September 2026

Responsible Artificial Intelligence Innovation

Insights from the Data Protection Commission's Supervision of AI (2021-2025)

An Coimisiún um Chosaint Sonraí
Data Protection Commission

Contents

Table of Figures

Glossary of Abbreviations

The following abbreviations are used throughout this report:

AI – Artificial Intelligence

DPC – Data Protection Commission

DPIA – Data Protection Impact Assessment

DPO – Data Protection Officer

EDPB – European Data Protection Board

EEA – European Economic Area

EU - European Union

GDPR – General Data Protection Regulation

LIA – Legitimate Interests Assessment

LLM – Large language model

LSA – Lead Supervisory Authority

Glossary of Terms

Term	Definition
Artificial Intelligence (AI)	Technology that enables computers and machines to simulate human intelligence, learning, problem solving, calculations, or ability to perform actions.
AI Opinion	EDPB Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models issued 19 December 2024.
Data subject	An identifiable natural living person who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, an online identifier or to one or more factors specific to their physical, genetic, mental, economic, economic or social identity.
Data Protection Act 2018	The Irish national law which gives further effect to the GDPR in and effect to the Law Enforcement Directive and which established the Data Protection Commission.
Data Protection Commission (DPC)	Ireland's data protection regulator.
Data Protection Impact Assessment (DPIA)	An assessment carried out by a data controller of the impact of envisaged processing operations on the protection of personal data where those processing operations are likely to result in a high risk to the rights and freedoms of natural persons.
European Data Protection Board (EDPB)	An independent European body, which contributes to the consistent application of data protection rules throughout the European Union and promotes cooperation between the EU's data protection authorities.
European Economic Area (EEA)	An area of free trade and free movement of peoples comprising the member states of the European Union, in addition to Norway, Iceland, and Liechtenstein.

General Data Protection Regulation (GDPR)	Lays down rules relating to the protection of natural persons, with regard to the processing of personal data and rules relating to the free movement of personal data. A principled piece of legislation which allows organisations of different sizes, and with different objectives, to decide on, and devise their own processes and procedures, to meet their respective objectives, once there is adherence to the basic principles as set out in data protection law.
Legitimate Interests Assessment (LIA)	An assessment that considers the reliance on the legal basis of Article 6(1)(f) that considers the legitimate interests of the controller or a third party, the necessity of the processing, balanced against the interests or fundamental rights and freedoms of the data subject impacted by the processing, in particular where the data subject is a child.
Large Language Model (LLM)	A type of generative AI language, model trained on text designed for natural language processing tasks, especially language generation.
Lead Supervisory Authority (LSA)	The supervisory authority of the main establishment or single establishment of a controller or processor who is competent for the cross-border processing, as well as the processing within their own jurisdiction, as per Article 56 GDPR.
Personal Data	Any information about a living person where that person is either identified or identifiable.
Processing	Doing something with the personal data, such as collecting, recording, organising, structuring, storing, adapting, altering, retrieving, consulting, using, disclosing by transmission, sharing/ making the data available, aligning or combining, restricting, erasing or destroying personal data.
Supervision	The act of monitoring, consulting, and supervising data protection compliance outside of other regulatory tools, such as complaints or inquiries.

Supervision Function	Refers to the broader use of Supervision by the DPC as a whole. It contains numerous units, each one focused on different sectors and controllers.
Technology Multinational Supervision Unit (TMSU)	A single unit within the DPC Supervision Function.

Executive Summary

As the EU Lead Supervisory Authority for many of the world's leading technology companies with European Headquarters located in Ireland, the Data Protection Commission (DPC) occupies a unique regulatory vantage point at the intersection of rapid technological evolution and fundamental rights protection. This report presents the insights gathered by the Data Protection Commission's Technology Multinational Supervision Unit within the DPC's Supervision Function regarding the supervision of Artificial Intelligence products and services between 2021 and 2025.

Distinct from formal inquiries, the DPC Supervision Function operates a quasi-regulatory sandbox which is a voluntary, engagement-driven model, allowing the DPC to influence product design at the pre-processing stage. The objective is to ensure that a data protection by design and default approach is embedded into product design at the earliest juncture of product development, and that significant risks and harms are identified before the launch of new products and services into the European market.

The report details a significant increase in AI-related engagements, driven in part by the rapid development of Generative AI in recent years.

Key findings highlight that innovation and rigorous data protection are not mutually exclusive. Through a large number of engagements involving different AI such as Large Language Models and recommender systems, the DPC Supervision Function has secured significant improvements in data protection compliance in the areas of lawful basis, transparency, data minimisation, and the protection of children. The report demonstrates a number of critical areas of focus for the DPC, particularly the application of Legitimate Interests as a legal basis for AI training, and the necessity of robust Transparency for what is novel technology, often involving opaque and complex processing operations.

While the majority of AI engagements resulted in recommendations being issued, the report also demonstrates the DPC's readiness to urgently intervene where risks to individuals' rights are unsatisfactorily mitigated.

Ultimately, this report serves as both a record of the DPC's actions in regulating emerging technologies such as AI, and as guidelines for controllers, signifying that early regulatory engagement is the most effective path to sustainable, responsible, and data protection and privacy-centric innovation.

An Introduction to DPC Supervision

As part of our work regulating under the General Data Protection Regulation (GDPR), the Data Protection Act 2018, ePrivacy Directive, and Law Enforcement Directive, the Data Protection Commission (DPC) conducts 'Supervision' of controllers. The DPC has a broad mandate for the purpose of monitoring and enforcing the aforementioned regulations, to ensure the rights and freedoms of individuals.

Supervision is one of the tools the DPC uses from its regulatory toolbox. It forms one prong of the DPC's mission to uphold the consistent application of data protection law, and to drive compliance with data protection legislation.[1] In line with the values of the DPC, Supervision provides fair, expert, independent, forward looking, results-driven engagement and guidance with stakeholders on upcoming products and services, while encouraging responsible innovation and data management practices.

Supervision also plays a key role in enhancing the DPC's technological foresight, by obtaining information on rapidly evolving technologies, like Artificial Intelligence (AI), to ensure the DPC is equipped to effectively regulate the technologies of the future. Given that the processing of data by AI can be very opaque, the proactive approach of Supervision has been a valuable tool in knowing what processing is occurring, and in delivering changes and improvements to benefit individuals.

The DPC's Supervision Function operates completely separately from other regulatory functions in the DPC, such as those involved in complaints or inquiries.[2] Rather than looking at infringements of the GDPR that are alleged to have occurred in the past, the Supervision Function engages with controllers on a voluntary basis, primarily before a product or service launches.

The Supervision Function does not carry out full compliance audits, but focuses on red line issues, high risks, and likely harms. Supervision implements a risk-based approach, prioritising services and products which have the potential to have a

[1] As per DPC's Regulatory Strategy 2022-2027: Mission: Upholding the consistent application of data protection law through engagement, supervision and enforcement, and driving compliance with data protection legislation.
[2] The DPC has Supervision Units servicing all sectors of the Irish economy.

higher impact on individual's lives. The DPC's goal at pre-processing stage is to ensure that controllers innovate responsibly, mitigate identified harms and risks to individuals, and appropriately consider individuals' rights by balancing their interests, rights, and freedoms, against the controllers' interests.

In effect, Supervision operates quasi-regulatory sandboxes with every issue assessed. The DPC does this work in cooperation with members of the European Data Protection Board (EDPB). Where Supervision identifies issues of concern, the DPC will engage with our peer EU regulators and relay the collective concerns of both the DPC and EU regulators to the controller, with the aim of ensuring sufficient GDPR guardrails are implemented to protect society's interests in the form of data subjects' rights across the EU/ EEA. Where appropriate, the DPC Supervision Function also consults with data protection authorities outside of the EU, on common areas of concern.

This work provides the DPC with insights into the processing carried out by controllers in our jurisdiction and across Europe, as well as providing the DPC with the opportunity to guide controllers towards compliance, by making recommendations prior to features or products being launched. The positive relationship between legal certainty and fostering innovation has been well established. The DPC's Supervision approach is an invaluable resource for firms on the cutting-edge of technology seeking more legal certainty and information-sharing in a fast-paced eco-system.

Supervision of AI

The EU's AI Continent Action Plan strategy aims to deliver both broad adoption across the economy and to ensure that "*AI is trustworthy and aligned with EU values*."[3] Although AI adoption in the EU has increased from 13.5% to 20%[4] in 2025, adoption is fragmented, and there is early evidence which indicates that increases in the rate of AI adoption tend to increase the risk of digital security breaches. [5] Notably in 2025, 52.7% of EU firms which considered using AI technologies and chose not to, cited data protection and privacy concerns as a reason why.[6] This suggests that questions related to data protection and privacy will be paramount for further AI development and diffusion, and in its consultative and awareness-raising capacity, the DPC can be of assistance.

The DPC has facilitated responsible, data protection and privacy focussed innovation through its engagement driven strategy on AI. While AI impacts all sectors, this Insights document is based on the engagements of the Technology Multinational Supervision Unit (TMSU), whose remit is to engage with the large multinational tech companies based in Ireland.[7]

The DPC is in a unique position in Europe due to its position as the Lead Supervisory Authority (LSA) for the many technology companies main established in Ireland. This means the DPC is the primary data protection regulator of those companies based in Ireland. The DPC engages with these controllers on a daily basis, while also working in close cooperation with the other supervisory authorities within the EDPB.

As LSA, the DPC has engaged intensively with many of the leading technology companies at the forefront of AI developments, in particular concerning the use of personal data to train LLMs in the EU/EEA. This has resulted in a number of

[3] Commission, 'AI Continent Action Plan' (Communication) COM (2025) 165 final
[4] In 2025, 19.95% of EU enterprises used AI technologies. Eurostat (2024) Use of artificial intelligence in enterprises. Retrieved 2025, from https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Use_of_artificial_intelligence_in_enterprises
[5] Kergroach, S., & Héritier, J. (2025). Emerging divides in the transition to artificial intelligence. *OECD Regional Development Papers*, 147th ed. doi:https://doi.org/10.1787/7376c776-en
[6] Eurostat 'Use of artificial intelligence in enterprises' (Online data code: isoc_eb_ai, 2024) doi:https://doi.org/10.2908/ISOC_EB_AI Retrieved 2025, from https://ec.europa.eu/eurostat/databrowser/view/isoc_eb_ai__custom_20332450/default/table>
[7] Please note, all statistics in this document relate to the TMSU in the DPC's Supervision Function

companies implementing improvements and additional data protection safeguards prior to launch.

In particular, from 2021 through 2025, the Supervision Function has engaged in supervising the creation, training, and deployment of AI by companies, including Airbnb, Apple, Deepseek, Google, LinkedIn, Meta, Microsoft, OpenAI, Pinterest, TikTok, X (formerly Twitter), and others. The DPC has supervised a range of AI including LLMs, age assurance, facial recognition, recommender systems, personalisation, agents, and more. The DPC recognised that the opportunity existed to foster responsible innovation by engaging with controllers in the EU concerning the use of personal data to create and deploy AI in the EU.

Between 2021 and 2025, the DPC engaged with controllers on the creation, launch, and deployment of approximately 180 AI products and services, and assessed thousands of pages of briefings, risk assessments, technical and organisational measures, and compliance documentation in relation to AI. Through the Supervision Function's work, the DPC has developed the skills and knowledge to provide educated, expert, and concrete recommendations and guidance to controllers at the forefront of AI development.

Through Supervision, the DPC has provided guidance on data protection matters and encouraged responsible innovation of AI in Ireland and Europe. This has resulted in a number of companies implementing improvements and additional data protection safeguards prior to launch in the EU.

While acknowledging that aspects of the application of GDPR to this fast-changing technology remain complex, through Supervision, the DPC has ensured that controllers understand that innovation and technological developments can work simultaneously with reducing the risk of harm to individuals.

Latest figures show that from 2021 to 2025, the DPC Supervision Function has seen a tenfold increase in AI-related engagements with large technology controllers.[8]

[8] Preliminary reporting stated that the Supervision Function saw over a six-fold increase in the number of AI-related engagements. See Joint Committee on Artificial Intelligence Deb 24 February 2026

Figure 1- DPC Supervision AI Engagements by Year (2021-2025)

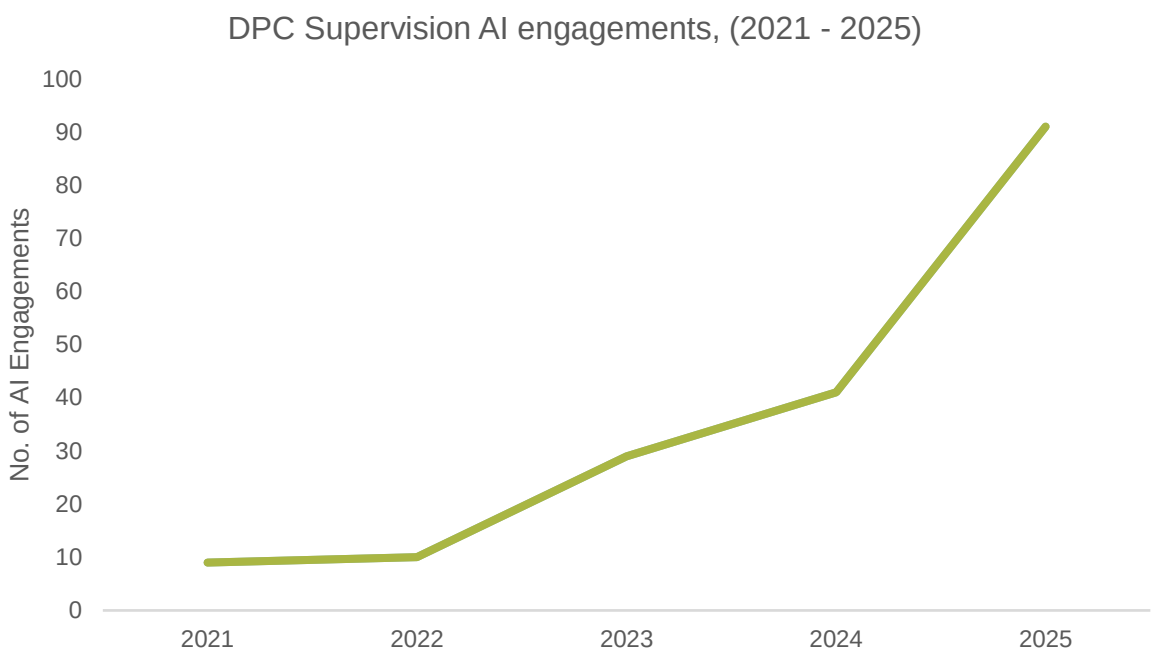

Note: Data sourced from internal TMSU Statistics including the full year of all AI engagements (January to December) in 2025.

To put this into perspective, in 2021, 1 out of every 35 engagements undertaken by the Supervision Function were related to AI. By 2025, AI engagements accounted for around 1 out of every 4 DPC Supervision engagements with large technology controllers.

Figure 2 - DPC Supervision AI Engagements (AI and Total)

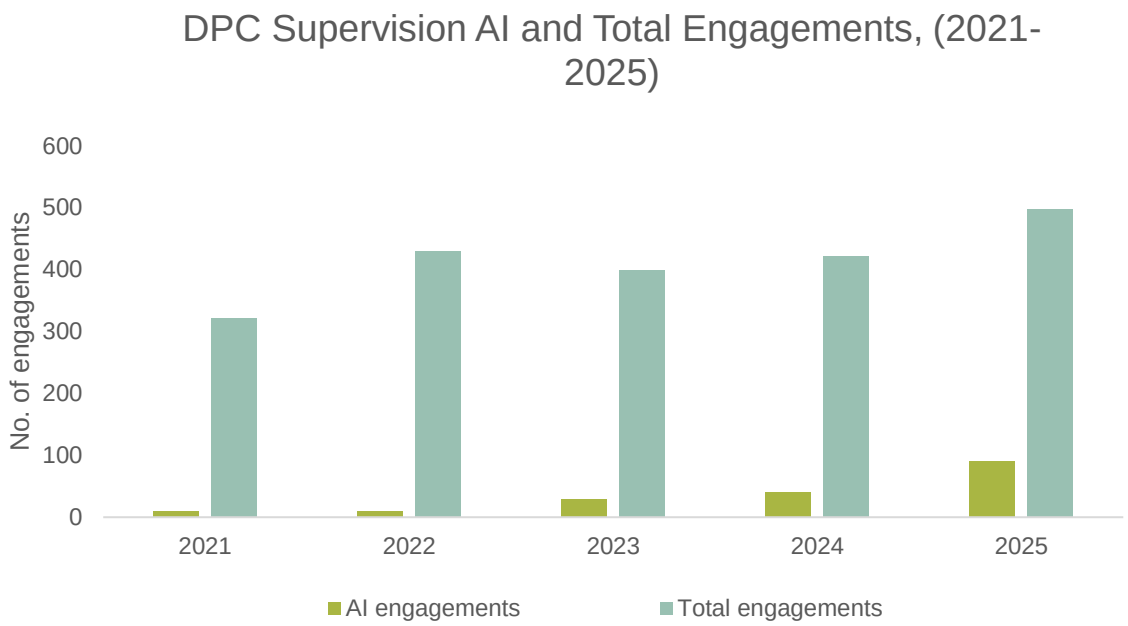

Note: Data sourced from internal TMSU statistics. Totals reflect all engagements with controllers and other stakeholders as reported in the DPC Annual Report.

Since its introduction to the market, Generative AI engagements with the Supervision Function have increased year on year; 25 in 2023, 36 in 2024, 82 in 2025. In total, during this time period, the Supervision Function had 143 engagements with controllers on Generative AI and 37 non-generative AI engagements, for a total of approximately 180 AI engagements. It is worth noting that any analysis which only considers the number of engagements should also consider that AI engagements typically involve more complexity and a larger workload per engagement.

Figure 3 - DPC Supervision AI Engagements (YOY Growth 2022-2025)

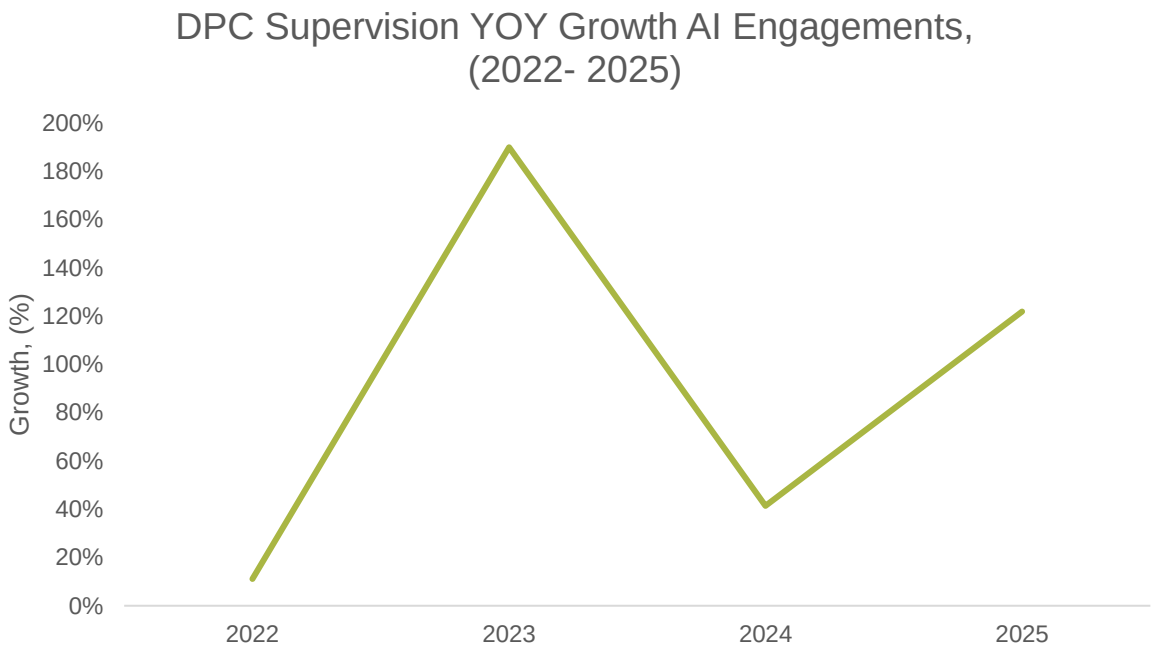

*Note: Data sourced from internal TMSU Statistics *Note the 2025 growth rate is calculated using the full year 2025 count of engagements.*

In 2023, there was a spike in the number of AI related engagements as generative AI was launched.

Figure 4 - Most Frequent AI Technology Type over a Five-year Period in DPC Supervision Engagements

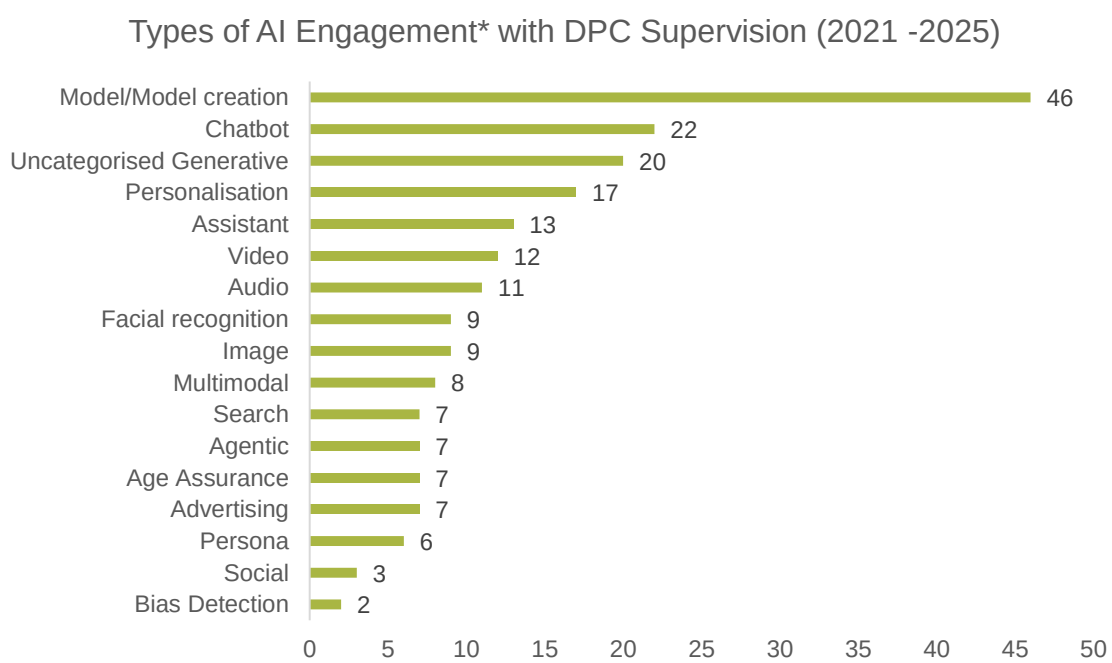

*Note: Data sourced from internal TMSU Statistics. *There is some crossover in types of AI counted for one AI engagement (for example, model creation and age assurance).*

Figure 5 - DPC Supervision Generative AI Engagements by Technology Type

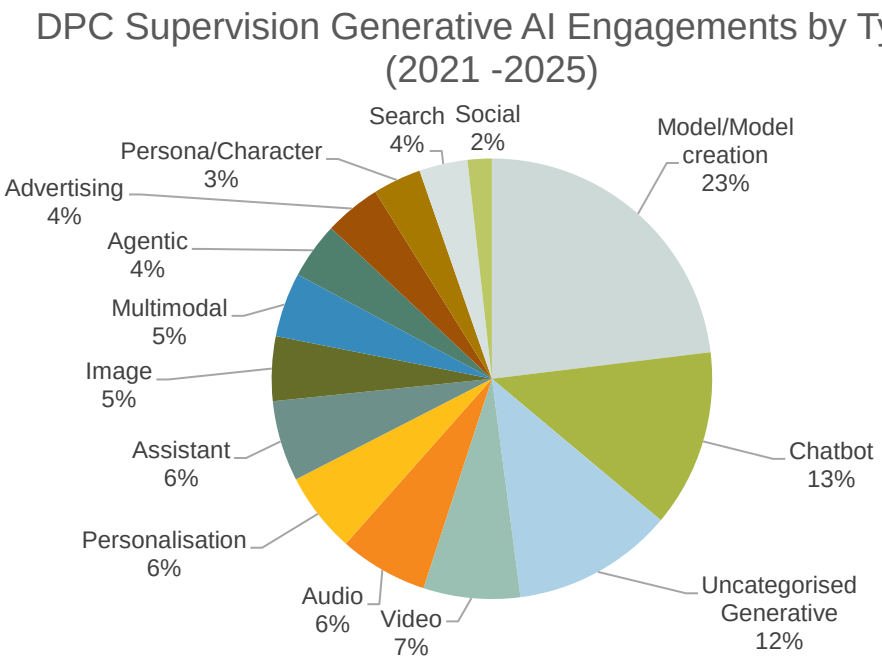

Note: Data sourced from internal TMSU Statistics

Figure 6 - DPC Supervision Non-Generative AI Engagements by Technology Type

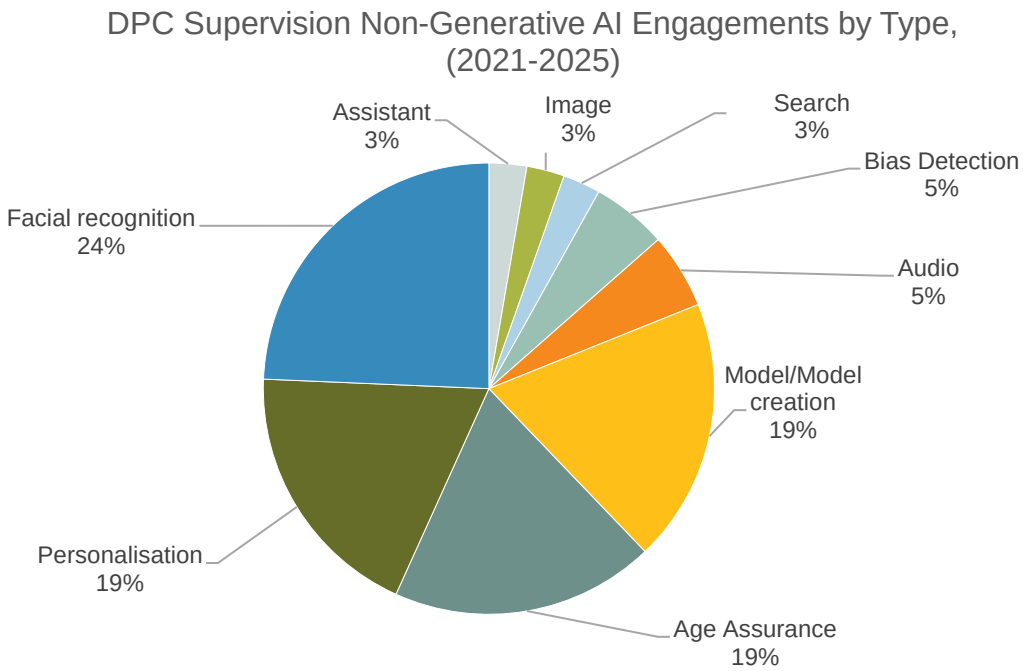

Note: Data sourced from internal TMSU statistics

Generative AI brought an explosion of AI engagements to the DPC. Despite generative AI engagements only commencing in 2023, generative AI was involved in 79% of all AI Supervision engagements between 2021 and 2025. Non-generative AI, which includes most age assurance and facial recognition engagements, only accounts for 21% of all engagements during the same period.

What does Supervision engagement entail?

The DPC is conscious that there is no one-size fits all approach for many aspects of data protection compliance under the GDPR, and this is true for Supervision engagements as well. Engagements with the DPC's Supervision Units vary, depending on the needs of the controller or the DPC, or due to the complexity of the AI or AI system being assessed. The DPC has used the full range of regulatory functions as part of its engagements, ranging from making recommendations, requesting controllers pause a proposed launch, and obtaining court orders to cease processing, via urgency action powers.

At a high level, Supervision involves:

- An open and honest conversation regarding a controller's intended processing or the capabilities of launches, products, services, etc. that will process personal data;
- Questions by the DPC and answers by the controller;
- As part of the assessment process, engagement with peer EU regulators to identify issues of concern.
- Recommendations (where necessary) by the DPC that indicate areas for improvement or point out likely issues for compliance;
- Follow-up by the controller on the DPC's recommendations; and
- Ongoing monitoring by the DPC post-launch.

Supervision will assess a product or proposed processing from a high-level, high-risk perspective, but not to the level of legal certainty that would be required for an inquiry. As such, **the DPC will never approve or state that a process, product, or service is compliant with the GDPR through Supervision**.

For engagements involving AI, this may include an assessment of a Data Protection Impact Assessment (DPIA), Legitimate Interests Assessment (LIA), and other risk assessments related to the product/service. It may include the monitoring of the technical performance of key mitigations or required processes, such as the operation of an objection form for compliance with GDPR Article 21.

The Supervision engagement process is usually conducted on a voluntary basis, as opposed to the more formal prior consultation under Article 36 of the GDPR. Controllers often give the Supervision Function a level of access that may not be afforded elsewhere, with the knowledge that providing this access will ultimately benefit them, as well as demonstrating cooperation with the DPC as LSA, as required under Article 31.

Benefits to the Supervision Model

In advance of the processing of personal data, Supervision advocates for the mitigation of harms, which protects the rights of individuals. Supervision also identifies and prevents noncompliant processing from taking place, while also supporting responsible innovation.

New products, services, and applications which process personal data that have benefited from Supervision engagement with the DPC are more likely to be compliant with the GDPR. This has the effect of potentially:

- Reduced risks and harms for individuals;
- Reducing the number of complaints;
- Avoiding investigations or inquiries under Article 58 GDPR;
- Avoiding potential urgency actions; and
- Providing regulatory guidance and a level of legal certainty for controllers.

These benefits directly relate to the Strategic Goals of the DPC, including safeguarding individuals, promoting data protection awareness, bringing clarity to stakeholders, and driving compliance.

Through Supervision, the DPC has achieved notable results for individuals by encouraging responsible innovation, but the benefits of Supervision are not just to the individual. Developers and deployers of AI technologies can take advantage of Supervision's 'informal sandbox' engagement. While a Supervision Unit will not categorically state that a product or service is compliant with the GDPR, the risks to controllers who proceed with processing based on the recommendations of the DPC are lower, having the certainty that they have considered the risks to individuals and GDPR requirements having had input from the regulator.

Supervision engagement with organisations includes the DPC being provided with confidential or commercially sensitive information. By satisfying the DPC's questions, including queries from other EU supervisory authorities, and providing access to documentation and technology, a controller can gain a level of confidence that its internal risk assessments, design decisions, and technical and organisational measures, are meeting regulatory expectations.

Controllers are also provided valuable guidance on what steps can be taken or areas can be improved to better meet the requirements of the GDPR. Being able to understand the expectations of the DPC as LSA working in cooperation with the EU/EEA's data protection authorities can assist in developing a controller's understanding of their obligations, with respect to particular products and processes, which is in line with the Helsinki Statement[9] aims.

The DPC does not validate that a particular feature or product is 100% compliant. However, controllers can significantly benefit from its views and insights. By adhering to the DPC's recommendations and mitigating specific harms and red line issues it has identified, controllers are well positioned to achieve compliance prior to a product or service launch. This can result in fewer complaints and reduce the risk of an inquiry under Article 58 of the GDPR.

[9] European Data Protection Board 'The Helsinki Statement on enhanced clarity, support and engagement' 3 July 2025 available at <https://www.edpb.europa.eu/our-work-tools/our-documents/statements/helsinki-statement-enhanced-clarity-support-and-engagement_en>

Supervision Recommendations and Analysis

The DPC informs organisations of its views in a number of ways, most often in the form of recommendations. Recommendations are made where the DPC has formed the view that the proposed/planned processing may not be GDPR compliant, should the product or service launch as currently designed. The Supervision Unit also uses recommendations to promote changes in corporate culture towards compliance, encourage good privacy practices by controllers, and instil a data protection by design and default mindset in line with the DPC's strategic goals.

As the Supervision Unit primarily engages with controllers pre-launch, recommendations may not be reflective of the actual launched product or service or the processing involved. This is because controllers have the opportunity to change or make improvements, prior to putting the product or service on the market.

Example recommendations include:

- The DPC recommends that clear retention periods or criteria to determine such periods, be developed and provided to users with regard to the transcription of audio data;
- The DPC recommends that the Privacy Policy language be adapted for children, in order to fully comply with Article 5(1)(a) and 13(1)(d) of the GDPR;
- The DPC recommends that additional transparency notices should be published prior to the creation of the datasets that will be used for training and testing of the models;
- If the controller intends to continue to process any personal data of users who object to the processing, the controller must demonstrate compelling legitimate grounds prior to launching in the EU/EEA;
- The DPC recommends a critical re-evaluation and adjustment of the consent flow, to ensure that it meets the requirements of 'freely given' consent under Article 7(4) of the GDPR;
- The DPC recommends that the controller develop and apply measures to remove, pseudonymise, anonymise, or otherwise protect full names of private individuals from the dataset prior to any AI training; and

- The DPC recommends that the controller take immediate steps to avoid collecting or processing personal data from pages, groups and elsewhere on the platform that would have a likely prevalence of special category data, or data capable of revealing special category data.

Recommendations are based on full, robust, analysis of the technology and law, based on the DPC's understanding of the processing. Such analysis is conveyed to controllers and can include:

- Stating the DPC's opinion that a controller's legitimate interests balancing test has not been met;
- Stating that the DPC questions the controller's ability to rely on a chosen legal basis;
- Informing a controller that risks have not been mitigated to a sufficient degree, or that risks have not been fully considered; and
- Informing a controller that the compatibility assessment of processing purposes has not been met.

The DPC's recommendations are non-binding in nature, however where redline issues are identified, the Supervision Function can indicate that the DPC may consider other regulatory powers if the controller decides not to implement the recommendations.

In 2024 and 2025 combined, 59% of all AI products or services that launched in the EU were issued with recommendations.

It should be noted that a single recommendation may cover several aspects of the GDPR. For example, a recommendation to avoid training on a specific range of personal data may also touch on data minimisation, purpose limitation, and the legal basis of a processing.

The median number of recommendations issued for an AI engagement was 3. For a single engagement, the highest total number of recommendations made was 18, with the lowest being 0. This further demonstrates that the Supervision Function engages in an 'as and when needed' risk-based approach but will engage deeply where multiple challenges to compliance are present.

Figure 7: Top Five Cited Data Protection Issues in DPC Supervision AI Recommendations.

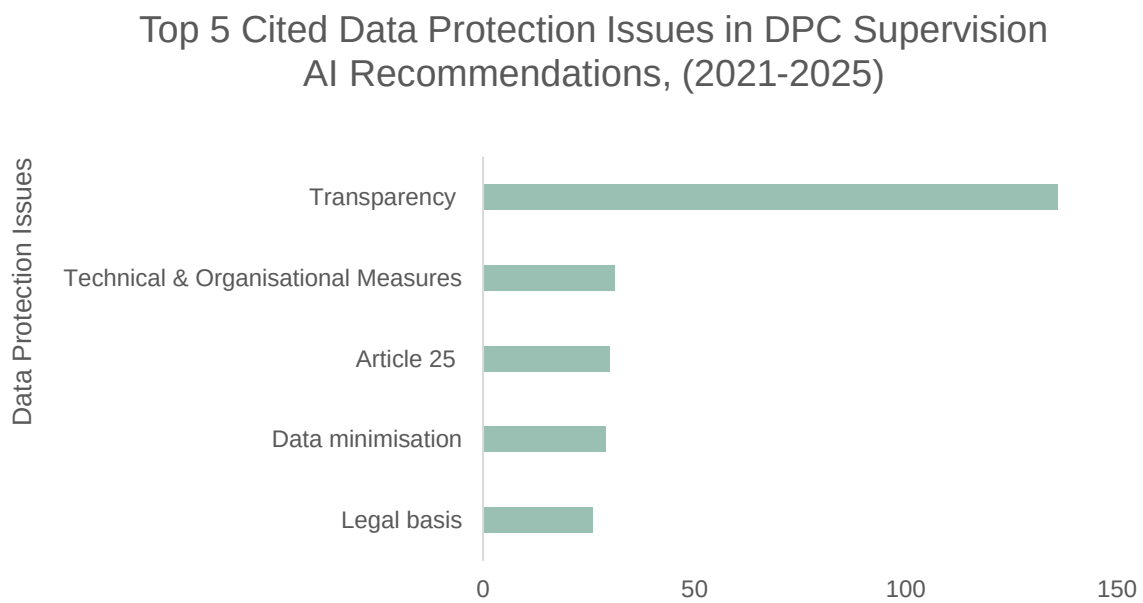

Note: Data is sourced from internal TMSU Statistics.

Although the specific nature of the Supervision Function's engagement with controllers often leads to citing particular issues more frequently in recommendations, all data protection issues are reviewed during DPC engagements, and controllers should create products and services with that in mind.

Figure 8 - Data Protection Issues Cited in DPC Supervision AI-related Recommendations

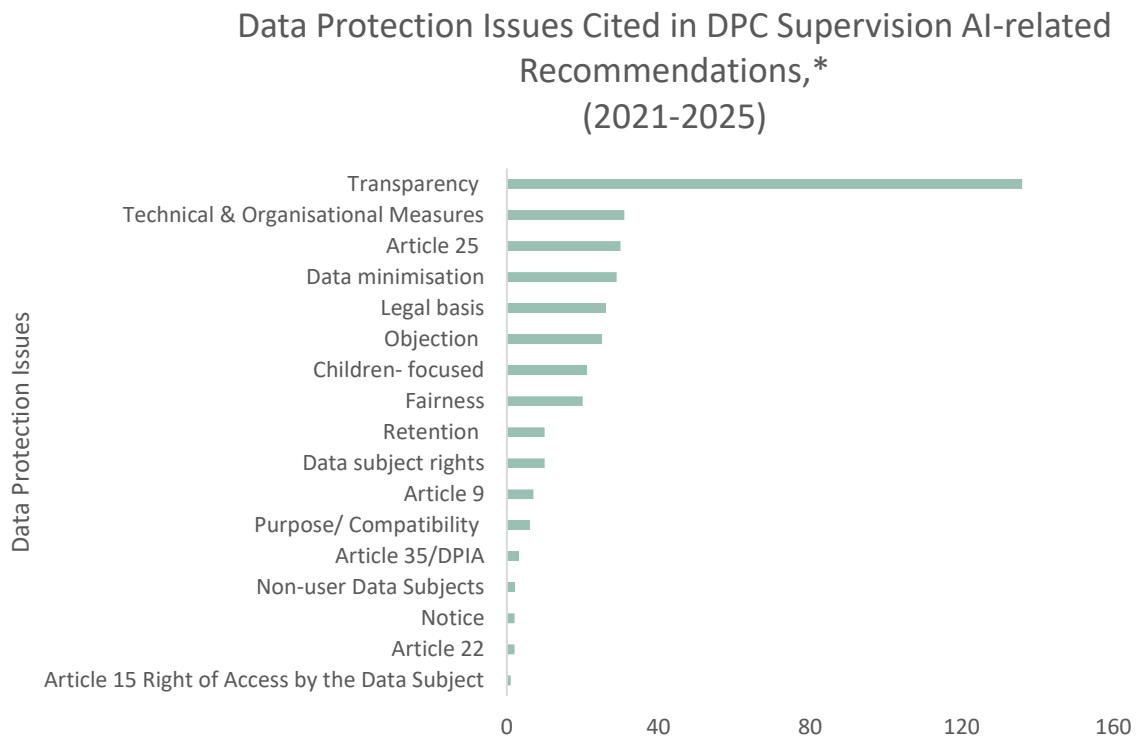

*Note: Data is sourced from internal TMSU Statistics *Please note that a single recommendation may cover multiple data protection issues, (for example transparency and objection).*

Good Engagement

The most effective engagements come with open, honest, considered and detailed information, at an early stage.

The DPC has had engagements with companies at all stages of AI development and deployment. The Supervision Function has made itself available at the planning stage of AI projects, providing feedback on planned projects that facilitate engineers building with data protection by design and default in mind, as well as considering post-launch deployments, making recommendations, and raising concerns where necessary.

However, the majority of the engagements with companies have happened shortly before the anticipated launch. In general, the DPC considers that a minimum 8-week lead time is necessary to enable the DPC to engage substantively. Artificial

intelligence, AI systems, and the resulting documentation, are complex, with many technical and data protection issues to consider. It is rare for the Supervision Function to have no questions when faced with new technology or products. Early engagement with the Supervision Function provides time for the DPC to assess the documentation, and for controllers to both fully address any outstanding questions, and implement any recommendations made prior to launch.

In some instances, the DPC has requested a voluntary pause by the controller before processing begins. A request to pause may happen for a number of reasons, including that the DPC, following consultation with its peer EU/ EEA regulators, has concerns around compliance, usually where significant issues have been identified, or simply that there has not been sufficient time for the engagement to come to its natural conclusion. However, out of all AI related engagements from 2021 through 2025, 1 launch in 2022, 1 launch in 2023, 4 launches in 2024, and 3 launches in 2025, were delayed or cancelled due to interventions by the DPC.

95% of AI engagements with the DPC were able to launch without delay due to DPC engagement. Whilst it is never the intention of Supervision to unduly delay the launch of a product, engagements can take time, and controllers should therefore engage early and often.

The Supervision Function considers that contacting the DPC as early as possible is the best policy to prevent delays and provide sufficient time for both the DPC and controllers to have a successful engagement prior to launch.

Case Study

One controller, in respect of a particularly complex matter, wanted to ensure that it considered all data protection aspects before launching a new AI model. The controller contacted the DPC in 2021 about creating an AI model from user personal data. The controller provided a DPIA in relation to its planned training of the AI model and the DPC provided feedback at the design stage. The next iteration of the DPIA had implemented the DPC's feedback and had also addressed the fine tuning of the AI model, resulting in significantly less personal data being processed. The DPC provided further recommendations on the design of the project in relation to transparency, data subject rights, and retention. Sometime after the initial contact with the DPC on the issue, the controller provided a DPIA on the proposed deployment of the AI model. The DPC made additional recommendations on data minimisation, proportionality, necessity, and the right to object. The controller further revised the DPIA based on DPC feedback. The controller undertook a pilot of the AI model and launched the feature after acting on the feedback from the DPC. Through the consideration and implementation of DPC feedback and recommendations at all stages, the controller was better able to demonstrate data minimisation, fairness, and compliance with the GDPR. The DPC considers that this was an effective engagement, resulting in iterative, data protection-focused innovation that mitigated detrimental impacts on individuals.

Key Takeaways

Early engagement with the DPC can ensure that foundational design choices that meet the GDPR's requirements, resulting in a more compliant final product. Supervision can engage at all stages of a development cycle, providing guidance to controllers that may prove beneficial in developing AI that is centred on the principles of data protection by design and default.

What happens if Supervision finds evidence of noncompliance?

Supervision's early input into proposed processing activities can be very effective in terms of being the 'canary in the coal mine' of data protection. Supervision is able to identify issues that may not arise through complaints or other investigations, which only happen after the fact. This is particularly useful, given the 'invisible' nature of AI processing, which may mean individuals are not aware of the processing, and are inhibited in the exercise of their rights.

Through effective engagement, the Supervision Function has been able to identify issues, and drive controllers to bring processing into compliance.

Case Study

Through Supervision, three instances were identified where controllers intended to train AI on user personal data, without adequately informing their users. The use-cases of the AI were very different, but in all three cases, the DPC recommended that the controller include both a notice and a notification period. The DPC based its recommendations on the fact that it would be unfair and not transparent were the processing to occur without notification. In all instances, the controller took the recommendations on board. This had the practical effect of informing individuals of the processing, as well as giving them a notice period wherein the individual could exercise their data subject rights, including the right to object.

Key Takeaways

Articles 5(1)(a), 12, and (as appropriate) Articles 13 and 14 of the GDPR place an obligation on controllers to provide a clear notification of the intended use of personal data for AI training, as well as a reasonable notice period wherein individuals can educate themselves and exercise their data subject rights. Such a notice period may also be necessary when a user's personal data becomes subject to AI training at a later date, for instance, upon the user reaching the age of 18.

Where necessary, the DPC can exercise its powers under section 134 of the Data Protection Act 2018. Section 134 allows the Commission, where it considers that there is an urgent need to act in order to protect the rights and freedoms of data subjects, to request an order from the Irish High Court suspending, restricting or prohibiting the processing of personal data concerned, or the transfer of such data to a third country or international organisation.

Case Study

In July 2024, the DPC became aware that key mitigation measures were not in place with regards to Twitter[10] training its Grok AI model using user personal data. The DPC immediately requested that Twitter cease processing, but Twitter refused. Having determined that there existed a real risk to the rights and freedoms of data subjects, and that there was an urgent need for the DPC as LSA to act, the DPC made an application to the High Court under Section 134 of the Data Protection Act 2018 on 08 August 2024. Subsequent to a hearing of the substantive proceedings, Twitter undertook to cease processing user data that had been collected during the period of time when the mitigations were not in place, and to delete the datasets that included such data.

Key Takeaways

Where the DPC determines that the processing or intended processing of personal data presents a real risk or harm to the rights and freedoms of data subjects, the DPC can and will act to prevent harm.

[10] Twitter's European operations, formerly known as Twitter International Unlimited Company (TIUC), were rebranded as X Internet Unlimited Company in April 2025.

The DPC can commence inquiries where noncompliance is identified concerning ongoing processing of personal data. This can result in a formal inquiry under Article 58 of the GDPR. The DPC utilised this power in relation to a number of AI engagements. At the time of writing, the following are the subject of an ongoing inquiry:

- Google, with regards to obligations to undertake an assessment under Article 35 prior to processing the personal data of EU/EEA data subjects associated with its foundational AI model PALM 2[11]
- X Internet Unlimited Company, with regards to the training generative artificial models, in particular the Grok Large Language Model;[12]
- X Internet Unlimited Company, with regards to processing of personal data of EU/EEA data subjects, including children, using generative artificial intelligence functionality associated with the Grok LLM within the X Platform.[13]

By using a variety of regulatory tools, the DPC balances the harder and softer regulatory approaches to enforce data protection compliance in an appropriate, proportionate manner which aligns with our regulatory objectives. This supports the overall mission of the DPC to regulate fairly and proportionally.

[11] Data Protection Commission 'Data Protection Commission launches inquiry into Google AI model' 12th September 2024 <https://dataprotection.ie/en/news-media/press-releases/data-protection-commission-launches-inquiry-google-ai-model>
[12] Data Protection Commission 'Data Protection Commission Announces commencement of inquiry into X Internet Unlimited Company (XIUC)' 11th April 2025 <https://dataprotection.ie/en/news-media/latest-news/data-protection-commission-announces-commencement-inquiry-x-internet-unlimited-company-xiuc>
[13] Data Protection Commission 'Data Protection Commission opens investigation into X (XIUC)' 17th February 2026 https://dataprotection.ie/en/news-media/press-releases/data-protection-commission-opens-investigation-x-xiuc>

Advocating Guidance and Harmonised EU Approach

The benefits of Supervision go beyond influencing better compliance and responsible innovation by controllers. In line with the DPC's strategic goal of bringing clarity to stakeholders, through the proactive Supervision approach, the Supervision Function was able to identify legal questions and gaps in the alignment of EU data protection regulators. The DPC has been able to leverage the insights gained through Supervision to advocate for consensus and consistency in interpretation across the EU, through guidance from the EDPB in relation to AI.

The DPC Supervision Function regularly engages with the EDPB and peer regulators on AI matters, sharing information and insights about controllers and their compliance with the EDPB through cooperation mechanisms and in meetings. The Supervision Unit often holds 'flash meetings' early in engagements where the DPC considers the matter to be high profile or of specific importance. This sharing of information helps inform our peer regulators and build consensus at EDPB level on matters of interpretation and importance, which the DPC then uses when communicating recommendations to controllers.

The DPC has focused on establishing consensus on matters central to the training and operation of AI models in Europe, as well as following up with controllers deploying AI to ensure that they know the requirements for such processing under the GDPR. This approach ensures safeguards and privacy enhancing measures are present in AI products ahead of an EU-wide roll out while providing guidance on how innovation and data protection / privacy can work hand in hand.

Example

In the summer of 2023, it became clear to the DPC that there were a number of differing opinions, interpretations, and understandings when it came to processing personal data for the training of AI. The DPC began sharing information and insights from engagements with controllers with other EDPB supervisory authorities, as well as advocating for discussions on points of

difference through EDPB subgroups and the ChatGPT Taskforce (now Generative AI Enforcement Taskforce). The insights gained from engaging with controllers who were training and deploying generative AI helped inform discussion at EDPB level, and this ultimately became the driving force behind the request by the DPC for an Article 64(2) Opinion on the legal basis for AI training. The DPC engaged the EDPB to provide a formal GDPR Opinion, to achieve Europe-wide regulatory harmonisation and clarity on a number of key AI model training and deployment related questions.

The EDPB Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models ('AI Opinion'), issued in December 2024, provided general criteria that the DPC and all other EU Supervisory Authorities should take into account when assessing compliance of the processing of personal data for the development and the deployment of AI models. It clarified when an AI model could be considered anonymous and that Article 6(1)(f) was a possible legal basis for training AI. This brought consensus and regulatory clarity from the EDPB. The DPC has since used this opinion to provide consistent guidance to all controllers who engage with Supervision, delivering certainty to controllers that the positions communicated by the DPC have been agreed upon across Europe.

Key Takeaway

The DPC continues to advocate for and seek regulatory consistency across the EU in line with the DPC's strategic goals. The EU-wide harmonised approach achieved through the AI Opinion allows the DPC to provide clear, consistent advice to controllers deploying AI models in Ireland and across Europe. This drives regulatory consistency, a pillar of the DPC's Strategic Strategy, while enabling and encouraging data protection compliant innovation.

In some cases, lack of legislative, legal or other clarity in relation to specific aspects of the application of data protection requirements to AI, also mean that Supervision will not make a recommendation on a particular issue.

Example

Bias is a highly relevant and important issue when it comes to AI and algorithms. However, the ability to identify if a dataset is representative and avoids bias may require the processing of highly sensitive or special category data. This would be prohibited without the data subject's consent as per Article 9(1) of the GDPR unless an exception applies. Notwithstanding the general requirement that all personal data processing under the AI Act should be done in compliance with the GDPR, Article 10(5) of the AI Act allows special category data to be exceptionally processed for the purpose of bias detection and correction. Until the Digital Omnibus on AI was published in late 2025, it was unclear if this was intended to be an exception under 9(2)(g) of the GDPR or a new standalone exception to allow processing for bias detection, and so the DPC refrained from taking a position on the AI Act's 10(5) exception, merely alerting controllers to the possible bias considerations and reserving the right for the DPC to take action or reassess matters, either through Supervision or an inquiry under Article 58 of the GDPR, once consensus or further guidance is available. The DPC's position on processing for bias detection may change in the future, now that the Digital Omnibus on AI has provided further clarity in the subject matter.

Key Takeaway

Where matters are unclear, Supervision can identify potential issues for controllers, the DPC, and the EDPB, that need further clarification. The issues identified can assist in driving priorities for guidance or areas for consensus.

While the DPC may refrain from taking a position, it is not required to do so. The DPC has taken positions on aspects of AI even in the absence of consensus. The question of special category data in AI training datasets was one where the DPC took a clear position in providing guidance and direction to controllers.

Example

The existence of special category data in a dataset comes with risks to data subjects, but it can also provide benefits. It is widely known that a training dataset that is not representative is more likely to produce biased results. Accordingly, the inclusion of some special category data may have beneficial results for individuals and may be required for the creation of non-biased AI. Acknowledging that challenges in law and practice existed, the DPC advocated for stronger safeguards and the usage of de-identification, pseudonymisation, or anonymisation measures on datasets used for AI training. When the DPC made such recommendations, no regulatory guidance required such measures. While the DPC is aware of the possibility that de-identification processes may not fully mitigate possible risks, from a risk-based approach the DPC considers that such efforts are more likely to produce AI that minimise harms and risks to individuals, while still retaining sufficient diversity and information to avoid bias. That approach is now being followed in the draft Digital Omnibus on AI, which has been welcomed by the EDPB.

Key Takeaway

The DPC can provide guidance even in the absence of regulatory consensus. Controllers looking to develop AI should consider implementing measures to de-identify, pseudonymise or anonymise datasets used for AI training.

Proactive and Ongoing Monitoring

Supervision involvement does not end once a product launches. Data protection compliance does not stop once an AI model has been created. AI models can be trained compliantly, yet their deployment can still be noncompliant. The DPC is conscious that compliance can become an issue in the deployment phase as well, and Supervision engages in ongoing monitoring of higher risk systems post launch.

On the emergence of mainstream generative AI models, the DPC took a proactive approach to early engagement with controllers. In 2023, the Supervision Function developed an AI Questionnaire and proactively sought responses from controllers as to their future plans and developments regarding the creation of generative AI. This proactive approach gave the DPC early insights into the product roadmap of likely future developments. It also put controllers on notice, encouraging them to engage with the Supervision Function on generative AI projects at an early stage.

Starting in 2023, the DPC has also requested reports from controllers creating and deploying LLMs. These reports typically require the controller to gather live data to evaluate:

- the performance of the LLMs;
- the performance of safeguards including technical and organisational measures;
- reactions of actual individuals or the wider public; and
- other real-world issues that may impact on the privacy of the AI or AI system.

Some reports have also included a further assessment when an AI model's training was refreshed. All reports have required a controller to re-evaluate the original risk assessments, DPIA, and LIA (where applicable) against the real-world data collected, and identify shortcomings, and where necessary put in place new mitigating measures to ensure compliance continues after deployment. Real-world data provides clear indications where measures and safeguards fail or are inadequate, or verification that the processes are working as intended. A number of controllers have indicated that these reports are helpful to them, but they have also served to provide evidence of data protection compliance.

These monitoring reports have served both as sources of information regarding ongoing processing, while providing the DPC with opportunities to re-evaluate the processing in light of facts, not hypotheticals. Where reports have identified additional issues, if not already remediated by the controller, the DPC Supervision Function has made additional recommendations and corrections.

Case Study

Most controllers have identified issues that could be improved through the report review process. In one report, the controller was able to identify additional datapoints to exclude from AI training and expand their filtering, targeted improvements to transparency information, and additional transparency for younger users. From user reports during the period covered by the report, the controller was able to identify bugs and correct errors with outputs of the LLM. Other features were fine-tuned further, additional work done to identify and block hallucinations, and to implement safety guardrails to prevent incorrect responses. Overall, the report provided the controller the opportunity to take account of real-world data and feedback to both improve their product as well as their compliance with the GDPR.

Key Takeaways

Monitoring live launches and treating the data protection assessments as living documents can demonstrate responsible innovation and assist to develop safer, more compliant products.

However, the DPC has not relied solely on controllers' assessments post-implementation when monitoring. DPC Supervision Function proactively monitors the live LLMs, chatbots, and AI services, and re-engags with controllers where the DPC has identified issues with GDPR compliance.

Case Study

In January 2025, through ongoing monitoring, the Supervision Function identified that the opt out for processing personal data for AI training for a particular chatbot was not available to non-subscribers using mobile devices, although access to the chatbot was available to all users. As such, the DPC was concerned that many mobile users of the chatbot would not be presented with the option to opt out of the processing of personal data for AI training. The Supervision Function immediately contacted the controller to identify and fix the issue. Upon investigation, it was found that the issue was the result of a bug, and that approximately 7.7 million mobile users in the EEA could potentially have been impacted by the bug affecting the availability to opt out. A technical fix for the solution was enacted by February 2025, allowing all users to access the opt out. The controller committed to not training AI on any data collected during the period when the opt out was unavailable, and the DPC continues to monitor the situation.

Key Takeaways

Compliance does not end with a feature or product's launch. It is the responsibility of controllers to continue to monitor the implementation of their products as well as how the technical and organisations measures are functioning to ensure compliance. The DPC will continue to monitor post-launch and take action where infringements or issues arise.

Good Documentation and Common Pitfalls

The DPC encourages proactive engagement by controllers. The DPC expects those companies who do engage to provide in depth and, where possible, complete information. GDPR compliance documentation plays a vital role in allowing controllers and the DPC to assess and manage the risks to individuals in relation to AI and other processing activities. Specific implementation details matter throughout the lifecycle of an AI product or service. Detailed compliance documentation assists controllers with demonstrating compliance with the GDPR as per the accountability principle found in Article 5(2). The documentation should demonstrate the choices and decisions made by the controller.

Engagements with the DPC where controllers have engaged with the DPC early, meaning 8 weeks or earlier, with good documentation have been predominantly able to launch without delay due to DPC engagement. AI, AI systems, and the documentation to evaluate data protection risks on training AI, can be complex and take time both to prepare, assess, and be revised for a compliant launch.

Case Study

One controller engaged with the DPC Supervision Function in relation to AI training on user interaction data, i.e. prompts or responses from a chatbot. The controller provided a DPIA and LIA to the DPC. When assessing the documentation, there was a high level of detail, and it was organised such that likely questions from the DPC were immediately answered next to relevant information. For example, statements that raised questions on retention had retention measures footnoted or in the next paragraph, making it easy for the DPC to understand the context and measures involved. The documentation contained high-quality writing, with clear legal analysis based on existing law and guidance, and dense, but specific details of the implementation. There was not a lot of repetition, or an over-use of referencing other sections. It clearly documented design choices by

the controller in relation to data protection, as well as providing the Data Protection Officer's (DPO's) contributions. The quality of documentation meant that the DPC had a very clear understanding of the proposed processing, mitigations in place, and analysis by the controller. Engagement with the DPC did not result in a delay for the controller.

Key Takeaways

Good documentation facilitates speedy regulatory engagements. Organisations should anticipate and proactively address questions likely to arise from the Regulator. Organisations should ensure compliance documentation provides quality information that is in-depth featuring clear, articulated analysis and documented design decisions. Where designated, the DPO must be involved in creating and documenting a DPIA.

Controllers who have engaged with the DPC making high level statements concerning their state of compliance without evidence or reasoning to back up those assertions, have struggled to complete engagement with the DPC in time for their planned launch date. In such cases, the DPC has requested that the launch be delayed or paused to ensure that the DPC can thoroughly assess compliance prior to any AI training beginning. delayed or paused to ensure that the DPC can thoroughly assess compliance prior to any AI training beginning.

Case Study

Two controllers informed the DPC of their intention to begin training AI on public personal data. The sources of public personal data differed between controllers. However after assessing the documentation provided, in both cases the DPC still had many unanswered questions regarding the processing and each controller's measures to mitigate risks to individuals, as well as concerns relating to the processing. Such concerns included that the controllers' documentation was lacking or insufficiently clear to the DPC and our peer regulators to the extent that the DPC was not satisfied that the controllers had effective or sufficient measures in place to minimise the data protection risks. As the Supervision Function felt it necessary to have more time to engage with each controller to understand the proposed processing and thereafter assess compliance, the DPC asked both controllers to pause the launch of their proposed AI training. In both cases, the controllers provided updated documentation of better quality, which clarified and responded to the DPC's concerns. As a result, both controllers proceeded with their planned processing without the requirement to pause.

Key Takeaways

Where poor documentation is provided, engagements with the DPC can become protracted, and can delay launches. Poor documentation may also raise doubts as to the effectiveness and completeness of the organisation's compliance readiness programme. Where controllers provide clear, in-depth documentation, engagement with the DPC is more straightforward and faster.

The DPC considers that any training of a new AI model with personal data will necessarily involve a new technology[14] and likely large quantities of data.[15] Therefore, **the DPC considers that the training of a new AI model will likely require a Data Protection Impact Assessment.[16]** While AI as a technology has existed for many years, the scale and impact of current AI developments significantly differ today from what was on the market a few years ago. Any newly created AI has the potential to have risks, biases, and unknown impacts on individuals when deployed. Therefore, it is of utmost importance that the risks to individuals are taken into consideration prior to processing beginning.

[14] While AI as a category may not be 'new' the creation of a new AI or AI model may still be considered a new technology that may have unknown risks that will trigger the need for a DPIA.
[15] Two criteria for when a DPIA is mandatory as per WP29 Guidelines on Data Protection Impact Assessment (DPIA) and determining whether processing is "likely to result in a high risk" for the purposes of Regulation 2016/679, WP248 rev.01, adopted 04 October 2017
[16] Data Protection Commission 'Data Protection Commission launches inquiry into Google AI model' (12 September 2024) <https://dataprotection.ie/en/news-media/press-releases/data-protection-commission-launches-inquiry-google-ai-model#PaLM>

DPC Data Protection Observations on AI Technologies

Transparency

The DPC is conscious that transparency relating to the creation and deployment of AI is of utmost importance. This is particularly important as many individuals remain unfamiliar with the technology or the implications of their personal data being processed for the creation of AI models. The nature of AI often renders the data processing involved invisible, which makes transparency regarding the training, creation, and deployment of AI a crucial component. Without transparency, individuals may not to know how or that they can exercise their data protection rights.

Transparency issues have been present in 72% of all recommendations made by the DPC Supervision Function. The DPC experienced reticence on the part of some controllers to inform users about the processing of personal data by AI. In some cases, this was for reasons the controllers believed were legitimate. For example, AI that is intended to detect fraud or illegal activities could potentially be evaded by bad actors if they learn of the existence of the AI or if they can figure out how it works. The DPC has emphasised to controllers that no amount of post-processing transparency will relieve a controller of its obligations to provide transparency prior to processing under Articles 5, 12, 13, and 14 of the GDPR.

The novel nature of how AI is now being deployed may require additional transparency by controllers, as individuals may be encountering the technology for the first time. Users less familiar with a technology are much less likely to be aware of how data is processed by the technology or the risks that such processing may entail. The DPC considers that the onus is on controllers who deploy AI technologies to educate individuals on the data protection implications of their products and services.

Case Study

Two controllers intended to launch AI Agents that rely on capturing screenshots of a virtual browser. Given the novelty of this technology, the Supervision Function considered that the controllers should be transparent with users about the nature of taking such screenshots, particularly regarding the risk that personal data may be captured. The DPC recommended that the controllers enhance transparency, tailor warnings specifically to the use of agents and provide greater clarity on the risks associated with this type of data processing.

Key Takeaways

Novel technology may require additional transparency to not only explain how to use the technology but also educate the individual on the data protection implications and risks.

In relation to LLMs, the DPC considers that transparency should always be provided both in the creation and deployment phase. Recital 39 of the GDPR states "*[n]atural persons should be made aware of risks…"* in relation to processing of personal data. The DPC has found this to be an area that many controllers do not adequately consider when designing transparency notices.

Case Study

In 2023, the DPC was conscious that many users were unaware of how generative AI worked, how generative AI processed personal data, and the risks to processing such data. In assessing one controller's documentation, the DPC Supervision Function was concerned that, despite the planned transparency notice, users who were confronted with this new technology would not fully understand the risks involved with the AI chatbot processing their personal data. The DPC recommended that the controller create transparency notices explaining the privacy risks, to assist in educating individuals of their rights and how a generative AI chatbot's processing intersects with data protection concerns. Since then, the DPC has made similar recommendations to 5 other controllers developing LLMs.

Key Takeaways

The DPC considers that controllers have a duty to not only explain that processing is occurring but also explain the data protection risks that may exist in relation to the processing of personal data by AI and LLMs, particularly where new technology is concerned.

Specifically for AI training, the DPC is conscious that AI training presents particular challenges when removing personal data from a model at a later time. Therefore, transparency prior to training is of the utmost importance.

The DPC considers that where there is a direct relationship with the individual, controllers have an obligation to ensure they provide individuals with transparency information prior to their data being processed. Where users do not log in to a platform or service regularly, this may mean that controllers should consider issuing transparency information external to their platform, such as through the use of email. Whilst the DPC notes the use of blog posts or general online posts, it is unclear how

many average users of a service actually see or would regularly access such transparency information. Controllers should use as many appropriate communication tools as necessary to ensure that individuals have optimal transparency, notice, and information.

Transparency is also a key element of the balancing test and the creation of user expectations needed for reliance on the legitimate interests legal basis. Additionally, lack of transparency can contribute to a user's consent not being fully informed. Failure to provide sufficient transparency can therefore undermine the validity of a controller's legal basis.

Transparency needs to be directed to the individual and be data subject centric. The DPC does not consider it appropriate to rely on transparency notices that are difficult to find, not directed to users, or that are intended for features that individuals do not actually engage with. Controllers cannot use indirect methods as the primary vehicle to provide information to those individuals.

Case Study

A controller's transparency information was addressed directly to verified users, in a section of their website that was unlikely to be accessed by non-verified users, despite the information being pertinent to all users. The DPC recommended that the controller develop and publish easily accessible transparency information that was designed for and addressed to all individuals that may be impacted by the processing.

Key Takeaways

Transparency information should be addressed to and be accessible for all individuals impacted by AI processing of personal data to ensure compliance with the GDPR.

Another area that is consistently lacking is transparency to individuals who do not use a service, but whose data may be processed. These 'non-users' remain data subjects under the GDPR and the DPC has made several recommendations to various controllers that additional transparency measures be implemented for these data subjects or that existing transparency information is moved to an accessible place for non-users.

Legal Basis

All processing of personal data must have a legal basis under Article 6 of the GDPR. It is the responsibility of the controller to determine which legal basis is most appropriate. However, where the DPC identifies concerns regarding a controller's chosen legal basis, the DPC will intervene and issue recommendations where appropriate. While the legal basis under Article 6 featured in only 7% of the data protection issues raised in recommendations, these recommendations were often expansive and of critical importance to a controller's compliance.

The EDPB AI Opinion in December 2024 set out the legitimate interests of a controller as a possible legal basis for AI training. The DPC has applied the guidance in the Opinion when assessing the development of LLMs and other AI systems.

The DPC Supervision Function undertakes detailed analysis of all claims to rely on Article 6(1)(f) in relation to AI. This involves analysis of LIAs by controllers to determine if a controller has likely met the three-step test[17] or if the rights and freedoms of data subjects outweigh the controller's legitimate interest.

When assessing LIAs, the DPC has taken into consideration all aspects of the processing including, but not limited to:

- Whether there is sufficient and appropriate transparency, particularly how such transparency addresses the reasonable expectations of a data subject;
- Ability and ease of an individual to exercise their rights, and any friction created by the implementation of data subject rights request forms;

[17] For more information, please see Data Protection Commission 'Guidance Note: Legal Bases for Processing Personal Data' (December 2019) <https://www.dataprotection.ie/sites/default/files/uploads/2020-04/Guidance%20on%20Legal%20Bases.pdf>

- The stated legitimate interest(s) by the controller;
- The necessity of the processing for the purpose(s) identified;
- The compatibility of the processing with the purposes for which the data was collected (if applicable);
- Whether there are any less intrusive ways to meet the goals of the processing and why the controller has not chosen a less intrusive option, if available;
- Mitigation measures put in place to limit the impact of the processing on data subjects; and
- Consideration of the balancing test conducted by the controller such that the controller's legitimate interest is not outweighed by the rights and freedoms of the individuals concerned.

Where the DPC has identified issues that would undermine a controller's legitimate interest assessment, these concerns have been communicated to the controller. This allows controllers the opportunity to put additional measures in place or to adjust their planned processing of personal data to satisfy the legitimate interests three-step test.

Case Study

In May 2024, Meta publicly announced their intention to begin training AI on user personal data. Immediately, difficulties with users attempting to exercise their right to object and confusion around the transparency notices became apparent to the DPC Supervision Function through our ongoing monitoring. Users struggled to understand what the processing would entail and how to object to it. The Supervision Function engaged with Meta, as well as our colleagues in the EDPB, to establish the full impact of the transparency and objection form issues. After assessing the 2024 LIA provided by Meta, the DPC still had significant questions and concerns regarding the processing and Meta's ability to rely on Article 6(1)(f) of the GDPR as a legal basis. The DPC was of the opinion that Meta's implementation of transparency notices combined with multiple objection form issues (including difficulties in accessing and completing an objection) fundamentally undermined Meta's reliance of the legitimate interests lawful basis for processing personal data. Prior to the processing commencing for training the LLM, the DPC requested that Meta pause the proposed AI training, which Meta complied with. The DPC also requested updated documentation.[18]

In February 2025, Meta re-engaged with the DPC and provided a revised LIA and DPIA. The rollout of transparency notices and the objection form did not encounter the same issues that occurred in 2024. Having obtained the views of the other EU/EEA supervisory authorities, the DPC made a number of additional recommendations which Meta implemented. Through this engagement, the DPC concluded that risks to individuals would be significantly reduced.

The DPC's conclusion was based on the implementation of the Supervision Function's recommendations, the additional measures Meta had committed to (such as filtering of data and de-identification prior to training), and the additional transparency measures introduced over the previous year which gave individuals an extended period of time in which to exercise their right to object, in advance of the processing of personal data. While not determining that Meta had complied with the GPDR, these additional measures had a significant impact on the balancing assessment required under Article 6(1)(f), to the extent that the DPC concluded that there were no grounds justifying further injunctive regulatory action at that time.[19]

Key Takeaways

It is essential that controllers carry out thorough due diligence ensuring that all aspects of the design and notifications associated with a product or service have been assessed, risk/harm mitigation measures are implemented, and all aspects of the proposed processing meet the requirements of the GDPR.

Related to the legal basis for the processing of personal data is the question of reasonable expectations of a data subject in the context of the legitimate interests test and the principle of purpose limitation found in Article 5(1)(b) of the GDPR. While the standard rule is that data should be collected for a specified, explicit, and

[18] EDPB Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models, adopted 18 December 2024
[19] Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications) [2002] OJ L 201/37

legitimate purpose and should not be further processed for other purposes, Article 5(1)(b) of the GDPR provides for the possibility for further processing provided it is not done in a manner incompatible with the original purpose(s). Article 6(4) of the GDPR provides for a compatibility assessment to be conducted to determine whether a purpose is compatible. Recital 47 of the GDPR emphasises that the existence of a legitimate interest needs careful assessment including whether a data subject can reasonably expect at the time and in the context of the collection of the personal data that processing for that purpose may take place.

Public discourse concerning AI, in particular regarding generative AI, would suggest that many individuals have only in recent years become aware of AI in any substantive sense. That being the case, it is likely that few would expect their personal data to be processed for or by AI, in particular beyond direct engagement with an AI model or system i.e. for further model training. Therefore, the DPC has assessed Compatibility Assessments carefully where a controller seeks to use personal data for AI training that may have been collected in the past. In order to rely on Article 6(1)(f) of the GDPR as a lawful basis for processing, a controller must, as part of the balancing test, take account of the reasonable expectations of data subjects. Such expectations must be understood in light of the circumstances at the time of collection and the context in which the data is collected, and any analysis must reflect the realities and reasonableness of those expectations.

Case Study

A controller intended to use all public user posts, comments and articles collected from the inception of the platform, for training their generative AI models. One of the risks identified in the DPIA by the controller was that data subjects might not have expected their data to be used for AI training when they first provided it. The controller provided a compatibility assessment, and argued that the usage of personal data did not constitute a new purpose than for what it was originally collected, namely to '*enable and improve*' its services. The Supervision Function requested the historical privacy policies

covering the proposed processing period as transparency from the date of collection is relevant to determining data subjects' reasonable expectations. On review of the privacy policies, it appeared that there was no consistent equivalent terminology until at least 2014 and later. There was an arguable case that the purposes from 2014 onwards were compatible with the proposed AI training, however prior to 2014, purposes of processing were confined largely to the improvement of cookies or failed to have any reference to improvement of service. The DPC recommended that the controller should not process any personal data collected prior to, at the very earliest, 2014, for the purposes of training generative AI, on the basis that individuals were not afforded sufficient predictability regarding the future use of their data to form or sustain reasonable expectations. The controller took this recommendation on board and reduced the relevant period of processed data to 2014 to present.

Key Takeaways

As past collections of data are unlikely to have been conducted with users expecting their data would be used for AI training, the DPC considers that using historical data will heighten the need for enhanced transparency. Where historical data is amongst a controller's intended dataset(s), any analysis cannot disregard historical conditions or transparency provided by the controller, as they form the basis upon which expectations were shaped. The DPC will carefully assess compatibility assessments based not only on the arguments within, but evidence provided, to determine the reasonable expectations of data subjects.

While not within the original scope of this report, during the first half of 2026 the DPC has been presented with the opportunity to consider and assess the use of private messages for AI training, such as private communications between users via a messaging service. Given that there is an inherent expectation of privacy to private communications, the **DPC considers it is unlikely that a controller will be able to demonstrate that data subjects expect the usage of their private messages or other private data for AI training.**

As highlighted in the EDPB Opinion on AI, the mere fulfilment of transparency requirements under the GDPR may not be sufficient to consider that data subjects can reasonably expect a certain processing. The DPC considers it insufficient to demonstrate that individuals are merely aware of the possibility of their private messages being used for AI training. Controllers will have to demonstrate conclusively that the inherent expectation of privacy does not exist or that the controller has sufficiently informed individuals, including third parties, to the degree that they will expect their private communications will be used for AI training in order to rely on legitimate interests as a legal basis.

Additionally, the necessity prong of any legitimate interests assessment will need to clearly explain why private message data is necessary for the AI training, particularly where alternative sources of data are available (such as synthetic data or publicly available sources) that may be suitable for training the particular AI model considering its purpose(s). Any necessity analysis must clearly address why other data sources are inappropriate and why the private messages are the only suitable data that will allow the controller's purposes to be fulfilled. In any event, given the nature of the personal data concerned, the DPC considers that in order to rely on legitimate interests for processing private messages for AI training, a controller will have to meet a very considerably higher threshold when justifying the use of private messages than if the controller were to process publicly available information. It is the DPC's assessment that this threshold will be very difficult to achieve in the context of the use of private messaging data for AI training purposes.

The DPC notes that GDPR may not be the only relevant legislation where private messages are processed for AI training. The ePrivacy Directive requires confidentiality of communications to be guaranteed, and that measures should be

taken in order to protect the confidentiality of private communications. The DPC considers that making private messages available for AI training would significantly erode the trust and confidence of users that their privacy will not be at risk, which would undermine the aims of the ePrivacy Directive. Controllers will have to demonstrate that any usage of private messages for AI training does not breach the ePrivacy Directive, as well as the GDPR.

Right to Object

Where a controller relies on Article 6(1)(f) of the GDPR as a legal basis, they are required by Article 21 of the GDPR to provide data subjects with the possibility of objecting to the processing. Article 21 and the right to object featured across 7% of data protection issues raised in DPC recommendations. However, often the objection form and the relevant recommendations covered multiple AI engagements by the same controller as controllers regularly provided a general way for individuals to object, rather than bespoke or separate objection solutions for each AI product or service.

With the current state of the art, once an AI model is trained it can be very difficult to remove any training data that might be part of the model. This can mean that any individual whose personal data is used to train the model may suffer a loss of control of their data and may not be able to fully exercise their data subject rights under Article 15 through 22 of the GDPR. Given these challenges, there may be no effective way for users to exercise their Article 21 rights if not provided the opportunity to opt out or object prior to the processing beginning.

The DPC considers that the design of an objection form, its clarity, accessibility, ease of use, and the ability for an individual to understand if their objection has been honoured is of paramount importance in ensuring the rights of data subjects and also demonstrating that the fundamental rights of data subjects do not outweigh the legitimate interests of a controller.

Where a controller is processing data collected from the individual for AI training, a controller should ensure individuals have sufficient notice prior to the processing beginning, with clear, easy-to-use ways to exercise their Article 21 right to object. Particularly where the controller has a relationship with the individual, the DPC

considers that controllers have an obligation to ensure such individuals have an easy to use and accessible process by which to object.

The DPC considers that any chilling effect on the exercise of data subjects rights is unacceptable and any friction to access an objection form or method should be minimised to the greatest extent possible. Such a chilling effect can be related to the transparency provided to individuals. For example, it is the DPC's view that the phrase '*if your objection is honoured…*' is misleading where the controller's intention is to provide an unconditional opt out, and would likely contribute to a chilling effect on the exercise of their rights.

To ensure that individuals can object, the DPC has made a number of recommendations to controllers that provide clear guidance on what a controller should do. This includes, but is not limited to:

- Objection forms or ways to object should be easy to use and easily accessible by data subjects;
- Being clear on the scope, extent, and duration of an objection;
- Ensuring that all users, regardless of what type of device is used, have received the transparency notices regarding this processing and the ability to opt out prior to their personal data being used for training;
- Where applicable, making clear that any objection to AI training is treated as an unconditional and automatically honoured opt-out;
- Transparency on the notice period provided should allow users the opportunity to educate themselves on the implications of the processing and allow them an opportunity to exercise their right to object prior to their data being processed;
- The initial transparency design should ensure that the data subject is not likely to overlook or have difficulties reading the information on objections contained within; and
- Ensuring users who have made objections are aware of the status of their request.

Additionally, the DPC considers it insufficient to provide an objection form at the time the processing that AI training begins. Individuals must be provided with sufficient

notice of the proposed processing with access to a way to object to the processing before any AI training begins. This must be accompanied with sufficient transparency to understand the processing. To do otherwise would likely be in conflict with the fairness and transparency principles found in Article 5(1)(a) of the GDPR.

The GDPR does not distinguish the right to object based on whether an individual provided their personal data directly to the controller, or whether the controller obtained the data by other means, i.e. 'third party data'. Therefore, controllers must ensure that individuals have the ability to object, even if the controller did not receive the personal data directly from the individual.

> **Key takeaways**: Objection forms or opt outs should be easily accessible, easy to use, with sufficient transparency and notice for data subjects to understand the processing and its implications.

Data Minimisation

Article 5(1)(c) of the GDPR requires that only personal data that is adequate, relevant and limited to what is necessary for the purposes should be processed. This principle of data minimisation has obvious challenges when the purpose is the creation of a Large Language Model, which by its nature requires a large amount of data. It is currently unclear from the state of the art what constitutes 'enough' data for the creation of LLMs. Nevertheless, the DPC has made recommendations where obvious issues with data minimisation were identified.

While the mitigations to avoid unnecessary, irrelevant collection and processing of personal data have varied between controllers, the DPC encourages the following to minimise the personal data used when creating LLMs or any AI (although this is not an exhaustive list):

- Pseudonymisation or anonymisation of a dataset;
- De-duplication of a dataset;
- Reducing the timeframe of collection (either at the time of data collection or the historical period of collection);
- Avoiding training on the person data of individuals under the age of 18;

- Avoiding sources with a likely prevalence of sensitive personal data, special category data or data capable of revealing special category data; and
- Avoiding sources of personal data where the user may not be aware of the processing, such as dormant accounts.

In one instance, a controller had identified that certain personal data points were to be removed from one dataset as they were not necessary but did not apply the same standards or removals to all datasets being used for the same processing purpose. Where a controller's own risk assessment identifies a personal data point as unnecessary or of sufficient risk to be removed, the DPC has recommended that such data be removed from all datasets.

Data minimisation is not just important when creating AI, but remains relevant when an AI application is being used by an individual. Controllers should ensure that the data minimisation principle is also applied when considering deployment of AI. In some cases, design choices can lead to processing more personal data than is necessary.

Case Study

Two controllers took similar design choices, resulting in audio continuing to be recorded after the last interaction by the user, in one case for up to one hour. In the case of the longer timeframe, the controller explained that the one-hour timeframe was the maximum capacity of the AI. From the use cases described by the controllers, it was clear that these AI were likely to be used in hands-free situations, similar to other non-AI voice assistants on the market. There was no wake word, and the controllers appeared to heavily rely on the default notification options available on the device's operating system to alert users that recording was still taking place. In both cases, the DPC recommended that the timeframe for recording be reduced as well as the inclusion of an audio reminder that the AI was still listening to notify the user. By providing users with more appropriate reminders that the AI application was listening and reducing the timeframe for the AI to listen, this assisted the controllers in better demonstrating compliance with the principles of data minimisation and transparency.

Key Takeaways

Controllers should ensure that any design choices on the length of time of collection be justified as necessary and provide individuals with appropriate information and reminders. Simply because it is technologically possible to process personal data for a period of time does not mean that doing so complies with the principles of data minimisation. Existing guidance on similar use cases in relation to data protection and GDPR principles, generally, can be applied to comparable AI use cases.

Some organisations have stated that the onus is on individuals to limit the personal data they share, and that any data shared can be processed. This argument is clearly against the CJEU jurisprudence, in cases such as C-73/07 *Satamedia*, and the DPC considers that requiring an individual to self-censor will very likely impact negatively on that individual's rights and freedoms.

AI and Children

The protection of children and other vulnerable groups remains an important focus under the DPC's Regulatory Strategy. Through Supervision, the DPC has identified a number of issues that are particularly relevant for those under 18 years of age.

As children have different capacities depending on their age, finding an appropriate level of transparency for different age cohorts has been challenging for controllers. AI terminology may be unfamiliar to children and the usage of language associated with AI may prove difficult for some. For example, AI terminology such as 'hallucinations' or 'machine learning' may not be understandable to a young person.[20]

The DPC has made a number of recommendations, including:

- Enhanced information at onboarding for child users;
- Clearer transparency using plain English, explaining terms, or otherwise making AI information more readily accessible to younger individuals; and
- Provision of transparency that is directed to children on AI and its risks.

Particularly for generative AI, the DPC has had concerns that it may be difficult for children to understand the risks posed or that children may have a greater likelihood to trust an AI that could provide incorrect information, particularly in the case of generative AI hallucinations. For six controllers, in relation to seven different AI products/services, the DPC has recommended they provide greater transparency to children on data protection risks and ongoing in-app transparency to support children in learning about tools to check the veracity of generative AI outputs.

[20] Although the DPC notes that such terminology may also be difficult for some adults to understand.

The majority of engagements regarding the development of AI models, and particularly generative AI models, have seen controllers choose not to process the personal data of children for AI training. If the controller has chosen to avoid processing children's data for AI training, the DPC has recommended that the controller not process personal data from any account where the controller is unsure if the user is a child or not.

Where a controller intends to process historical data for AI training, users who transition to adulthood may find themselves confronted with their personal data being used for AI training, which was not part of their experience before. Given this may constitute a very significant change to the way their data is processed, controllers must give greater consideration to these individuals and their rights.

Case Study

Three controllers intended to implement AI training on user personal data. As part of the implementation, adults were given a notice period, transparency, and an opportunity to object to the processing. The personal data of children would not be used for AI training, but once a user became 18, the data on their account would then be used for AI training. The DPC recommended that users who transition to 18 are provided with a notice period and transparency related to AI training, specific to such users, allowing the individuals the opportunity to educate themselves on the implications of the notice, the new way their personal data will be processed, and allow them an opportunity to exercise their right to object prior to their data being processed.

Key Takeaways

Where an individual transitions to being an adult, the DPC considers that controllers have an obligation to extend them a notice period, transparency, and opportunity to object prior to processing of their personal data where such processing was not occurring when they were under the age of 18.

Automated Decision Making

In In some instances, AI may involve automated decision-making. Where a decision produces legal effects or similarly significantly affects an individual, data subjects have a right not to be subject to such a decision based solely on automated processing. Controllers need to ensure that they design and deploy their AI systems and processes with data protection front and centre. This assists in demonstrating compliance with data protection by design and default, and is imperative to avoiding prohibited automated decision-making.

Case Study

One controller sought to address disruptive behaviour of users in voice communications via an AI solution. The decision made by the AI solution could result in a user being banned from the service. The DPC considered that such a decision fell within the scope of the prohibition on automated decision-making. The controller argued that compliance with the prohibition in Article 22 GDPR was achieved because the AI system was triggered by another user, and there was an appeal process that involved a human review. The DPC recommended that the controller review the AI system and, in the context of Article 22 of the GDPR, consider if and where in its processes further human review was needed. In response, the controller made significant changes to the process, ensuring human involvement before any ban was imposed.

Key Takeaways

The DPC considers that triggering an AI system by a human or an appeals process with human intervention does not necessarily ensure compliance with the obligation not to subject a data subject to a prohibited decision under Article 22 of the GDPR based solely on automated processing. All phases of deployment must be considered independently to ensure data protection compliance.

Looking to the Future

While this document has focused on 2021-2025, the work of the DPC continues in areas on the cutting edge of AI. Engaging with AI technology firms has provided the DPC insight into not just what is currently on the market, but what is coming.

Agentic AI is growing rapidly, and the DPC has engaged with 5 separate controllers on their Agentic AI launches in the EU. While the DPC has already issued recommendations on a number of Agentic AI engagements, the common themes across controllers were not yet clear at time of writing beyond a continuing lack of transparency and clarity on how the AI processes personal data how the AI works and its potential impacts on users' rights.

As AI adoption becomes more diffuse, various use cases have created new ways to process personal data. In some circumstances, it appears that data protection law is being used as a replacement for moral or ethical social norms in the use of AI, even where data protection may not be the most appropriate regulation or law to address such matters. For example, image and video generation improvements in generative AI has continued apace, including varying capabilities for the generation of non-consensual intimate or sexualised images. Other legal regimes, such as harassment, harmful communications, defamation, or laws governing the distribution of intimate images may provide better reliefs to victims than data protection law. Nevertheless, while the DPC's jurisdiction does not necessarily extend to all use cases, the DPC has and will continues to act where there are concerns that personal data is being processed and questions exist as to whether such AI capabilities comply with the GDPR.[21]

Additionally, this document has not dealt with the Supervision of AI under the AI Act or other legislation. Conscious of the ongoing efforts in the EU for simplification,[22] the

[21] Data Protection Commission opens investigation into X (XIUC) | 17/02/2026 | Data Protection Commission

[22] European Commission, *Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL amending Regulations (EU) 2016/679, (EU) 2018/1724, (EU) 2018/1725, (EU) 2023/2854 and Directives 2002/58/EC, (EU) 2022/2555 and (EU) 2022/2557 as regards the simplification of the digital legislative framework, and repealing Regulations (EU) 2018/1807, (EU) 2019/1150, (EU) 2022/868, and Directive (EU) 2019/1024 (Digital Omnibus),* COM(2025) 837 final, Brussels, 19.11.2025; European Commission, *Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL amending Regulations (EU) 2024/1689 and (EU) 2018/1139*

DPC remains live to the possible changes in law and will adjust recommendations on AI in future, as needed.

as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), COM(2025) 836 final, Brussels, 19.11.2025.

Conclusion

The DPC's Supervision of AI seeks to achieve the DPC's strategic goals, that is to *Regulate consistently and effectively, Safeguard Individuals and promote data protection awareness, Prioritise the protection of children and other vulnerable groups,* and *Bring clarity to stakeholders and Support organisations and drive compliance.* Key findings of this insights report highlight that rigorous and ongoing data protection assessments complement and deliver responsible AI, helping to guide the design, development, deployment, and use of AI. Through engagements involving different AI such as LLMs, the DPC has secured significant improvements in data protection compliance in the areas of lawful basis, transparency, data minimisation, and the protection of children.

The DPC continues to engage with controllers on a daily basis on various issues including AI matters. Through Supervision, the DPC is able to take a risk-based approach to regulatory interventions which drives compliance with the GDPR resulting in data protection, privacy-centric AI products which mitigate potential harms and risks to individuals. Controllers who engage early, with well-considered compliance documentation can benefit from Supervision. While Supervision does not 'green-light' or 'approve' projects, controllers can gain insights into the DPC's assessment of their proposed processing and potentially avoid escalation pathways such as urgency proceedings before the Irish High Court and inquiries.

The DPC is conscious that EU simplification measures, in the form of both the Digital Omnibus and AI Omnibus, which have had, and may have further changes for both controllers and data subjects. The DPC will continue to apply the GDPR and any legislative amendments arising from the Omnibus procedures; treating all controllers fairly, encouraging compliance, and protecting the rights of data subjects.

Appendix A: Key Takeaways

For ease of reference, please find a list of all key takeaways:

- Early engagement with the DPC can ensure that foundational design choices meet the GDPR's requirements, resulting in a more compliant final product
- Supervision can engage at all stages of a development cycle, providing guidance to controllers that may prove beneficial in developing AI that is centred on the principles of data protection by design and default
- The DPC considers that Articles 5(1)(a), 12, and (as appropriate) Articles 13 and 14 of the GDPR place an onus on controllers to provide a clear notification of the intended use of personal data for AI training, as well as a reasonable notice period wherein individuals can educate themselves and exercise their data subject rights. Such a notice period may also be necessary when a user's personal data becomes subject to AI training at a later date, for instance, upon the user reaching the age of 18
- Where the DPC determines that the processing or intended processing of personal data presents a real risk or harm to the rights and freedoms of data subjects, the DPC can and will act to prevent harm
- The DPC continues to advocate for and seek regulatory consistency across the EU in line with the DPC's strategic goals. The EU-wide harmonised approach achieved through the AI Opinion allows the DPC to provide clear, consistent advice to controllers deploying AI models in Ireland and across Europe. This drives regulatory consistency, a pillar of the DPC's Strategic Strategy, while enabling and encouraging data protection compliant innovation
- Where matters are unclear, Supervision can identify potential issues for controllers, the DPC, and the EDPB that need further clarification. The issues identified can assist in driving priorities for guidance or areas for consensus
- Controllers looking to develop AI should consider implementing measures to de-identify, pseudonymise or anonymise datasets used for AI training
- Monitoring live launches and treating the data protection assessments as living documents can demonstrate responsible innovation and assist to develop safer, more compliant products for use

- Compliance does not end with a feature or product's launch
- It is the responsibility of controllers to continue to monitor the implementation of their products as well as how the technical and organisations measures are functioning to ensure compliance
- The DPC will continue to monitor post-launch, and take action where infringements or issues arise
- Good documentation facilitates speedy regulatory engagements
- Organisations should anticipate and proactively address questions likely to be asked by the DPC.
- Organisations should ensure compliance documentation provides quality information that is in-depth featuring clear, articulated analysis and documented design decisions
- Where designated, the DPO must be involved in creating and documenting a DPIA
- Where poor documentation is provided, engagements with the DPC can become protracted and delay launches. Poor documentation may also raise doubts as to the effectiveness and completeness of the organisation's compliance readiness programme
- Where controllers provide clear, in-depth documentation, engagement with the DPC is more straightforward and faster
- The DPC considers that the training of a new AI model will likely require a Data Protection Impact Assessment
- Novel technology may require additional transparency to not only explain how to use the technology, but also educate the individual on the data protection implications and risks
- The DPC considers that controllers have a duty to not only explain that processing is occurring, but explain the data protection risks that may exist in relation to the processing of personal data by AI and LLMs, particularly where new technology is concerned
- Transparency information should be addressed to and be accessible for individuals to ensure compliance with the GDPR
- It is essential that controllers carry out thorough due diligence ensuring that all aspects of the design and notifications associated with a product or service

have been assessed, risk/harm mitigation measures are implemented, and all aspects of the proposed processing meets the requirements of the GDPR

- As past collections of data are unlikely to have been conducted with users expecting their data would be used for AI training, the DPC considers that using historical data will heighten the need for enhanced transparency
- Where historical data is amongst a controller's intended dataset(s), any analysis cannot disregard historical conditions or transparency provided by the controller, as they form the basis upon which expectations were shaped. The DPC will carefully assess compatibility assessments based not only on the arguments within, but evidence provided, to determine the reasonable expectations of data subjects
- The DPC considers it is unlikely that a controller will be able to demonstrate that data subjects expect the usage of their private messages or other private data for AI training
- The DPC considers that any chilling effect on the exercise of data subjects rights is unacceptable and any friction to access an objection form or method should be minimised to the greatest extent possible
- Objection forms or opt outs should be easily accessible, easy to use, with sufficient transparency and notice for data subjects to understand the processing and its implications
- Controllers should ensure that any design choices on the length of time of collection be justified as necessary and provide individuals with appropriate information and reminders
- Simply because it is technologically possible to process personal data for a period of time, does not mean that doing so complies with the principles of data minimisation
- Existing guidance on similar use cases in relation to data protection and GDPR principles, generally, can be applied to comparable AI use cases
- Where an individual transitions to being an adult, the DPC considers that controllers have an obligation to extend them a notice period, transparency, and opportunity to object prior to processing of their personal data where such processing was not occurring when they were under the age of 18

- The DPC considers that triggering an AI system by a human or an appeals process with human intervention does not necessarily ensure compliance with the obligation not to subject a data subject to a prohibited decision under Article 22 of the GDPR based solely on automated processing.
- All phases of deployment must be considered independently to ensure data protection compliance.

Appendix B: Statistics Reference

For ease of reference, all figures are reproduced here in Appendix B.

Figure B1

DPC Supervision AI Engagements by Year (2021-2025)

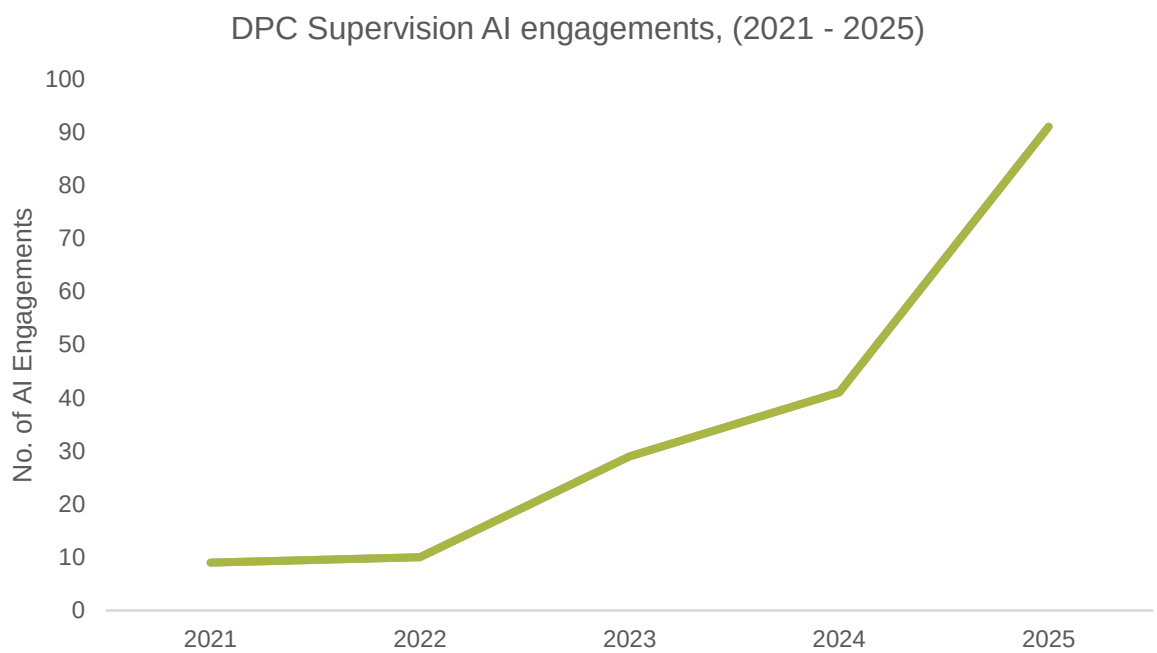

Note: Data sourced from internal TMSU Statistics including the full year of all AI engagements (January to December) in 2025.

Figure B2

DPC Supervision AI Engagements (AI and Total)

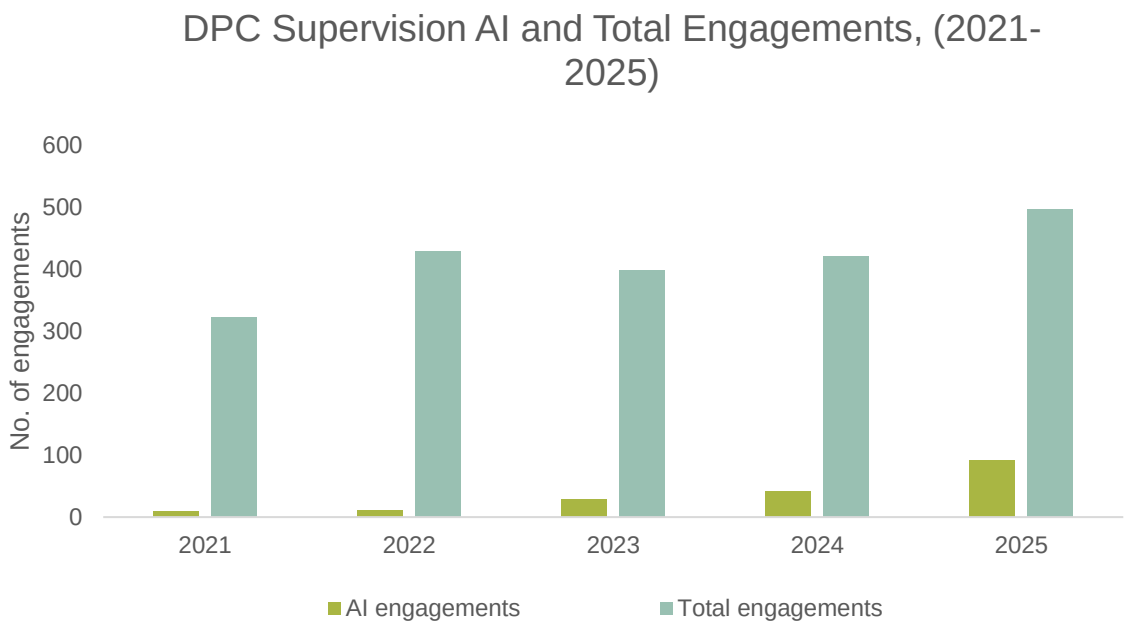

Note: Data sourced from internal TMSU statistics. Totals reflect all engagements with controllers and other stakeholders as reported in the DPC Annual Report.

Figure B3

DPC Supervision AI Engagements (YOY Growth 2022-2025)

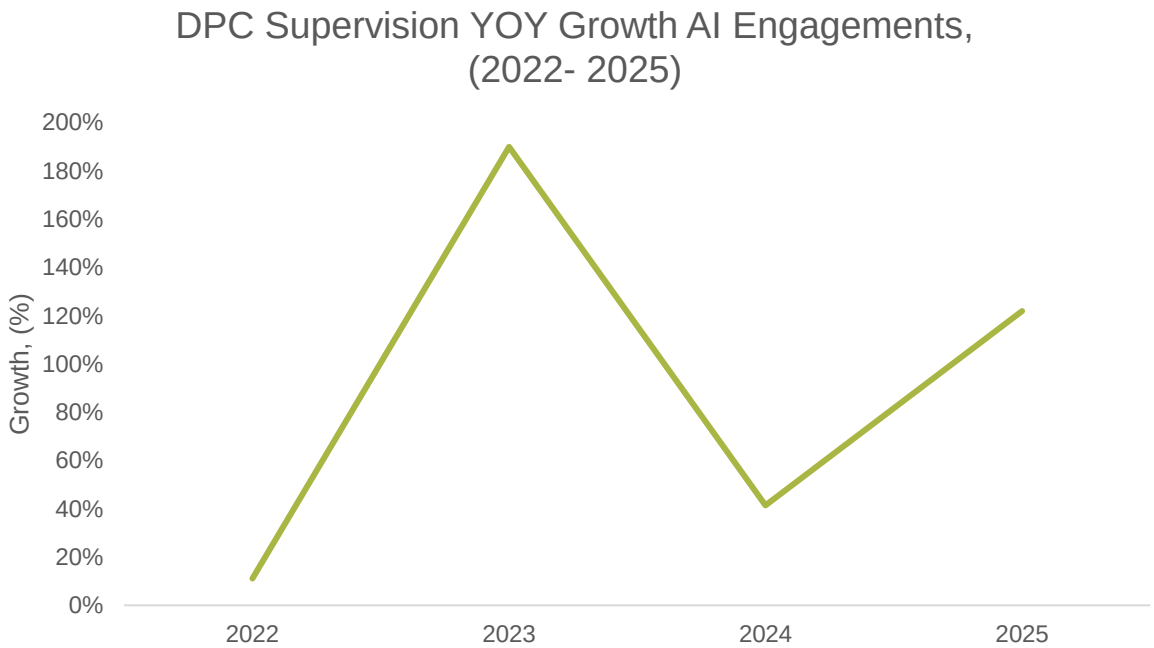

*Note: Data sourced from internal TMSU Statistics *Note the 2025 growth rate is calculated using the full year 2025 count of engagements.*

Figure B4

Most Frequent AI Technology Type over a Five-year Period in DPC Supervision Engagements

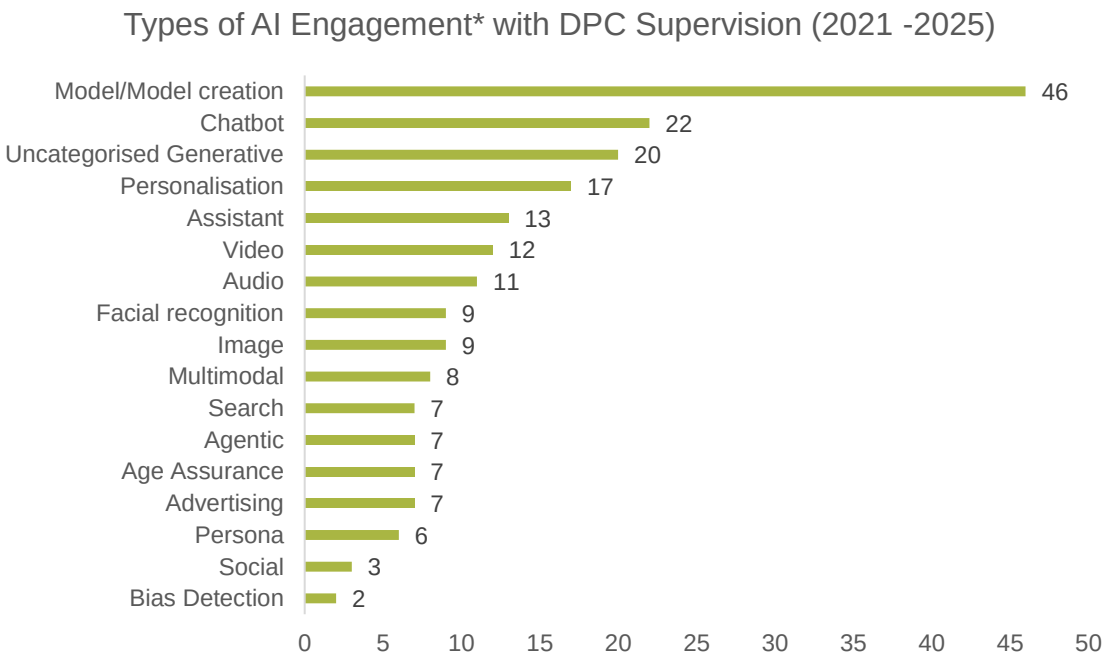

*Note: Data sourced from Internal TMSU Statistics. *Please note there is some crossover in types of AI counted for one AI engagement (for example, model creation and age assurance).*

Figure B5

DPC Supervision Generative AI Engagements by Technology Type

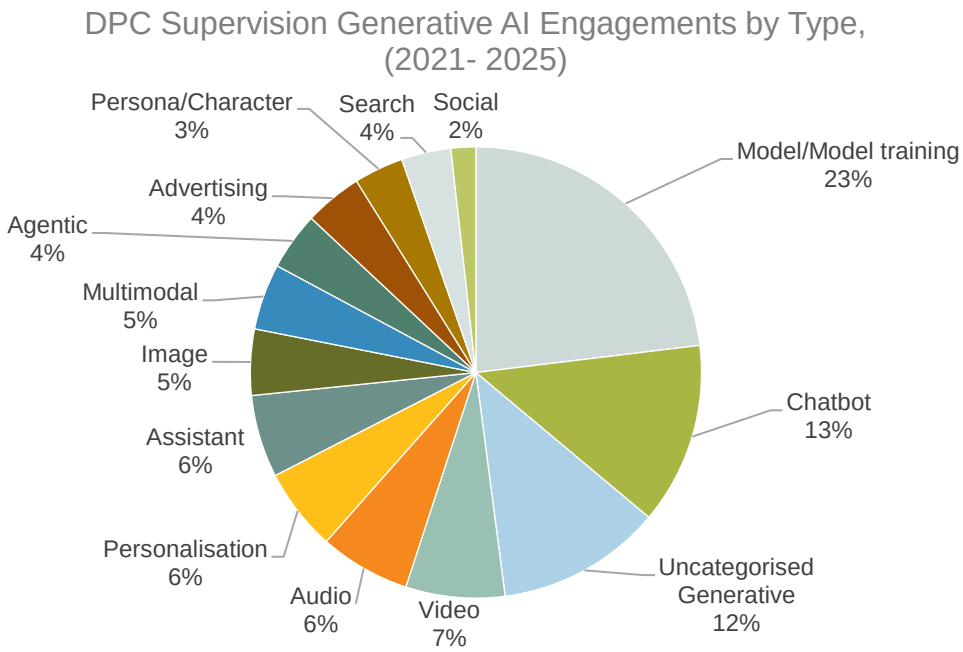

Note: Data sourced from internal TMSU Statistics

Figure B6

DPC Supervision Non-Generative AI Engagements by Technology Type

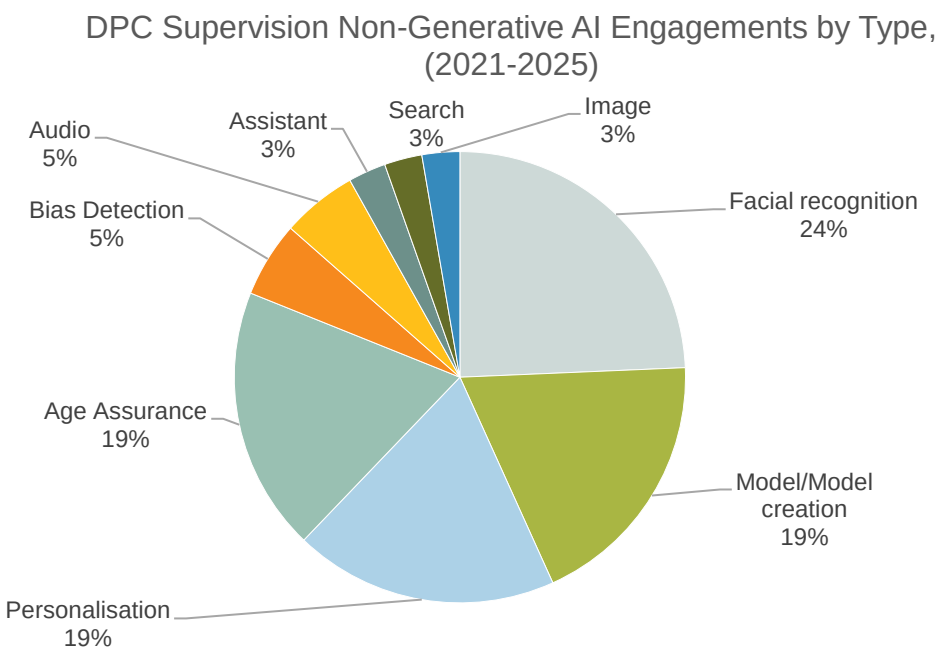

Note: Data sourced from internal TMSU statistics

Figure B7

Top Five Cited Data Protection Issues in DPC Supervision AI Recommendations.

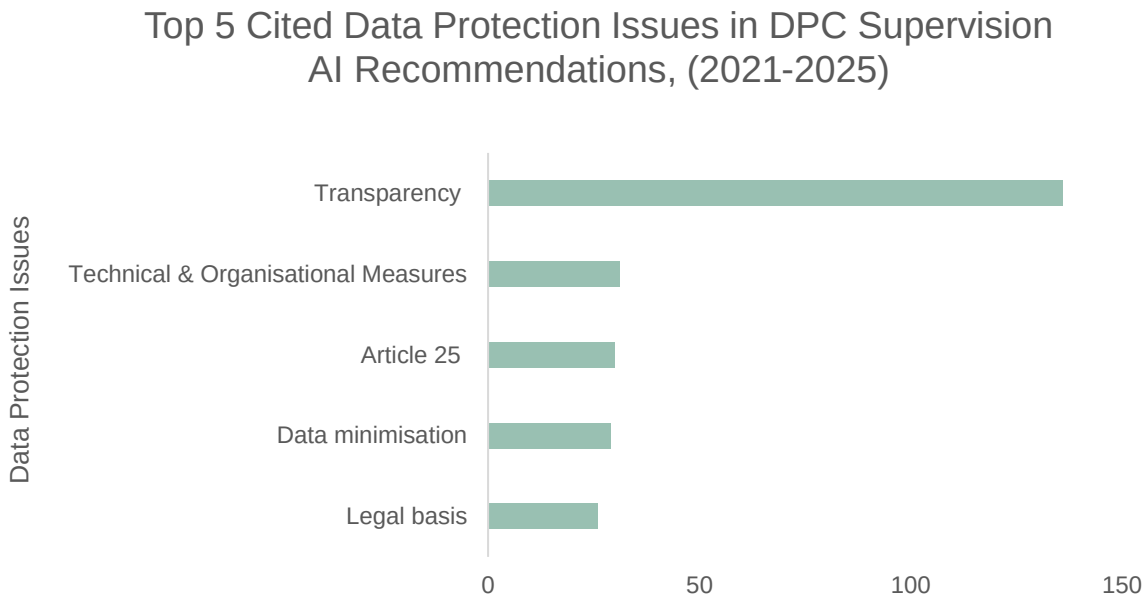

Note: Data is sourced from internal TMSU Statistics.

Figure B8

Data Protection Issues Cited in DPC Supervision AI- related Recommendations

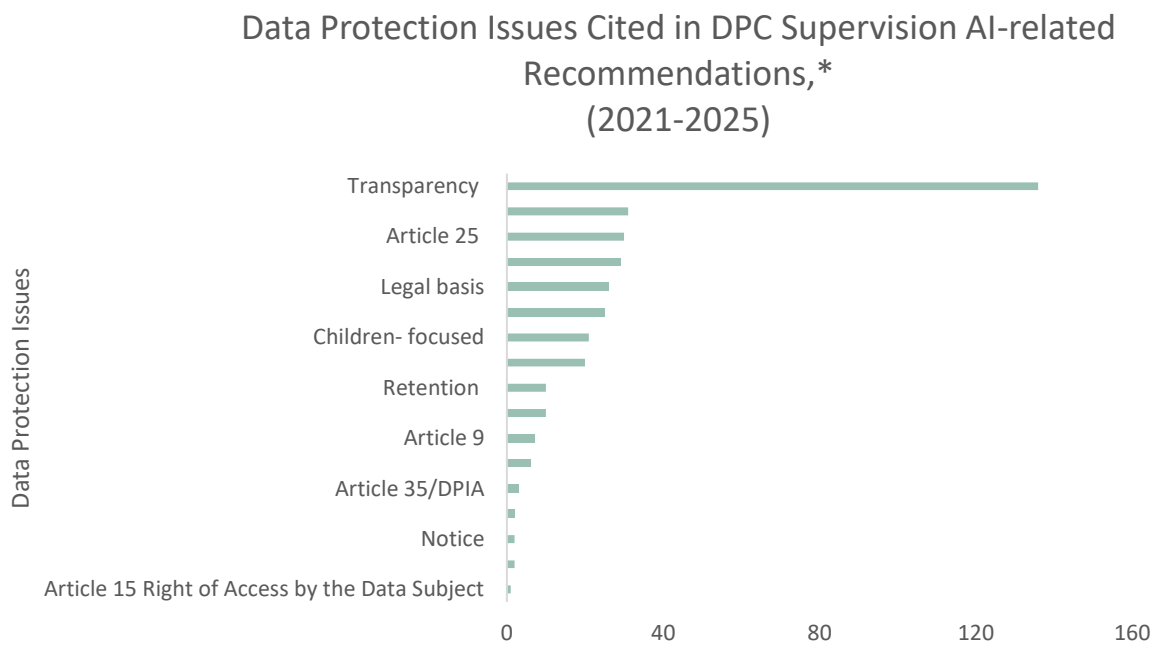

*Note: Data is sourced from internal TMSU Statistics *Please note that a single recommendation may cover multiple data protection issues, (for example transparency and objection).*