

Children's Rights and AI Oversight – 5Rights Position Paper on the EU's Artificial Intelligence Act

A DUTY TO REGULATE AUTOMATED DECISION-MAKING SYSTEMS THAT IMPACT ON THE SAFETY OF CHILDREN

December 2021

Introduction

The Draft Artificial Intelligence Act exemplifies the EU's commitment to ensuring technology is developed and used in conformity with Union values and with a high level of protection for health, safety and fundamental rights – with particular attention for children. It builds on the global consensus that digital services that interact with or otherwise affect children must be designed with them in mind. Childhood is a time of experimentation and personal growth, and while no environment is entirely risk free, technology environments in which children socialise, learn and play must be designed in a way that minimises risk and gives children the privacy, safety and security to which they are entitled.¹

Central to the digital world is artificial intelligence, commonly referred to as AI. AI is not a standalone or fixed technology, but plays a part in automated decision making (ADM) systems and many other data-driven features common across digital services. Automated systems shape the experiences of children and young people in the digital world, both as a result of their direct engagement, for example receiving friend/follower or content recommendations on social media, and from systems that they may not interact with directly, for example, automated-decision making used to allocate welfare funding.

Much emphasis is put on the challenges of regulating new and emerging technologies, but AI is not new. The term 'AI' was coined in the 1950s to describe the science and engineering of machines that can make automated choices based on specific criteria using given information. In many ways the word 'intelligent' is used to give humans confidence in the efficacy and authority of machine-made choices. Since then, huge advances in the application of AI and greater availability of data have led to more sophisticated, data-driven decision making.

Systems that use AI are still human-made with specific objectives, design goals, chosen inputs, a set of rules by which information is given importance or weight, and a combination of outcomes and outputs. At each of these stages, automated decisions

¹ Cf. UNCRG General comment No. 25 on children's rights in relation to the digital environment

are made that are often imperceptible to those they impact, particularly if they are a child.

Automated decision-making sits behind features that are ubiquitous across digital products and services that interact with or otherwise effect children. It can support children to navigate the online world and the mass of content available, and help them to identify activities and outcomes that are useful or beneficial to them. But there are also many situations when automated decision-making systems undermine their rights or put them at risk. For example;

- 75% of the most popular social networking sites make automated friend recommendations.² These introduce users based on the data profiles the platform has built of them, irrespective of their age, which has been found to enable predators to contact children.³
- Misinformation is spread and amplified by automated systems designed to promote content that is most likely to engage users, irrespective of its veracity or potential to harm. In 2020, vaccine misinformation alone was estimated to be worth up to \$1bn to the largest tech companies.⁴
- Automated nudges encourage users to make in-app purchases or engage with gambling-style features. In 2019, British children alone spent €320 million on loot boxes and other in-app purchases.⁵
- Vast amounts of harmful material is 'suggested' or 'recommended' to users on social media feeds, including material promoting self-harm or suicide behaviours, disordered eating and pornography.⁶

The draft AI Act offers a vision of what a responsible digital world looks like. However, the absence of ex-ante risk assessments to determine which AI systems are high-risk, in particular for children, is a critical weakness that will leave children exposed to a wide range of systems that negatively impact on their rights. In addition, it will not be possible to meet one of the objectives of the Bill (to prohibit "practices that have a significant potential to manipulate persons through subliminal techniques beyond their consciousness or exploit vulnerabilities of specific vulnerable groups such as children or persons with disabilities in order to materially distort their behaviour in a manner that is likely to cause them or another person psychological or physical harm") unless there is a clear duty for the regulator to investigate algorithms on behalf of children, and an agreed standard by which to assess them. A requirement for services to conduct ex-ante risk assessments and regulatory duty to investigate AI systems would ensure the risks of such systems to children are identified, eliminated, mitigated or effectively managed. Such requirements could equally be applied to ensure the full

² https://5rightsfoundation.com/uploads/But_How_Do_They_Know_It_is_a_Child.pdf

³ <https://www.thetimes.co.uk/article/instagram-sends-predators-to-private-accounts-of-children-as-young-as-11-wqvmjc2df>

⁴ The Anti-Vaxx Industry: How Big Tech powers and profits from vaccine misinformation, Center for Countering Digital Hate

⁵ Young People Losing Millions to Addictive Gaming – REPORT, Safer Online Gambling Group, August 2019.

⁶ <https://5rightsfoundation.com/uploads/Pathways-how-digital-design-puts-children-at-risk.pdf>

implementation of the risk management of AI systems likely to interact with or otherwise affect children described in Article 9.

The four-step AI oversight model set out in this paper is platform neutral and can be applied across different sectors, including but not limited to, health, education, financial services, public services, entertainment, B2B services and social media. It can also be applied to different parts or features of a service, including advertising, content recommendation, moderation and redress. The adoption of this model would give clarity to businesses in fulfilling their safety duty to children and power to the regulator to inquire, analyse and assess whether a system is conforming to requisite standards. When new risks and harm are revealed, they can act as an early warning, especially when that harm was an unintentional by-product of an automated decision-making process optimised for another purpose.

This short paper builds on the work of many in the international community, notably the Centre for Data Ethics and Innovation,⁷ UNICEF⁸, IEEE⁹, Ada Lovelace Institute¹⁰ and Council of Europe¹¹. We are grateful for their expertise and recognise that this practical application of their work could not have been done without their thoughtful and detailed insights.

Thanks are due also to Dr. Rebekah Tromble, Associate Professor in the School of Media and Public Affairs and Director of the Institute for Data, Democracy, and Politics at George Washington University. Dr. Tromble developed the four-step model articulated in this report.

5Rights is committed to building the digital world young people deserve. That world is one in which they share the benefits of digital engagement as participants, citizens and consumers, and in which businesses respect and uphold their existing rights and respond to their needs and evolving capacities - automatically.

⁷ In November 2020, the Centre for Data Ethics and Innovation conducted a [review into bias in algorithmic decision making](#) and made recommendations to the government and regulators designed to produce a step change in the behaviour of organisations making life changing decisions on the basis of data.

⁸ UNICEF's Draft Policy Guidance on AI for Children is designed to promote children's rights in government and private sector AI policies and practices, and to raise awareness of how AI systems can uphold or undermine these rights. The policy guidance explores AI and AI systems, and considers the ways in which they impact children. It draws upon the Convention on the Rights of the Child to present foundations for AI that upholds the rights of children.

⁹ The IEEE (Institute of Electrical and Electronics Engineers) has a [global initiative](#) on the ethics of autonomous and intelligent systems. Its aim is to move from principles to practice with standards projects, certification programs, and global consensus building to inspire the ethically aligned design of autonomous and intelligent technologies.

¹⁰ Ada Lovelace Institute are [developing tools](#) to enable accountability of public administration algorithmic decision-making, such as a typology and a public register.

¹¹ <https://www.coe.int/en/web/artificial-intelligence>

Definitions

- **Artificial intelligence** (AI) describes when machines are able to mimic the problem-solving and decision-making capabilities of the human mind.
- **Machine learning** is a branch of artificial intelligence (AI) that employs computational algorithms to detect patterns in—and learn iteratively from—data, generating output with minimal human intervention and improving over time.
- An **algorithm** is a sequence of instructions or set of rules designed to complete a task or solve a problem.
- **Automated decision-making** is the process of making a decision by automated means, without any human involvement. These decisions can be based on factual data, as well as on digitally created profiles or inferred data.¹²
- **Algorithmic bias** is commonly used to describe an automated system that produces results that discriminate against or disadvantages groups of people (for example, based on age, disability, gender, or race).
- **Algorithmic fairness** is an automated system that produces results that do not discriminate against nor systematically disadvantage groups of people: it also seeks to ensure that automated systems do not violate rights, exploit vulnerability, manipulate, nor withhold information in a way that impairs one’s ability to make informed choices.
- **A child** is a person under the age of 18.¹³
- **Document analysis** is a research method involving the review and interpretation of written materials such as emails, legal records, and meeting notes, designed to gather evidence on the topic being studied and answer specific questions.
- **Code analysis** is a way of assessing how an algorithm is structured and how it might function in practice without actually executing the program, allowing errors or vulnerabilities to be detected.
- **Variables** are individual items in a dataset being analysed, for example age, gender and location.

¹² <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/automated-decision-making-and-profiling/what-is-automated-individual-decision-making-and-profiling/#:~:text=Automated%20decision%20making%20is%20the,created%20profiles%20or%20inferred%20data.&text=an%20online%20decision%20to%20award%20a%20loan%3B%20and>

¹³ Article 1 of the United Nations Convention on the Rights of the Child states “a child means every human being.”

Regulating AI in products and services that impact on children

The four-step model set out below offers a mixed method approach to algorithmic oversight. It describes how organisations and regulators can evaluate each element of an automated decision-making process, from the goals, inputs, implementation and outcomes, to ensure that applications of AI meet the established rights and needs of children, as set out in:

- The EU Charter of Fundamental Rights and United Nations Convention on the Rights of the Child to which all EU Member States are signatories. General comment No. 25 on children's rights in relation to the digital environment is the authoritative document which sets out the relevance of the convention to the digital world.¹⁴
- The General Data Protection Regulation (GDPR) and associated guidance from EDPB and national Data Protection Authorities, including the Age Appropriate Design Code in the UK¹⁵ and Ireland's Fundamentals for a Child-Oriented Approach to Data Processing,¹⁶ which set out the standards providers of digital products and services must meet in relation to children's data.
- The effective application of existing EU policies and strategies, laws and protections that pertain to children, such as the EU Strategy on the Rights of the Child,¹⁷ the EU Strategy for a Better Internet for Children,¹⁸ the Audio-Visual Media Services Directive¹⁹, and EU consumer and product safety laws, as well as national government guidance may also be relevant when considering the impact of AI systems on children.²⁰

Any assessment of automated decision-making systems must have the flexibility to uncover harms that are currently unknown or not anticipated. It must also allow for potential improvements or benefits to be identified so that they might be shared and used to guide best practice.

1 Understand the design goals

Aim:

Algorithms are formulated with a purpose and intended outcomes. In assessing the fairness and appropriateness of algorithms, it is important for the regulator to

¹⁴ General comment No. 25 (2021) on children's rights in relation to the digital environment.

¹⁵ Age Appropriate Design Code

¹⁶ https://www.dataprotection.ie/sites/default/files/uploads/2020-12/Fundamentals%20for%20a%20Child-Oriented%20Approach%20to%20Data%20Processing_Draft%20Version%20for%20Consultation_EN.pdf

¹⁷ <https://eur-lex.europa.eu/legal-content/en/TXT/?uri=CELEX%3A52021DC0142>

¹⁸ <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:52012DC0196>

¹⁹ <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32010L0013&qid=1632209662952>

²⁰ For example, the German Youth Protection Act.

understand the original intent and goals of its creators and how those goals evolved over time by asking the following questions:

- a. What was/were the problem(s) or challenge(s) those designing the algorithm set out to address?
- b. What was/were the intended outcome(s)?
- c. Why was this product, feature or process considered necessary?
- d. Who was involved in defining the problem(s) and desired outcome(s)—including internal and external stakeholders? What was their role in shaping the understanding of the problem(s) and desired outcome(s)?
- e. How and why did any of these things change over time?

Method:

- Undertake interviews with stakeholders
- Analyse information - product development documents and internal communications such as emails and meeting notes in which the algorithmic product was discussed.

2. Consider the data inputs

Aim:

Every algorithm contains a series of inputs — data points and variables that can be thought of as the “ingredients” of the algorithm. Unfair, discriminatory or biased outcomes are often the result of problematic data (“garbage in, garbage out”). It is therefore essential that any framework intended to examine algorithmic fairness assess the quality and appropriateness of the data used to build and train the algorithm, by asking the following:

- a. What features (variables) did the algorithm’s designers want to include as inputs and why?
- b. Were they able to include those features? Did they have to settle for proxies and/or exclude some features altogether and why?
- c. What dataset(s) was/were used as input(s) for building, training, and testing the algorithm?
- d. Were other datasets considered? For training/testing? For final implementation? If not, why not?
- e. If so, what were the perceived advantages and disadvantages, strengths and weaknesses of this/these datasets compared to other options?
- f. Were multiple datasets and/or features tested? If so, how were they evaluated? And why were the final datasets/features selected?
- g. Who had input into these decisions, and what was their role in the process?

Methods:

- Undertake interviews with stakeholders.
- Analyse information - product development documents and internal communications

- Code analysis
- Data sample analysis

3. Assess the model selection and execution

Aim:

If data inputs are the “ingredients” of an algorithm, the mathematical model and its parameters offer the instructions for how to put the algorithmic recipe together. They lay out how the inputs should be combined, at what point and in what amount, as well as the ways in which those inputs might be altered or transformed. Careful scrutiny of the model and the assumptions it is built upon is needed to assess its appropriateness. Note that such scrutiny is possible even with machine learning algorithms. The questions to consider as part of this scrutiny include:

- a. What is the mathematical formula/model applied?
- b. Why was this model selected?
- c. What assumptions are built into this model?
- d. Did those designing or implementing the algorithm deviate from any of the assumptions built into the model? If so, how and why?
- e. Within the model, what is being optimised for? How is this optimisation carried out (e.g., how are the various features weighted)?
- f. How and when was the model tested and changed/updated?
- g. When changes were made, what were the reasons for making those changes?

Method:

- Undertake interviews with stakeholders.
- Analyse information - product development documents and internal communications
- Code analysis
- Implementation experiments (e.g., running independent tests on real or synthetic data, including on platform).

4. Identify outputs and outcomes

Aim:

After an algorithm is launched, it will generate certain outputs. It is important to examine these outputs to reveal whether the model performs as intended. However, at this stage, it is also important to look at the actual *outcomes* – the *real world impacts* generated by the algorithm/s and its uses.

The previous three steps help to determine why and how something went wrong what elements of the design and implementation results in discrimination, disadvantage, exploitation, manipulation, or rights violations. However, the output (step four) is likely

to be the first place that harm is identified and indicated that the four step process is necessary.

Many observers note that algorithms are not autonomous, neutral entities. They are designed by people, with all the biases, blind spots, and other foibles associated with being human. It is therefore crucial to examine the *interplay* of technical features on the one hand, business decisions, and human interactions on the other. The regulator will need researchers and investigators with training in the social sciences as well as computer scientists to conduct such assessments.

Below we lay out the three lenses through which to examine algorithmic outputs and outcomes. First, we describe assessments of the ways in which relevant companies interpret outputs and outcomes, as well as their techniques for mitigating perceived harms. Second, we outline a broad approach for considering how users interact with and are impacted by algorithms. Finally, we discuss broad approaches to uncovering impacts on society as a whole.

1. Companies

Aim:

To examine how either the company that designed the algorithm or companies that make use of those algorithms evaluate outputs and outcomes.

Questions:

- a. What model outputs (variables) does a company use internally? (I.E., What outputs matter to them and why?) In what ways do they use these outputs?
- b. What is the internal process for evaluating the performance of an algorithm? What standards are applied? What metrics are applied? By whom?
- c. What is the internal process for determining whether an algorithm should be changed? Who is involved in this process? Who makes final decisions and how?
- d. What, if anything, is the company doing to assess larger impacts on users and society?
- e. If such assessments occur, are they ad hoc or systematic?
- f. What techniques and methodologies are used for such an assessment? What standards and metrics are applied? Who is involved in this process and how?
- g. Are changes ever made to algorithms on the basis of such assessments? What is the process for doing so? Who is involved in this process? Who makes final decisions and how?

Method:

- Interviews
- Document analysis
- Code analysis

2. Users

Aim:

To assess whether users' reasonable expectations for how they interact with and what they expect from an algorithm align with the actual outcomes, and whether any harms (either perceived by the user or not) accrue.

Questions:

- a. What, if anything, do users understand the algorithm to be doing? Are they even aware that an algorithm is involved? If they are, do they perceive specific advantages and disadvantages to the algorithm?
- b. What do users expect from the algorithm? Are outcomes aligned with those expectations?
- c. Is the algorithm creating disparities between users and non-users and/or between different types of users?
- d. Is the algorithm limiting user choice(s)? If so, in what ways? And what are the consequences (positive or negative) of those limitations?
- e. Does the algorithm directly or indirectly exploit user vulnerabilities?
- f. Does it directly or indirectly manipulate users?
- g. Does it violate users' rights or contribute in any way to the violation of those rights?

Method:

- User surveys and interviews
- (Controlled) experimental user studies

3. Societal impacts

Aim:

To understand the social, financial, environmental and human impacts of automated-decision making systems.

Questions:

- a. Is the algorithm contributing directly or indirectly to social harms? If so, in what ways? And to whom? Is the harm caused by certain features of the algorithm? Can these harms be mitigated by changes to the algorithm? Can these harms be mitigated without causing harm to others?
- b. Is the algorithm benefitting certain members of society? If so, are those benefits accrued fairly and equitably?
- c. Is the algorithm benefitting society as a whole? If so, in what ways? Can those benefits be amplified or expanded?
- d. Are there "best practice" lessons to be learned from the design and implementation of this algorithm?

Method:

- A variety of social scientific and humanistic research designs

Ex-ante risk assessments and a regulatory duty to investigate

Children cannot be expected to understand or take action against automated decision making or algorithmic unfairness. It is unlikely that they have the developmental capacity, the knowledge or the resource to understand the subtle, cumulative or even acute nudges and impacts those automated systems have on their online experience. In fact, many children do not understand that an algorithm could be responsible for introducing them to a 'suggested friend', nor do they have the tools to prevent an onslaught of automated harmful material.

The AI Act would benefit from requiring ex-ante risk assessments in order to determine which AI systems are high-risk, with special consideration for the impact on children being a core part of this process. Barring this, AI systems that are likely to be accessed by or impact on children should be considered high-risk by default.

Then, the AI Act must give the regulator – whether the national market surveillance authorities alone or together with the Board and the Commission – not only the powers to interrogate automated systems but create the expectation that they will be actively analysing automated decision-making systems and algorithms of services that impact on children – a duty to investigate.

In order to fulfil this duty, the regulator(s) must have the expertise, resource, and processes in place to scrutinise the design goals, data inputs, model selection and outputs and outcomes of algorithms. Where there is evidence or outputs that indicate such systems are discriminating against or systematically disadvantaging children or violating their rights, the regulator(s) should set out a mandatory course of action for compliance.

While transparency is a key component of the four-step process set out above, decades of research show, transparency alone can result in layers of obfuscation and does not always result in a better systems or more positive outcomes. The value of transparency lies not in the availability of information itself, but in the way it allows for scrutiny and accountability. A duty for the regulator(s) to undertake the four steps on automated decision-making systems that impact on children would deliver that accountability.

Companies often use commercial sensitivity as a defence to usurp transparency reporting requirements. On the whole, this should be resisted, and where there are legitimate commercial sensitivities, the regulator(s) must have the power to maintain private oversight.

Proposed amendments for the Draft Artificial Intelligence Act

Title I – General Provisions

Article 3 – Definitions**Addition**

“child” is any person under the age of 18.

Title II - Prohibited Artificial Intelligence Practices**Article 5.1**

(b) the placing on the market, putting into service or use of an AI system that exploits any of the vulnerabilities of children or a specific group of persons due to their age, physical or mental disability, in order to materially distort the behaviour of a person pertaining to that group in a manner that causes or is likely to cause that person or another person physical or psychological harm;

Title III – High-Risk AI Systems**Chapter 1 – Classification of AI Systems as High Risk****Article 6 - Classification rules for high-risk AI systems**

2. AI systems likely to interact with or impact on children²¹ shall be considered high-risk.

Chapter 5 – Standards, Conformity Assessment, Certificates, Registration**Article 41– Common specifications**

Where harmonised standards referred to in Article 40 do not exist or where the Commission considers that the relevant harmonised standards are insufficient or that there is a need to address specific safety or fundamental right concerns, the Commission may, by means of implementing acts, adopt common specifications in respect of the requirements set out in Chapter 2 of this Title. Those implementing acts shall be adopted in accordance with the examination procedure referred to in Article 74(2).

The Commission shall adopt common specifications setting out how risk management systems should give specific consideration to interaction with or impact on children.

Title VI – Governance**Chapter 1 – European Artificial Intelligence Board****Article 58 – Tasks of the Board**

The Board shall provide statutory guidance in relation to children’s rights, applicable law and minimum standards for the evaluation of automated decision-making systems to meet the objectives of this Regulation pertaining to children and to investigate the design goals, data inputs, model selection, implementation and outcomes of such systems.

Title VIII - Post-market monitoring, information sharing, market surveillance**Chapter 3 – Enforcement****Article 65 - Procedure for dealing with AI systems presenting a risk at national level**

1. AI systems presenting a risk shall be understood as a product presenting a risk defined in Article 3, point 19 of Regulation (EU) 2019/1020 insofar as risks to the health or safety or to the protection of fundamental rights of persons are concerned.

²¹ For an AI system to be 'likely' to interact or impact on children, the possibility of this happening needs to be more probable than not. Whether an AI system is likely to interact with or impact on children will depend upon whether the content and design of the system is likely to appeal to children, and any measures in place to restrict or discourage their access to the service.

1a. When AI systems are likely to interact with or impact on children, the precautionary principle shall apply.²²

2. Where the market surveillance authority of a Member State has sufficient reasons to consider that an AI system presents a risk as referred to in paragraph 1, they shall carry out an evaluation of the AI system concerned in respect of its compliance with all the requirements and obligations laid down in this Regulation. When risks to the protection of fundamental rights are present, the market surveillance authority shall also inform the relevant national public authorities or bodies referred to in Article 64(3).

Where there is sufficient reason to consider that that an AI system exploits the vulnerabilities of children or violates their rights intentionally or unintentionally, the market surveillance authority shall have the duty to investigate the design goals, data inputs, model selection, implementation and outcomes of the AI system and the burden of proof shall be on the operator or operators of that system to demonstrate compliance with the provisions of this Regulation.

The relevant operators shall cooperate as necessary with the market surveillance authorities and the other national public authorities or bodies referred to in Article 64(3), including by providing access to personnel, documents, internal communications, code, data samples and on platform testing as necessary.

Where, in the course of its evaluation, the market surveillance authority finds that the AI system does not comply with the requirements and obligations laid down in this Regulation, it shall without delay require the relevant operator to take all appropriate corrective actions to bring the AI system into compliance, to withdraw the AI system from the market, or to recall it within a reasonable period, commensurate with the nature of the risk, as it may prescribe. The corrective action can also be applied to AI systems in other products or services judged to be similar in their objectives, design or impact.

Conclusion

The tech industry is worth over \$5 trillion to the world economy.²³ It is central to children's lives and life outcomes. Algorithmic oversight is critical if the next generation of digital technologies, products and services are to pay more than lip service to the needs of consumers and citizens, particularly children.

The four-step model of algorithmic oversight will reveal the goals, inputs, implementation and outcomes of algorithms and automated decision-making systems. This kind of transparency will support a change in corporate behaviour that meets the expectations of parents and uphold the rights of children. By giving the regulator a duty to interrogate automated decision-making systems on behalf of children, and service providers a clear process by which it will be done, the risks to children from automated decision-making systems can be reduced – by default and design.

²² Applying the Precautionary Principle for technology that may impact children and young people ensures that Child Online Safety is considered at an early stage. UNESCO's World Commission on the Ethics of Scientific Knowledge and Technology (COMEST) put forward a 'working definition' of the Precautionary Principle: 'When human activities may lead to morally unacceptable harm that is scientifically plausible but uncertain, actions shall be taken to avoid or diminish that harm. Morally unacceptable harm refers to harm to humans or the environment that is threatening to human life or health, or serious and effectively irreversible, or inequitable to present or future generations, or imposed without adequate consideration of the human rights of those affected.'

²³ <https://www.statista.com/statistics/507365/worldwide-information-technology-industry-by-region/>

There is no silver bullet to fix all the ills of the digital world or to guarantee children will be safe from harm, either through regulation or technological development. But the argument that regulation and accountability stifle innovation or impose limits on a child's freedom in the digital world is simply untrue.

It is in the interests of all parties to have a more equitable and trustworthy system of oversight that allows growth and innovation, but which reduces negative outcomes for children.

To do nothing is no longer an option.