Centre for Information Policy Leadership
HUNTON

Legitimate Interests for Data in AI Training - The DPO Perspective

December 2025

CIPL
Centre for Information Policy Leadership
HUNTON

Legitimate Interests for Data in AI Training – The DPO Perspective

Contents

Introduction

Artificial intelligence (AI) has become a defining force of technological and economic change, driving innovation across sectors and challenging existing legal and regulatory frameworks. Governments around the world are positioning themselves to harness the transformative potential of AI by investing in AI infrastructure, skills, and innovation-friendly frameworks. In Europe, for example, Mario Draghi identified critical structural challenges holding the region back from global leadership, which included fragmented and complex regulatory frameworks and limited access to high-quality data.[1] Similar national strategies elsewhere[2] underline the global momentum behind AI as an engine for sustainable growth.

Yet this momentum underscores a central tension at the heart of AI development: AI systems often depend on access to large volumes of data, while data protection frameworks often globally place strict conditions on the collection, use and re-use of personal data. Transformative AI technologies, including large language models (LLMs) are already subject to a wide range of data protection laws and principles.[3] Governments, regulators and corporate compliance functions are therefore increasingly focused on how to interpret and apply data protection principles to enable responsible AI development and, in particular, how to enable broad collection, use and re-use of data. These legal discussions are taking place against a general recognition of the significant value these technologies bring, and how their continued responsible development and use have the potential to deliver broad benefits across all sectors and geographies, supporting economic growth and societal advancement worldwide.

The DPO Forum @CIPL

The Centre for Information Policy Leadership (CIPL)[4] has been spearheading work to interpret the legal basis for processing personal data under data protection laws[5] as well as apply data protection laws and principles to AI technologies since 2019.[6] Most recently, CIPL has published a paper applying data protection principles to generative AI, highlighting the tensions and the interplay between such technologies and globally adopted data protection principles.[7] At the same time, many data protection authorities (DPAs) around the world have been considering these same topics in a number of consultations and regulatory guidance, including the EU's European Data Protection Board (EDPB), which published an Opinion on AI Models in December 2024,[8] and the UK's Information Commissioner's Office (ICO), which consulted on a series of positions on generative AI and data protection[9], alongside many other international regulators.[10]

Building on CIPL's prior thought leadership on AI and recent regulatory guidance, CIPL has convened a select group of Data Protection Officers (DPOs) from leading companies – known as the 'DPO Forum @CIPL'. This initiative enables DPOs to exchange insights, discuss best practices, analyse emerging challenges, and contribute informed, practice-oriented perspectives to the broader policy and regulatory discourse.

This paper distils discussions with the DPO Forum @CIPL and explores the practical application of the legitimate interest (LI) legal basis[11] for processing personal data in the context of AI. We recognise that applying LI to AI training can vary significantly based on the technology in question, the stage in the AI lifecycle, and the type or source of the personal data. We have deliberately kept the reference to AI training wide enough, so that we can rely on examples from all of these important factors. AI training in this paper therefore includes everything from narrow, task-specific models to broad-scope foundation or general-purpose models. The legal basis analysis and safeguards will differ substantially for each and they will be dealt with separately throughout the paper, where relevant.

The objective of this paper is threefold:

a) To underscore that LI is often the most appropriate ground for data processing in AI training, enabling responsible and accountable innovation;

b) To share DPOs' practical experiences and considerations of the specific compliance steps in applying the LI legal basis, including the Legitimate Interest Assessment (LIA) in the context of training AI; and

c) To contribute to the development of future best practices in the implementation of LI and to inform the stakeholders, including regulators, of the expectations of success criteria in the application of LI for data processing in AI development and deployment.

This paper is structured into two parts:

1. Section 1 contains a set of key considerations and potential solutions, which may arise for organisations and their DPO offices when seeking to rely on LI as a legal basis for processing personal data for AI development and deployment.

2. Section 2 methodically examines some of the key compliance steps required in the process of choosing and relying on LI legal ground for data processing for AI training, including the three-step LIA. We look at compliance steps through the contextual lens of AI development and highlight how compliance obligations can be met through a risk-based, accountable approach.

While DPO Forum @CIPL is comprised of DPOs appointed under the General Data Protection Regulation (GDPR), this paper and its findings are relevant for organisations and their DPOs and privacy counsels in all countries whose data protection laws include LI as a ground for data processing, or similar legal concepts.

Proposed considerations for organisations

1. **Recognise that AI training is not one single activity**. Address each stage of the AI lifecycle, including data acquisition, pre-processing, model training, fine-tuning, evaluation and monitoring, in a way that reflects its specific context and risks.
2. **Differentiate between types of AI models and contexts**. Clearly identify whether the model is narrow or has a sector-specific use case, or whether it is a general-purpose AI model. This will impact the proportionality, balancing and safeguards of the LIA.
3. **Define a specific and evidence based legitimate interest**. The interest must be concrete, not abstract. When describing the benefits of the processing, be as precise as possible and ensure each claim is evidenced.
4. **Consider the impact of inaction**. Document the potential harms of not processing, as well as the risks of processing.
5. **Integrate a taxonomy of potential harms**. Use structured taxonomies of potential harms to ensure the LIA is repeatable. Quantify or categorise those impacts to ensure consistent internal decision-making and to demonstrate accountability to regulators.
6. **Consider integrating a taxonomy of benefits**. Alongside your harms taxonomy, consider establishing an internal taxonomy of benefits arising from the AI training, using this framework to articulate the positive impacts of the processing.
7. **Embed monitoring and periodic review**. Embedding monitoring and review into the AI training lifecycle ensures the LIA remains a living and adaptive process. As not all risks can be foreseen before training, periodic review allows organisations to identify emerging impacts, adjust safeguards, and keep risk management aligned with real-world outcomes.
8. **Integrate LIAs with other risk assessments**. Coordinate LIAs with other risk assessments (for instance, data protection impact assessments and other fundamental rights impact assessments) to ensure consistent treatment of risks and rights across all areas of the business.
9. **Assign clear roles and responsibilities**. Assign clear roles across all relevant teams. Establish clear escalation routes to DPO, where relevant, for high-risk or novel cases.
10. **Adopt a forward-thinking, risk-oriented culture**. View the LIA not as an obstacle but as an enabler of responsible innovation.

Proposed considerations for regulators

1. **Recognise that AI training is not one single activity**. Address each stage of the AI lifecycle, including data acquisition, pre-processing, model training, fine-tuning, evaluation and monitoring, in a way that reflects its specific context and risks.
2. **Provide clear interpretative guidance on applying LI to AI training**. Clarify how the three-stage test applies to different AI use cases. Consistent guidance would support convergence across jurisdictions and reduce regulatory uncertainty.
3. **Endorse proportional and contextual balancing tests**. Encourage risk-based and contextual approaches that recognise both the benefits and the harms.
4. **Support wider adoption of taxonomies**. Promote the development of shared taxonomies of harms and benefits across regulators.
5. **Clarify expectations for reliance on Article 14(5)(b) GDPR**. Issue guidance on 'disproportionate effort' in AI contexts, with practical examples.
6. **Promote harmonised approaches across supervisory authorities**. Coordinate between DPAs to ensure consistent interpretation of LI for AI training.
7. **Recognise the role of safeguards and accountability**. Endorse effective organisational and technical safeguards, such as PETs, and organisational accountability, which can effectively reduce risk and tilt the

balance in favour of organisations relying on LI.

8. **Address all relevant fundamental rights**. Encourage organisations to undertake holistic balancing that considers the full range of fundamental rights, alongside privacy and data protection.
9. **Acknowledge the limits of pre-launch assessments**. Accept that not all risks related to AI training can be foreseen before deployment, and promote ongoing monitoring and post-launch review to ensure continued compliance and accountability.
10. **Encourage evaluation of societal benefits**. Provide guidance on how organisations can systematically identify and document the broader benefits of AI training, to ensure a balanced and evidence-based assessment.

Part 1 – Why LI can be the right legal basis

AI development and the application of the GDPR

Under the GDPR, organisations must identify a legal basis for processing personal data.[12] As the EDPB recently reinforced in its AI Opinion,[13] there is no hierarchy amongst the legal bases, and organisations have to determine the most appropriate legal basis for their processing context.[14] LI is one of the six available legal bases, designed to accommodate new, beneficial and legitimate uses of personal data, including those not foreseeable at the time of collection, provided such uses do not unduly prejudice individuals' rights.

When relying on LI, an organisation should conduct a three-stage test, also known as a legitimate interests assessment.[15] The LIA identifies that (i) the LI is pursued by the controller or by a third-party ('the purpose test'); (ii) the processing is necessary to pursue the LI ('the necessity test'); and (iii) the LI is not overridden by the interests or fundamental rights and freedoms of individuals ('the balancing exercise').[16]

There is a growing consensus among data protection practitioners, lawyers, experts, and increasingly, regulators[17], that among the GDPR's six legal bases, LI is often the most appropriate legal basis to support responsible AI innovation, while safeguarding fundamental rights.[18] Its flexibility, contextual balancing, and built-in safeguards make it well-suited for dynamic and evolving processing activities such as AI model training. Increasingly, regulators such as the EDPB and the ICO have acknowledged that LI may serve as the sole legal basis for certain AI training activities (for example, web-scraped or other third-party data processing),[19] but they also set a high threshold, requiring organisations to demonstrate necessity, proportionality, and robust safeguards,[20] including transparency requirements.[21]

LI has historically been misunderstood and under-utilised by both regulators and organisations: it has erroneously been considered a last resort that should be reserved for exceptional cases.[22] As a result, organisations are frequently hesitant to rely on it, turning instead to other legal bases (consent, for instance) that may be seen as less controversial or more certain. Reluctance persists due to: (1) perceived regulatory uncertainty and heightened scrutiny; (2) the high bar and complexity of the LIA; (3) continuing scepticism from some regulators about LI's appropriateness for AI training.

Despite these concerns, LI offers significant advantages that make it one of the most appropriate legal bases for training AI for the following reasons:

1. **Built-in accountability and risk mitigation**. LI is effectively grounded in organisational accountability, enabling a robust level of protection for individuals. The burden is on the organisation to assess the risks relating to the data processing activities and to define mitigations. In the context of AI training, this creates a structured framework for managing risks, while enabling responsible data use, and enhancing protection for individuals without stifling innovation. This only serves to highlight the in-depth accountability that the LI represents, especially in comparison with other legal bases.

2. **Supporting compatible re-use**. LI also provides a specific mechanism to reuse personal data held by organisations for purposes that are compatible with the original purposes. This flexibility could be vital in AI training, where iterative refinement and retraining are routine. By contrast, consent requires re-obtaining permission for each new use, even where compatibility exists, which hampers scalability. This is critical to enable organisations to responsibly use their own (first-party) data for a new but compatible purpose, while maintaining safeguards for individuals.

3. **Access to publicly available data with safeguards**. LI also creates the opportunity to utilize publicly available data to train AI, where the balancing test is met and appropriate safeguards are in place. This is particularly relevant for training AI models, where access to diverse and representative datasets, such as those gathered

via web-scraping, can improve model accuracy, fairness, and mitigate bias. The ICO has acknowledged that, in the context of web-scraped data, LI is often the most appropriate legal basis when models are developed in ways that align with appropriate safeguards to minimise risk.[23]

4. **Individuals' right to object**. Individuals retain a meaningful safeguard, as organisations must be prepared to assess and, where appropriate, honour such objections, unless they demonstrate compelling overriding interests.[24]

5. **Third parties**. LI expressly refers both to the legitimate interest of the controller or that of "a third party". This means organisations may consider not only their own legitimate interests, but that of others, as a lawful basis for the training of AI using personal data.

Reluctance to use the LI legal ground for data processing could result in a loss of opportunities stemming from the use of new technologies and a risk aversion to progress projects that may be beneficial to individuals, organisations and society. It decreases focus on the LI of third parties and restricts the scope of organisational considerations. It could lead to undermined competitiveness and echoes the concerns raised in Mario Draghi's report on European innovation. By contrast, LI is a critical enabler for organisations of all sizes, offering the flexibility needed to support responsible data use in fast-moving, data-driven environments, while safeguarding privacy through the implementation of appropriate safeguards.

However, as highlighted in CIPL's paper *Rethinking Sensitive Data in the Age of AI*[25], the processing of sensitive data to train the AI model is not permissible relying solely on LI. Where such data is present in the training dataset, an appropriate condition for processing under Article 9 or 10 GDPR must be identified and applied. Ultimately, in order to facilitate the responsible use of sensitive data in AI systems, re-visiting existing laws may be necessary to address the legal gaps that limit the responsible use of sensitive data in AI development, as the European Commission has proposed in its Digital Omnibus[26].

To realise the full potential of AI innovation in a rights-respecting manner, greater regulatory clarity and dialogue around the application of LI will be essential. This paper aims to contribute to this conversation by highlighting key considerations based on practical DPO perspectives on the lawful use of data for AI training.

Part 2 – Key considerations when relying on LI

LI for first-party data and third-party data

The application of LI to AI training requires different considerations depending on the source of the personal data. A key distinction arises between first-party and third-party data:

- **First-party data**: Collected and created directly by an organisation from its users, customers, employees, within the context of an existing relationship.

- **Third-party data**: Collected and created indirectly by an organisation, from third parties and often from publicly available sources. In the context of generative AI, the scraping of publicly available data from the web is a common practice for developers to gather large and diverse datasets.

LI can be a viable legal basis for use of both first- and third-party data, but the balancing analysis will likely differ.

Unduly impeding an organisation's ability to process its own (first-party) personal data for AI training may hinder responsible AI model development and broader innovation. The existence of a pre-existing relationship between the individual and the organisation does not limit an organisation's ability to rely on LI to use existing data for a new purpose of AI training.

Contextual consideration: Using first-party data to train a GenAI model.

In some cases, an organisation may only be able to provide a generative AI model in a certain language if it processes first-party data, because third-party sources do not provide sufficient or accurate data. In these situations, the organisation should be able to rely on LI, upon successful completion of the LIA.

Considering other legal bases

Although it may arise that ultimately LI is the most appropriate (or only viable) legal basis for AI training, organisations should first consider which legal basis is the most appropriate by going through the list under Article 6 GDPR. In practice, the focus of such consideration will likely be on consent, contractual necessity and LI, as the remaining legal bases (legal obligation, public interest, and vital interests) apply only in very narrow, or clearly defined circumstances.

A. Consent

In theory, organisations could adopt a consent-first strategy, in order to reduce perceived regulatory risk. However, this approach overlooks the legal and practical limitations of consent in many modern data processing scenarios.[27] In the context of AI training, relying on consent is often neither feasible nor appropriate for the following reasons:

1. **Practical feasibility reduces the viability and integrity of the data set**. Obtaining consent for large and complex training datasets is unmanageable at scale, particularly for foundation or generative AI models. Refusals may fragment datasets, leading to incomplete, biased, or unreliable training data, which undermines accuracy, fairness, and non-discrimination. Re-obtaining consent for iterative retraining or compatible reuse is similarly impractical and restricts innovation.[28]

2. **Validity concerns**. In AI training contexts involving employment, educational institutions, public services, or other asymmetrical relationships, power imbalances may render consent invalid.[29] In addition, the opacity and technical complexity of AI training makes it difficult for individuals to fully comprehend the processing, undermining informed consent.[30]

3. **Weaker accountability and consent fatigue**. Once consent is obtained, organisations may treat it as sufficient without maintaining ongoing justification or oversight or implementing appropriate safeguards. Moreover, the overuse of consent requests can result in consent fatigue, reducing interactions to a formality and eroding the value of genuine, informed choice, effective transparency and empowerment for individuals.[31]

4. **Limitations on reuse**. When relying on consent, organisations are restricted in their ability to reuse personal data for compatible purposes. In AI training, where iterative development and refinement are key, this rigidity limits lawful data reuse that would otherwise be permissible under LI.

Consent, while important in specific and clearly defined contexts, is not a universally appropriate legal basis for the training of AI systems. A more nuanced and context-driven application of the lawfulness principle is essential.

B. Contractual necessity

Contractual necessity applies only where processing is objectively necessary to perform a contract with the individual, and the contract cannot be performed without the processing.

In the context of AI, contractual necessity may be appropriate where the developer or deployer of AI has a direct relationship with the individual, for instance:

- The training itself is integral to delivering the contracted service to the individual;
- The processing is necessary to meet contractual obligations or legitimate user expectations (for example, using AI to ensure safety or security on a platform, such as age verification, where no less intrusive means are available).

Contextual consideration: Contractual necessity (from CNIL How-to Sheets[32])

A text editor software company offers an automated and personalised mail generation service, to which the user contractually subscribes, and for which the editor collects the data of the users of this service. The data processing for this personalisation service may be considered, subject to its specific characteristics, necessary for the performance of the contract.

Contextual consideration: Safety model training

A company offers a smart home security service to consumers under a subscription contract. The service includes an AI-powered camera system that detects unusual activity and alerts the homeowner. To deliver the contracted service safely and effectively, the provider must train and periodically retrain the AI model on user footage to minimise false alerts, detect new threat patterns, and prevent system failures that could compromise household safety.

Because this model training is integral to maintaining the promised level of safety and performance, it may be regarded as necessary for the performance of the contract. The processing is limited to what is required to ensure the AI system functions as described in the agreement - i.e., to keep the home monitoring service accurate and reliable for the user.

However, the contractual necessity legal basis, while important, remains narrowly interpreted by both the Court of Justice of the European Union (CJEU) and DPAs including the EDPB. For instance, in the case of *Meta Platforms and Others v Bundeskartellamt* the CJEU adopted a strict and literal approach to contractual necessity, stating it must be "objectively indispensable" for core contractual purposes and that there are no other workable or less intrusive alternatives.[33] Furthermore, in many AI training contexts it is difficult to establish a direct contract with all data subjects whose data are processed. Even where such a contract exists, the necessity threshold remains high, the processing must be objectively necessary to perform the contractual service, not simply envisaged in the terms. Finally, and as under LI, the processing of sensitive personal data to train the AI model is not permissible on contractual necessity and an Article 9 GDPR condition should be identified and applied.[34]

C. Legal obligation

The legal obligation basis may be appropriate for AI training, where processing is necessary to meet a specific statutory or regulatory requirement. This legal basis is interpreted narrowly and applies only where the obligation is clearly set out in law. While most AI training activities are voluntary and innovation-driven, certain use cases may fall within mandatory compliance frameworks, for example, anti-fraud, anti-money laundering, safety or child protection requirements.

As the EU's digital rulebook continues to evolve, obligations under instruments such as the AI Act, the Digital Services Act (DSA), or sectoral safety laws may increasingly require organisations to train or retrain AI systems to meet defined regulatory standards. In such contexts, reliance on the legal obligation basis could become more relevant, particularly where AI training is an essential component of complying with statutory duties or safeguarding public interests.

Contextual consideration: Anti-money laundering obligations

A financial institution is required under anti-money laundering (AML) and payment services legislation to monitor transactions for fraud, money laundering, and terrorist financing. To meet these obligations, the bank develops and trains AI models that can detect suspicious payment patterns and anomalies across vast volumes of personal transaction data.

In this scenario, the processing of personal data for training the fraud-detection AI may be considered necessary to comply with a legal obligation imposed on the bank, under applicable financial regulations. In such a case the purpose of the AI training is directly linked to fulfilling statutory duties, rather than optional product improvement.

Intersections with other data protection principles

Lawfulness is only one of several core data protection principles under the GDPR. Individual protections are not derived solely from the legal basis selected, but from the full application of remaining GDPR provisions and its broader framework of rights and obligations. Regardless of the legal basis relied upon, organisations must apply all the data protection principles in a risk-based and proportionate manner for all data processing in the training and development of AI. These include fairness, transparency, data quality, security, accountability and the exercise of individual rights.

Focusing solely on the choice of legal basis risks encouraging overly restrictive or conservative approaches to AI, which could hinder innovation and societal development. A holistic application of the GDPR, one that integrates lawfulness with proportional safeguards across the lifecycle of AI development, is essential to building trust, meeting regulatory expectations, and enabling responsible innovation.

- **Accuracy and fairness**. The accuracy principle requires organisations to ensure that personal data is accurate[35] and, where necessary, kept up to date. This obligation applies throughout the AI lifecycle, including during training and output stages.[36] At the training stage, a representative dataset is central: without diverse and balanced inputs, models risk generating biased or misleading results. Accuracy is closely tied to fairness, which ensures organisations do not process personal data in a way that is "unjustifiably detrimental, unlawfully discriminatory, unexpected or misleading" to the individual.[37] AI systems that are trained on diverse and representative datasets are more likely to have accurate and non-discriminatory outputs, which enhances fairness. For narrow, domain-specific models, accuracy and fairness may be measured in relation to a defined task or dataset, and risks may be easier to test and isolate. For foundation or general-purpose models, on the other hand, ensuring accuracy and fairness may be more complex as the training datasets may be larger.

Contextual consideration: AI in recruitment

An AI recruitment tool trained only on CVs from one region or demographic may systematically favour candidates from that group. By contrast, training the model on a diverse and representative dataset that includes applicants across genders, ethnicities, ages, and educational backgrounds reduces the likelihood of biased recommendations. The broader dataset leads to more accurate predictions of candidate suitability and produces fairer outcomes.

  These datasets often include sensitive personal data, which require additional protections under the GDPR. Currently Article 9/10 GDPR conditions do not provide sufficient grounds for processing sensitive data solely to mitigate bias or discrimination, leaving a recognised legal lacuna. As further discussed in CIPL's *Rethinking*

Sensitive Data in the Age of AI,[38] the AI Act rightly acknowledges the significance of processing sensitive personal data[39] for bias mitigation, by providing high-risk AI developers the ability to process sensitive data for bias detection and correction,[40] subject to certain safeguards such as deletion once testing is complete. However, this exemption does not extend to non-high-risk AI, limiting its potential impact. A European Parliamentary Research Service report highlighted this very issue, ultimately concluding that "a reform of the GDPR or further guidelines on its interplay with the AI Act might help address these issues".[41] Complexity is further increased by divergent definitions of personal data and sensitive data across different jurisdictions. For instance, under the California Consumer Privacy Act (CCPA) as amended by the California Privacy Rights Act (CPRA), sensitive personal data includes financial and precise geolocation data which goes beyond the classification under the GDPR.

Fairness also depends on the reasonable expectations of individuals. Meaningful transparency can act as a bridge between fairness and privacy. If individuals have been informed of the processing, it is more likely to be in their reasonable expectations and thus contribute to overall fairness. Consistency with previous actions can help inform what is considered fair or within the reasonable expectations of an individual.

Fairness should also be understood in a broader sense: sometimes the absence of processing may itself lead to unfair outcomes. For example, not providing security or anti-fraud measures for financial payments and transactions may not be fair to individuals using the service. Hence, organisations and regulators should consider what the impact on individuals is, if they do not carry out the processing in question.

Ultimately, ensuring compliance with the accuracy and fairness principles can also enhance the quality of AI training and development, alongside meeting data protection obligations and fostering wider accountability. It will also be more likely to tip the balancing test in favour of the organisation, as it reduces the impact on individuals.

Contextual consideration: Measuring accuracy in a user-based product or service

Accuracy in a language model can be measured by how well it serves users across different languages. If training data is primarily in English, the model may perform poorly for speakers of Hindi, Spanish, or smaller European languages. By sourcing diverse multilingual data, developers can improve accuracy, ensure the model is fully representative, and deliver fairer outcomes for users worldwide.

Contextual consideration: Measuring and defining fairness

The four-fifths rule, originally established by the U.S. Equal Employment Opportunity Commission (EEOC) in its Uniform Guidelines on Employee Selection Procedures (1978), is a long-standing benchmark for detecting potential discrimination in hiring and selection practices. Under this rule, the selection rate for any protected group should be at least four-fifths (80%) of the group with the highest selection rate, and if it is any lower this may indicate a discriminatory or unfair practice.

It can also be leveraged to guide fairness in relation to AI and privacy, helping developers identify potential bias in AI outputs, although compliance with the four-fifths (80%) threshold does not guarantee that an AI system is fair or lawful, so it should be used in conjunction with other fairness assessments.

- **Data minimisation and proportionality**. The data minimisation principle requires that personal data be limited to what is necessary for the specified purpose. However, as further explained in CIPL paper *Reconciling AI with*

the Data Minimization Principle: Bridging the Innovation and Privacy Gap,[42] in the context of AI, what is "necessary" may often involve large and diverse datasets.[43] Depending on the type of technology or processing activity, limiting data indiscriminately may undermine model quality, inclusivity, and fairness. For instance, for a narrow model used in a sectoral application, it may not be necessary to train the AI on large and indiscriminate data sets, whereas a foundation model with a broad scope may need much larger and more diverse datasets.

Interpreted proportionately, data minimisation allows the use of data volumes appropriate to the benefits and risks of a model, a point emphasised by the EDPB in its AI Opinion.[44] The key questions for organisations are: how much data is necessary for the specific type of AI technology or processing activity, whether the processing is proportionate at each stage of development, and how privacy-enhancing technologies and privacy-preserving technologies (PETs/PPTs) can help reduce risks without undermining model quality.

When relying on LI, data minimisation operates as a key safeguard during the balancing test, particularly in high-risk AI use cases. Applied well, it can refine AI training, improve model efficacy and ultimately support accountability. Crucially, data minimisation cannot mean using less data than is required, but using the data that is necessary and appropriate to deliver accurate, fair, and trustworthy AI models.[45]

- **Safety and security**. The security principle obligates organisations to protect data confidentiality, integrity, and availability. For AI, security is not just a legal obligation but a business imperative: all AI models must be robust and trustworthy, in order to be commercially viable. Security measures themselves can lead to increased data volume and sensitivity, which must be weighed against its impact on the individual. That is a fine balance for organisations to strike, and it must evolve as threats to security and the technological state of the art change. Effective governance requires close cooperation between DPOs, chief information security officers and AI engineers. Properly implemented security and privacy should be seen as mutually reinforcing principles of trustworthy AI models and demonstrating robust security controls can serve as a tangible safeguard in the LIA, which helps tilt the balancing test in the organisation's favour.

  Safety, although not an explicit GDPR principle, is embedded in its risk-based approach and reinforced through other laws such as the AI Act and DSA. It relates to avoiding harmful or discriminatory outcomes, ensuring predictable and reliable system behaviour and mitigating risks to individuals. In practice, ensuring safety may make training AI on personal data, including sensitive data, essential as it can enable the AI to detect harmful content or unsafe information in the context of online safety or child protection. For example, foundation AI models that interact directly with individuals may need to be trained to identify and respond to users that demonstrate distress or may need additional support. The increasing use of AI technologies to provide safety must still be proportionate and carefully safeguarded, but it demonstrates how safety considerations often justify broader use of personal data.

  The type of AI technology can impact the safety or security risk, and change the safeguards and mitigations involved. Narrow use case models may have well defined organisational environments which can enable strict access controls or other controls, whereas general purpose models may be open to a wider variety of security attacks such as model inversion or prompt-injection attacks, and additional measures may be required.

  Both security and safety may expand the scope of processing but function as safeguards that reduce overall risk and strengthen the balancing test in favour of the organisation.

Contextual consideration: Sensitive data (from CIPL paper *Rethinking Sensitive Data in the Age of AI*[46])

An individual is communicating with a personalised therapy chatbot, the purpose of which is to provide mental health support, personalised therapeutic exercises, and cognitive behavioural therapy techniques to users. The AI powering the chatbot needs to ensure safety under both data protection and other content moderation and safety laws. They will need to process personal data, potentially including sensitive health data, in the training

stage of the AI to ensure that the AI is able to recognise harmful content, such as suicidal thoughts and suggestions. This could include: mental health conditions, emotional state, user's self-reported feelings, and chat transcripts. Such observations require a nuanced perspective broader than explicit assertions, which require an increased amount of data, including sensitive data.

Taken together, the process of weighing data protection principles with each other, in the context of AI training, illustrates the unique challenges raised by disruptive technologies to our existing legal interpretation. When interpreted flexibly and through a risk-based lens, the other principles do not create tension with privacy; they operationalise it. Accuracy and fairness prevent discrimination, minimisation ensures proportionality, and security and safety mitigate harm. The task for organisations is to demonstrate in their LIA how these principles interact, how risks have been addressed, and how the benefits to individuals and society justify the processing. This flexible, contextual reconciliation of principles underscores why LI is well-suited to AI training: it embeds accountability and requires organisations to confront and resolve these complex trade-offs transparently.

In a data-driven world, powered by AI, there is an urgent need to re-examine the application of some key data protection principles to AI technology and how these principles fit together in the ever-changing world.

Identifying the legitimate interest: risks and benefits

At the heart of the LIA is the requirement to identify a legitimate interest. Certain purposes are incompatible with the fundamental rights and values underpinning EU law, and cannot constitute a legitimate interest. For example, it can never be in an organisation's legitimate interest to use web-scraped child sexual abuse material (CSAM) to train a model that is used to generate synthetic pornography.

In identifying a legitimate interest, organisations must not only ensure that the purpose pursued is not incompatible with the essence of the right to privacy and data protection, but also recognise that these rights often need to be balanced against other fundamental rights and interests protected under EU law.[47] No single fundamental right is absolute. Recital 4 of the GDPR makes clear that data protection must be balanced against other rights, including the freedom of expression and information,[48] freedom from discrimination,[49] or freedom of thought and autonomy.[50]

Accordingly, an LIA should not only assess the risks to individuals, but also the *benefits* of the processing, both for the individual and for other third parties (or even society at large). In some contexts, not acting could itself produce harm, for instance, by entrenching bias, limiting fairness, or withholding life-saving innovation. Thus, benefits are not merely positive externalities; they can represent the avoidance of harm to other rights and interests.

In the context of AI training, these benefits can be significant. Training AI models on diverse datasets may be essential for identifying and mitigating bias, thereby advancing the right to non-discrimination leading to more accurate and inclusive services. At a societal level, AI trained on diverse datasets can help reduce systemic bias, enable scientific research, and drive innovation that supports economic and social development. The ICO has explicitly recognised that "the legitimate interests of the public" may form part of the balancing test, reinforcing that societal gains can legitimately outweigh limited impacts on individual privacy, where robust safeguards are in place.[51] In such cases, upholding privacy in absolute terms could undermine the protection and promotion of equally vital rights.

In practice, responsible organisations must approach data processing in a way that gives due weight to all relevant rights and interests, but this must be both precise and well-evidenced. Broad or non-specific arguments made at this stage about the specific model training will likely not satisfy this part of the test, unless these spillover benefits are tangible and evidenced. A well-reasoned, contextual balancing of rights ensures that the full spectrum of fundamental rights is respected and advanced in the digital age. In this vein, moving towards a taxonomy of benefits, as well as a taxonomy of harms, could be a useful resource for organisations seeking to emphasise these effects in their LIA.

A forward-looking LIA should also consider the risks of inaction. Failing to process data may withhold opportunities for individuals and society, such as life-saving technologies or tools that promote fairness and accessibility. Explicitly documenting these potential negative consequences strengthens the legitimacy of the interest pursued and supports the proportionality analysis within the LIA.

Clear, harmonised guidance from regulators on how to incorporate such forward-looking assessments can help organisations make confident, well-reasoned decisions and reduce regulatory uncertainty, when relying on the legitimate interests basis.

By explicitly weighing individual and societal benefits alongside risks, the LIA ensures a more complete assessment of competing interests. This structured balancing underscores why LI is a particularly suitable basis for AI training: it embeds accountability, while enabling innovation that can deliver value not only to organisations, but to individuals and the society as a whole.

Understanding and measuring risk and harm

The LIA is a real example of a risk-based approach that has been embedded in the GDPR, as it enables organisations to assess not only the benefits of AI training, but also the risks of harm to individuals. This is one of the most complex aspects of the LIA, as the GDPR provides no common methodology for identifying or quantifying potential harms, and regulators apply different approaches. In the context of AI training, risks and harms are dependent on factors such as the type of AI involved, the source of the data or the context of the processing. Some consideration and calculation of risks and harms are present in other provisions of GDPR, such as in data protection impact assessments (DPIAs), assessing potential personal data breaches, and building and implementing accountability programs. Beyond the GDPR, the AI Act requires a Fundamental Rights Impact Assessment, which also implies carrying out a risk assessment similar to DPIA and LIA. There will be an increasing need to industrialise and integrate all these processes of risk assessment in a repeatable and measurable process.

A first challenge is the absence of an agreed taxonomy of potential harms. Data protection harms are subjective and often difficult to identify and quantify in a repeatable and consistent process. Organisational risk management guidance from DPAs provides a good starting framework for understanding data protection harms, such as the ICO's taxonomy of harms,[52] and CIPL has long argued for an overarching agreement across regulators on the taxonomy of harms and the framework for assessing their likelihood and severity.[53] While technologies often evolve too rapidly for a fixed taxonomy of harms, developing a shared, principles-based understanding remains valuable. A common framework can evolve alongside innovation, offering flexibility while promoting consistent, risk-based decision-making.

In practice, organisations often rely on structured tools such as risk scoring, whereby severity and likelihood are assigned a numerical value to enable comparison across processing activities. While such methods promote consistency, it can often be hard to capture and measure risk in strict numerical terms.

Mitigations and safeguards

The effectiveness of LI depends on the presence of meaningful safeguards that reduce risks to individuals. These measures ensure that the flexibility afforded by LI does not compromise individual rights, and they demonstrate accountability by showing that risks have been identified, understood, and addressed throughout the AI lifecycle.

Safeguards can be technical, organisational, or procedural and will vary depending on the nature of the AI system. They typically include measures such as robust governance, access controls, transparency, data quality checks, and PETs. What matters is that they are proportionate to the risks created by the training activity and reflect the specific characteristics of the model and the data sources involved.

Different AI technologies present different risk profiles and therefore require different types of safeguards. Broad, foundation models trained on large-scale third-party datasets may require stronger upstream controls, such as dataset filtering, quality screening, and provenance checks, given the size and heterogeneity of the data and the absence of a direct relationship with the individual. By contrast, narrow or sector-specific AI systems, typically trained on curated first-party datasets, may rely more heavily on organisational safeguards, such as stricter access controls, clearer purpose limitation, and targeted retention practices.

Safety-critical or high-risk systems may require additional layers of testing, monitoring, and human oversight, while systems that rely on sensitive personal data will require enhanced protections regardless of their type. These differences illustrate the need for safeguards to be tailored and proportionate: a foundation model trained on open-web data will require extensive dataset filtering and transparency around training sources, whereas a sector-specific customer-service model may require more focused controls, such as role-based access, robust retention limits, and fairness testing relevant to its task.

In all cases, the appropriate combination of safeguards will materially influence the LI balancing exercise by reducing risks to individuals and demonstrating responsible, accountable model development.

Transparency

Transparency is not only a cornerstone data protection principle, but is also a recognised safeguard and, as such, is intimately linked with the balancing test. Meaningful information about how personal data is used for AI training helps align processing with the reasonable expectations of individuals. Transparency also supports the balancing test by mitigating the risk to individuals, informing them about processing, such as the right to object on justifiable and specific grounds.

AI developers should be transparent with individuals and the public, when they are processing personal data for AI training purposes and bring the people on the journey of accepting and adopting AI tools with meaningful user-centric transparency, contributing to continued digital trust when individuals understand the value and benefits of the technology. The type of transparency required will depend on the type of AI and source of the data involved, as direct transparency is most feasible for AI models trained on first-party data, as there is a direct link between the individual and the organisation. For AI models trained on third-party data, transparency is more likely to be published through higher level means, including public notices, because individual notification is often impossible.

Regulators are increasingly emphasising the need for improved transparency and, in particular, where the AI is trained on third-party data and, as such, there is no direct relationship with the individual. They often find that basic transparency has not been met with respect to organisations relying on LI to use third-party data to train their AI models.

Where organisations train AI models using third-party data, Article 14 GDPR allows for an exemption to providing transparency where such information would involve a disproportionate effort, provided that appropriate safeguards, including making the information publicly available, are implemented. A proportionate approach to Article 14 should focus on the principle of meaningful transparency, rather than exhaustive individual notification.

However, there needs to be more discussion and constructive engagement between regulators, industry and user interface experts to agree on the appropriate level of and mode of delivery of transparency that is required for AI model training.

Weighing interests in the balancing test

An LIA requires organisations to weigh their legitimate interest against the impact on individuals' rights and freedoms.[54] The GDPR outlines a set of data protection principles, but does not specify how to reconcile them when

they interact with each other or other fundamental rights and organisations are required to assess these on a case-by-case basis.

The challenge for DPOs is that the GDPR does not provide a methodology for reconciling these principles, when they appear in tension. Moreover, individual DPAs interpret and apply these principles differently. Moving towards a more standardised way of weighing interests and rights when applying the GDPR in practice, especially in relation to reoccurring trade-offs in the context of AI training, would be useful.

Although it is considered that tensions arise between data protection principles in the context of AI training, or between those principles and other legitimate interests, such as innovation or efficiency, in practice they function as mutually reinforcing pillars of trustworthy AI training and development. When applied contextually and interpreted through a risk-based lens, these principles do not conflict with the right to privacy, but instead serve to operationalise it. Ensuring accuracy promotes fairness and reduces harm; robust security protects the confidentiality and integrity of personal data; and data minimisation strengthens accountability, while supporting proportionality. For AI systems these principles should not be understood as simple trade-offs, but as interdependent safeguards that, collectively, enable lawful, ethical, and socially beneficial innovation.[55] However, this requires a flexible and forward-thinking interpretation of the principles that both protect fundamental rights and also promote innovation.

Reasonable expectations of individuals

The balancing test also depends on whether the processing is within the reasonable expectations of individuals. These expectations are not static and an individual's reasonable expectations will also evolve over time, as their understanding and acceptance of technologies and related data uses changes. Technological evolution is at an inflection point, where society is grappling with innovative data processing activities that were previously unexpected and unforeseen.

AI use cases and their ensuing benefits have become more commonplace in individuals' everyday lives and awareness of the technology is increasingly demonstrated and examined in the media and broader culture. DPOs should assess expectations in context, taking into account the product or service, whether it is well-established, and the degree of public debate around AI. For historic data or less transparent uses, stronger mitigations (enhanced notice, opt-outs, or targeted communications) may be needed to bring processing within reasonable expectations.

Digital literacy initiatives, as reinforced by Article 4 of the AI Act, will play an important role in raising awareness of users, customers, employees and the general public about data use and the important role of data in the modern economy and society, and in transformative technologies, such as AI.

The assessment of the reasonable expectations of individuals may vary depending on the relationship the organisation has with the individual. For AI training on third-party data, the organisations are less likely to be able to directly communicate information with the data subject and, therefore, there may be a higher difficulty associated with transparency. As emphasised by the EDPB in their AI Opinion, this is not a "blanket exemption" for large scale processing, and organisations must be able to demonstrate why individual notification is disproportionate.

Organisations training AI on third-party data may seek to make the transparency information publicly available, if it would require disproportionate effort to deliver to the individual. Datasets containing first-party data are less likely to be able to rely on this provision, due to the close relationship with the individual.

LI and wider risk assessments

Given the number of risk assessments to be performed under an increasing body of laws, including for an LIA, there is a real need to harmonise these processes. This will be essential in order to ensure repeatability and consistency of the process, to reduce compliance burden on business teams, and to take away the subjective, human variables

in conducting the assessments. At the same time, organisations will need to balance the need for operational efficiency in carrying out risk and impact assessments, with the need to preserve human review and escalations for high-risk processing and use of data. Any tools, to be effective, need to be flexible enough to be context-specific and not be solely reliant on a quantitative measurement. Regular reviews and updates to the scoring methodology will also be required to maintain its relevance and accuracy.

Table 1: Risk assessment in the EU's digital framework

Law	Type of Assessment	Relevant Articles	Scope
GDPR	Data Protection Impact Assessment (DPIA)	Art. 35 GDPR	Required where processing is "likely to result in a high-risk" to individuals' rights and freedoms (e.g., large-scale special category data, profiling, systematic monitoring).
	LIA (*not formally in GDPR, but practice under Art. 6(1)(f)*)	Art. 6(1)(f) GDPR, Recital 47	Structured balancing test: purpose, necessity, balancing. Often documented as an LIA.
AI Act	Fundamental Rights Impact Assessment (FRIA) for certain AI systems	Art. 29a (for deployers of high-risk AI in public sector and certain private contexts)	Must assess impact on fundamental rights, including risks of discrimination, bias, and exclusion.
	Risk Management System (continuous risk assessment & mitigation)	Art. 9 AI Act	Obligates providers of high-risk AI to establish a risk management framework across lifecycle.
Digital Services Act	Risk Assessments by Very Large Online Platforms (VLOPs) and Very Large Online Search Engines (VLOSEs)	Art. 34 DSA	Assess and mitigate systemic risks, including disinformation, electoral integrity, protection of minors, public health and safety. Must be updated annually.
	Independent Audit of Risk Management	Art. 37 DSA	External audits to verify compliance with systemic risk assessments and mitigation measures.
NIS2 Directive	Cybersecurity risk management measures	Art. 21 NIS2	Requires entities to assess cybersecurity risks and adopt appropriate technical and organisational measures.

The multiplication of risk assessments across various digital disciplines and compliance areas will make it imperative for organisations to consider an integrated and holistic approach to risk assessment, as well as harmonising this process. There is a need to ensure cohesion between these risk assessments, which may in some cases be happening independently of each other. Pre-existing and concurrent risk assessments should be considered and leveraged where appropriate. Such a holistic approach ensures a consistent and harmonised approach to risks and fundamental rights across the business.

Part 2: How to conduct an LIA for AI training

Part 2 of this paper outlines the step-by-step process of conducting an LIA, drawing on the best practices shared by DPOs in the DPO Forum @CIPL. Its purpose is to provide practical guidance that any organisation can apply when seeking to rely on LI as the legal basis for AI training.

However, conducting an LIA is a highly contextual exercise that will depend on the interest pursued and the data and AI involved. For instance, large data-driven organisations may take a more resource intensive approach, which may be less feasible for smaller organisations or start-ups.

Potential considerations before you start the LIA

There are a number of considerations that could be useful before the LIA is formally started.

- **An LIA is a living document**. An LIA should be approached as a dynamic process rather than a one-off document. It should be revisited if circumstances change, with versioning and monitoring built in to ensure continued relevance. Templates and processes could also evolve with regulatory guidance and case law.

- **An LIA is a risk-based process**. The depth of the assessment should be proportionate to the potential risk involved in the AI training, with more extensive analysis required where the processing poses a higher risk to individuals.

- **LIAs are often cross-disciplinary procedures**. Whilst the DPO office plays a significant role in providing guidance and support in an LIA, it is not solely responsible for the process. The responsibility for conducting the LIA may be split between multiple compliance functions. The extent to which it is involved can depend on the significance of the processing. The responsible parties will likely involve or escalate high-risk processing activities, if they are identified, to the DPO office.

I. Scope and facts

Before beginning the LIA, it may be useful to understand the scope of the AI activity and the key facts.

- Define the AI training objective, lifecycle stage and deployment context.

- Carry out data mapping, including source, categories, volumes, access, locations/transfers.

- Identify other relevant information and design well-informed, targeted questions.

- Identify stakeholders across the organisation and establish efficient methods for finding and consolidating relevant information.

II. Legal bases review

Organisations should consider the legal bases under Article 6 GDPR, considering why each is suitable or not. This should not be treated as a tick box exercise; it is an opportunity to ensure legal and operational alignment and to spot subtle differences which could ultimately affect whether LI can be relied upon.

- Carefully consider the most appropriate legal basis and, where appropriate, document why other legal bases are not suitable to rely on. Due to the novelty of the data processing involved in AI training, the consideration of why alternative legal bases are not appropriate is crucial, in order to set a precedent and inform future AI development.

- Documentation could support both internal governance and accountability, as well as provide external stakeholders with confidence that the assessment has been carried out in a thorough, transparent, and responsible manner.

- Documentation should follow the risk-based approach. Where processing carries more risk or is more impactful, the documentation will likely be more evolved or complex.

- Final determination to rely on legitimate interests can only be made once the LIA is completed.

- If the processing is considered to present a high-risk to individuals, a DPIA will be required to be conducted in tandem, and the processing will be escalated to the DPO office. Even if an organisation doesn't consider a DPIA at this stage, the outcome of the LIA may trigger the requirement for a DPIA to be subsequently carried out. Alternatively, organisations may choose to carry out a full-scale DPIA instead, in which case a separate LIA may not be required.

- It may be useful to establish a timeline for completing the assessment that is flexible enough to allow for the possibility of identifying an alternative legal basis, if the LIA yields negative results.

Stage 1: Purpose test

As the first stage of the LIA, organisations are required to identify the legitimate interest they are pursuing and demonstrate that it is specific, real, and lawful.

- Demonstrate that the processing amounts to a legitimate interest. This interest can be the interest of the controller, third-party, society, or group of individuals.

- LI often includes everyday, routine and established business practices[56] such as fraud prevention, cyber security and direct marketing (Recital 47 GDPR). In the context of AI models, an LI may include developing the service of a conversational agent to assist users, developing an AI system to detect fraudulent content or behaviour or improving threat detection in an information system.[57]

- Identify benefits to demonstrate that the processing is in the pursuit of a legitimate purpose, and will contribute to the balancing of these benefits with the risks during a later stage. This can be a commercial purpose,[58] a benefit to society, a third-party, or an individual interest. In the context of AI models, ensuring AI models are trained on diverse and representative data can bring broader public benefits. Organisations may want to identify a taxonomy of benefits, as in the example below, to better inform their risk assessments.

- Analyse what the impact would be if they did not go ahead with the processing, and as part of this consider the 'loss of opportunity' to society.

Contextual consideration: Taxonomy of benefits

<u>Benefits to the individual</u>

Improved accuracy: AI systems can provide more reliable and precise results, reducing the likelihood of errors that inconvenience or disadvantage individuals. For example, users experience fewer false positives in content moderation or more accurate assistive technologies.

Enhanced accessibility and inclusion: AI models can support equal access to products, services, and information by accommodating diverse languages, accents, and disabilities. For example, improved speech

recognition or accessible digital interfaces, such as visual accessibility features for low-vision or blind individuals, providing practical assistance to allow more people to fully participate online.

Increased personal safety and security: AI can help protect individuals from harm, fraud, and exposure to unsafe or misleading content. For instance, more effective detection of scams, impersonation, or harmful material enhances users' sense of security.

Improved fairness: AI systems trained on representative datasets can support more equitable and consistent decisions, reducing bias or discriminatory outcomes. This means individuals are treated fairly across automated processes, such as recruitment or service access.

Benefits to Society

Advancement of scientific and technical knowledge: AI training can contribute to collective innovation and scientific understanding across sectors such as medicine, environmental research, and linguistics, supporting social and economic progress.

Enhanced public safety and resilience: AI systems that detect and predict threats can improve responses to public health, environmental, or security challenges, strengthening societal resilience.

Economic growth and productivity: The use of AI in products and services can drive innovation, support new industries, and improve productivity across the economy, benefiting workers and consumers.

Improved civic engagement and information access: AI that summarises, translates, or simplifies complex information can help individuals engage more easily with public institutions and democratic processes.

Advancement of education and skills: AI training can support the development of personalised learning tools, adaptive educational systems, and accessible content, improving educational outcomes and promoting lifelong learning opportunities across diverse communities.

Preservation and promotion of linguistic and cultural diversity: Training AI models on diverse languages and cultural data can help preserve underrepresented languages, improve translation accuracy, and strengthen access to information for minority communities, supporting inclusivity in the digital sphere.

Social and environmental sustainability: AI trained on diverse datasets can identify patterns that inform more sustainable practices, from optimising energy use to advancing environmental research, supporting broader societal goals aligned with public interest and sustainability.

Stage 2: Necessity test

As the second stage of the LIA, organisations must assess whether the processing of personal data is necessary and proportionate to achieve the stated purpose, and confirm that no less intrusive means are reasonably available. If the processing is not necessary, LI cannot be relied upon.

- Show the personal data is necessary and proportionate for the purpose of the legitimate interest pursued. This determination is context specific and will vary considerably case-by-case, with high-risk processing requiring more scrutiny than low-impact uses. As time progresses, teams may build up a repository of evidence to help demonstrate that the processing is necessary for the stated purpose.

- Consider whether the processing activity will allow the pursuit of the purpose.[59]

- Consider whether there is a less intrusive way of pursuing this purpose, taking into account the data minimisation principle,[60] for instance, through the use of PETs/PPTs[61] such as federated learning, sampling, pseudonymisation, or differential privacy. The *EDPS v SRB* judgement confirmed that personal data is relative, rather than absolute. In AI training, data transformed or aggregated or pseudonymised before being shared with a model developer may not be personal data, if the recipient has no realistic means of re-identifying individuals using information.[62]

- Consider whether the data processing is proportionate by aligning safeguards to the expected risks and benefits.

- The CJEU is moving towards a strict interpretation of necessity, which includes the processing being indispensable or strictly necessary for the purpose. However, such a strict standard may not be workable in the case of training AI models. For instance, what is 'less intrusive' may not be able to be demonstrated without first conducting the processing.

- Training robust AI models often requires large volumes of representative, real-world data. While synthetic data or fully anonymised data may reduce privacy risks, such datasets may not provide the diversity, context, or nuance needed to train accurate AI models. In many circumstances the processing of such personal data will be necessary to train accurate, safe and secure AI models, as it could not be achieved through other means.

- In the context of AI training, the DPO will likely engage with product teams or technologists to understand the necessity of using personal data. It is important that these teams clearly articulate and evidence why the data is required for the intended purpose. Providing a well-reasoned justification supports the DPO office in fulfilling its oversight responsibilities and reinforces the organisation's accountability in line with data protection principles.

Contextual consideration: Where the necessity test could fail in the context of training AI

A multinational manufacturer uses an AI system to predict supply-chain disruptions and optimise production schedules.

To improve its accuracy, the company decides to train the AI on supplier-level performance data, which includes supplier contact information and key employee names (first-party data obtained through procurement systems), delivery timelines and delay patterns, and publicly available company credit ratings (third-party data).

The stated legitimate interest is improving operational efficiency and avoiding stock shortages through better forecasting.

The necessity test could fail here because:

1. Personal data is not needed for the stated purpose, as the AI model's goal is to forecast delivery reliability and demand fluctuations which could be achieved through anonymised supplier identifiers rather than named individuals.
2. The training dataset includes communications metadata, which is disproportionate to the stated purpose.
3. Less intrusive alternatives are available, including the use of aggregated delivery records and historical performance data stripped from personal data.

Stage 3: Balancing test (risks, transparency, wider safeguards, right to object)

At the final stage of the LIA, organisations are required to determine whether the interests or fundamental rights and freedoms of the individual outweigh the identified legitimate interest of the organisation, or a third-party in processing the personal data in question. If the former is the case, the organisation would be unable to rely on the LI legal basis.

I. **Risks and harms**

- Identify the potential risk to individuals. Consider the likelihood that each risk will occur and the severity of the impact if it does. Risks include those relating to privacy and confidentiality, fairness and discrimination, individual autonomy and expectations, security, and safety. The more 'sensitive' or private the personal data the more it is likely to present a risk to the individual. Risks may also include wider economic or social harms, such as loss of opportunity or exclusion from benefits including privacy/confidentiality, fairness, discrimination, security, and safety.

- Understand whether the individuals impacted by the processing include children, vulnerable groups, or minorities.

- If the processing is likely to result in a high-risk to the individual then a full scale DPIA is required.

II. **Contextual factors**

- Take into consideration the context in which the personal data was collected and is now being processed, as the likelihood and severity of the impact can be influenced by the nature of the personal data, the context of the processing and any further consequences that the processing has.

- Assess the nature of the relationship between the individual and the organisation, including whether the data was originally public. In the context of AI training, first-party data collected directly from the individual may carry different expectations than third-party data.

- Ensure processing is also within the reasonable expectations of individuals. This is an objective test which requires the processing to be within the reasonable expectations of the average individual, and not any particular individual.

- This also includes a consideration of the wider context of the processing (for instance the context and relationship between the individual and the organisation, the nature of the product or service and the context or source of the data collected).[63]

- Although a contextual analysis will always be necessary, if the processing is listed in transparency documentation that has been made publicly available to individuals, it may be more likely to be within the reasonable person's expectations, as emphasised by the EDPB.[64]

III. **Safeguards and mitigations**

- Where risks or harms are identified, implement measures to reduce or eliminate the risks.

- Consider whether there are already safeguards in place, and to what extent they are effective.

- Transparency is a key safeguard to mitigate the risks posed by AI, as well as demonstrating strong accountability practices. This includes clear and accessible information about how data is being used in AI training.

- Provide individuals with ways to exercise their rights, including the right to object. The right to object is not an absolute right and allows individuals to object to processing, carried out under LI, unless the controller demonstrates a compelling legitimate ground, which overrides the individual's interest.

- Consider use of PETs/PPTs as risk mitigation measures, especially where risks cannot be mitigated by other means.[65] This should also be done in a proportionate and risk-based manner, appropriate for the actual risk. Examples in the context of AI training may include strict data access, limitations on categories or sources of data, data use limitations, security measures, retention schedules, data minimisation (including anonymisation and pseudonymisation),[66] signing non-disclosure agreements, making documentation available in secured platforms with 'view only' access, using filters to restrict access to private information, and internal access controls.

- Apply organisational controls, including retention limits, audit logging, vendor due diligence, and review cycles.

- Build safety and security into the design of AI models.

> Contextual consideration: safeguards for AI training
>
> In order to develop an accurate cancer detection model, large amounts of data are required to be collected from multiple hospitals and medical centres. However, pathology images are highly sensitive and cannot necessarily be freely shared, without introducing privacy and security challenges[67].
>
> Through the use of federated learning and trusted execution environments in combination, healthcare institutions are able to securely share their data to jointly train a model. Federated learning allows model training to take place locally, with model updates aggregated and sent to a global model. A trusted execution environment is used to provide further security by encrypting model updates during transmission, thus preventing unauthorized parties from intercepting and accessing the data.

> Contextual consideration: transparency as a mitigation in AI training
>
> An organisation operates an online retail platform and collects customer messages and feedback to provide customer support. It intends to use those interactions to train and improve an AI model that supports its virtual assistant.
>
> When a user creates an account, a message appears to inform them that their personal data may be used to help train an AI model.
>
> The notice has a second layer which links to more detailed privacy information. The link also directs users to the privacy control centre where they can exercise their data subject rights.

Example of Safeguards		
Organisational	Technical	Procedural
• Strong data governance	• Input/output filtering	• Transparency

framework ● Data Protection by Design and Default	● Anonymisation or pseudonymisation ● PETs/PPTs such as differential privacy, federated learning, secure multi-party computation ● Encryption or trusted data environments ● Data security	● Right to object ● Storage limitation

IV. **Loss of opportunity consideration**

Consider the risks of not processing by documenting what individuals and society may lose if AI training does not succeed. For example, withholding training data may prevent improvements in model offering, accuracy, fairness, accessibility, or safety. Explicitly recording these opportunity costs helps demonstrate that the balancing test has been applied comprehensively.

While individuals retain the right to object to processing based on legitimate interests, this right is not absolute. Under Article 21 GDPR, controllers may continue processing where they can demonstrate a compelling legitimate ground, which overrides the individual's interest.[68] An objection should be carefully evaluated against the broader opportunity and benefit that the processing seeks to achieve, for example, improvements in safety, fairness or accessibility. Controllers should document not only how they assess and respond to objections but also how the potential harm, if rejecting an objection, is balanced against the wider legitimate interests in continuing the processing.

After the LIA: ongoing monitoring and review

Regular monitoring is a core component of organisational accountability and can help organisations verify that data processing continues to meet the conditions under which it was initially justified. This also allows for the early detection of new risks presented by the processing itself, as well as new risks presented by the evolving nature of data processing in the advancing technological era.

The frequency and depth of post-LIA monitoring is determined on a case-by-case basis and depending on the level of risk involved. Furthermore, regulators are increasingly expecting evidence of ongoing risk assessments and not just pre-launch analysis. If there is a corresponding DPIA, monitoring and review may be conducted in a combined process.

Bibliography

[1] European Commission, 'The Future of European Competitiveness: A Competitiveness Strategy for Europe' (Draghi Report), 2024, https://commission.europa.eu/document/download/97e481fd-2dc3-412d-be4c-f152a8232961_en.

[2] Brazil launched its National Strategy for Artificial Intelligence (EBIA) in 2021 and the Brazilian Artificial Intelligence Plan (PBIA 2024–2028) to promote responsible AI, innovation and competitiveness; China's New Generation Artificial Intelligence Development Plan (2017) aims to make China a global AI leader by 2030, supported by subsequent AI governance frameworks; and the United Arab Emirates adopted its National Strategy for Artificial Intelligence 2031 to integrate AI across government and industry.

[3] Examples include the EU General Data Protection Regulation (Regulation (EU) 2016/679) (GDPR); Singapore's Model Artificial Intelligence Governance Framework (2020); China's Interim Measures for the Management of Generative Artificial Intelligence Services (2023); and Canada's Directive on Automated Decision-Making (2019).

[4] The Centre for Information Policy Leadership (CIPL) is a global privacy and data policy think tank within the Hunton law firm that is financially supported by the firm, 85+ member companies that are leaders in key sectors of the global economy, and other private and public sector stakeholders through consulting and advisory projects. CIPL's mission is to engage in thought leadership and develop best practices for the responsible and beneficial use of data in the modern information age. CIPL's work facilitates constructive engagement between business leaders, data governance and security professionals, regulators, and policymakers around the world. For more information, please see CIPL's website at www.informationpolicycentre.com. Nothing in this document should be construed as representing the views of any individual CIPL member company or Hunton. This document is not designed to be and should not be taken as legal advice.

[5] CIPL Paper, 'How the "Legitimate Interests" Ground for Processing Enables Responsible Data Use and Innovation', 2021; CIPL Paper, 'Legitimate Interest Paper - CIPL Examples of Legitimate Interest Grounds for Processing of Personal Data', 2017; CIPL Paper, 'Transparency, Consent and Legitimate Interest Paper - Recommendations for Implementing Transparency, Consent and Legitimate Interest under the GDPR', 2017.

[6] CIPL, 'Artificial Intelligence and Data Protection in Tension', 2018.

[7] CIPL, 'Applying Data Protection Principles to Generative AI: Practical Approaches for Organizations and Regulators' (Discussion Paper), 2024.

[8] EDPB, 'Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models', 2024, https://www.edpb.europa.eu/system/files/2024-12/edpb_opinion_202428_ai-models_en.pdf.

[9] ICO, 'Consultations Series on Generative AI and Data Protection', 2024, https://ico.org.uk/about-the-ico/ico-and-stakeholder-consultations/2024/09/ico-consultation-series-on-generative-ai-and-data-protection/.

[10] Other jurisdictions around the world have published key guidance on the interplay between AI and data protection, including: South Korea PIPC 'Personal Information Processing Guide for the Development and Use of Generative AI', 2025; Brazil ANPD, 'Preliminary Study on Generative AI', 2025; Singapore IMDA, 'Advisory Guidelines on the Use of Personal Data in AI Recommendation and Decision Systems', 2024.

[11] Article 6(1)(f) GDPR outlined that processing necessary for the legitimate interests of the controller or a third-party can be processed lawfully, except where such interests are overridden by the interests of fundamental rights of individuals.

[12] Article 5, Article 6 GDPR.

[13] AI Opinion, supra n.8.

[14] As echoed in the AI Opinion, supra n.8.

[15] The three-part Legitimate Interests Assessment originates in CJEU case law interpreting Article 7(f) of the Data Protection Directive (see, e.g., C-465/00 Österreichischer Rundfunk, 2003; and C-13/16 Rīgas satiksme, 2017), but was codified as a practical methodology in WP29 Opinion 06/2014.

[16] CJEU, judgment of 4 July 2023, Case C-252/21, Meta v. Bundeskartellamt para. 106; CJEU, judgment of 11 December 2019, Case C-708/18, Asociaţia de Proprietari bloc M5A-ScaraA, paragraph 40. See also EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0 adopted on 8 October 2024, paragraph 12 and ff.

[17] Elsewhere regulators are actively engaging on this issue. For instance, Brazil's National Data Protection Authority (ANPD), in its recent consultation on AI and data protection, is exploring the use of legitimate interests as a potential legal basis for AI training.

[18] For instance, the ICO states that "Legitimate interests remains the sole available lawful basis for training generative AI models using web-scraped personal data based on current practices", supra n.9.

[19] ICO, Ibid.

[20] EDPB Opinion, supra n.8.

[21] ICO, supra n.9.

[22] The EDPB has sought to clarify this in their most recent guidelines on legitimate interests, explicitly stating that legitimate interests should not be seen as a last resort legal basis, European Data Protection Board, 'Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR', 2024, https://www.edpb.europa.eu/system/files/2024-10/edpb_guidelines_202401_legitimateinterest_en.pdf.

[23] ICO, supra n.9.

[24] See CIPL's report on individual rights in AI (forthcoming, Fall 2025).

[25] CIPL, 'Rethinking Sensitive Data in the Age of AI', 2025.

[26] European Commission, 'Digital Omnibus Regulation Proposal', 2025, https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-regulation-proposal

[27] CIPL, 'The Limitations of Consent as a Legal Basis for Data Processing in the Digital Society', 2024.

[28] EDPB, 'Law & Compliance in AI Security & Data Protection: Training curriculum on AI and data protection (Support Pool of Experts, delivered by Dr Marco Almada)', 2024, https://www.edpb.europa.eu/system/files/2025-06/spe-training-on-ai-and-data-protection-legal_en.pdf.

[29] CNIL, 'Ensuring the lawfulness of the data processing – Defining a legal basis' (AI How-To Sheet 4), 2024, https://www.cnil.fr/en/ensuring-lawfulness-data-processing-legal-basis.

[30] EDPB Supra n.25.

[31] European Data Protection Board, 'EDPB: cookie pledge initiative should help protect fundamental rights and freedoms of users', 2023, https://www.edpb.europa.eu/news/news/2023/edpb-cookie-pledge-initiative-should-help-protect-fundamental-rights-and-freedoms_en.

[32] CNIL, 'Relying on the legal basis of legitimate interests to develop an AI system', 2024, https://www.cnil.fr/en/relying-legal-basis-legitimate-interests-develop-ai-system.

[33] Meta Platforms Inc and Others v Bundeskartellamt.

[34] CIPL, ibid, n.25.

[35] Data protection accuracy can be distinguished from model efficacy (or statistical accuracy) in AI, which measures the accuracy of the outputs of the AI system in relation to the expectation. To comply with the principle of data protection accuracy, the outputs do not need to be 100% accurate, but model efficacy is one part of meeting the obligations under this principle. ICO, 'What do we need to know about accuracy and statistical accuracy?', https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/what-do-we-know-about-accuracy-and-statistical-accuracy/.

[36] ICO ibid; AEPD, 'Artificial Intelligence: accuracy principle in the processing activity', 2023, https://www.aepd.es/en/prensa-y-comunicacion/blog/artificial-intelligence-accuracy-principle-in-processing-activity.

[37] EDPB, 'Report of the work undertaken by the ChatGPT Taskforce', adopted 23 May 2024, https://www.edpb.europa.eu/system/files/2024-05/edpb_report_chatgpt_taskforce_en.pdf; EDPB, 'Guidelines 4/2019 on Article 25 Data Protection by Design and by Default', Version 2.0, adopted 20 October 2020, https://www.edpb.europa.eu/system/files/2020-10/edpb_guidelines_201904_dataprotection_by_design_and_by_default_v2.0_en.pdf; Article 29 Working Party, 'Guidelines on transparency under Regulation 2016/679', revised and adopted 11 April 2018, endorsed by the EDPB on 25 May 2018, https://ec.europa.eu/newsroom/article29/items/622227/en.

[38] Sensitive data, or Special Category Data, requires organisations to identify a condition under Article 9 as well as a legal basis under Article 6 GDPR.

[39] CIPL, supra n.25.

[40] Article 10(5) AI Act.

[41] European Parliament, 'Algorithmic discrimination under the AI Act and the GDPR', 2025, https://www.europarl.europa.eu/thinktank/en/document/EPRS_ATA(2025)769509.

[42] CIPL, Reconciling AI with the Data Minimization Principle: Bridging the Innovation and Privacy Gap, 2025.

[43] CIPL, supra n.7.

[44] EDPB, supra n.8, paragraph 64.

[45] CIPL, supra, n.39.

[46] CIPL, supra n.25.

[47] EU, 'Charter of Fundamental Rights of the European Union', 2012, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A12012P%2FTXT.

[48] Ibid, Article 11.

[49] Ibid, Article 21.

[50] Ibid, Article 10.

[51] ICO, 'What is the "legitimate interests" basis?', https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/legitimate-interests/.

[52] ICO, 'Theory of harm in data protection', [n.d.], https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/theory-of-harm-data-protection.

[53] CIPL, 'Risk, High-risk, Risk Assessments and Data Protection Impact Assessments under the GDPR', 2016.

[54] EDPB Opinion, supra n.8, paragraph 61.

[55] CIPL supra n.7.

[56] The UK has gone further to recognise instances that amount to a legitimate interest by the introduction of a new legal basis 'recognised legitimate interests' and has provided a further list of processing activities likely to pass the purpose part of the LIA.

[57] EDPB, supra n.8, paragraph 69.

[58] CJEU, Case C-621/22 (Koninklijke Nederlandse Lawn Tennisbond v Autoriteit Persoonsgegevens), in which the Court held that processing personal data for a controller's purely commercial interest may, in principle, be justified under Article 6(1)(f) GDPR, provided that the three-step interest test is satisfied.

[59] CJEU, Case C-524/06 Heinz Huber v Bundesrepublik Deutschland, 16 December 2008, paragraph 66.

[60] CJEU, Case C-73/16 Peter Puškár, 27 September 2017.

[61] CIPL, 'Privacy-Enhancing and Privacy-Preserving Technologies: Understanding the Role of PETs and PPTs in the Digital Age', 2023.

[62] However, the threshold remains high: pseudonymisation alone does not automatically remove data from the GDPR's scope. Controllers must be able to demonstrate that re-identification would be practically or legally impossible using means reasonably likely to be employed by the model provider.

[63] EDPB Legitimate Interests Guidelines.

[64] EDPB AI Opinion, paragraph 94.

[65] CIPL, 'Privacy-Enhancing and Privacy-Preserving Technologies in AI: Enabling Data Use and Operationalizing Privacy by Design and Default', 2025.

[66] CIPL, supra n.39.

[67] Duality Technologies, 'Secured Collaborative AI for Oncology Research', https://dualitytech.com/use-cases/healthcare-industry/collaborative-ai-for-oncology-research/.

[68] Article 21 GDPR provides a qualified right to object for individuals if an organisation relies on legitimate interests. The organisation would have to stop processing the personal data, unless they can demonstrate a compelling legitimate ground for the processing, which overrides the interests, rights, and freedoms of the individual.