Centre for Information Policy Leadership
HUNTON

Reconciling AI with the Data Minimization Principle: Bridging the Innovation and Privacy Gap

December 2025

Contents

Reconciling AI with the Data Minimization Principle: Bridging the Innovation and Privacy Gap

CIPL makes the following recommendations for organizations and regulators to effectively operationalize data minimization in AI systems:

Adopt a contextual and flexible interpretation of data minimization: Organizations and regulators should apply data minimization proportionally, focusing on necessity, balancing risks and benefits, acknowledging AI-specific needs, tailoring measures across AI lifecycle stages and permitting socially beneficial secondary purposes

Implement a structured necessity test framework: Organizations should assess data needs using a practical, step-by-step framework that defines purpose, considers less intrusive alternatives, justifies data volume, scope, and categories, limits retention, and ensures accountability and review

Embed strong safeguards and accountability: Organizations should establish robust governance, cross-functional collaboration, and technical measures, including privacy-enhancing technologies, to responsibly manage and minimize personal data use

I. Introduction and Background

While artificial intelligence (AI) technologies are not new, the advent and uptake of generative AI (genAI) have prompted regulators and policymakers to renew their focus on their governance. In the context of privacy and data protection, this has sparked a debate on how data protection principles apply to AI, what new risks these systems may present, if any, and how to address them. At the same time, there is an increasing understanding of the economic and societal value of AI and a push to ensure its continued development.

This paper examines how data minimization should be interpreted in the AI context, building on the Centre for Information Policy Leadership's (CIPL) previous work on AI governance and organizational accountability.[1] While data minimization remains a central safeguard across global privacy frameworks, its purpose is to ensure that personal data is limited to what is necessary and proportionate for a legitimate aim – not to mandate data reduction at the expense of effectiveness or utility.

a. The Global Landscape of Data Minimization

In Europe, data minimization has long been a core principle of data protection law. The former Data Protection Directive (DPD) required that personal data be "not excessive" in relation to the purposes for which it was collected,[2] building on the OECD Data Protection Principles[3] and Council of Europe Convention 108[4]. The GDPR later refined this language, requiring that personal data be "limited to what is necessary," a linguistic rather than a substantive shift.[5] The GDPR standard does not forbid the use of large datasets where they are necessary for the purpose.

In the United States, data minimization has also emerged as a key area of focus in privacy policymaking in state legislatures and Congress. Some US state privacy laws impose a duty of data minimization, limiting the collection of personal data by businesses to what is adequate, relevant, and reasonably necessary in relation to the purposes for which such data is processed, as disclosed

to the consumer.[6] However, the State of Maryland has adopted – and other states and Congress have considered and, to date, rejected – a stricter standard of necessity: what is reasonably necessary and proportionate to provide or maintain a specific product or service requested by the consumer to whom the data pertains.[7]

Globally, data minimization is recognized as a core data protection principle, though its interpretation varies across jurisdictions:

- Brazil defines "necessity" as limiting processing to the minimum data required for the purpose, emphasizing proportionality and non-excessive processing.[8]
- In India, consent permits the processing of personal data strictly necessary for the specified purpose, and controllers must erase that data once the purpose has been achieved, unless retention is legally required.[9]
- Singapore describes data minimization as a good practice to reduce data protection and cybersecurity risks, encouraging organizations to collect and use only the personal data necessary to train AI systems,[10] and to delete such data once its purpose has been fulfilled.[11]
- Canada similarly requires that collection, use, and retention of personal data be limited to what is necessary for identified purposes, with the data deleted or anonymized once it is no longer needed.[12]
- Japan requires organizations to specify the purpose for which personal data is used and prohibits processing beyond what is necessary to achieve that purpose without the individual's consent.[13]

Together, these frameworks reflect a global consensus that data minimization entails both restricting unnecessary processing and ensuring purpose-bound, proportionate use of data to safeguard individuals' privacy.

b. Why it Matters for AI

Since their introduction to the market, genAI systems have seen widespread adoption by individuals and organizations around the world.[14] These systems rely on foundation models, also called general-purpose AI models. While all AI systems rely on data for training, foundation models require particularly large and diverse datasets. They are typically trained on a multitude of sources, such as publicly available data from the web, licensed data, and academic and industry datasets, to achieve a variety of purposes.

GenAI models are trained to recognize statistical relationships between words and other types of data, such as images, videos, and audio (i.e., datapoints), and make probabilistic predictions, often in response to user prompts. Some training data may qualify as personal data under many data protection laws. This raises questions about how principles such as data minimization apply, as personal data collected should be **adequate, relevant, and limited to what is necessary** for the intended purpose.

Beyond genAI, agentic AI often autonomously determines what data to process to achieve specific goals.[15] Such systems can pose particular challenges for data minimization, requiring appropriate measures to manage and mitigate these issues.

Strict limits on the collection, retention, or use of personal data, or requirements to remove it entirely, can conflict with what is required to build effective AI systems. An overly narrow interpretation of the data minimization principle and what data is "necessary" may ultimately limit the innovative potential of machine learning and the 'uptake of AI', a goal expressly articulated in the EU AI Act.[16]

Data minimization should instead be applied in a contextual and flexible manner rather than as a strict quantitative limit, focusing on the necessity and proportionality of data use in relation to legitimate purposes. A nuanced, risk-based application of data minimization can therefore reconcile innovation with privacy protection.

II. The Importance of the Data Minimization Principle

Data minimization is a foundational principle of responsible data governance.

Limiting data collection and retention helps protect individuals from unnecessary privacy intrusion, which is especially important in model training contexts where web-scraped data may contain sensitive or even illegal material, such as child sexual abuse material.[17] Limiting the volume and sensitivity of the data processed also significantly reduces the attack surface for cyber breaches.

Beyond risk mitigation, there are compelling business and social incentives in the context of data minimization, including reduced operational costs for larger data volumes, lower compliance burdens (e.g., fewer individual rights requests),[18] and reduced environmental costs associated with unnecessary storage and computation.

The objective of data minimization is to limit the unnecessary and excessive collection, processing, and retention of personal data. The focus of the data minimization principle is on what is *necessary for a legitimate purpose*, to limit the risks of misuse, unauthorized access, and data breaches (e.g., large datasets may be justified where they demonstrably reduce risks such as bias or improve fairness).[19] This principle protects individuals' privacy by ensuring more responsible data practices through preventing the collection of data as a "nice to have".[20]

When contemplating the data minimization principle in the context of AI, it is important to recall its purpose. Data minimization requires that personal data is adequate, relevant, and limited to what is necessary for the intended purpose.[21] Data minimization is not, however, about minimizing data for its own sake.[22]

Necessity further requires that processing proceed only where the objective cannot be achieved through less intrusive means.[23] If such alternatives exist, the processing cannot be deemed necessary.[24] Where the same result can be achieved with less personal data, for instance, then this processing would not be considered *necessary*.[25]

Necessity, therefore, links the processing of personal data to a specific and justifiable purpose. It does not inherently dictate a particular volume of data. A single item of personal data may already be excessive if it is not necessary for the intended purpose. Conversely, processing large volumes of personal data may be entirely appropriate, provided that each data point serves a legitimate and necessary purpose.[26]

Data minimization should not be viewed in isolation, but as part of a broader ecosystem of data protection safeguards that collectively prevent the unrestrained processing of personal data. Other

data protection principles establish boundary conditions that enable a more flexible and proportionate approach to data minimization.

In practice, this means embedding data protection by design and by default, assessing the compatibility of purposes, ensuring that every processing activity has a lawful basis (it is difficult to justify a legitimate interest in processing excessive amounts of data), implementing pseudonymization wherever possible, and linking data retention periods closely to demonstrable utility.

Figure 1: Defining Data Minimization

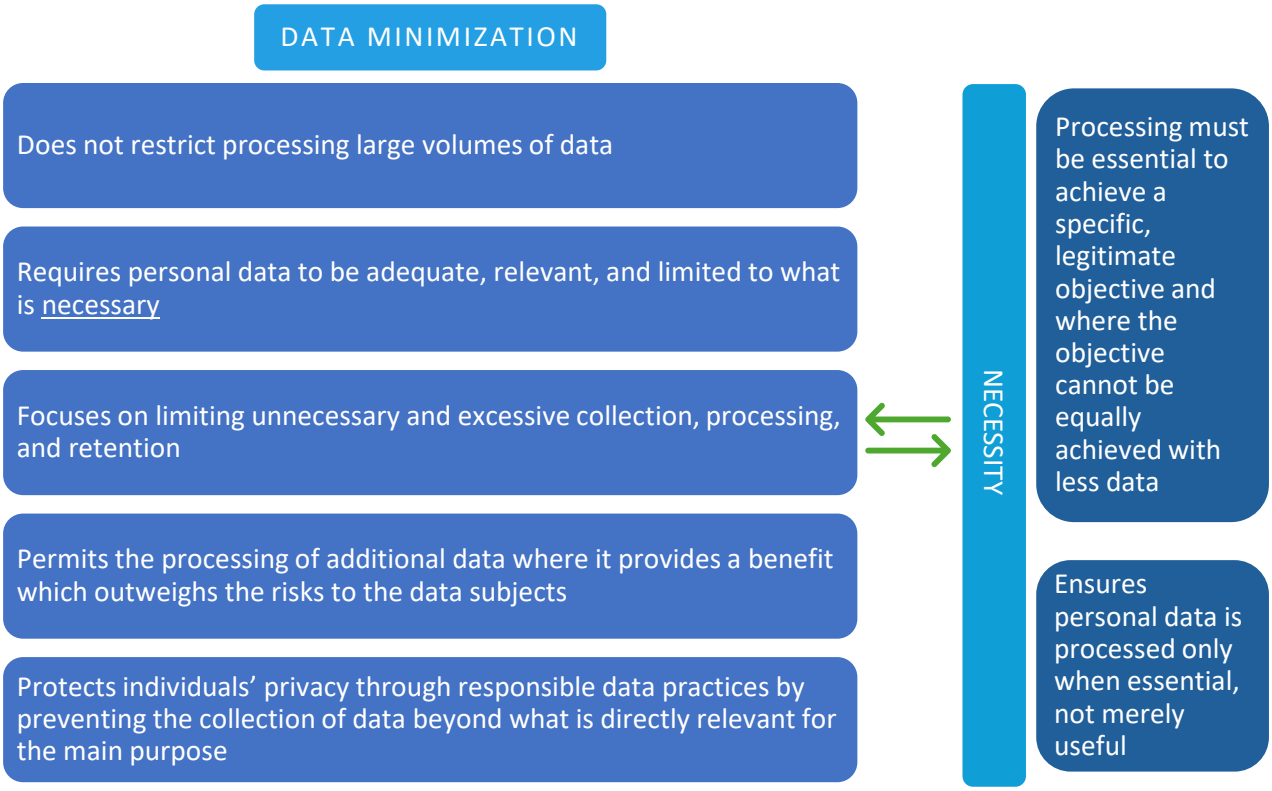

III. Interpreting the Data Minimization Principle in the Context of AI

The data minimization principle must be interpreted in the context of how AI systems are developed and operate. It does not prohibit the use of large datasets, but rather it requires that personal data be adequate, relevant, and limited to what is necessary for a legitimate purpose.

In the AI context, these boundaries shift across the lifecycle: pre-training often demands vast and diverse datasets; accuracy and fairness may require broad inputs, including for bias mitigation; and iterative model development can blur the line between what is "excessive" and what may later prove essential.

AI's evolving capabilities and broad use cases, therefore, require a contextual application of data minimization, grounded in purpose and risk, rather than a uniform expectation to reduce data in all circumstances.

Figure 2: Interpreting the Data Minimization Principle in the Context of AI

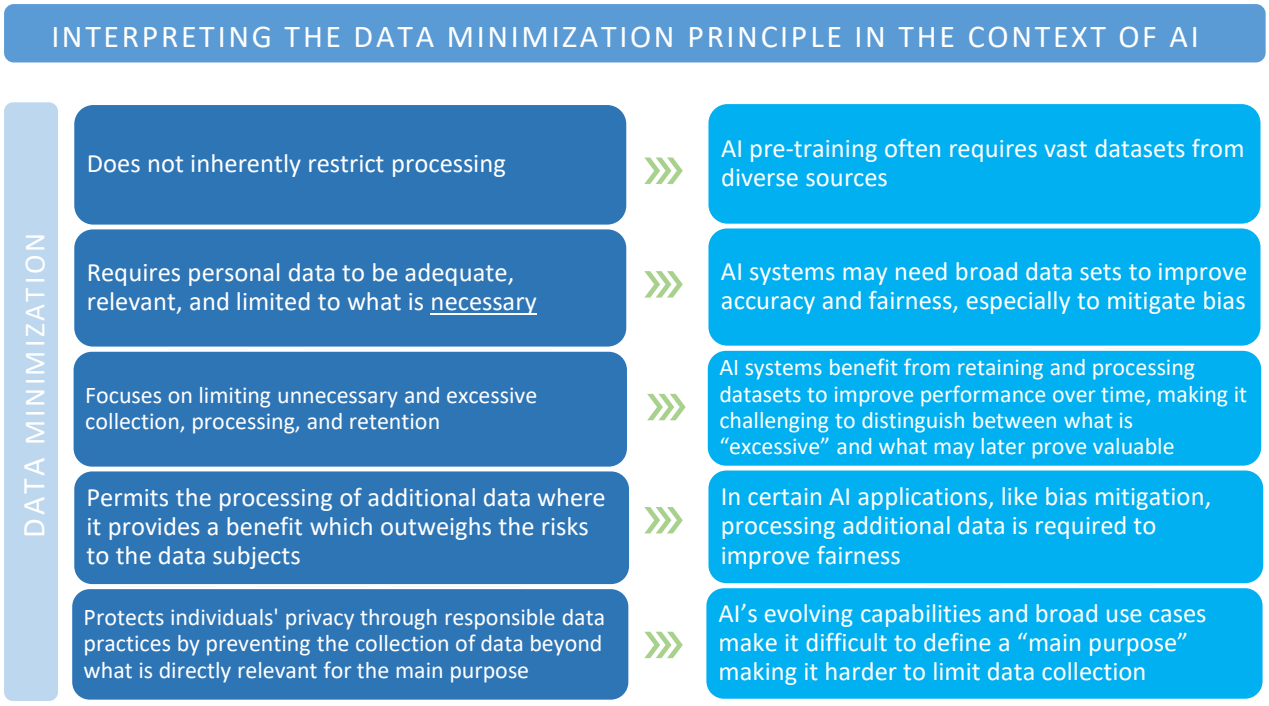

a. Applying Data Minimization across the AI Lifecycle

Practical implications of the data minimization principle should be considered in context and must distinguish between the different phases of the AI lifecycle. Each stage can be associated with specific minimization measures and safeguards.

The table below outlines these phases, illustrating the types of processing involved and the potential measures available to support compliance with the data minimization principle.

Table 1: Applying Data Minimization across the AI Lifecycle

AI Lifecycle Phase	Typical Processing Activities	Potential Minimization Measures
Data Collection/Sourcing	Gathering raw data from various sources	• Where possible, use synthetic or anonymized datasets • Apply data sampling or aggregation to reduce individual-level detail
Data Analysis and Pre-processing	Cleaning, labeling, deduplication, normalization	• Remove redundant or irrelevant features • Apply feature selection techniques • Use privacy-preserving pre-processing techniques (e.g., pseudonymization, differential privacy)
Model Training	Training base or foundation models	• Where possible, employ federated learning to avoid centralizing personal data • Minimize retention of training data post-training • Deploy other privacy-enhancing technologies (e.g., homomorphic encryption, trusted execution

		environments, secure-multi-party computation)
Model Fine-tuning/Alignment	Adapting base models to specific domains	• Use parameter-efficient fine-tuning • Consider differential privacy during parameter updates
Model Evaluation	Testing and validation	• Where possible, use synthetic or anonymized test sets
Deployment and Monitoring	Running and updating deployed systems	• Limit personal data retention • Implement on-device or edge processing

The pre-training stage often demands more or larger datasets to build accurate AI models. Here, a more flexible application of data minimization is justified, particularly if supported by appropriate technical and organizational safeguards. By contrast, the deployment phase typically may involve the processing of less data as the AI model interacts with users.

Model development is, however, an iterative process. Early stages are exploratory, and the relevance of particular data categories often becomes clear only as the model is tested, refined, and aligned. For this reason, assessments of what data is necessary should accommodate the fact that salience is not always predictable in advance.

The importance of personal data for AI development varies significantly according to the nature of the model. In some cases, personal data is deliberately collected to support functionality, security, accuracy, and bias mitigation. In others, it is captured incidentally from publicly available sources as part of broader efforts to assemble rich datasets.

Applying the data minimization principle, therefore, requires a contextual, risk-based assessment of what personal data is genuinely necessary for each stage and purpose, recognizing that relevance is not always fully predictable during early stages of AI development.

b. Determining Necessity

Interpreting what is "necessary" must follow this contextual analysis. The data minimization principle neither permits broad, indiscriminate collection nor requires controllers to avoid large datasets where they are genuinely needed. For example, the CNIL rejected a pharmaceutical company's request to process medical records from the entire active patient population of the medical centers participating in the clinical research project (most of whom did not have prostate cancer), deeming the scale of data collection disproportionate and unnecessary for the stated purpose.[27]

On the other hand, for certain AI applications, it may be necessary to collect significant amounts of personal or sensitive data to improve accuracy or quality, mitigate bias, and improve fairness. For example, in hiring and recruitment, it may be necessary to collect and fine-tune on information such as race, gender, or educational background to ensure that the AI tool is not biased in candidate selection.[28] In this instance, using more personal data (and potentially even special category data) for training the model is necessary to create an accurate and fair AI system.[29] Necessity should be interpreted to include the processing of large volumes of personal data for a legitimate purpose, such as the proper functioning of a task-specific model, mitigating bias, or improving fairness.

A recent paper from the Hamburg and Schleswig-Holstein data protection authorities stresses that the principle of data minimization should not be understood merely as a command to "use less

data," but as a requirement to limit processing to what is necessary for a lawful and legitimate purpose.[30] In the AI context, it outlines that the processing of large datasets to mitigate harmful, biased outcomes is not a violation of the principles, but essential to fulfil the duty to protect data subjects. The accuracy principle creates a positive duty to protect individuals from harm caused by flawed systems, and processing substantial datasets may be necessary to fulfil that duty. Similarly, Article 10(3) of the EU AI Act demands that training, validation, and testing datasets be sufficiently representative, complete, and statistically sound – requirements which may necessitate large, diverse, and detailed datasets.[31]

The following table presents a practical framework to help organizations operationalize the concept of data minimization and determine what data is necessary for specific AI purposes:

Table 2: Necessity Test for Data Minimization in AI Systems[32]

Step	Objective	Suggested Questions
1. Define the specific purpose and how success will be measured	Ensure the data used is directly linked to a clearly defined, measurable purpose	• What is the specific purpose of the AI system? • Is personal data necessary for the purpose? • What metrics define success (accuracy, precision, fairness, robustness)?
2. Consider less intrusive alternatives	Demonstrate that less intrusive alternatives have been considered and were insufficient	• Which less intrusive alternative methods were considered (e.g., subsampling, feature reduction, privacy-enhancing technologies, synthetic data)? • Why were alternative methods insufficient?
3. Show that the chosen amount of data is necessary	Evidence that the volume, scope, and categories of data are required to achieve the defined purpose, and that additional data would not materially improve performance	• How does reducing the dataset (size, variables, or sensitivity) affect outcomes? • Does collecting more data provide proportionate improvements, or do returns flatten?
4. Limit data retention to functional necessity	Retain data only as long as it is needed	• Is data retention justified (e.g., for model retraining or compliance needs)? • Are retention periods proportionate?

			• What deletion/anonymization measures are in place?
5.	**Implement accountability and review procedures**	Maintain accountability, transparency, and periodic review of decisions	• Who is responsible for conducting, approving, and monitoring necessity assessments? • Is there a structured review process as AI models and their data use evolve?

c. Data Minimization and Purpose Limitation

Data minimization is closely tied to the principle of purpose limitation, which requires a specified and legitimate purpose and prevents "just in case" processing.[33]

AI systems often develop in ways that make it difficult to specify every beneficial purpose at the outset. This is especially true for general-purpose AI, which is designed to support a wide range of functions that may only become apparent as the model matures. A purpose that individuals may reasonably expect and value may therefore not be foreseeable at the training stage, even though it remains connected to the system's broader, legitimate objectives.

While purpose limitation remains important in the era of AI, it should not be absolute; data should be reused for socially valuable innovation, provided this is done ethically and with safeguards to protect privacy and rights. In the context of the European Health Data Space, for example, it was recognized that future beneficial uses of health data cannot always be fully defined at the point of collection.[34] A similar flexibility should extend to AI development, where it is often impossible to anticipate every potential model iteration or application.

Purpose limitation should prevent harmful misuse while allowing compatible, beneficial uses. As with data minimization, it should be interpreted differently at different stages of the AI lifecycle. In particular, the use of personal data for training a foundation model should be treated as a distinct processing purpose, separate from the use of data to build, deploy, or refine individual applications built on top of that foundation model. Model developers should clearly define and justify the purposes of training and pre-deployment processing, including the types of personal data used. This distinction ensures that data subjects' rights and expectations are respected at each stage, preventing the repurposing of personal data beyond the original, transparent intent.

Accountability measures – including transparency, compatibility assessments, and risk evaluations – help ensure that new or expanded uses do not create or increase risks. Purpose limitation should be interpreted with enough flexibility to allow legitimate, socially beneficial secondary uses while still maintaining safeguards against misuse.[35]

d. Data Minimization and Storage Limitation

Storage limitation requires personal data to be kept only for as long as it remains necessary for the purpose for which it is processed. In AI systems, this assessment is complex because model development is iterative. Training is not a one-off event; it continues throughout the lifecycle of the model.

It is often difficult to know in advance which data will ultimately prove relevant for model training or pre-training. Given that the AI domain is characterized by identifying unexpected patterns, defining which data is relevant or sufficient at the outset may, in many cases, be in conflict with the very purpose of the development of an AI model, especially considering that combining more data can lead to deeper insights.[36]

Similarly, the relevance of available or historical data may only emerge over time. For example, traffic data can be invaluable in predicting congestion patterns or directing new infrastructure construction. In medical care, historical patient records may later reveal patterns that improve early disease detection or personalized treatment recommendations. Data that initially appears unnecessary may become essential for improving accuracy, mitigating bias, or strengthening robustness and security.

As models evolve, developers may need to retain and reuse the same training data to protect against bias or set guardrails, for instance, and to preserve the robustness, accuracy, quality, and security of developed models. Retention needs may therefore narrow or expand as models develop and as general-purpose or task-specific objectives change. A rigid approach to deletion may undermine performance or prevent the correction of harmful outcomes.

For these reasons, storage limitation must be applied with flexibility in the AI context. Controllers should document, for each data category, why retention is necessary (e.g., impacts on model performance or fairness) and how long it remains justified. Retention periods should be proportionate, evidenced, and subject to periodic review as the model matures and risks evolve.

e. Emerging Regulatory Positions

Data protection authorities are increasingly assessing how data protection and innovation, including AI, can co-exist when data protection principles are interpreted in a forward-thinking, risk-based, and proportionate way by providing useful guidance on common data protection principles for AI developers and deployers.

Table 3: Regulatory Guidance on Data Minimization in AI

Regulatory Body	Guidance
European Data Protection Board (EDPB)[37]	• AI model development and deployment must use personal data that is adequate, relevant, and necessary. This includes data used to prevent bias or errors, if clearly required for the stated purpose • Compliance with data minimization depends on the specific processing activity • Different stages of AI development or deployment may involve the same or different processing activities • The volume of personal data used must be evaluated against less intrusive, equally effective alternatives • The amount of personal data processed must be proportionate to the legitimate interest pursued
European Data Protection Supervisor (EDPS)[38]	• Data controllers must restrict the collection and processing of personal data throughout the AI system's lifecycle, ensuring it is not used indiscriminately and only when no suitable

	alternatives, such as synthetic or anonymized data, can achieve the same goal • If personal data processing is necessary, staff involved in AI development must apply technical measures to minimize its use at all stages of model creation and deployment • AI models should be trained using high-quality, well-labelled, and curated datasets containing only the personal data required for the intended purpose, supported by strong data governance and regular reviews. • Using more data does not automatically improve AI performance; this requires high-quality datasets, careful design, supervised training, and regular monitoring
Information Commissioner's Office (ICO)[39]	• Data minimization requires clearly stating what personal data is adequate, relevant, and limited, based on the AI system's use case • More data should not be processed just because it might be useful later on; it must be necessary for the purpose • Data should only be kept for as long as needed • Retention periods must be proportionate, balancing organizational needs with the impact on individuals' privacy
Commission Nationale de l'Informatique et des Libertés (CNIL)[40]	• Data minimization does not prevent the use of large volumes of data • Data should be carefully selected and cleaned for training to avoid the unnecessary processing of personal data • The quantity of data needed for training must be accurately estimated and proportionate to the processing purpose • The training phase can require a large amount of data to develop an AI system and explore the potential of machine learning • Data minimization means only using personal data useful for the development of the AI system and applying technical measures to limit collection • Training data may be retained for extended periods if justified and secured with appropriate safeguards
Hamburg Commissioner for Data Protection and Freedom of Information (HmbBfDI) and the Schleswig-Holstein data protection authority (ULD)[41]	• Data minimization is not a command to use less data, but requires limiting processing to what is necessary for a lawful and legitimate purpose • Processing large datasets to mitigate harmful, biased outcomes stemming from AI systems is not a violation of data minimization, but essential to protect data subjects • The accuracy principle creates a duty to protect individuals from harm caused by flawed systems, therefore allowing data processing when necessary to uphold fairness, reliability, and data integrity in line with data minimization

IV. SUFFICIENT SAFEGUARDS

At the same time, organizations must embed strong data protection management programs throughout their operations. This includes a clear commitment to: (1) adhere to applicable data protection laws, (2) uphold accountability for their data processing activities, and (3) implement

technical measures to ensure the data minimization principle is considered at every stage – from the initial design and training of AI models to their ongoing deployment and use.

Accountability in this context requires organizations to proactively demonstrate how they operationalize data protection principles, including data minimization, through structured governance frameworks. Strong data protection management programs should include clear roles and responsibilities, regular data protection impact assessments, internal audits, and documented decision-making processes. These programs also benefit from cross-functional collaboration, involving legal, technical, and ethical expertise to ensure that privacy considerations are fully embedded into AI development workflows. By implementing such governance measures, organizations can ensure that data minimization is not treated as a one-time exercise, but as an ongoing responsibility that evolves alongside the AI systems themselves.

Organizations should also integrate both technical solutions and strategic design choices to reduce the collection and use of personal data without unduly compromising model performance. Privacy-enhancing technologies (PETs) are increasingly recognized by data protection authorities as essential tools in achieving this balance.[42] Organizations should also consider the implementation of red teaming and automated risk measurement pipelines to inform their mitigation strategies. These evaluations should be documented and ultimately shared to foster industry-wide learning.

The UK's ICO, for example, acknowledges that PETs such as perturbation (including adding noise via differential privacy), federated learning, and the use of synthetic data can help to minimize the processing of personal data during the training phase.[43] The ICO recommends that, where possible, these techniques be applied prior to data processing or use, as part of a privacy-by-design approach, to better mitigate risks to individuals.[44]

For the inference phase, other approaches, including converting personal data into less human-readable formats, performing inferences locally on users' devices, and employing privacy-preserving query techniques, can be effective.[45] For example, converting raw data into abstract representations, such as feature vectors, can reduce direct visibility of personal data during processing.[46] Local inference involves running AI models locally on users' devices, limiting the need to send personal data over networks.[47] Finally, privacy-preserving query techniques can be used to obtain predictions without fully revealing the user's data. These strategies help protect personal information during inference and can work alongside training-phase measures to support data minimization.

France's CNIL emphasizes the importance of AI system design and model selection in achieving data minimization.[48] The CNIL states that when multiple technical approaches can achieve the same outcome, preference should be given to the method that requires the least personal data and presents the lowest privacy risk.[49] This principle extends not only to model architecture but also to training protocols and deployment strategies. Techniques such as decentralized training (e.g., federated learning) and encrypted computation (e.g., secure multi-party computation and homomorphic encryption) are promising approaches that allow models to be trained without exposing underlying data. However, they acknowledge that these methods may carry limitations in terms of computational cost and maturity, and so recommend monitoring the development of these techniques.

Of course, PETs are not one-size-fits-all solutions: they are highly use-case specific, and to maximize both utility and protection, they may need to be used in combination. PETs do have trade-offs, and improper application can lead to unintended consequences, such as degraded model performance or ineffective de-identification.[50]

In CIPL's 2025 report on PETs in AI, we examine how these tools are being implemented to address privacy concerns across the AI lifecycle while also unlocking new opportunities for data use.[51]

V. CONCLUSION

The principle of data minimization should not be understood as a mandate to reduce data volume indiscriminately, but rather as an obligation to ensure that personal data is collected and retained only when demonstrably relevant to the specific purpose pursued. It should be interpreted flexibly to allow personal data use for machine learning purposes, where the AI systems are socially beneficial and respect data protection rights.

CIPL makes the following recommendations for organizations and regulators to effectively operationalize data minimization in AI systems:

Adopt a contextual and flexible interpretation of data minimization: Organizations and regulators should apply data minimization proportionally, focusing on necessity, balancing risks and benefits, acknowledging AI-specific needs, tailoring measures across AI lifecycle stages and permitting socially beneficial secondary purposes

Implement a structured necessity test framework: Organizations should assess data needs using a practical, step-by-step framework that defines purpose, considers less intrusive alternatives, justifies data volume, scope, and categories, limits retention, and ensures accountability and review

Embed strong safeguards and accountability: Organizations should establish robust governance, cross-functional collaboration, and technical measures, including privacy-enhancing technologies, to responsibly manage and minimize personal data use

Applying data minimization in AI calls for a contextual approach that recognizes the need for large and diverse datasets while placing clear boundaries on their use. This balanced interpretation enables the development of powerful, socially beneficial AI systems without compromising the protection of individual rights.

ABOUT CIPL

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.

---

[1] See CIPL Report, "Applying Data Protection Principles to GenAI: Practical Approaches for Organizations and Regulators," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_applying_data_protection_principles_genai_dec24.pdf. See also CIPL Report, "Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_building_accountable_ai_programs_23_feb_2024.pdf. See also CIPL, "Ten Recommendations for Global AI Regulation", available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_ten_recommendations_global_ai_regulation_oct2023.pdf. See also CIPL Report, "Hard Issues and Practical Solutions," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_second_report_-_artificial_intelligence_and_data_protection_-_hard_issues_and_practical_solutions__27_february_2020_.pdf. See also CIPL/Hunton Andrews Kurth Legal Note, "How the GDPR Regulates AI," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl-hunton_andrews_kurth_legal_note_-_how_gdpr_regulates_ai__12_march_2020_.pdf. See also CIPL Report, "Artificial Intelligence and Data Protection in Tension," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_first_ai_report_-_ai_and_data_protection_in_tension__2_.pdf. See also CIPL "Response to CNIL How-To Sheets on the Development of Artificial Intelligence Systems", available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_response_to_cnil_consultation_on_ai_-_second_series-c.pdf. See also CIPL "Compilation of Responses to UK ICO GenAI Consultations," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_-_compilation_of_uk_ico_generative_ai_responses.pdf. See also CIPL "Response to the National Institute of Standards and Technology (NIST)'s Request for Comment on the Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipls_response_to_nists_ai_rmf_gai_profile.pdf.

[2] Directive 95/46/EC, Article 6(c).

[3] OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, referred to as the Data Quality Principle, available at https://www.oecd.org/en/publications/2002/02/oecd-guidelines-on-the-protection-of-privacy-and-transborder-flows-of-personal-data_g1gh255f.html.

[4] Council of Europe Convention 108, Art. 5(4)(c), available at https://www.coe.int/en/web/data-protection/convention108-and-protocol.

[5] See also de Terwagne in Kuner, Bygrave, Docksey, GDPR, Art. 5, p. 317: "The CJEU has required a "strict necessity" in some cases like *Digital Rights Ireland*, *Tele2*, and *Schrems II*, where the Court invalidated measures such as broad data retention and the EU-US Privacy Shield for failing to meet the strict necessity threshold".

[6] Followed by Colorado, Connecticut, Delaware, Florida, Indiana, Kentucky, Minnesota, Montana, Nebraska, New Hampshire, New Jersey, Oregon, Rhode Island, Tennessee, Texas and Virginia.

[7] Maryland Online Data Privacy Act of 2024 (MODPA), Section 14-4607(B)(1)(I); See CIPL Report, "Data Minimization in the United States' Emerging Privacy Landscape: Comparative Analysis and Exploration of Potential Effects," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_data_minimization_us_privacy_landscape_aug24.pdf.

[8] Brazilian Data Protection Law (LGPD) (As amended by Law No. 13, 853/2019), Article 6.

[9] The Digital Personal Data Protection Act, 2023, Articles 6(1), 9(7) and 12(3).

[10] Advisory Guidelines on use of Personal Data in AI Recommendation And Decision Systems, available at https://www.pdpc.gov.sg/-/media/files/pdpc/pdf-files/advisory-guidelines/advisory-guidelines-on-the-use-of-personal-data-in-ai-recommendation-and-decision-systems.pdf.
[11] Personal Data Protection Act 2012, Article 25.
[12] Personal Information Protection and Electronic Documents Act (PIPEDA), Section 5.
[13] Act on the Protection of Personal Information (Act No. 57 of 2003), Article 15-16.
[14] OpenAI's ChatGPT now has nearly 800 million weekly active users (Shubham Singh, "ChatGPT Statistics 2025", 5 June, 2025, available at https://www.demandsage.com/chatgpt-statistics/), Microsoft's Copilot Studio is already being used by more than 230,000 organizations, including 90% of Fortune 500 companies, to build AI agents and automations (Frank X Shaw, "Microsoft Build 2025: The age of AI agents and building the open agentic web", 19 May, 2025, available at https://blogs.microsoft.com/blog/2025/05/19/microsoft-build-2025-the-age-of-ai-agents-and-building-the-open-agentic-web/), McKinsey reports that more than 78% of companies are now using gen AI in at least one business function (McKinsey & Company, "The state of AI: How organizations are rewiring to capture value", March 12, 2025, available at https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai).
[15] See CIPL Report, "Agentic AI: Fostering Responsible and Beneficial Development and Adoption," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_agentic_ai_fostering_responsible_development_adoption_oct25.pdf.
[16] Regulation 2024/1689, Article 1, Recitals 1, 2, 3, 165, 176.
[17] Davey Alba and Rachel Metz, "Large AI Dataset Has Over 1,000 Child Abuse Images, Researchers Find", *Bloomberg*, 20 December 2023, available at https://www.bloomberg.com/news/articles/2023-12-20/large-ai-dataset-has-over-1-000-child-abuse-images-researchers-find.
[18] See CIPL's report on individual rights in AI (forthcoming, 2026).
[19] Prakhar Ganesh et al., The Data Minimization Principle in Machine Learning, 2024, available at https://arxiv.org/abs/2405.19471.
[20] Gemma Galdon Clavell et al., Auditing Algorithms: On Lessons Learned and the Risks of Data Minimization, 2020, available at https://dl.acm.org/doi/10.1145/3375627.3375852.
[21] GDPR, Article 5(1)(b).
[22] As the UK ICO puts it "*either 'process no personal data' or 'if we process more, we're going to break the law'* (Information Commissioner's Office, Guidance on AI and Data Protection, 15 March, 2023, available at https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/how-should-we-assess-security-and-data-minimisation-in-ai/#whatdataminimisation).
[23] *Meta Platforms Inc v. Bundeskartellamt* (C-252/21) EU:C:2023:537 (04 July 2023).
[24] ibid.
[25] See Recital 39 GDPR, which specifies that personal data should only be processed if the objective cannot reasonably be achieved through other, less intrusive means; See also European Data Protection Board (EDPB), Guidelines 2/2019 on the processing of personal data under Article 6(1)(b) GDPR in the context of the provision of online services to data subjects.
[26] Commission Nationale de l'Informatique et des Libertés (CNIL), AI and GDPR: the CNIL publishes new recommendations to support responsible innovation, 2025, available at https://www.cnil.fr/en/ai-and-gdpr-cnil-publishes-new-recommendations-support-responsible-innovation.
[27] ibid.
[28] Marvin van Bekkum and Frederik Zuiderveen Borgesius, Using sensitive data to prevent discrimination by artificial intelligence: Does the GDPR need a new exception? (2023) p. 3, available at https://www.sciencedirect.com/science/article/pii/S0267364922001133.
[29] The Information Commissioner's Office (ICO) notes that to assess and address discrimination risks in AI systems, that special category data may be required. Information Commissioner's Office, What about fairness, bias and discrimination?, 2023, available at https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/how-do-we-ensure-fairness-in-ai/what-about-fairness-bias-and-discrimination/.
[30] Thomas Fuchs et al., The Bridge Blueprint, available at https://datenschutz-hamburg.de/fileadmin/user_upload/HmbBfDI/Datenschutz/Informationen/EN-Bridge-Blueprint-v.0.9.pdf.

[31] Article 12 of the EU AI Act requires high-risk AI systems to record and trace its operations, implying the retention of extensive logs and metadata, potentially increasing the volume of personal or sensitive information processed.
[32] In practice, these steps may not follow a strictly linear sequence, and their answers may vary across the AI lifecycle.
[33] GDPR, Article 5(1)(b).
[34] Recital 61 of the European Health Data Space allows the secondary use of healthcare data for the purposes of research, innovation, policymaking, education, safety, regulatory activities and personalized medicine. The GDPR itself also provides a foundation for this flexibility through the concept of purpose compatibility, which allows further processing if it is compatible with the original purpose in GDPR, Article 6(4).
[35] For example, CIPL encourages the pragmatic and staged approach taken by the ICO when applying purpose limitation to general-purpose AI, who allow for broader data processing when training foundation models, and then limits the volume that is deemed necessary as the purpose for which a model is being deployed becomes clearer. See ICO, Generative AI second call for evidence: Purpose limitation in the generative AI lifecycle, available at https://ico.org.uk/about-the-ico/what-we-do/our-work-on-artificial-intelligence/generative-ai-second-call-for-evidence/.
[36] International Working Group on Data Protection in Telecommunications, Working Paper on Privacy and Artificial Intelligence, 64th Meeting, 29-30 November 2018, Queenstown (New Zealand), p. 9.
[37] EDPB, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models, available at https://www.edpb.europa.eu/system/files/2024-12/edpb_opinion_202428_ai-models_en.pdf.
[38] European Data Protection Supervisor (EDPS), Generative AI and the EUDPR: Orientations for ensuring data protection compliance when using Generative AI systems, available at https://www.edps.europa.eu/system/files/2025-10/25-10_28_revised_genai_orientations_en.pdf.
[39] ICO, supra note 22.
[40] CNIL, supra note 26; CNIL, Development of AI systems: the CNIL's recommendations to comply with the GDPR, 22 July, 2025, available at https://www.cnil.fr/fr/developpement-des-systemes-dia-les-recommandations-de-la-cnil-pour-respecter-le-rgpd; CNIL, AI: ensuring GDPR compliance available, 21 September, 2022, available at https://www.cnil.fr/en/ai-ensuring-gdpr-compliance.
[41] Thomas Fuchs et al., supra note 30.
[42] See CIPL Report, "Privacy-Enhancing and Privacy-Preserving Technologies: Understanding the Role of PETs and PPTs in the Digital Age," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl-understanding-pets-and-ppts-dec2023.pdf.
[43] ICO, supra note 22.
[44] While these methods can be effective in reducing data processing, the ICO emphasizes that they must be deployed thoughtfully. As perturbation adds controlled noise to data to protect privacy, it must be carefully balanced to avoid harming model accuracy; federated learning trains models locally and shares only updates, yet these updates can still leak personal information, requiring careful risk assessment; and poorly designed synthetic data may leak unique traits that risk re-identification, underscoring the need for thorough model testing and assessments of the synthetic data.
[45] ICO, supra note 22.
[46] However, conversion to feature vector is not an absolute guarantee of privacy, with several studies highlighting the tendency of data leakage from text embedding, see for example, John X Morris et al., "Text Embeddings Reveal (Almost) As Much As Text", available at https://arxiv.org/pdf/2310.06816.
[47] However, running AI models on-device can be challenging due to device or model constraints, as on-device memory is limited, thereby restricting the user to a narrower selection of small models.
[48] CNIL, Taking into account data protection when designing the system, 7 June 2024, available at https://www.cnil.fr/en/data-protection-when-designing-system#:~:text=To%20ensure%20that%20the%20development,the%20original%20publication%20in%20French.
[49] Article 88(c) of the European Commission's Digital Omnibus Regulation Proposal references that technical and organizational measures should be used to operationalize the data minimization principle during the stage

of selection of sources and the training and testing of an AI system or model, available at https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-regulation-proposal.

[50] See CIPL PETs Report, supra note 41.

[51] See CIPL Report, "Privacy-Enhancing and Privacy-Preserving Technologies in AI: Enabling Data Use and Operationalizing Privacy by Design and Default," available at https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_pets_and_ppts_in_ai_mar25.pdf.