A Warm Body in the Loop

Rethinking Human Control of AI in EU Tech Regulation

Sebastian Schwemer & Ida Koivisto

18.09.2025 doi: 10.59704/aa18974494405499

Brussels has recently signalled a shift in its approach to technology regulation, focusing on simplification through various Omnibus packages (Commission 2025). In the digital context, besides the stated goal of cutting "red tape," these packages offer an opportunity to reconsider the foundations of human involvement regulation in EU legal instruments in the age of AI.

This post examines human-AI interaction in EU technology regulation, in the context of the EU's Artificial Intelligence Act (AIA), the General Data Protection Regulation (GDPR), and the Digital Services Act (DSA). It questions whether such human involvement is meaningful or merely symbolic. While EU regulation consistently inserts humans "in the loop" as safeguards against the risks of automation, these mechanisms are often at risk of being formalistic, providing the appearance of oversight rather than its substance. We argue that for human involvement to offer genuine protection, it needs to be based on clearer conceptual foundations.

The changing role of humans

Over the last two decades, regulating human-machine interaction has become a central, albeit subtle, theme in EU tech law. It appears in areas such as robotics, AI regulation, content moderation, automated decision-making (ADM) in the public sector, and in the governance of drones and financial markets, to name a few.

Since the Code of Hammurabi, one of the earliest written legal codes, law silently relies on human labor: a civil servant granting administrative benefits, a judge ruling cases or a lawyer giving legal advice. Today, however, humans are increasingly supported or supplanted by machines (read: AI). The human monopoly of applying the law is questioned.

What role is left for humans, then? Is it simply to control automation? Keeping a human involved is the choice of the EU legislator. Across the GDPR (and its predecessor from 1996), the DSA, and the AIA, the primary safeguard against automation risks – such as threats to health, human rights, and safety – is precisely to keep a human in the loop (Laux, Wachter and Mittelstadt 2024).

"Human in the loop" refers to a system or a process where human involvement is an important but not the exclusive part of an operation. In law, this entails a new hybridization of machines and human legal labor. Perhaps paradoxically, this both represents and reproduces the underlying human-machine dichotomy found in technology regulation (Koulu & Koivisto, forthcoming 2025). The assumption is that adding humans into the loop can provide safety, legitimacy, and trust in machine operations that might otherwise feel daunting and unpredictable.

Yet, this regulatory design raises a crucial question: what exactly does human involvement achieve, and under what conditions does it work? Let us look at a few examples from recent EU technology regulations.

- Human oversight: Human oversight means that humans must be able to supervise AI systems. While it now clearly features as a legal norm (most prominently in Art. 14 AI Act), it also functions as a legal principle and interpretative tool (visible in recital 27 AIA which refers to the AI High-Level Expert Group principles). Research has shown, however, that human oversight often falls short. Humans can all too easily become mere rubber stampers, unable – or perhaps unwilling – to critically evaluate automated decision-making, particularly when systems are complex or highly specialized. The

AI Act, for instance, does not adequately recognize this limitation, instead presuming that humans possess almost superhuman capacities, assigning them extensive responsibilities such as monitoring, detecting, interpreting, intervening, and ultimately deciding when to stop a system (Article 14 AIA). This raises concerns that oversight may, in practice, become an empty shell, especially when humans are tasked with monitoring systems they do not fully understand or control.

- Human intervention: Human intervention can serve as a legal backup if full or semi-ADM fails (C-634/21). In Art. 22 GDPR, a data subject has the right not to be subject to automated decision-making with legal or similar effects. However, the right has many exceptions. For instance, ADM may be permitted under national law, in which case legal safeguards – including the right to human intervention (recital 71) – must apply. The assumption is that if ADM fails, a human can step in to solve or mitigate the problem. Yet it remains unclear what gives human intervention its legitimating effect: does mere presence imply accountability, or is it largely symbolic with little corrective power?
- Human review: For the better part of the last decade, online platforms have performed algorithmic content moderation. This means large-scale AI-based micro-legal decision-making on whether content is illegal or against their terms of service, where human content moderation plays a decreasing role (Guardian 2025). Nonetheless, human elements are still introduced through two provisions. Both the infamous Art. 17 of the Directive on copyright in the Digital Single Market, and Art. 20 DSA, put a human – i.e. "appropriately qualified staff"– in the content moderation loop, albeit only ex post. While the initial decision may be automated, a human case handler is required at the redress or complaint stage. Although this provision ensures a degree of procedural fairness, it remains unclear whether such post hoc review is sufficient to correct errors or prevent harm already caused by automated removals.
- Limits of automation: Some legislation does not explicitly mandate human involvement but implies its importance, suggesting that certain tasks should not be fully automated. For example, in EU platform law, a key question is how far content moderation can be automated. The requirement for human review is especially significant, as it implies that everything before the review stage may be automated (Quintais et al 2022). At the same time, the DSA encourages internet intermediaries to take voluntary measures to address illegal content, which may include automation. By contrast, both the GDPR and the AI Act explicitly exclude certain kinds of automation, indicating that some tasks must remain under human control (cf. Huq 2020) or be disallowed entirely (cf. AI Act, Art. 5). In this way, the acceptable scope of automation is always defined in relation to the scope of human involvement.

What we must inquire

The European Court of Justice has consistently noted that the necessity for safeguards is "all the greater where the interference [with fundamental rights] stems from an automated process" (see, e.g., C-401/19, para. 67). As we can see from the examples above, humans are called to oversee, intervene and reassess automated processes which otherwise might seem uncontrollable. Is human involvement the panacea to create trust in machines in law, although humans are flawed in numerous ways and can sometimes only serve the role of "warm bodies" in the process (Crootof et al 2023)? To move beyond regulatory lip service, we must ask deeper conceptual and empirical questions: what makes human oversight legitimate, and under what conditions does it work?

First, we must confront the – naïve – assumption of human exceptionalism that silently underpins EU tech regulation. This is reflected in the trust in fellow humans (in the loop), whereas machines come across as unknowable and unpredictable. Yet, this assumption

risks creating blind spots: it may entrench flawed practices simply because they are human made, while ignoring improvements offered by well-designed automation. Humans might be flawed, but at least they are flawed in a familiar and foreseeable way. Is unquestioned human exceptionalism the way to go or should we approach it critically?

Second, more attention needs to be devoted to locating human agency in automated processes. Where in a fully or semi-automated decision-making process is human involvement necessary? And why is it necessary in the first place? What can serve as a benchmark; e.g. decision quality and error rates (Schwemer 2024) or is it about something else entirely? And what is the standard for that human? Not all human input is meaningful. The quality, timing, and expertise of that input must be assessed empirically, rather than assumed effective by default.

Toward conceptual clarity

A concrete way forward in legal research is to systemise the scattered provisions on human involvement in order to establish a shared vocabulary and conceptual basis (Schwemer 2021). Such efforts are especially timely as the EU lawmaker is revisiting parts of the digital acquis. Therefore, a legal bird's-eye view might be especially useful.

It may be worthwhile to explore whether consolidating human involvement provisions in the AI Act, DSA, GDPR could provide clearer definitions of human oversight, intervention and review to rely on, which could then be applied in use-case or sector-specific legislation.

Notably, all this could be done without immediately addressing the challenging normative questions of whether, when, where and how humans should remain involved. Namely, a more fundamental question inevitably lingers. Do problems arising from the human-in-the-loop stem from a lack of a systematic approach, or do they reflect a deeper issue of modern law: the difficulty of regulating non-human entities without human intermediaries? Systematization alone cannot answer this; basic research is needed.

ABOUT THE AUTHORS

Sebastian Schwemer
Sebastian Schwemer is Professor of Law and Technology at BI Norwegian Business School. He is also Associate Professor at the Centre for Information and Innovation Law (CIIR) at the University of Copenhagen, where he previously served as Director. He advised the European Commission on the DSA proposal.
Ida Koivisto
Ida Koivisto is Professor of public law, specializing in the digitalization of public administration at the University of Helsinki.

SUGGESTED CITATION

Schwemer, Sebastian; Koivisto, Ida: A Warm Body in the Loop: Rethinking Human Control of AI in EU Tech Regulation, Verfassungsblog, 18.09.2025, doi: 10.59704/aa18974494405499.

Licensed under CC BY-SA 4.0.

Langzeitarchiviert bei IntRechtDok unter: doi.org/10.17176/20250919-104620-0.