Centre for Information Policy Leadership

HUNTON

CIPL EU Simplification Project– Workshop I Report

Recommendations on the Digital Omnibus on AI Regulation

About This Initiative

As part of its project **"Towards a Coherent Digital Rulebook"**, the Centre for Information Policy Leadership (CIPL) is convening a series of technical workshops focused on the EU Simplification agenda and the forthcoming Digital Omnibus packages. The objective of this initiative is to provide constructive, practical, and grounded recommendations to legislators to ensure that the EU digital regulatory framework remains coherent, innovation-friendly, legally certain, and globally competitive.

The first workshop was dedicated to the proposed AI Act Omnibus. The workshop brought together organisations and technical experts to examine how targeted amendments under the AI Omnibus could simplify implementation of the law while preserving the Act's objectives of supporting a well-functioning EU internal market by boosting trustworthy, human-centric AI and innovation, while protecting health, safety, fundamental rights, and the environment from AI harms.

Participants focused on ways to ensure that the AI Act remains technologically neutral and future-proof, operationally workable and enforceable, globally competitive, and consistent with existing EU digital legislation.

The recommendations below reflect considerations emerging from the roundtable on practical improvements that would enhance legal certainty, reduce fragmentation, and strengthen the Act's effectiveness without weakening its safeguards.

1. Protect the AI Act's Long-Term Effectiveness by Preserving Technological Neutrality

The current debate regarding **"Agentic AI"** illustrates a broader legislative challenge: the understandable desire to respond directly to emerging technological developments by appending increasingly granular definitions. However, this risks departing from one of the foundational principles of effective AI regulation: technological neutrality. AI capabilities are evolving rapidly, as illustrated by the shift from generative AI to increasingly agentic systems since the Act's adoption. It will not be feasible to update the framework to reflect every new technical development. Maintaining a technology-neutral approach remains the most robust and future-proof option to ensure that emerging developments are captured within the existing framework. Creating a distinct legal category for "Agentic AI" is unnecessary and may create unintended complexity.

Existing Definitions Are Sufficient

The existing definition of AI systems already captures systems that take autonomous or semi-autonomous actions. Conceptually, if an AI system performs an action, it necessarily involves a form of decision-making. Since decision-making is embedded in the current definition, introducing a separate category for "agents" does not appear to add legal clarity or substantive differentiation.

Risk of Definitional Fragmentation

Introducing "Agentic AI" risks fragmenting the definitional architecture of the Act. Developers could be required to assess systems simultaneously as "models," "agents," and "AI systems," creating overlapping regulatory categories without clear added value. Given the link between definitions and substantive obligations across the AI Act, ambiguity at the definitional level can cascade into uncertainty regarding risk classification and potential prohibitions.

Risk of Legislative Obsolescence

Specifying technologies that reflect a particular moment in the innovation cycle risks accelerating legislative obsolescence. The AI ecosystem evolves rapidly; embedding specific technological paradigms in primary legislation may require continual revision and may, in the reverse, inadvertently raise questions about whether or not future technologies are included.

> **Recommendation:** Maintain technological neutrality and avoid introducing new technology-specific definitions unless demonstrably necessary to address clearly distinct risk profiles.

2. Calibrate Article 5 Prohibitions to Avoid Overbreadth and Ensure Legal Coherence

There is no question that the generation and dissemination of non-consensual sexualised imagery must be effectively addressed. However, it is important to recall that the AI Act has been designed to be primarily a **product safety framework** governing the development and placing on the market of AI systems, rather than a legal instrument designed to regulate specific forms of unlawful conduct.

Existing Frameworks Already Apply

In many cases, the harm at issue stems from malicious human intent and criminal behaviour. The established principle that what is illegal offline should be illegal online remains central. Sectoral frameworks, including national criminal law and instruments such as the Digital Services Act, already provide more targeted mechanisms to address the creation, dissemination, and removal of illegal content.

Risk of Overbroad Provisions

It is also important to recognise that tools capable of manipulating visual or audio content have existed for many years, including widely used professional image and video editing software. The mere technical capability to alter or generate content does not, in itself, indicate harmful use. Broad provisions targeting systems that can manipulate content risk capturing a wide range of legitimate creative and professional tools that were never intended to fall within the scope of prohibitions.

Where regulatory intervention is considered, a clearer distinction should therefore be maintained between general-purpose tools with legitimate uses and systems deliberately designed or marketed for harmful exploitation, including features specifically enabling the generation of non-consensual sexualised imagery. Without such differentiation, there is a risk of blurring the boundary between product safety legislation and conduct-based regulation, while diverting compliance resources towards interpretative uncertainty rather than effective harm mitigation.

> **Recommendation:** Ensure that practices under Article 5 are narrowly and precisely calibrated to address clearly defined harmful practices, while preserving legal certainty for legitimate general-purpose tools.

3. Provide Fixed and Predictable Implementation Deadlines to Ensure Legal Certainty

The debate between fixed implementation dates and deadlines linked to the availability of harmonised standards reflects a broader tension between flexibility and legal certainty. Organisations consistently emphasise that they require **fixed calendar milestones** to structure internal governance, budgeting, and technical deployment cycles.

A "moving target" approach, whereby implementation is triggered by the availability of standards, is out of touch with the realities of corporate compliance planning. Even where standards are delayed, a fixed end-date provides clearer internal coordination than open-ended conditional triggers.

- ### Minimum Viable Window

  A **12-month window** was considered the minimum technically viable period for meaningful integration. A six-month window was broadly viewed as unrealistic for complex enterprise systems and integrated software stacks.

- ### ISO-Certified Organisations

  This transition period is particularly critical for organisations already certified under international ISO frameworks. These companies must "map" existing standards to the specific requirements of the EU high-risk regime, a process that requires structured engineering and governance work.

- ### Consequences of Insufficient Lead Time

  Without sufficient lead time, organisations face unrealistic expectations to redesign integrated systems within tight deadlines.

> **Recommendation:** Provide fixed implementation dates and ensure a minimum 12-month transition window following finalisation of relevant harmonised standards.

4. Ensure Registration Requirements Are Proportionate and Effective

The proposal to exempt providers from registration requirements where they have concluded that their system does not qualify as high-risk reflects the logic of EU product safety legislation. Under the EU's New Legislative Framework, compliance is primarily ensured through internal conformity assessment, technical documentation, and post-market surveillance rather than through broad ex ante registration obligations. In this context, mandatory registration of systems that providers have already assessed as not falling within the high-risk category risks introducing limited additional oversight while creating unnecessary administrative burdens.

1. National authorities may face significant constraints in meaningfully reviewing potentially large volumes of registrations. Registries may therefore become largely formalistic tools, offering limited supervisory value.
2. Public registers may provide only partial transparency while still imposing substantial compliance costs. Registration alone does not demonstrate that effective risk management processes are in place.

As a positive example, the GDPR also moved away from previously existing comparable notification and registration obligations, recognising that these mechanisms often created administrative burdens without materially improving compliance outcomes. By contrast, structured and harmonised self-assessment templates would provide far greater value for accountability, enabling demonstrable, and auditable risk management processes.

> **Recommendation:** Focus on meaningful accountability documentation and targeted supervision rather than blanket registration mechanisms that may generate administrative volume without added value.

5. Enable Practical and Lawful Access to Data for Effective Bias Mitigation

The Omnibus proposal to clarify the conditions under which sensitive data may be processed for bias detection—including for AI systems that are not classified as high-risk—is a welcome step towards improving the practical implementation of fairness testing requirements under the AI Act.

However, the discussion highlighted a persistent structural challenge relating to data availability. In particular, many organisations do not collect special category data as part of their core operations. While the proposal does establish a clearer legal basis for processing such data for the purpose of bias detection and mitigation, this clarification alone does not resolve the underlying issue where companies do not hold the relevant demographic data in the first place.

In practice, meaningful bias testing may require access to representative demographic datasets sourced from third parties. Such datasets are often difficult to obtain at a sufficient scale. As a result, organisations may continue to face a structural data scarcity challenge when attempting to conduct robust fairness assessments, even where the legal basis to process such data has been clarified. Addressing this issue will be important to ensure that bias detection and mitigation requirements remain operationally feasible and capable of delivering the intended policy objective.

Furthermore, the proposal to require the immediate deletion of sensitive data following bias "correction" raises practical and technical concerns. Bias detection and mitigation are not one-off exercises but ongoing, iterative processes. AI systems may evolve over time, and model performance can drift as data or deployment contexts change. Effective fairness governance, therefore, requires the ability to conduct periodic reassessment and continuous monitoring. An obligation to delete sensitive data immediately after an initial correction risks undermining the capacity to verify whether mitigation measures remain effective over time.

Recommendations:

Facilitate Lawful Data Access

Consider ways to facilitate lawful access to appropriate demographic datasets for fairness testing, including from third-party providers.

Allow Retention for Ongoing Monitoring

Allow the retention of sensitive testing data for as long as strictly necessary to support ongoing bias detection, monitoring, and mitigation, recognising that fairness assessments are continuous rather than one-off exercises.

Maintain Safeguards and Proportionality

Ensure that safeguards and proportionality remain central to such processing.

6. Strengthen and Harmonise Regulatory Sandboxes to Support Responsible Innovation

Well-designed regulatory sandboxes and other modern regulatory tools can foster responsible innovation in a collaborative environment between an organisation and regulators. Participants welcomed the updated provisions on regulatory sandboxes but stressed the need for greater EU-wide coordination.

To maximise the effectiveness of sandboxes:

Sandbox outcome reports should benefit from mutual recognition across Member States.

The development of specialised "centres of excellence" should be promoted, enabling Member States to build sectoral expertise (e.g., healthcare, scientific research).

Sandbox exit reports should be fully recognised within conformity assessment procedures, creating tangible incentives for participation.

Proportionate DPA Involvement

Participants noted that, given the general applicability of the GDPR, mandatory DPA involvement in all sandbox cases may create avoidable bottlenecks. Particularly where AI applications rely predominantly on non-personal data or the processing of personal data is not central to the sandbox project, a **"consultative trigger" model** may be more proportionate, reserving formal DPA engagement for cases raising significant GDPR complexity.

> **Recommendation:** Strengthen cross-border mutual recognition and better integrate sandbox participation into conformity assessment processes and ensure the AI Act provides other lighter-touch innovation tools.

7. Reinforce Institutional Coordination and Mutual Recognition to Prevent Market Fragmentation

A key opportunity in the AI Act Omnibus process lies in strengthening coordination mechanisms similar in spirit to the GDPR's One-Stop-Shop. Without effective mutual recognition and cooperation among Member States, there is a risk of fragmented supervision, where a single AI system could face parallel scrutiny from multiple authorities with differing interpretations. This would undermine the objective of a coherent Digital Single Market. Enhanced regulatory cooperation and mutual recognition are therefore essential to ensuring consistency and legal certainty.

Moreover, as the AI Office assumes increased responsibilities, particularly for general-purpose AI, it must be adequately resourced to provide timely guidance, benchmarks, and interpretative clarity. Unintended centralisation without operational capacity risks creating bottlenecks and uncertainty.

Reinforce Cross-Border Recognition

Reinforce mechanisms that ensure decisions and conformity assessments in one Member State are respected across the EU, thereby safeguarding the integrity of the Single Market.

Resource the AI Office Appropriately

Ensure that the AI Office is appropriately resourced to fulfil its mandate and provide technical and interpretative leadership.

Looking Ahead

CIPL looks forward to continuing to engage constructively with EU lawmakers, the Commission and Member States to ensure that the AI Act remains coherent, future-proof, and practically enforceable within a dynamic technological landscape.

Explore CIPL resources on AI

- Learning from Practice: Designing Effective Regulatory Sandboxes
- Reconciling AI with the Data Minimization Principle: Bridging the Innovation and Privacy Gap
- Legitimate Interests for Data in AI Training - The DPO Perspective
- Agentic AI: Fostering Responsible and Beneficial Development and Adoption
- AI Act Article 4: AI Literacy Best Practices and Recommendations for Practitioners
- Key Takeaways - Inaugural CIPL EU AI Act Forum: Setting the Direction of Implementation
- Applying Data Protection Principles to Generative AI: Practical Approaches for Organizations and Regulators
- Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework
- CIPL Recommendations on Adopting a Risk-Based Approach to Regulating AI in the EU