stating type and purpose of collection and linking to Privacy Policy. No further verification of parental holder of responsibility
Technical online identifiers used for internal operations purposes only (analytics, contextual advertising, personalisation, security) Country-level geo-location information	Low	Casual games site with no registration, only contextual advertising	Processing on **Legitimate Interest** basis. No age check required.	Processing on **Legitimate Interest** basis. Parental consent not required. n/a
No data collection	None	Corporate website for marketing purposes, no advertising, no trackers	No age check required.	Parental consent not required. n/a

All of the above is of course subject to the prerequisite that the ISS meet the transparency requirements, in particular when it comes to notices children can understand.

Example 1: educational website that finances itself primarily through advertising.

If advertising is delivered only contextually and no cross-domain tracking is allowed, then this represents Low sensitivity and would not require age verification or parental consent.

Publisher would have to ensure all technology and advertising partners are aware of child-directed nature of site and is responsible for guaranteeing that they are not collecting technical online identifiers that could be used to profile users. Social media plugins would not be allowed.

Example 2: mobile social application that enables chat, connecting with friends, sharing content under real names.

Use of real names, open text chat and the ability to connect with strangers make this High Sensitivity, eg a service that requires age verification and/or verified parental consent.

Example 3: virtual world that allows interactions between anonymous avatars.

Provided measures are in place to prevent disclosure of personal information (eg filtering out real names or phone numbers in unmoderated channels or chat rooms), then this represents Low sensitivity, with no verification or parental consent required.

Example 4: voice-based virtual assistant, or Internet-connected toy.

Given that audio files are likely to be stored and analysed in the cloud, and it is not technically feasible to filter out personal information in moderation, this represents High sensitivity and should require both age verification and Verified Parental Consent.

If the service provider can demonstrate that it is using any collected audio files solely for purposes of transcribing a command, and immediately deletes the audio files thereafter, we may consider this case Medium sensitivity, requiring only a simple opt-in + Direct Notice to parents.[10]

b. *LOOK AHEAD:* Data minimisation via portable, universal verification

The experience of COPPA in the U.S. has helped establish several effective methods for obtaining parental consent and verifying the identity of the parent. Many of these will be useful in applying the GDPR's Article 8 requirements for verifiable parental consent.

However, there is one crucial difference between COPPA and the GDPR that significantly limits the applicability of the parental verification methods most widely in use today—the principle of data minimisation.

The most common methods of verifying parental consent under COPPA are: completing a credit card transaction or checking a parents' ID card or number against a national database. Both of these methods require the collection of personal data, even when the data processing activity that triggered the consent requirement is far lighter (such as an email for password reset).

Second, the verification requirement also introduces a significant blocker in the accessibility of services. Our experience in powering the consent flows for many apps and websites shows that for every 100 children that request access, between 70 and 90 never get it (depending on the verification method). In other words, the hurdle imposed by seeking parental consent and then verifying that parent's identity is so great, the vast majority drop out. This effect directly contradicts another key principle of EU law, which is to ensure universal access and not to unduly prevent children from benefiting from innovation in digital services.

[10] We recommend as a best practice following the recent guidance from the U.S. Federal Trade Commission (FTC) on how virtual digital assistants can comply with COPPA: https://www.ftc.gov/news-events/press-releases/2017/10/ftc-provides-additional-guidance-coppa-voice-recordings

For these reasons, we believe the industry should be encouraged to promote a strategy that includes **universal, portable parent verification**, one in which the parent's identity has to be verified just once, and where that verification is reusable across many services that would like to obtain consent under Article 8.

Such sharing of a verification status would mean parents do not repeatedly have to upload sensitive data to verify they are indeed the guardian of their child. At the same time, it would provide sufficient audit trails to meet the verification requirements.

There are multiple ways to make this work, but one approach is as follows:

1. The first time a parent signs up to go through a parental consent process (triggered by publisher 1), they provide their email and verify their identity however they want to (credit card, identity card or other)
2. Publisher 1 adds a simple entry to a permission-based central database, registering the parent's used e-mail address as well as the level of verification performed (email+, credit card, identity card, etc.)
3. The second time the same parent goes through a parental consent process (triggered by publisher 2), they merely provide their email address and verify the ownership of this email address by clicking on a link that has been sent to them.
4. In the background, publisher 2 places a simple API request to the central database described in step 2, realises the parent has already verified their identity and operates under this assumption.

Note that this process needs to be accompanied by transparent notices, informing the parent of the actions they have performed and the permissions they have given. This communication should both be instant (when a change in permission configuration occurs) as well as periodical (informing the parent of the current state of permissions).

The benefits of such a centralised verification approach are clear:

1. Significantly reduced personal data collection from parents
2. Higher conversion rates for children and hence greater access to new digital services

We believe there is scope for the regulator to encourage the creation of such an industry initiative by setting out the standards and legal frameworks to make sharing of verification status an acceptable method under GDPR. We would welcome working with the ICO and other bodies to develop such a concept further.

4. MONETISATION & ADVERTISING

Principle: **advertisers should be able to reach kids, and publishers should be able to monetise their kids' audience effectively, without collecting any personal data.**

By far the largest category of personal data collection from children (without consent) occurs in conjunction with digital advertising. As children browse the web, use their favourite apps as well as many services designed for adults (eg Facebook) and social media services like YouTube, their every action is recorded by a web of 'adtech' companies building profiles for marketing purposes.[11] Much of this data collection is out of the scope of the regulations designed to protect the data privacy of children, because the sites and app they are browsing are not caught by the regulation, or are social media platforms where a child is logged in by a parent, or are being accessed through shared devices.

We believe that designing the age-appropriate infrastructure—or kidtech—for zero-data monetisation of kids' websites and apps, plus improved age verification technologies accessible to advertisers (see Section 1) should be a key component of the Age-Appropriate Design Code.

Second, unfiltered and unmoderated digital advertising is a major source of inappropriate content exposure, both on kids' sites and apps as well as on the many adult services, including Google search and YouTube, commonly used by children[12]. Even publishers who use the tools provided by leading ad networks to limit advertising categories to those considered 'safe' for children report continuing incidences of inappropriate ads being shown.

Any guide to designing age-appropriate services for children must include design of **zero-data ad delivery**, as well as design of **adverts themselves**, including guidance on calls-to-action and click destinations (landing pages), and labelling and transparency of ads to enable children to identify their commercial intent. Some of this is encapsulated in the ASA's Cap Code 05 Children and in the more recent Recognition of advertising: online marketing to children under 12, but we believe there is scope for more guidance for both advertisers and publishers.

a. Zero-data advertising delivery

Our research—updated with 2017 figures from Ofcom's 2017 *Children and parents: media use and attitudes* report—indicates that by the time they complete the age of 12, the typical UK child has had more than 78m data points collected by trackers included in digital ads. This is based on:

[11] See: Friend or foe? The rise of online advertising aimed at kids (*The Guardian*, 28 Feb 2014); Marketing online to kids in the age of GDPR poses new challenges (*Martech Today*, 4 Jan 2018); Adtech firms collecting 'vast amounts' of data on kids despite online regulations (*The Drum*, 13 Dec 2017)

[12] See: Displaying inappropriate ads in a children's game is a bad idea (but it'll happen anyway) (Gamerlaw, Jun 2015); ASA: Advertising to children: Inappropriate Advertising (*RPC*, 3 Oct 2016); YouTube 'urgently' removing ads from inappropriate videos of kids as big brands freeze ad spend (*The Drum*, 24 Nov 2017); YouTube pulls ads on 2 million inappropriate children's videos (*The Verge*, 28 Nov 2017)

- Time spent online per day by age category—ranging from 1.1 hrs (3-4 year olds) to 3 hrs (12-year olds) (*Ofcom*)
- An estimated average exposure of 11,250 ads per month for a typical user[13], equating to 61 ads per hour based on the average monthly Internet and social media use[14].
- Our own analysis of advertising tags in programmatic platforms where the average advert contains 15 trackers, and each tracker typically collects 14 data points.[15]

As a result, children using the internet in a typical way can expect to be exposed to between 377,000 and 992,000 trackers each year, collecting between 5,000,000 and 14,000,000 data points annually:

Age	No of hrs spent online per day±	Trackers per year	Data points collected each year
3	1.1	376,915	5,276,805
4	1.1	376,915	5,276,805
5	1.3	434,168	6,078,345
6	1.3	434,168	6,078,345
7	1.3	434,168	6,078,345
8	1.9	639,324	8,950,530
9	1.9	639,324	8,950,530
10	1.9	639,324	8,950,530
11	1.9	639,324	8,950,530
12	3.0	992,383	13,893,360
			78,484,125

There are two primary reasons why data collection from kids remains so prevalent:

1. The majority of sites and apps where kids spend their time are not in scope of kids' data privacy protection either because they have (thus far) avoided being classified as 'offered to children' or because they rely on an ineffective age gate mechanism; or,
2. They rely on industry signaling standards, such as the COPPA flag, which notifies other vendors involved in ad delivery that the ad request comes from a child and personal data

[13] Yes, There Are Too Many Ads Online (*HuffPost*, 9 Feb 2017)

[14] Adults spend almost 8 hours each day consuming media (*IPA*, 21 Sep 2017)

[15] SuperAwesome's *AwesomeAds* platform connects with major ad exchanges (such as Rubicon and AppNexus) and actively filters out tracking codes. The number of trackers per ad is based on data generated by AwesomeAds over the course of 12 months in 2017. Each tracker was found to collect—at a minimum—a unique persistent identifier (such as device ID, IP address, IDFA), user agent consisting of multiple data points about the device or browser configuration (operating system version, language settings, etc), a URL address, geolocation and likely one or more custom events specific to that tracker's purpose.

should not be collected. These signaling standards tend to be ineffective as most vendors simply pass on the signal along with the personal data.

What we suggest is that in the kids space, technological solutions should be employed that **guarantee** no data is able to be collected by any of those parties, such as to prevent any potential misuse.

Whilst in theory the GDPR's provisions on Automated Decision-Making and Profiling together with Recital 38 expressly prohibit this type of data collection for profiling or marketing purposes, in practice it is rampant. In order to reduce the level of personal data collected from children, advertisers must design **zero-data delivery** into their advertising strategies, and publishers must have access to **technology and to legal frameworks** to ensure they can monetise without allowing their child users to be profiled by third parties.

We make recommendations for both of these use-cases below.

(i) Designing zero-data ad campaigns

Designing age-appropriate advertising campaigns starts with the advertiser's data strategy, which must be differentiated from typical adult-targeted advertising. The biggest kids' brands, such as the toy companies or entertainment companies, already practice an approach which includes designing campaigns for data minimisation; using dedicated kid-safe campaign landing pages; making use only of contextual targeting approaches; relying on the fewest possible third parties to minimise the risk of data leakage; and, enforcing human moderation of every ad creative.

Our long-standing guidance to advertisers and their media agencies is to apply the **zero-data standard** that satisfies both US and EU data privacy regulations, namely:

1. Do not engage in behavioural or interest-based audience selection and any retargeting or remarketing strategies.
2. Do not work with or link the campaign to any third-party Data Management Platforms (DMPs), as these are designed to profile all users and do not have effective mechanisms to protect the data of children.
3. Brief your suppliers to select audiences based solely on context, by choosing sites or apps that are known to cater to the desired age demographic.
4. Whenever possible, use zero-data ad serving technologies and infrastructure, e.g.:
    a. ad serving platforms designed to not collect persistent identifiers, such as those built specifically for compliance with the kids' data privacy regulations;
    b. ad filtering technologies that actively strip trackers from ad tags and/or remove personal information from ad requests—this will protect you from configuration errors or breaches by partners.

5. Minimise the number of 'hops' in your advertising delivery chain, the number of tag wrappers used for a single ad, and the number of third-parties involved in delivery, targeting and measurement. Ideally create your own 'clean' ad tag (with only your required performance trackers, such as click counters) and deliver it directly to kids' sites/apps, or a specialised kids' ad network.
6. Require that every third-party supplier involved in delivering your campaign (DSP, exchanges, SSP, verification vendor) confirms in writing:
    a. their 'actual knowledge' that this is a kids' campaign;
    b. they understand their obligations under relevant data privacy laws; and,
    c. they will not profile users or share data related to any views of the campaign.
7. If accepting ad tags or trackers or any other code from partners, audit these to ensure they do not include 'piggy-back' or 'zombie' trackers.
8. If you are buying inventory programmatically,
    a. Bid only on sites/apps, not on user profiles; or buy only from guaranteed kid-safe sources[16]
    b. If any personal data (including unique identifiers) is passed back to your systems, discard it or ensure it is treated compliantly by:
        i. making no use of it for profiling/targeting;
        ii. ensuring it is not shared and no other party can access or use it;
        iii. having documented, clear processes in place to prove you are treating it compliantly; and,
        iv. keeping your privacy notice up-to-date to describe your approach.
    c. Regularly audit your delivery chain for potential data leaks back to your clients or partners.
9. Respect kids' publishers who have their own stringent policies on data collection and seek to protect themselves from 3rd party trackers.

(ii) Technologies and legal frameworks for publisher compliance

Our anecdotal evidence suggest that the majority of new digital content creation for children is funded by advertising, and that proportion is growing. As the number of digital subscription services multiply, it becomes more and more difficult for new kids' services to persuade parents to pay directly for quality content. For this reason we expect advertising to remain the primary funding source for kids digital content—whether for education or entertainment.

When responsible kids' publishers want to generate revenue compliantly, they face difficult choices. Few if any existing sources of advertising they can connect to—such as Doubleclick

[16] Note that the so-called 'COPPA flag' which many ad exchanges pass on from publishers cannot be relied on to identify or avoid kids' inventory, as it (a) is inconsistently applied by publishers, (b) often erroneously used by exchanges to block inventory as a whole, and (c) is purely a flagging mechanism that does not prevent personal data from being passed among parties.

(Google), OpenX or Facebook Advertising Network—are able or willing to guarantee (a) that personal information will not be collected, or (b) that no inappropriate ads will be shown, even if publishers self-declare as services directed to children.

For publishers' **monetisation strategies**, age-appropriate design means:

1. **Clearly defining (and segregating) your audience**—is it primarily under-16s or a mixed audience? If mixed, how will you segregate the adults from the children—using an age gate or sign-posting clearly which section is for which audience? For more on this, see *The GDPR-K Toolkit for Kids Publishers Part Two: Defining your audience*.
2. **Choosing a technology approach you can control**—will you operate your own ad serving platform or outsource monetisation to a partner? In either case, it is critical to set out and document your policy on how to protect your audience from personal data collection—whether by ensuring correct configuration of your ad serving platform, or contractually committing your outsourced monetisation partner. For more on this, see *The GDPR-K Toolkit for Kids Publishers Part Four: Safely monetise your site or app*.
3. **Sticking to your guns**. Advertisers and agencies will always push for more data collection and the use of additional third-party trackers, and generally will not know about the zero-data requirements of kids' publishers. Hold firm and insist on the best available technical and contractual measures to prevent personal data collection from your users, and to ensure age-appropriate content only. If your partner can't or won't provide appropriate guarantees, find another source of advertising revenue.

Thanks to the emergence of the kid-tech market[17], publishers are gradually finding more technical solutions to help them remain compliant whilst enabling monetisation and growth. But the legal protections available to them remain inadequate. In the runup to the GDPR enforcement deadline this year, many publishers were rushed into signing Data Protection Agreements that leave them with all the responsibility of obtaining consent from their users and few legal protections against unscrupulous data collectors.

We believe that kids' publishers need a legal framework that allows them to inform upstream advertising partners that they are child-directed, and that shifts the onus of ensuring (a) no PII collection, and (b) appropriate content to that partner. This is akin to the 'actual knowledge' concept under COPPA, which has proven reasonably effective.

As a framework, we propose empowering publishers with a standard contract template they can use with partners, along the following lines:

> We operate a [website/app] directed at children and young audiences (the "**Site**"). You deliver advertising and/or use trackers or pixels ("**Trackers**") to provide analytics, verification or other

[17] Won't somebody please think of the children? The 'kidtech' space is about to explode, led by European startups (*TechEU*, 25 Feb 2015)

services to advertisers. You wish to place advertisements or Trackers on our Site in a compliant manner, and so the parties agree the following:

1. We hereby notify You—and you hereby confirm your actual knowledge—that the Site is directed at children and young audiences, including those aged under 16, and are therefore subject to specific data privacy protections relating to children.
2. We agree to permit advertisements to serve on our Site, or Trackers to be included in advertising tags, provided You:
    a. use personal information (including persistent identifiers) collected solely for the purpose of providing aggregated anonymised performance metrics, analytics, verification or similar services to your advertiser clients;
    b. do not associate any data collected from users of our Site with any personally identifiable information from any source;
    c. do not use data collected to: (i) create user profiles; (ii) undertake behavioural targeting, re-marketing or re-targeting activity; (iii) attempt to locate or identify the same user on other sites or apps;

4. You grant us the right at any time during the term of this Agreement, upon request, to re-test and inspect Your ad tags or Trackers for privacy compliance. If an ad tag or Tracker is deemed to be non-compliant, at our sole discretion, you must immediately cease using and remove the non-compliant Tracker or ad tag.

b. Advertising Content

Thanks to the well-established self-regulatory framework of the ASA and CAP Code, children in the U.K. are among the best protected in the world when it comes to the design of advertising content. Chapter V of the CAP Code sets out the key principles that children are vulnerable and advertisers should avoid exploiting their credulity. This approach has been copied in other countries both independently and via the International Chamber of Commerce (ICC), and is widely regarded as global best practice.

However, as the digital markets continue to change, we believe some gaps have emerged which leave children exposed even when the CAP Code is fully applied. We recommend that guidelines for age-appropriate design of advertising extend the CAP Code to include some additional provisions, specifically:

1. **Avoid linking to non-child-directed sites**—many advertising campaigns directed at children, and running on child-directed sites, allow users to click out to a social media page, or an e-commerce site, or to the advertiser's home page. In nearly all cases, none of these sites are considered child-directed, and hence they are all likely to **(a)** collect PII from all visitors, and **(b)** contain content that is not age-appropriate. In particular, when children are using a shared device, they are highly likely to see advertising or content targeted behaviourally at a parent, which is often inappropriate, eg alcohol for example.

Child-directed adverts should either **(a)** always link to child-directed sites, ideally a landing page created specifically for this child-directed campaign—one with age-appropriate content, no third-party trackers, and only internal performance analytics; or **(b)** use a messaging 'bumper' to inform the child that it is leaving a safe site and going to an adult page for which it should seek permission from a parent. An alternative to the latter is to use a so-called parental gate—posing a math question or other challenge only an adult can answer—before transferring the user to an adult social media site or similar.

2. **Enforce human moderation of every advert**—the rapid growth of programmatic distribution of advertising means that millions of adverts are being served to users without any matching of the ad content to the context where it is being displayed. As has been widely reported, many advertisers do not fully know where their ads appear, and most publishers have little knowledge or control of what ads are being delivered to their sites by the major ad networks. Even when the relevant campaigns are in principle appropriate to children, it still means that—for example—11+ rated video games may be advertised on child-directed sites whose primary audience is 4-6.

   We recommend that children's adverts be required to be moderated *in relation to their delivery context* by a human being. In other words, a person on the advertiser side or publisher side should review each ad creative in full knowledge of where it will be served, in order to ensure CAP Code compliance and age-appropriateness to the site. Adverts and sites should be tagged by the moderator for the correct age bracket: pre-school, kid, tween, young teen.

c. Labelling & Transparency

Advertising campaigns that are appropriate for children must also comply with the GDPR's principle of transparency, which is primarily reflected today in the CAP Code's requirement that "marketing communications must be obviously identifiable as such".[18] When it comes to children, however, we believe that labelling 'commercial intent' is insufficient. The objective must be to:

1. Ensure the child understands that the advert has commercial intent,
2. Make clear which product or brand is being promoted, and,
3. Further inform the child (a) whether the ad's content including the click destination is safe, and (b) whether any personal data is being collected.

When the CAP Executive developed guidance in 2016 about enhanced disclosure for so-called 'integrated advertising', it mentioned research showing that most children from the age of 8 generally recognise the difference between advertising and editorial content (provided they are clearly separated). But given the continuing growth and extraordinary creativity of "native

[18] Recognition of advertising: online marketing to children under 12 (*ASA*, 28 Apr 2017)

advertising"[19], which is likely to further blur the lines, it is clear that a more rigorous approach to labelling kid-safe advertising is required.

In order to ensure children—and parents—can rely on a standardised approach to labelling and disclosure in advertising, we recommend that all children's ad campaigns be designed to include a watermark that identifies both its commercial intent and safety features for children. A useful and widely recognised example of this is the Safe Ad watermark that is automatically applied to every ad served through SuperAwesome's platform. The SafeAd mark signals to the user that:

1. This ad is not collecting any personal data (eg, it is not behaviourally targeted, nor collecting data for profiling purposes);
2. This ad has been reviewed by a human (a) for compliance with the CAP Code (or domestic equivalent), and (b) to ensure it is age-appropriate to the site the child is on.

The Safe Ad watermark can be applied across standard display, video and rich media advertising, as well as integrated or 'native' campaigns such as brand integrations into virtual worlds, product placements, or even influencer marketing. In all cases, it is able to meet the CAP Executive's requirement that disclosures be (a) interruptive, (b) prominent, (c) understandable by children, (d) link to media literacy materials.

See specific usage examples below:

Interactive video:

[19] Native Ad Spend Will Make Up Nearly 60% of Display Advertising in 2018 (*eMarketer*, 11 apr 2018)

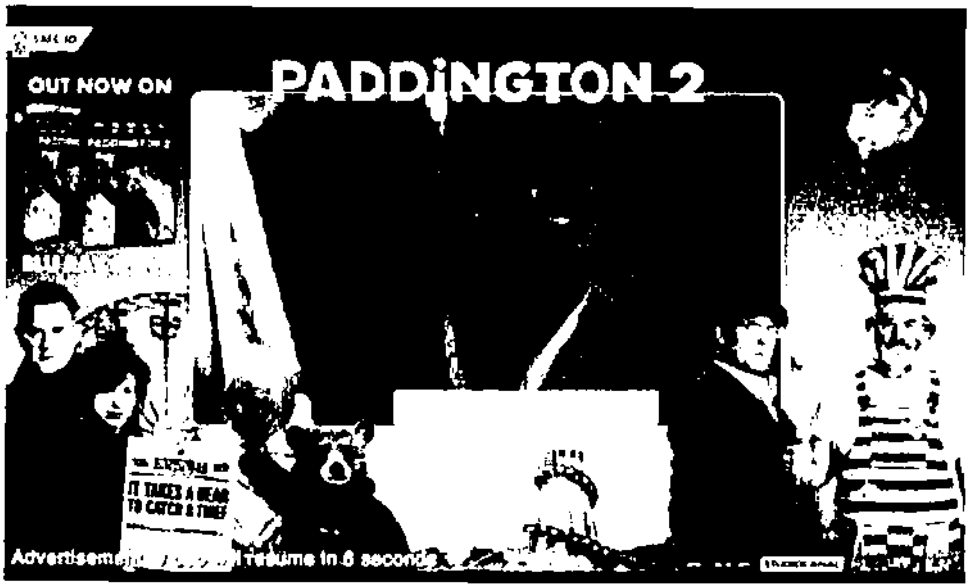
OUT NOW ON
PADDINGTON 2
IT TAKES A BEAR TO CATCH A THIEF
Advertisement ... resume in 6 seconds

Display (floor ad):

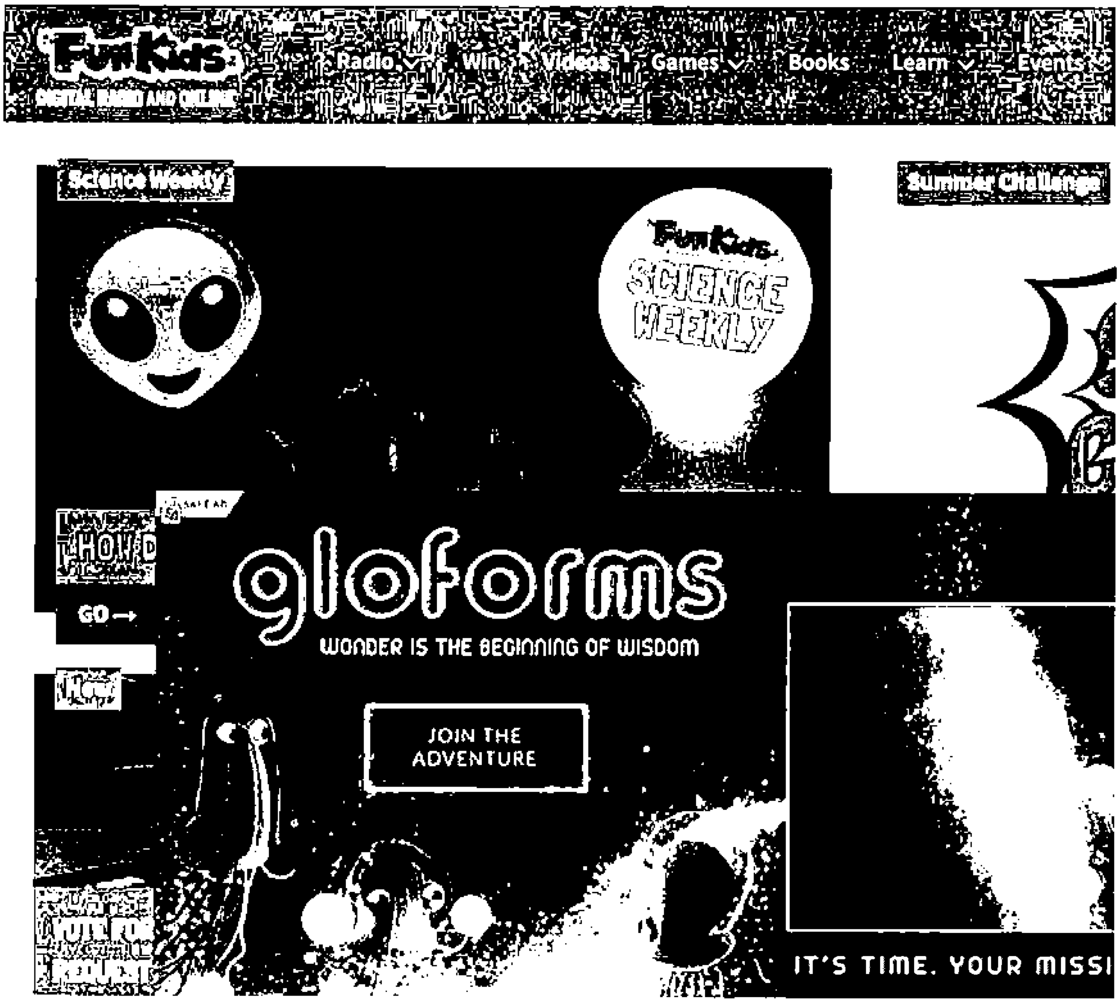

Mobile interstitial ad:

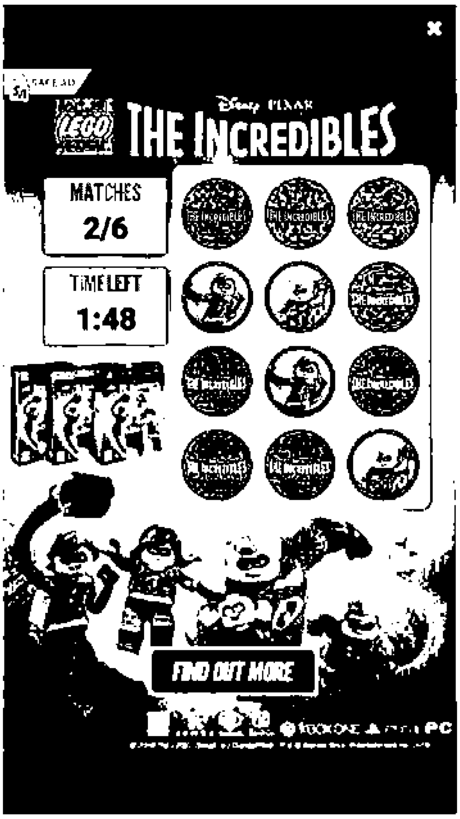

Virtual world integration:

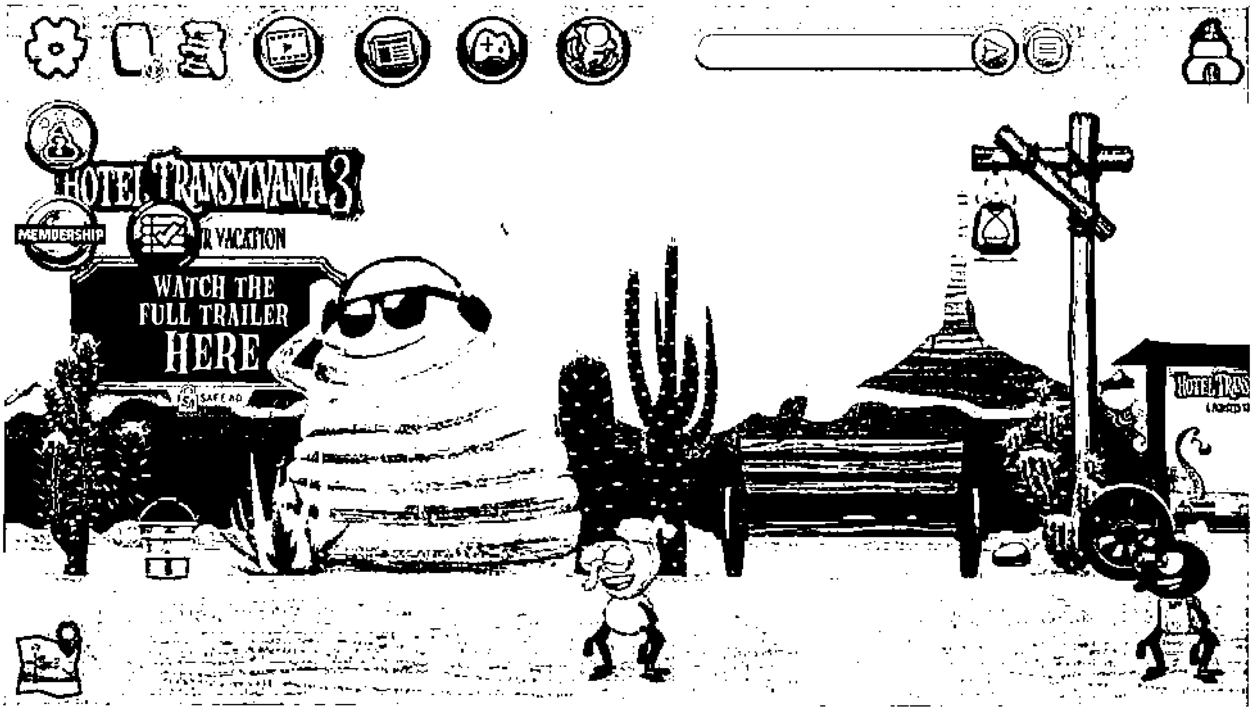

Native mobile ad (in-feed post):

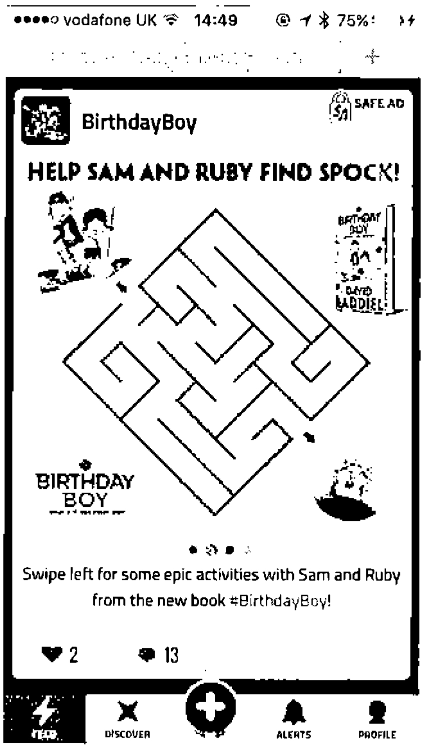

BirthdayBoy

HELP SAM AND RUBY FIND SPOCK!

Swipe left for some epic activities with Sam and Ruby from the new book #BirthdayBoy!

♥ 2 13

5. RESPONSIBLE DESIGN

Principle: **apps and websites should be designed with children's real-world behaviour in mind, to ensure kids can engage and interact with appropriate content in a safe and responsible way.**

Our PopJam platform for creative expression and sharing is now used by companies to safely engage with over 2 million kids monthly in the UK and Ireland, the US and Canada and Australia.The platform enables children (typically 6-11 years old) to create art with drawing tools, stickers, GIFs and photos (we don't allow selfies and photos can only be added when a user has 50 followers), take quizzes, watch videos and view, heart, share and comment on posts and create posts that others can view, heart, share and comment on. PopJammers can also add and follow friends and take daily creative challenges.

PopJam was designed from the ground up to provide a safe, moderated experience for children within a nurturing online community, whilst collecting the absolute minimum of data. It is certified as compliant with the Children's Online Privacy Protection Act (COPPA) in the US, and delivers an engaging experience without collecting any personal data as defined under the GDPR.

Furthermore, the PopJam platform embodies what we call 'Responsibility by Design', a series of functional design concepts to ensure a healthy environment for children. As we believe many of these concepts are universally applicable to kids' digital experiences and will be of interest to anyone looking for design guidance for a kids' app, we illustrate some of them below:

a. Closing hours

On PopJam we are closed between 11pm and 6pm. Kids can still open PopJam, and take quizzes and heart posts, but they cannot post or comment on posts.

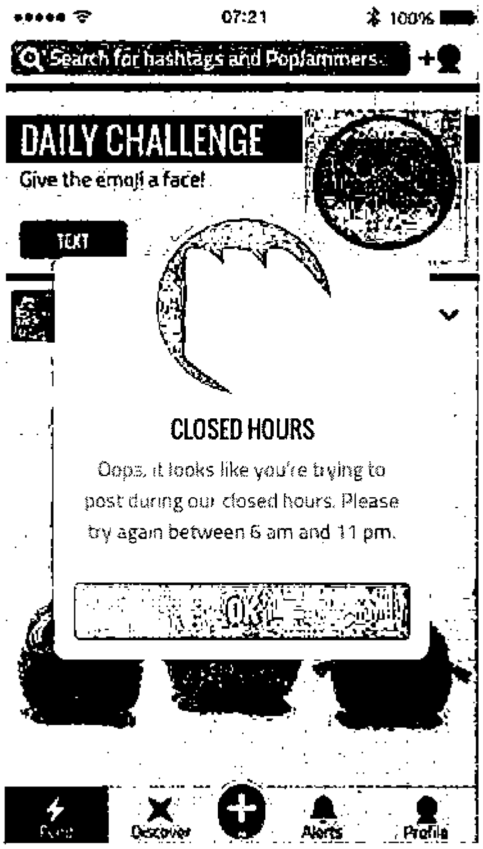

CLOSED HOURS

Oops, it looks like you're trying to post during our closed hours. Please try again between 6 am and 11 pm.

Our definition of closed is to keep kids away from any social elements. We know it is very unlikely that children are with their parents or supervised at this time, and our experience (see below) shows that after-hours social activity tends to contribute little that is positive to the community. This is a popular feature with parents, teachers, as well as charities and safeguarding groups.

The benefits of closing hours are:

- Reduction in naughty behaviour and reduced need for human moderation
- Improved safety for children, less likely to be exposed to predators
- Aligns with parents' offline bedtime restrictions
- Contributes to better health due to reduced exposure to screens late at night

Evidence

[redacted]

"We did see naughty posts late at night. It's sort of a no brainer in that kids who are up late are normally unsupervised and therefore assume we won't supervise them either—hence, naughtier behavior at night. Also professionally I can say that predators look for kids who are not supervised—kids who are up late and online. They know the risk of being caught is much lower when they find kids who are policing themselves online."

Research has shown that use of technology in the evening has a negative effect on sleep, especially when it comes to children.[20]

Clearly, government regulation has long ago recognised the need to align media consumption rules for children with bedtime, as exemplified by Ofcom's broadcasting rules and the watershed in particular.[21] This suggests children are accustomed to offline media restrictions, so extending these into the digital realm is a logical evolution.

What kids (PopJammers) say about closing hours

From 6 to 10th September 2018, we anonymously polled our users for their thoughts on closing hours.

Whilst we never collect personal data from our PopJammers, we do ask for their views and opinions on a range of topics. The views expressed below are from typical PopJam users, who are 7-12 year olds, 60/40 girl/boy split):

> "I honestly used to dislike the closing hours, but now I get why they do it. I think it's good to stop people from being on their phone all night."
>
> "Also sometimes they sorta remind me to do my homework so I don't mind closing hours unless I have a great idea for the dcentry!!"
>
> "The good things... it makes me shut off my device and have more good times with family like watching good movies or play cards!!"
>
> "I think it's good cause young kids shouldn't be encouraged to be on late but bad because people cannot socialise with other regions"
>
> "I think there's nothing good about closing hours. It isn't going to stop kids from using other apps during that time and it's the parents responsibility to make sure that the child is not using electronic devices late into the night."
>
> "One of the reasons why I don't like it is because we can't say 'happy new year' at twelve a clock in the morning on New Year's day and get people to say happy birthday to me as it's my birthday on New Year's Day!"
>
> "Omg I hate it I stay up so late but I can't go on PopJam until the morning"
>
> "It helps with younger kids going to bed but say if we were at a sleepover over and you want to show your friend something then it closes then it's really annoying also when you find a really amazing post but you can't like rejam or comment anything nice about it it's quite annoying."

[20] Bedtime Use of Technology and Associated Sleep Problems in Children (*Global Pediatric Health*, 27 Oct 2017)
[21] What is the Watershed? (*Ofcom*, 10 May 2013)

b. Time limits

We do not operate or enable individual time limits on PopJam. This is currently a topic of active discussion in the industry, and a number of initiatives are underway including Apple unveiling screen time controls, independent apps launching (such as Goozby), and Facebook and Instagram introducing time limit features.

But these approaches are fairly blunt instruments, place a significant burden on parents to decide on and proactively manage their children's exposure to screens. For these and other reasons, opinion among researchers remains divided on whether they are an effective tool.[22]

Rather than enabling screen time controls (and inviting parents to get more involved), PopJam's focus is on ensuring time spent in the app is 'quality time'—engaging in creative activities, within a positive community, in a safe environment. For us, responsible design means encouraging positive behaviours and avoiding features whose primary purpose is to extend screen time or encourage frequent returns to the app.

What kids (PopJammers) say about screen time limits

From 6 to 10th September 2018, we anonymously polled our users for their thoughts on screen time limits.

> "I think screen breaks are a very good thing. They help you limit your time online which helps your eyesight, mental health and physical health. You can definitely spend too much time online because being online for to long is bad for your eyes and health. During screen breaks I personally like to draw, do my homework, do sports, paint and sometimes just think."
>
> "When you have a screen break you tend to do things that are productive and often you might find things that you haven't used in a while and use them."
>
> "i don't like taking my eyes off fortnite ever i think its a bad thing"
>
> "Yes, I think screen breaks are really important, especially because of your brain. If your online 24/7 all day every day, that's bad for your brain. You're also not using your body and feeling what re@l life is like. Instead of spending so much time on your device and online, you can go outside and adventure. Or, you could spend time with your family and go on a trip, play with your siblings, read a book, go somewhere special, and just be in the re@l world."
>
> "No screen breaks are useless, the myth that our phone, tv and tablet screens hurt are fragile eyes is fake."
>
> "Guys i only have an hour a day now of screen time one like equals one prayer for me🙏🙏🙏🙏🙏🙏"

[22] See Sonia Livingstone: The Trouble with screen time rules (*LSE Blog*, 8 Jun 2017); Why the very ideas of 'screen time' is muddled and misguided (*LSE Blog*, 28 Feb 2018)

> "all my cousin ever does is play his ipad i tell him let's go play but he yells at me saying NO do you think he should take a break from screen time comment below."

c. No private chat

Early versions of PopJam (then called JellyChat) included a feature that allowed one-to-one conversations in the app (private chat). But this capability was removed in 2014 and since then it is our firm policy that there is no means of private communication in PopJam—all communications are public.

Private chat in children's digital services creates numerous problems. Because children assume it is fully 'private' (including out of sight of moderators), it encourages inappropriate behaviour that can easily degenerate into bullying. The very presence of a private chat feature may attract predators and enable a channel for grooming.[23] Any chat functions in kids' services must be moderated in real-time, which is a logistical and financial challenge that even the largest social media platforms have not been able to address.[24]

The benefits of not allowing private chat are:

- Improved safety for children, less likely to be exposed to predators
- Better behaviour from users, who know that all their communications are visible to the community
- It engenders an early understanding among children that there is no real privacy online
- It's very popular with parents and teachers

Evidence

[redacted]

"When exploring private chat at Mind Candy, I pulled all the content from private chats for a month or so and found the following "Hi" "Bye" "What's up" "later" was the majority of the private chat. The more lengthy and in depth private chats always led to troubled behavior including cybersex and solicitations for cybersex. Once I presented the team with the evidence, we shut down private chat immediately. We KNOW via data that nothing good comes out of private chat for under 13s."

If digital services choose to enable private chat, we recommend a number of responsible design elements:

[23] Roblox. What parents must know about this dangerous game for kids (*Family Zone*, 24 Feb 2017); Fortnite game craze is putting children at risk from online paedophiles, NCA warns (*The Telegraph*, 5 Apr 2018)

[24] Facebook tops list of sites used for online grooming (*NSPCC*, 16 Apr 2018)

- All chat should be moderated in real-time or near-real-time using both technology and human moderation
- Private chat channels should come with block and report functions that are clear and easy to use for kids
- Users should be informed very clearly (and reminded) of the difference between the public and private sections of the experience, and their risks

<u>What kids (PopJammers) say about private chat</u>

From 6 to 10th September 2018, we anonymously polled our users for their thoughts on private chat features.

> "What if u could accept them private chatting? And u can decline?"
>
> "If you have something to say to someone it shouldn't be very personal, if you need to say something to someone you should be able to say it in a public chat (so it ensures to us it isn't very personal), or get another app that it's sole purpose is a private chat ❤"
>
> "I'd say I like not having public chat, because if I did, my mom would make me delete Popjam 😅😂😂"
>
> "Good: people can't send you hurtful messages that only you can see bad: you can't privately talk to someone online if something serious happens"
>
> "A bad thing is if you need to have a private talk and u create a post there are multiple peeping toms who do online eavesdropping"
>
> "Some good things about private chat are, firstly, you can plan a surprise for people! Say someone has recently reached 1k, you and your friend can make some surprise Fanart! Secondly, you can talk to friends without anyone seeing, which is good because other people won't see, keeping everything secret."

d. Community and self-moderation

Effective moderation of user-generated content on kids' services has been a huge challenge for the industry for many years. At PopJam we have developed an approach that we believe to be extremely effective. Key to its success is the amalgamation of three approaches into a cohesive whole: AI-based technology, experienced human moderation, and community self-moderation based on trust.

<u>Trust</u>. On PopJam, children post art creations and comments in public channels. Each new user starts out with a low 'trust score', which means that all their posts are pre-moderated (checked before posting). Once a user has demonstrated that they are generally well-behaved, their trust score increases and their posts start to be post-moderated. Any single post that breaches the community guidelines will lower their trust score and return them to pre-moderation.

Technology. PopJam uses software tools for Optical Character Recognition (OCR), image recognition, and textual analysis to filter out posts that need to be reviewed by a human moderator. These tools not only look for inappropriate words and behaviour, but also for potential personal information being shared, such as phone numbers, or photos showing school uniforms.

People. Thanks to these technology and 'trust' filters, our experienced human moderators can focus specifically on those posts or users that have been escalated to them by the system. Their response tends to be constructive and community-focused, eg they will engage with users who misbehave and encourage them to reword their post, or reconsider their actions, ie giving that user an opportunity to repair their trust score. This approach has significant positive side effects within the community, generally raising the standard of behaviour throughout.

In addition, PopJam provides a number of tools for users to self-moderate, which should be part of any age-appropriate design code in relation to apps that allow user-generated content and interaction between users, such as:

- Kids can block others (they understand how to do it and what it means)
- Kids can report others (they understand how to do it and what it means)
- Kids can unfollow content they previously followed
- Kids can delete content they have created
- Kids can delete comments they have written
- Kids can delete comments or stickers under their comments

Evidence

[redacted]

"The overall rejection percentage of image content is around 5%. Rejection percentages of pre-moderated content is higher, as much as 11% for pre-moderated profile pictures. This shows we are catching attempted selfies before they are posted. The vast majority of deleted content is simply inappropriate (not dangerous)."

Example: Images from PopJam's in-house channel, PJHacks. This channel shows PopJammers how to use the features of the app. Here we have a GIF reminding users how to report content:

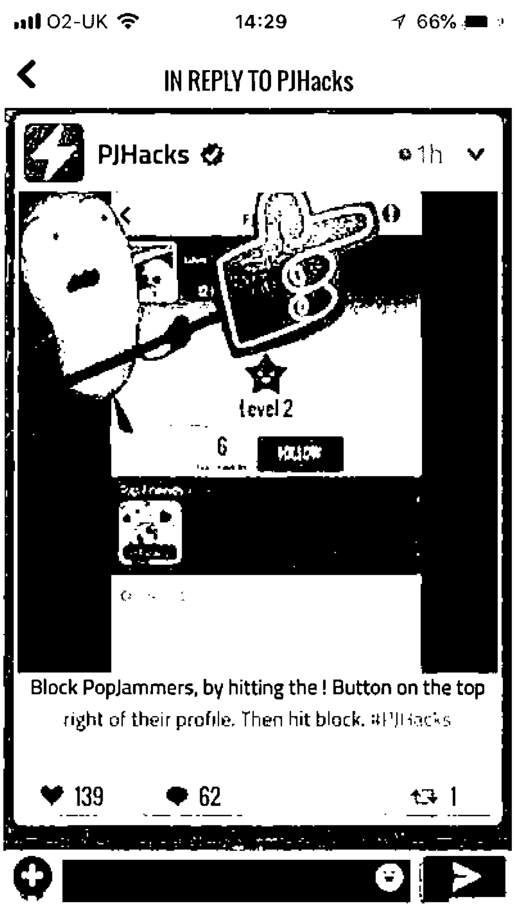

Block PopJammers, by hitting the ! Button on the top right of their profile. Then hit block. #PJHacks

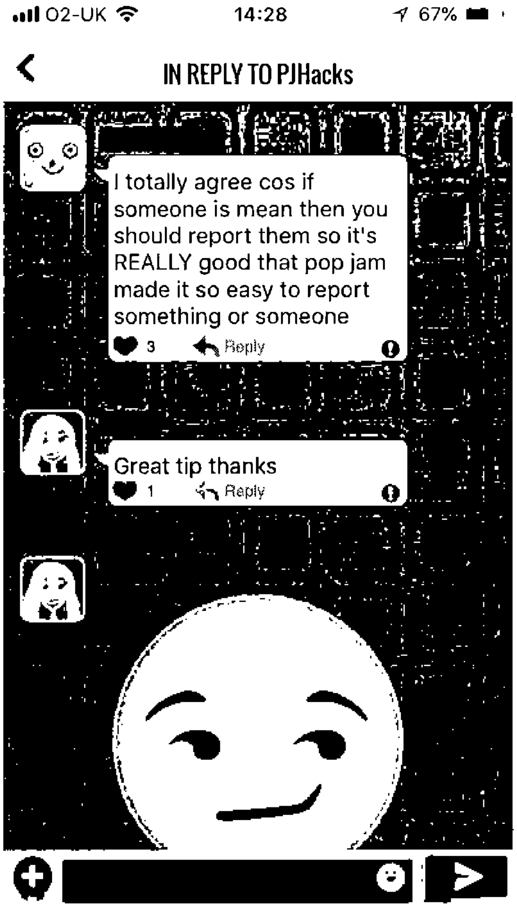

I totally agree cos if someone is mean then you should report them so it's REALLY good that pop jam made it so easy to report something or someone

Great tip thanks

<u>Example</u>: attempted negative language in our moderation tool, and how user self-corrected. Red text was caught by the AI and not featured on PopJam. The user changed saw their comment continually disappearing, so self edit until the suitable statement is accepted.

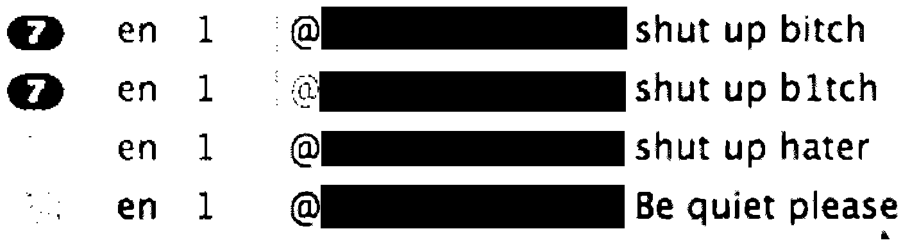

7	en	1	@[redacted] shut up bitch
7	en	1	@[redacted] shut up b1tch
	en	1	@[redacted] shut up hater
	en	1	@[redacted] Be quiet please

We constantly remind kids how to report and block users. Our PJHacks channel reminds kids about good community behaviour, fun tips for the platform, safety messaging and practical tips, such as how to block and report. This information is also available on:

https://www.popjam.com/faq
https://www.popjam.com/privacysummary

e. Promoting good behaviour

Promoting good behaviour is a key part of our content cycle, and an approach that we believe would be good practice for any child-directed app where users engage with each other.

An effort needs to be made to provide regular content to keep kids busy, leading to signposted activities and things to do everyday. Positive in-house content, from Safer Internet Day to regular themed weeks promotes a positive and enabled community.

Developers should be aware of the power of reward and recognition. For example, PopJam awards Neon Stars (a badge that appears on the player's profile) and shares Daily Challenge winners and runners-up.

Staff interact with active members and acknowledge good behaviour in posts and comments by hearting or commenting. Most behaviours are fine, and kids will naturally take small chances and risks to test new experiences. See below for some examples of in-house content promoting kindness and good digital citizenship in the app:

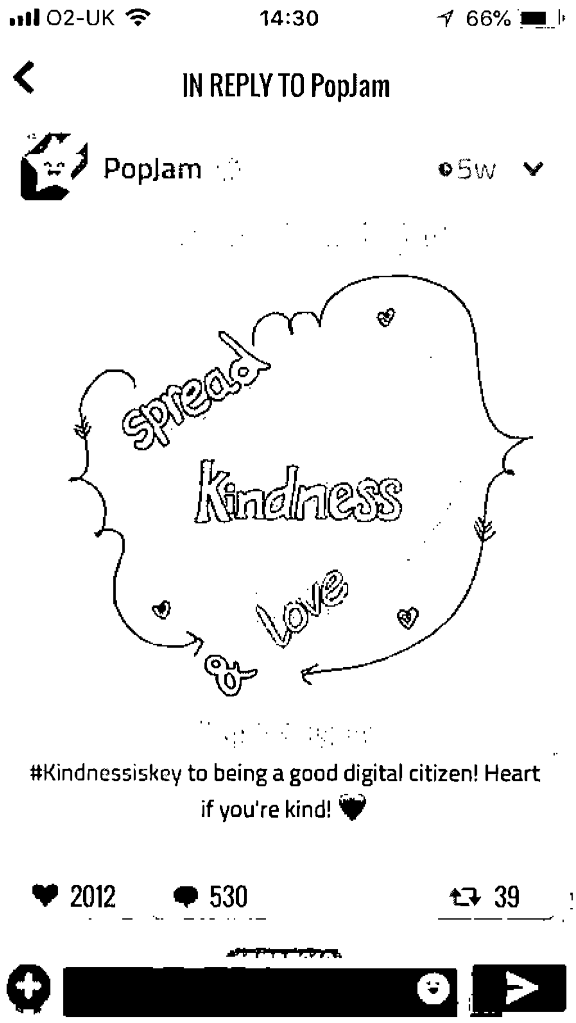

O2-UK 14:30 66%

IN REPLY TO PopJam

PopJam 5w

#Kindnessiskey to being a good digital citizen! Heart if you're kind! ♥

♥ 2012 530 39

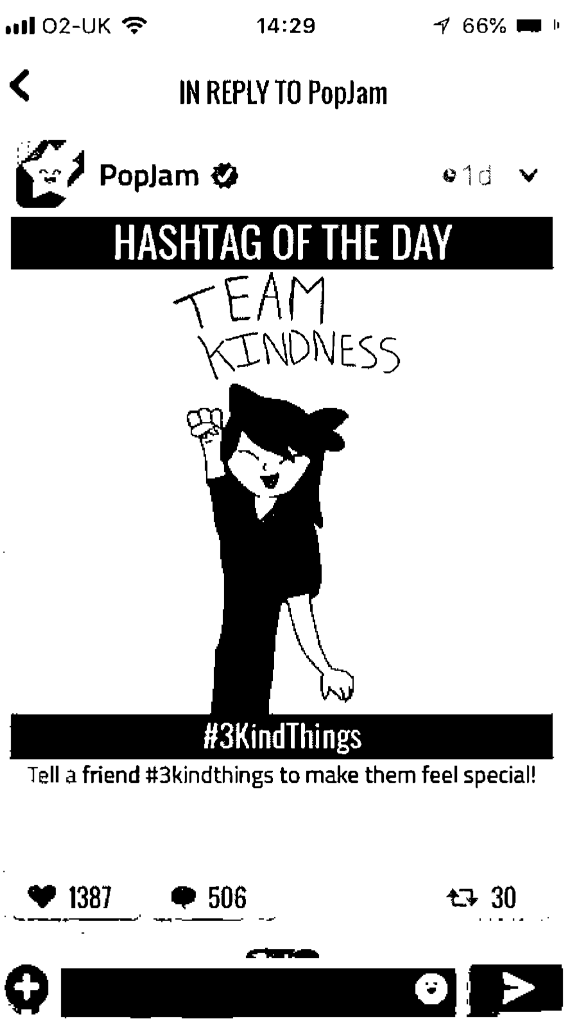

O2-UK 14:29 66%

IN REPLY TO PopJam

PopJam 1d

Tell a friend #3kindthings to make them feel special!

♥ 1387 506 30

f. What the kids are saying

One of the ongoing criticisms of child safety regulation and data privacy regulation is that the views of children are insufficiently considered when drafting new rules.[25] Thanks to PopJam we can gain unique insights into how children view internet safety and their experiences online. Some of these views are set out in the links below:

- SuperAwesome blog about kids top online fears.
- SuperAwesome blog about how kids think about online safety.
- SuperAwesome blog about kids creating positive change for Safer Internet Day.

[25] Among others: Sonia Livingstone on the GDPR: No more social networking for teens? (Better Internet for Kids, 31 Mar 2016)

And in our recent PopJam insights report (June 2018) we anonymously polled our users for their thoughts on what would make the internet safer, data which may be useful to the Commissioner:

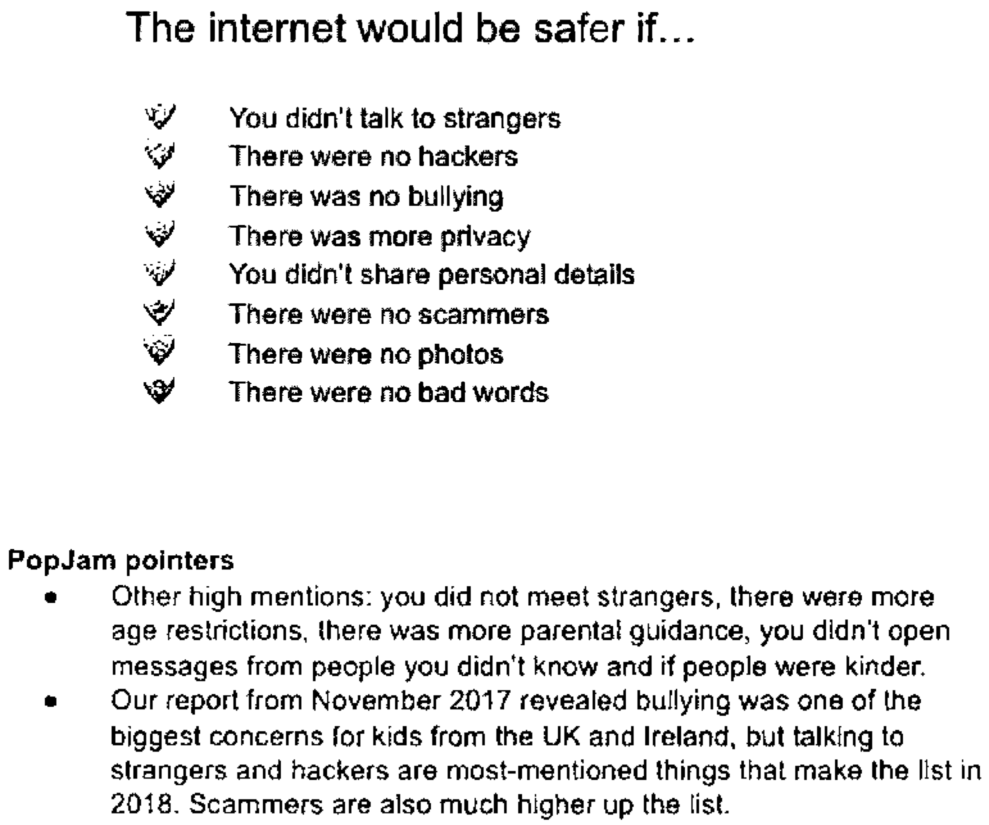

The internet would be safer if...

- You didn't talk to strangers
- There were no hackers
- There was no bullying
- There was more privacy
- You didn't share personal details
- There were no scammers
- There were no photos
- There were no bad words

PopJam pointers

- Other high mentions: you did not meet strangers, there were more age restrictions, there was more parental guidance, you didn't open messages from people you didn't know and if people were kinder.
- Our report from November 2017 revealed bullying was one of the biggest concerns for kids from the UK and Ireland, but talking to strangers and hackers are most-mentioned things that make the list in 2018. Scammers are also much higher up the list.

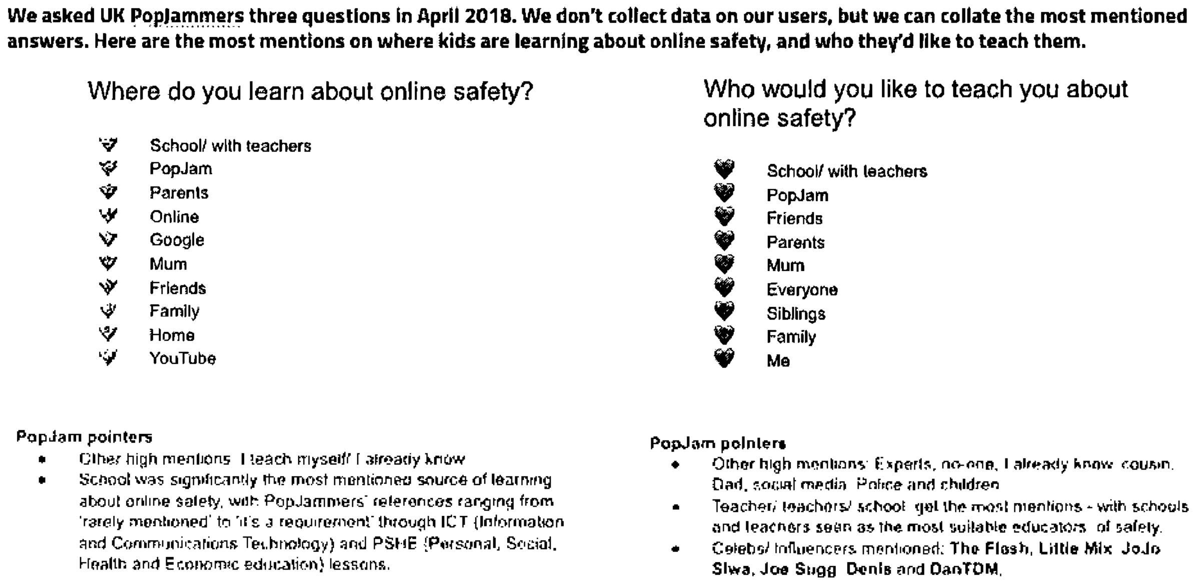

We asked UK PopJammers three questions in April 2018. We don't collect data on our users, but we can collate the most mentioned answers. Here are the most mentions on where kids are learning about online safety, and who they'd like to teach them.

Where do you learn about online safety?

- School/ with teachers
- PopJam
- Parents
- Online
- Google
- Mum
- Friends
- Family
- Home
- YouTube

PopJam pointers

- Other high mentions I teach myself/ I already know
- School was significantly the most mentioned source of learning about online safety, with PopJammers' references ranging from 'rarely mentioned' to 'it's a requirement' through ICT (Information and Communications Technology) and PSHE (Personal, Social, Health and Economic education) lessons.

Who would you like to teach you about online safety?

- School/ with teachers
- PopJam
- Friends
- Parents
- Mum
- Everyone
- Siblings
- Family
- Me

PopJam pointers

- Other high mentions: Experts, no-one, I already know cousin, Dad, social media Police and children
- Teacher/ teachers/ school get the most mentions - with schools and teachers seen as the most suitable educators of safety.
- Celebs/ influencers mentioned: **The Flash, Little Mix JoJo Siwa, Joe Sugg Denis** and **DanTDM.**

g. Additional considerations by age bracket

Each of the elements we've explored should be further considered in context—by age, platform and application.

Topic	Applicability	Age considerations
Closing hours	Should be applied against all platforms with easily understandable and accessible parental controls, so parents and guardians can set closing hours, time limits and any other restrictions. If parents or young teens and teens choose to ignore suggested closing hours options, then appropriate moderation should be in place to ensure a safe and positive experience.	**Pre-school (3-5)** Closing hours likely to be controlled by parents, but could be introduced to the product to help kids develop understanding of appropriate time spent online. **Kid (6-9)** Closing hours enforced by product or controlled by parents **Tween (10-12)** Closing hours enforced by product or controlled by parents **Young Teen** (13-15) Closing hours suggested by product **Teen (16-17**) Closing hours suggested by product
No private chat	To be applied as a standard requirement in platforms such as social media/ gaming/ video streaming where chat could be a feature.	**Pre-school (3-5)** No private chat or no chat functions at all **Kid (6-9)** No private chat (consider limited pre-selected phrases and icons or moderated chat experience) **Tween (10-12)** No private chat (consider limited pre-selected phrases and icons or moderated chat experience) **Young Teen (13-15)** Moderated private chat with blocked words and phrases, and reporting and blocking mechanisms in place. **Teen (16-17)** Moderated private chat with blocked words and phrases, and reporting and blocking mechanisms in place.
Parental controls	To be applied against all platforms with easily understandable and accessible parental controls, so parents and guardians can set closing hours, time limits and any other restrictions.	**Pre-school (3-5)** Experience fully controlled by parents. **Kid (6-9)** Parental controls if there are any elements of data collection, or if the product is not moderated or has report/ blocking mechanisms in place. **Tween (10-12)** Parental controls if there are any elements of data collection, or if the

		product is not moderated or has report/ blocking mechanisms in place. **Young Teen (13-15)** No parental controls. Information on how to talk to parents about issues. **Teen (16-17)** No parental controls. Information on how to talk to parents about issues.
Moderation & community	Applied to any platform or application where children can question, search, engage, write free text or interact, such as web, in-app and voice. It's equally important where chat and community applications are being used that children can experience communities with rules, moderation and positive journeys.	**Pre-school (3-5)** Full moderation in place. Word and phrases filters and blocking. Messages to parents/ kids about safety/ resilience/ citizenship/ kindness. **Kid (6-9)** Full moderation in place. Word and phrases filters and blocking. Messages to kids about safety/ resilience/ citizenship/ kindness. **Tween (10-12)** Full moderation in place. Word and phrases filters and blocking. Messages to kids about safety/ resilience/ citizenship/ kindness. **Young Teen (13-15)** Full moderation in place. Word and phrases filters and blocking. Messages to young teens about safety/ resilience/ citizenship/ kindness/ mental health. **Teen (16-17)** Full moderation in place. Word and phrases filters and blocking. Messages to kids about safety/ resilience/ citizenship/ kindness/ mental health.

6. *LOOK AHEAD:* INFRASTRUCTURE PRIVACY BY DESIGN

Principle: service providers should design their technology infrastructure so as to NOT create stores of personal data that are susceptible to misuse or to data breaches.

An often overlooked aspect of building kids' digital services is the backend technical infrastructure. Most service providers simply rely on the established technology tools, architectural standard practices, and cloud-based infrastructure services which were typically designed for the adult internet. All of these components are biased toward data collection, and our reliance on then makes it extremely difficult to minimise data collection downstream.

But there are, in fact, some fairly simple infrastructure design concepts that can dramatically improve the safety and data privacy compliance of kids' digital services, and we believe such concepts should be an integral part of any age-appropriate design code.

a. It's all anonymous... until it isn't

Some data protection laws allow publishers to collect personal information so long as it is within the scope of internal operations. Others such as GDPR allow the collection of personal information subject to the appropriate legal basis, with a self-assessment based on the sensitivity of the *intended* data usage rather than the *potential* data usage. This approach is problematic, because once collected, personal data lives in systems that were designed for making most use of it, making leakage to third parties and unintended uses almost inevitable.

The only fail-safe way to to protect user's data privacy is to embed 'privacy by design' in the infrastructure itself, to design the data storage paradigm in such a way that personal data *can't* be used outside its intended scope. Below we set out a technical design principle for data storage that could underpin such an approach.

In theory, the current system should work: if organisations have good intentions *and* they are technically adept *and* they did not suffer data breaches, then basing the data collection rules around their intentions seems sensible. The reality in most cases, however, is very different. The risks are compounded by the general public's and most companies' poor understanding of what "anonymisation" really means[26]. The truth is many digital profiles of identifiable people are built upon supposedly anonymous data that was assembled beyond the intentions of any one service provider—for example, using the common practice of 'fingerprinting'.[27]

b. The prevalence of fingerprint-based profile of children

The idea of anonymisation is solid. Anyone maintaining a database makes sure there is no personal data (in the classical sense -- no first and last names, addresses, etc.) in their system, and only allows storage of records by "anonymous identifiers" such as numbers or a string of characters that (in theory) are not attributable to an identifiable person.

The anonymity of the user is secure, at least within the confines of that single database. But by the time we are teenagers we have already interacted with hundreds of 'anonymous' data

[26] There is no such thing as anonymous data (*HBR*, 9 Feb 2015)

[27] A fingerprint is a set of individually anonymous data points that, when combined, form a uniquely identifiable pattern (and thus is unique to a particular user). An example would be the following: A user agent (internet browser name and version); an approximate location; a time of day when you were seen using an app; your favourite colour; your favourite sports team. All of the above information is anonymous, but when combined, it creates narrow filter that—when combined with a third-party dataset—can be used to identify a particular person. Fingerprinting is widely used in the digital content industry to build user profiles for advertising targeting purposes. For an illustrative test, see: Are you unique?

collecting services, and so there are hundreds of databases tracking us, constantly adding new variables to our profiles.

As these databases grow with more and more data points, they are statistically creating fingerprints of people, which are then easily correlated across other dataset, as illustrated below.

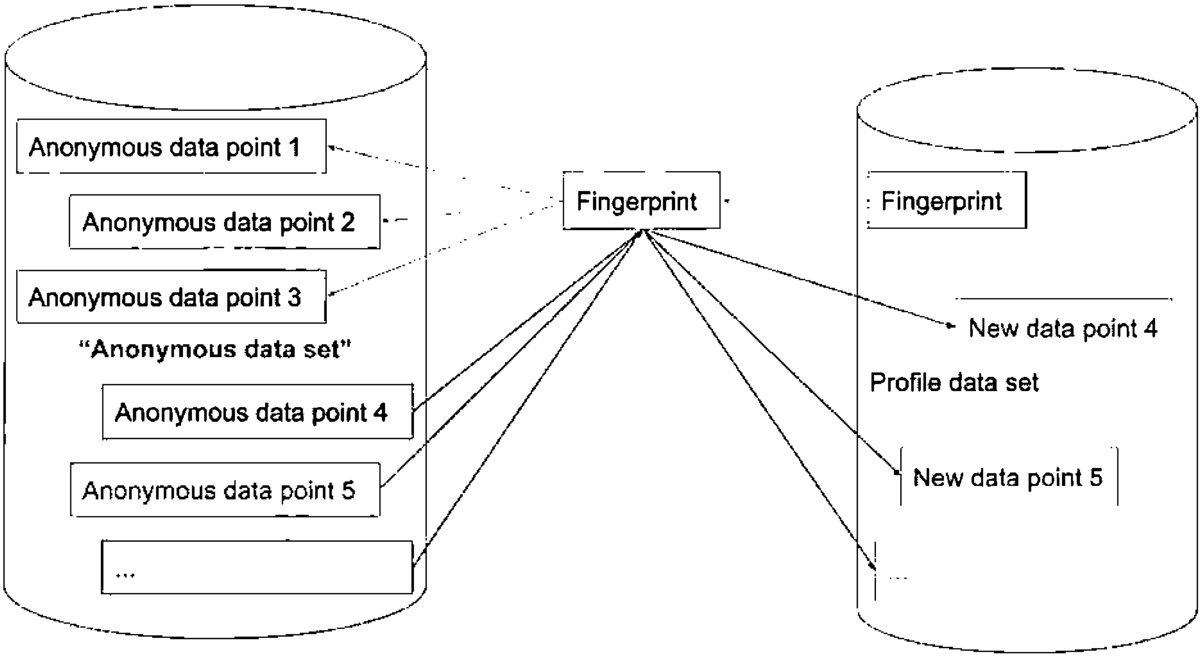

This means that what started out as legal, well-intentioned, anonymous data collection is now a collection of personal profiles that can be resolved back to a natural person including a child.[28] This is a problem for a number of reasons:

- Many organisations suffer data breaches[29], and the frequency of hacks is only increasing
- Data analysis technology is becoming more efficient (and humans understand less of it) [30], making the identification of fingerprints easier
- Organisations are too big to vet every single piece of data that is being stored, especially if it is not personal data (yet)
- There are few reliable technical methods for policing the actual usage of data once it has been collected

c. Designing privacy *into* the infrastructure

[28] Browser fingerprinting: Online anonymity elusive Australian researcher finds (*Computerworld AU*, 21 Jul 2016)

[29] See February 2018 Website Hacking Statistics (*Web Arx*, Feb 2018); and ICO warns on over-reporting of data breaches (*Out-Law.com*, 13 Sep 2018)

[30] The Dark Secret at the Heart of AI (*MIT Technology Review*, 11 Apr 2017)

The only way to eliminate the above risks altogether is to move beyond the commonly used approaches to storing data. In fact, it is entirely possible to design data storage architectures that do not rely on unique identifiers. We call this a move from *ID-based* to *zero-data-based.*

Rather than tying every single piece of data to a user ID and then building regular aggregation and reporting tools on top (which leaves the individual user records vulnerable to fingerprinting), our proposed approach is to design systems that do not correlate this data to any unique records in the first place. The simple way to do this is to create two, separate data stores—one for user IDs (if needed, for example, to enable log-in to a service), and the other for *events* (eg, what users do in the service). This allows the organisation to benefit from analysis of the events, without inadvertently creating an increasingly rich profile of a specific user.

In summary:

Traditional approach	**Fail-safe approach**
- Create a user - Link all events to this user - Potentially aggregate events daily for performance purposes - Write reporting / statistics tools / etc. on top of the events or on top of the aggregation	- Create user - Create events separately - Potentially add relevant individual data points to these events - Potentially aggregate your events daily for performance purposes - Write reporting / statistics tools / etc. on top of the events directly or on top of the aggregation

By not tying the unique user ID to every single event that is raised in the system, there is no single data store that can be used (or abused) to build fingerprintable profiles. Data privacy and enhanced security are effectively *built in* to the architecture, which means:

- The controller's initial data collection intention and legal basis is respected, as no unintended fingerprinting is possible;
- The consequences of a data breach are hugely mitigated as hackers will not be able to correlate the events and user data; and,
- Particularly in the case of children, there is no accumulation of hundreds of data profiles that could be used to identify them by the time they are teenagers.

Designing fail-safe privacy into the infrastructure in this way is not difficult, but it will require guidance and education from the regulator to help technology professionals shift away from their traditional approaches to database design. We would welcome the opportunity to assist in the formulation of such guidance, if this would be of use to the Commissioner.