CDT EUROPE

A Series on the EU AI Act

Pt. 2 – Privacy and Surveillance

May 2024

Authored by
***Laura Lazaro Cabrera**, Counsel and Director of the Equity and Data Programme, CDT Europe*

The AI Act clearly shows an ambition to protect human rights – and privacy features prominently as a right to be preserved in the deployment of artificial intelligence (AI) systems, with the acknowledgment that it is a right made vulnerable by certain types of AI. The Act not only explicitly highlights the applicability of existing EU law on privacy and data protection (Article 2(7)), but calls for the right to privacy and protection of personal data to be guaranteed throughout the entire lifecycle of the AI system (Recital 69).

Lastly, the Act creates clear obligations on high risk AI systems to ensure data governance that are consistent with data protection law priorities (Article 10). The multiple nods to privacy protections are a welcome – and indeed necessary – element of the AI Act, but numerous concerns about privacy and surveillance remain.

In this explainer, the second in our AI Act series, we delve into the key challenges posed by the AI Act in connection with the right to privacy, and more broadly on law enforcement uses of AI which raise surveillance concerns.

The National Security Carve-Out Issue

As civil society advocates observed throughout the AI Act negotiations, the AI Act's exemption for AI systems deployed for national security purposes is significant both in scope and in effect. Article 2(3) of the AI Act states that it shall not "affect the competences of the Member States concerning national security, regardless of the type of entity entrusted by the Member States with carrying out tasks in relation to those competences".

The fact that national security falls within the discretion of member states is not new. Indeed the Treaty on the European Union, the foundational text underpinning the creation of the EU, clearly states that "national security remains the sole responsibility of each member state". However, this does not mean that member states are free to manage national security as they please: the Court of Justice of the European Union has previously held, in the context of reviewing the applicability of privacy legislation, that the "mere fact that a national measure has been taken for the purpose of protecting national security cannot render the EU law inapplicable and exempt the Member States from their obligation to comply with that law". The AI Act exemption for national security, if read in a literal sense, would allow for intrusive and unethical technologies that would otherwise be illegal under this regulation to be created and deployed on grounds of national security.

This national security carve-out is both noteworthy and concerning. The AI Act defines the concepts of "law enforcement" and "law enforcement authorities" in ways that allow them to be construed broadly, and to potentially be protected as national security activities. For example, law enforcement is defined as the set of activities "carried out by law enforcement authorities or on their behalf for the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, including safeguarding against and preventing threats to public security". One example of an overlap would be the use of facial recognition capabilities to search for and identify a terrorism suspect at a large gathering or event, an activity which could be grounded in national security or law enforcement. A law enforcement entity could seek to avoid the restrictions imposed by the Act on real-time biometric identification carried out in public spaces simply by alluding to a national security purpose. This loophole can be dangerously interpreted and applied, opening the door to mass surveillance legitimised under the veil of "national security" activities.

As the AI Act gradually enters into application, it will be crucial for the European Commission to ensure that the overlap is not exploited to the detriment of fundamental rights across the EU.

Increased Scrutiny of Law Enforcement Use of AI

Despite the possible risk of conflation of national security and law enforcement activities, the Act does attempt to curtail the use of AI in law enforcement cases both by prohibiting specific uses of AI and by categorising all law enforcement uses of AI as "high-risk". This is an important step, although further clarification may be needed to better protect people's rights.

Prohibited Practices

No less than two of the eight practices prohibited by the AI Act directly tackle law enforcement activities: the use of real-time biometric identification (RBI) in publicly accessible spaces for the purposes of law enforcement, and the use of AI systems for making risk assessments predicting propensity towards criminal behaviour. However, these prohibitions are not absolute, and several exceptions threaten to undermine the rules.

The Ban on Real-Time Biometric Identification

The AI Act purports to ban real-time biometric identification (live facial recognition) carried out by law enforcement in publicly accessible spaces. In its definition of biometric identification, the Act describes the practice as the automated recognition of physical, physiological, and behavioural human features, including the face, eye movement, body shape, voice, prosody, gait, posture, heart rate, blood pressure, odour, keystrokes characteristics – all for the purpose of establishing an individual's identity by comparing biometric data of that individual to stored biometric data of individuals in a reference database. Facial recognition is currently far more advanced and broadly deployed than other remote biometric identification techniques, but a forward-looking definition that includes other biometrics is a smart measure to ensure rules apply as those identification technologies improve.

It is, however, noteworthy that the above definition is contained in a recital to the AI Act, as opposed to the section explicitly devoted to definitions under the Act. By contrast, the General Data Protection Regulation deals with the definition of "biometric data" directly in its definitions section, raising real questions as to whether recitals are the most suitable places to introduce key concepts that are likely to have a serious impact on fundamental rights.

To understand the extent and scope of the ban under the Act, the following key considerations must be taken into account:

- **The Act distinguishes between real-time biometric identification and post biometric identification.** While the real-time identification falls within the scope of the prohibition, "post" identification does not. By the Act's own definition, post biometric identification occurs when the biometric data has been captured and comparison and identification occur after "significant delay" (Article 3(43)). While the Act states that post biometric identification should not lead to indiscriminate surveillance, and there are some safeguards (including authorisation requirements in certain cases), post biometric identification still bears significant risks for human rights. Facial recognition is most broadly used by law enforcement in the "post" context for identification in investigations and other settings; absent strong safeguards there are harms to privacy, risk of misidentifications, and potential for abuse. Though this is explicitly discouraged in the Act, one could easily conceive of a scenario where law enforcement would seek to bypass the more stringent requirements applicable to real-time identification by collecting images at a protest or public event and postponing the comparison exercise in order to benefit from the lesser restrictions applicable to post biometric identification. The chilling effect on public assembly and expression would remain.

- **The prohibition on law enforcement RBI applies only to publicly accessible spaces.** The definition of "publicly accessible spaces" contained in the recitals explicitly excludes prisons and notably the border, a known site of human rights abuse. As such, the AI imposes no prohibition on RBI applied in migration control settings along the border, such as crossing points, despite the heightened vulnerability of people in such settings and the difficulty for civil society advocates or others to identify and remedy abuses.
- **The Act creates numerous exceptions to the RBI prohibition** which threaten to swallow the rule. The Act permits the use of RBI in three cases:
  - » To search for a missing person or a trafficking, abduction, or sexual exploitation victim;
  - » To prevent a specific, substantial, and imminent threat to individuals or a genuine and present/genuine and foreseeable threat of a terrorist attack; or
  - » To identify or locate a person suspected of having committed a crime listed in Annex II with a given degree of seriousness, established by reference to relevant member state penalties, which must be custodial sentences or detention orders of at least 4 years.

  These exceptions could be construed broadly: For example, the serious crime exception does not require an objective and verifiable link to an ongoing investigation, but only that identification is necessary in order to conduct a potential investigation for a serious offence. This means that the identification of the suspect may ultimately not lead to prosecution and may have negative impacts on the person.

  If any of the above scenarios apply, the prohibition is lifted and the RBI, instead of falling in the prohibited category of AI systems, is downgraded to the high-risk category and is accordingly subject to safeguards.

Safeguards Around Authorised Uses of RBI

Once RBI falls within any of the exceptions contemplated above, safeguards apply in the form of i) independent authorisation, ii) fundamental rights impact assessments, and iii) record-keeping. Law enforcement are required to obtain judicial or independent administrative authorisation prior to the deployment of RBI or, "in duly justified cases of urgency", 24 hours after deployment at the latest (Article 5(2)). The Act explicitly provides that the AI shall be limited to identify the identity of the specifically targeted individual, which overlooks the fact that the nature of scanning of biometric information in real-time is such that, even if an AI is looking for a particular individual, everyone's faces are scanned to establish a match.

In order to obtain approval, law enforcement must prepare a fundamental rights impact assessment, and ensure that the relevant AI system is registered in the non-public section of the EU database created for high-risk AI systems (Article 49). The authorising authority must consider whether RBI deployment is: i) necessary and proportionate to the objectives; and ii) whether it is strictly limited in time, geographic, and personal scope. If rejected, RBI must be stopped with immediate effect and all data as well as related outputs deleted and discarded.

In addition to the authorisation process, each deployment of RBI must be notified to the relevant market surveillance authority and the national data protection authority (Article 5(4)).

Crucially, member state legislation is needed for RBI deployment even under the circumstances authorised by the Act. Article 5(5) requires member states seeking to benefit from the exceptions to the prohibition on RBI to lay down in their national law the necessary detailed rules for the request, issuance, and exercise of, as well as supervision and reporting relating to, the abovementioned authorisations. Member states are free to introduce more restrictive laws on the use of remote biometric identification systems.

Safeguards Around Post Biometric Identification

Post biometric identification is subject to fewer safeguards than RBI. To start, the Act does not impose limits on the types of crimes the investigation of which could legitimately allow post biometric identification. In contrast to the Act's approach to real-time identification, which is only allowed in the context of criminal investigations when they concern a narrow set of offences listed in Annex II *and* only if these are subject to custodial sentences or detention orders of a minimum length of 4 years, the provisions on post biometric identification do not require that the investigated offence be of a given type or degree of seriousness.

There is an authorisation requirement in all cases except when an AI system is used for the initial identification of a potential suspect based on objective and verifiable facts directly linked to the offence. In all other cases, authorisation must be sought either prior to the deployment of the AI system or within 48 hours afterwards. While the AI Act does not explicitly state that a fundamental rights impact assessment is required for authorisation to be granted, our reading of the Act is that any use of post-biometric identification would require a fundamental rights impact assessment in any event as it would fall under the high-risk category of AI and concern a body governed by public law.

Similarly to the provisions around RBI, rejection of the authorisation request requires law enforcement to stop the use of the AI system with immediate effect and delete all related data. However, the reporting obligations are different from those applicable to RBI. Instead of requiring notification to the market surveillance authority and the national data protection authority of each use of post-biometric identification, the Act simply requires for the uses of post biometric identification to be documented and made available to these authorities upon request (Article 26(10)).

See chart on the next page.

	Real-time biometric identification by law enforcement in publicly accessible spaces	**Post-biometric identification**
Acceptable use cases	Criminal offenses, when offense is listed in Annex II and subject to custodial sentences of 4 years or more Targeted searches for crime victims or missing persons Prevention of threats to life or physical safety of individuals or a terrorist attack	Not restricted to specific use cases, but its use in the law enforcement context cannot be untargeted
Authorisation	Required prior to deployment or within 24 hours	Authorisation generally not required **except** for a targeted search of a person suspected or convicted of having committed a criminal offence, either prior to deployment or within 48 hours. Such authorisation is however not required for the “initial identification” of a potential suspect based on “objective and verifiable facts” linked to the offence
Fundamental rights impact assessment	Required for authorisation to be issued in all cases	Not required
Registration in high-risk AI database	Required for authorisation except in situations of urgency	Not required
Notification to regulators	Required for each individual use to the market surveillance authority and data protection authority	Not required for each individual use, but deployers must make annual reports to market surveillance authorities and data protection authorities (DPAs)
Notification to individuals	The Act explicitly states that the notification procedure laid out in pre-existing legislation (Directive 2017/680) shall apply	For law enforcement uses, the Act explicitly states that the notification procedure laid out in pre-existing legislation (Directive 2017/680) shall apply. In all other cases, notification is required under the Act when biometric identification is used to make decisions about individuals or assist in making decisions about them
Further domestic legislation	Required for states wishing to use RBI for law enforcement purposes	Not required

The Ban on AI Used for Criminal Profiling

The Act purports to prohibit AI systems which make risk assessments to assess or predict the risk of a person to commit an offence, where it is based solely on profiling of the person or assessing their traits and characteristics. However, as with RBI, a crucial exception exists: such AI systems are allowed if used to support "the human assessment of the involvement of a person in a criminal activity" which itself must be based on verifiable facts and linked to a criminal activity. In other words, the prohibition only applies to AI systems engaging in predictive policing if these risk assessments are made in a vacuum. If used to support a human assessment of criminal propensity, the prohibition is lifted and a predictive AI system may be used, subject to the safeguards for high-risk uses in Annex III.

There is a real concern that the exception to this prohibition may be exploited for predictive policing purposes, i.e. law enforcement uses. However, another equally concerning use of the exception would be in the field of administration of justice, in the context of predicting recidivism of a suspect or defendant. As some have pointed out, it would be possible for some of the existing risk assessment tools deployed to assess potential for recidivism to be seen as providing a supportive role in recidivism risk assessments, and therefore as falling within the exception.

As CDT outlined in an earlier explainer in this series, a key flaw with the AI Act provisions is that providers of AI are allowed to assess their AI as being outside the high risk category. Crucially, however, this is not possible where an AI system undertakes profiling, as this is the only type of high-risk AI from which a derogation cannot be obtained (Article 6(3)).

Law Enforcement Uses Categorised by Default as High-Risk AI

Annex III of the AI Act explicitly categorises as high-risk several types of AI deployed in the law enforcement context, ranging from the use of tools such as polygraphs and similar tools to AI predicting the risk of a person becoming a victim or offender.

The high risk categorisation is crucial under the Act, as it engages other provisions enabling the governance and oversight of high-risk AI systems. As mentioned above, this starts with the requirement to undertake a fundamental rights impact assessment in every case where the deploying entity is a body governed by public law, which squarely includes law enforcement. Additionally, law enforcement must register the AI system in the EU database created by Article 49, take appropriate technical and organisational measures to ensure compliance (Article 26 (1)), monitor the operation of the high-risk AI generally, and stop using the system altogether if it presents a risk to the health or safety or fundamental rights of persons under Article 79(1) (Article 26(5)).

While the above safeguards are welcome, there are key limits that obstruct public oversight and transparency. To start, the database in which law enforcement authorities must register their high-risk AI is non-public and only accessible to the Commission. Additionally, there are real limits to the extent to which an AI tool used in the law enforcement context may present as such to a person directly exposed to it. The AI Act creates only limited transparency obligations on providers and deployers of AI. By virtue of the fact that law enforcement departments frequently rely on technology provided by third parties, they are likely to be considered "deployers" under the Act. However, the Act allows deployers not to disclose that they are using an AI for biometric categorisation and emotion recognition purposes where they are permitted by law to detect, prevent or investigate criminal offences, subject to appropriate safeguards (Article 50 (3)). These safeguards are not further detailed by the Act. Similarly, deployers of AI systems generating deepfakes are not required to disclose that the content is artificially generated if the use of such systems is authorised by law to detect, prevent, investigate, or prosecute criminal offences (Article 50(4)).

Conclusion

Despite the AI Act's efforts to closely monitor and regulate AI uses by law enforcement, there are various aspects which lead to genuine concerns that not all harmful uses of AI by law enforcement have been effectively guarded against. To mitigate these risks where possible, all relevant actors under the AI Act will need to take action.

The Commission will be in a unique and crucial position to assess high-risk AI systems deployed by law enforcement by virtue of its exclusive access to the non-public section of the EU database where these will be listed. It will have to ensure that it is exercising oversight proportionate to its unparalleled level of knowledge and access. Market surveillance authorities and data protection agencies must ensure that they thoroughly review reported instances of use of real time biometric identification, and that they proactively request information on uses of post biometric identification. Lastly, market surveillance authorities must pay close attention to, and comprehensively review, fundamental rights impact assessments undertaken by law enforcement in relation to high-risk uses of AI.

Find more from the CDT Europe team on the EU AI Act at *cdt.org*.

*The **Center for Democracy & Technology (CDT)** is the leading nonpartisan, nonprofit organization fighting to advance civil rights and civil liberties in the digital age. We shape technology policy, governance, and design with a focus on equity and democratic values. Established in 1994, CDT has been a trusted advocate for digital rights since the earliest days of the internet. The organization is headquartered in Washington, D.C. and has a Europe Office in Brussels, Belgium.*