Article | Received 19 March 2025; Accepted 6 June 2025; Published 20 June 2025
https://doi.org/10.55092/let20250004

Negotiating ethical and technological ought in AI governance law: a SCOT analysis of GDPR Articles 22/25 and EU AI Act Recital 27

Nabila Khwaja* and Amal Robay

Weill Cornell Medicine-Qatar, Ar-Rayyan, Qatar

* Correspondence author; E-mail: nwk4001@qatar-med.cornell.edu.

Highlights:

- Applies SCOT theory to show how legal norms like consent and autonomy are socially constructed in AI regulation.
- Analyzes GDPR Articles 22/25 and AI Act Recital 27 as contested sites of ethical and technological negotiation.
- Explores the tension between technological 'ought' (efficiency, scalability) and ethical 'ought' (autonomy, dignity, oversight) in AI regulation.
- Challenges procedural views of consent, advocating for a more substantive, justice-oriented approach.
- Argues for participatory governance frameworks that embed moral agency and human rights into AI law.

Abstract: The use of artificial intelligence (AI) in high-risk domains like healthcare and human subject research raises critical ethical tensions particularly between 'technological ought' that prioritize efficiency and 'ethical ought' which focuses on autonomy, informed consent, and the principle of Respect for Persons. This article applies the Social Construction of Technology (SCOT) theory to analyze how these tensions are negotiated within legal instruments such as GDPR Articles 22 and 25, and Recital 27 of the EU AI Act. We explore how consent and autonomy core expressions of the Kantian principle for Respect for Persons (PRP) are socially constructed through interactions among regulators, developers, and civil society. SCOT reveals that legal protections are not static but reflect competing visions of accountability, transparency, and moral agency. Recital 27 of the EU AI Act, by exempting research and development applications, illustrates how anticipatory governance can be selectively applied, privileging innovation while potentially sidelining early-stage ethical safeguards in opaque domains like genomic diagnostics. We argue that meaningful consent, sustained human oversight, and the ethical commitment to respecting persons are essential to uphold ethical ought in AI system development. This article asks: how do legal norms around autonomy and consent become contested, negotiated, and stabilized through socio-technical processes in AI regulation, and how might this reshape our understanding of moral agency in law?

Keywords: SCOT theory in AI governance; GDPR Articles 22 and 25; EU AI Act Recital 27; ethical and technological 'ought' and moral agency in algorithmic regulation

1. Introduction

The use of artificial intelligence (AI) in high-risk domains like healthcare and human subject research raises critical ethical tensions, particularly between technological 'ought' that prioritize efficiency and ethical 'ought' to ground in autonomy, informed consent, and the principle of Respect for Persons (PRP). Recent interdisciplinary scholarship[1] increasingly emphasizes that both AI's socio-economic impacts and its legal governance are socially constructed outcomes of dynamic negotiations among technological, institutional, and ethical actors [1]. Building on this understanding, this article applies the Social Construction of Technology (SCOT) framework to analyze how key legal protections specifically Articles 22 & 25 of the General Data Protection Act (GDPR) and Recital 27 of the European Union Artificial Intelligence Act (EU AI Act) are shaped, contested, and stabilized within evolving AI governance landscapes. SCOT reveals that legal norms like consent, autonomy, and oversight are not fixed mandates but are fragile achievements, reflecting competing visions of accountability, transparency, and moral agency. Against this backdrop, we argue that sustaining meaningful consent, substantive human oversight, and respect for persons is essential to uphold ethical 'ought' in AI systems. The precautionary principle, a core pillar of European Union (EU) AI governance, advocates for proactive regulatory intervention to prevent irreversible harm before ethical safeguards are fully established [2,3]. This approach is evident in frameworks such as the General Data Protection Regulation (GDPR) [4]. EU AI, Act which shapes AI's ethical and legal trajectory through strict compliance requirements [5,6]. While these regulations aim to protect fundamental rights and mitigate risks, they may also be seen as prohibitive to innovation by imposing burdensome constraints that slow technological progress. The innovation principle promotes advancements while managing risks through minimal, case-specific regulations rather than broad restrictions. It assumes technology is generally beneficial and argues that market forces and existing laws can often mitigate risks without stifling progress. Unlike the precautionary principle[2], which supports anticipatory regulation in the face of uncertain harms, the innovation-oriented approach favors reactive intervention only in cases of demonstrable risk. This balance aims to prevent overly restrictive measures that could stifle AI development. While Recital 27 of the EU AI Act & Articles 22 & 25 of the GDPR, provide critical safeguards to regulate automated decision-making through data protection by design, such provisions may inadvertently pose barriers to innovation in complex fields like genomics [7]. Recital 27 of the EU AI Act introduces a partial counterbalance by exempting AI systems used exclusively for research and development from the regulation's scope. This exemption, while intended to foster innovation, also

[1]Approaches informed by ethics, law, and science and technology studies (STS) increasingly reject the notion that AI systems and their regulation are purely technical or objective. Instead, they foreground how design, deployment, and legal governance are shaped by normative judgments, institutional interests, and culturally embedded assumptions. What counts as 'risk,' 'harm,' or even 'autonomy' is not given by technology itself but constructed through legal classifications, ethical frameworks, and socio-political negotiation.

[2] In the context of EU laws, the precautionary principle serves as a foundational approach to regulatory decision-making in areas of environmental and technological risk. It guides policymakers to take preventive measures even when scientific evidence is uncertain, shifts the burden of proof to actors proposing potentially harmful actions (such as industries or developers), encourages the exploration of safer alternatives, and emphasizes inclusive public participation. Enshrined in Article 191(2) of the Treaty on the Functioning of the European Union (TFEU).

constructs a regulatory threshold that distinguishes between experimental and high-risk operational uses raising questions about how legal boundaries are drawn and justified in emerging scientific domains like genomic research. which is developing rapidly [8,9]. A pertinent example is the use of polygenic risk scores (PRS), where AI-driven models assess individual susceptibility to conditions like cancer or cardiovascular disease [10]. Although PRS promises personalized medicine and early intervention, it also surfaces ethical and legal concerns related to autonomy, consent, and data governance. The accuracy of AI predictions depends on training data, which often lacks genetic diversity, leading to biased risk assessments that may disproportionately misclassify individuals from underrepresented populations. Additionally, informed consent becomes complex, as patients may not fully understand how their genomic data will be used over time, particularly as AI models evolve. While GDPR Articles 22 & 25 mandate transparency, data protection, and the right to explanation, enforcing these safeguards remains challenging due to the lack of algorithmic transparency in many black-box models. A notable example is the Dutch childcare benefits scandal, where thousands of families were wrongly accused of fraud by a semi-automated system. Although an individual signed off on the final decisions, the overwhelming reliance on algorithmic outputs rendered the oversight ineffective demonstrating the legal and ethical limits of Article 22 in practice [11]. Beyond genomic research, the complexity of machine learning (ML) and deep learning (DL) models in AI-driven research presents additional regulatory challenges. The black-box nature of many AI systems obscures decision-making processes, making it difficult to ensure fairness, assess risks, and enforce ethical compliance [12]. Regulatory bodies, including Institutional Review Boards (IRBs), often lack domain expertise in AI, leading to gaps in oversight and inconsistencies in ethical evaluations [13]. Addressing these challenges requires a multi-faceted approach, including explainable AI (XAI) methodologies, pre-submission transparency reports, and interdisciplinary collaboration between AI researchers and regulators. While SCOT offers a valuable lens for understanding how legal norms like consent and autonomy are socially constructed, it has been critiqued for insufficient attention to structural power and normative commitments. By focusing on micro-level interactions and avoiding ethical stances, SCOT risks underplaying the impact of systemic inequalities and failing to challenge unjust power dynamics [14]. This is particularly relevant in AI governance, where the supposed neutrality of automated systems can obscure deeply embedded biases and unequal effects[3] [14,15]. In the context of Respect for Persons (PRP), such critiques highlight that formal rights like opting out of automated decisions may be inadequate if broader social structures prevent individuals from exercising those rights meaningfully. To remain relevant in high-risk technological contexts like AI, SCOT benefits from integration with PRP as well as other bioethical principles, elevating human autonomy and dignity. Recent interdisciplinary scholarship has emphasized that AI is not merely a technical domain, but a site where broader societal inequalities are reproduced, contested, and negotiated. For instance, applications of Margaret Archer's social realism highlight how AI technologies and governance structures are shaped by dynamic interactions between social context, agency, and structural power [1]. This article uses the SCOT framework to show how legal norms like consent, autonomy, and accountability are not fixed safeguards but socially constructed through

[3] See Tuba Bircan, 'Unmasking Inequalities of the Code: Disentangling the Nexus of AI and Inequality'. The authors argue that the inequalities produced by AI are not inevitable but socially constructed outcomes, shaped by agency, structural conditions, and embedded power dynamics. Drawing on Margaret Archer's social realism, they highlight how AI systems reflect ongoing interactions between social structures and human actors, reinforcing the need for governance approaches that address these deeper socio-technical roots

contested socio-technical processes. It asks how such norms are negotiated in AI governance, particularly through evolving interpretations of GDPR Articles 22 and 25 and Recital 27 of the EU AI Act.

2. Introducing SCOT: interpretive flexibility in law

While a comprehensive review of all scholarly discussions on Article 22 & 25 of the GDPR or Recital 27 of the EU AI Act is beyond the scope of this article, we have selected a subset of influential contributions that directly inform our SCOT-based analysis of autonomy, consent, and human oversight in AI governance. We prioritize analyses that exposes the ambiguity of key terms and show how legal meaning is shaped by institutional, technical, and normative negotiations rather than fixed formalism. SCOT theory emerged in response to the limitations of earlier theories of technological change, which were often rooted in linear, deterministic models that treated technology as an autonomous force. Its founders Bijker, Hughes, & Pinch argued that such models obscure the complex negotiations and power dynamics involved in shaping technological artefacts [16,17]. SCOT theory challenges the notion that technology evolves independently, instead emphasizing *interpretive flexibility*[4] the idea that different social groups assign different meanings to the same technology. Developed in the 1980s by Pinch and Bijker, SCOT presents technology design as an open, socially influenced process shaped through negotiation, contestation, and context-specific interactions [18]. For example, Pinch & Bijker's classic study of the bicycle illustrates how technological design is not inevitable, the eventual adoption of bicycle safety over the high-wheeled 'ordinary' model resulted from negotiations among diverse social groups. Technological change is shaped by human values, choices, and power dynamics not by inevitable progress. SCOT challenges deterministic views by highlighting how technologies including legal and regulatory systems are shaped through social negotiation. This is particularly relevant for AI governance in human subject research, where frameworks like the GDPR & the EU AI Act reflect contested values around autonomy and consent. Articles 22[5] and 25[6], for example, are not neutral rules but socio-technical artifacts shaped by tensions between innovation and ethical safeguards rooted in the principle of PRP a foundational bioethical concept tied to human dignity and moral worth [19]. The meaning and expression of PRP vary across cultures and contexts [20]. Historically, bioethics has emphasized PRP particularly through the lens of autonomy. This article treats PRP not as a fixed ethical norm but as a contested construct shaped through legal instruments, institutional practices, and technological design. Using SCOT, we analyze how PRP's expressions like consent and autonomy are negotiated in contexts such as in Articles 22 & 25 GDPR. This perspective underscores the interpretive flexibility of legal protections and highlights the importance of anticipatory governance in aligning ethical commitments with evolving AI systems.

The SCOT framework challenges technological determinism the idea that AI evolves autonomously and inevitably improves society. In human subject research, this assumption undermines accountability and human agency, especially in healthcare, where opaque algorithms can obscure informed consent and

[4] Interpretive flexibility refers to the idea that the meaning and function of a technology are not fixed or inherent but are shaped by different social groups who assign varying interpretations and uses to the same artifact. The concept originates from the Social Construction of Technology (SCOT) framework developed by Pinch and Bijker in the 1980s.

[5]Article 22 GDPR prohibits solely automated decisions with legal or significant effects, unless based on consent, law, or contractual necessity, and subject to safeguards like human review and the right to contest.

[6] Article 25 GDPR addresses the *effects* of automated decisions on individuals, Article 25 concerns the obligations of data controllers to proactively embed data protection into system design, i.e., Data Protection by Design and by Default (DPbD).

intensify privacy risks. As AI systems increasingly repurpose sensitive data beyond its original intent, enforcing data protection rights under the GDPR becomes more difficult. The growing autonomy of these systems complicates accountability for harm, underscoring the need for governance frameworks that embed transparency, human oversight, and ethical accountability across the entire AI research lifecycle [21]. Existing regulatory frameworks, including the Helsinki Declaration [22], and Belmont Report [23], often fall short in addressing the complexities of AI. This is particularly concerning in biomedical research involving big data, such as genomics, where AI processes sensitive personal data without clear mechanisms for consent withdrawal or effective oversight [24]. Additionally, AI-driven research tools could evolve independently of human oversight, there is a risk of where responsibility for harm or bias is displaced onto 'the system' rather than accountable researchers or institutions.

Regulatory lag intensifies the challenge of governing opaque and unpredictable AI systems, as traditional human subject protections struggle to keep pace particularly in jurisdictions like the United States [25]. In predictive healthcare and genomic research, AI design embeds assumptions about fairness, risk, and consent underscoring the need for ethical oversight throughout the lifecycle, as highlighted by SCOT-informed analyses. Developers influence AI through choices about data selection, model architecture, and risk thresholds, directly affecting research participant [26,27]. Therefore, AI used in human subject research must be critically examined through both ethical and legal lenses to ensure that it upholds established bioethical principles such as PRP [28]. However, SCOT has also faced important critiques particularly around its lack of attention to structural power and normative commitments. Critics have argued that by focusing primarily on micro-level interactions and refusing to adopt an ethical stance. It has been argued that SCOT can understate the influence of systemic inequalities and fail to challenge unjust power dynamics [29]. This is especially relevant in AI governance, where claims of neutrality can mask structural bias and inequality. In the context of PRP, formal rights like opting out of automation may be inadequate if broader social conditions limit individuals' ability to act on them. While SCOT usefully reveals how consent and autonomy are socially constructed, it must be paired with ethics-focused analysis that centers power, justice, and human dignity. In doing so, SCOT exposes how law mediates ethical 'ought' through uneven regulatory processes, underscoring that PRP in AI requires more than procedural compliance it demands ongoing scrutiny of power within socio-technical systems.

3. Article 22 GDPR as a socio-technical artifact: a SCOT analysis of interpretive flexibility

SCOT's concept of *interpretive flexibility* is particularly useful for understanding how Article 22 GDPR is mobilized, contested, and operationalized across legal, technical, and institutional contexts. The Schufa[7] ruling raised critical ambiguity about whether algorithmic outputs even when subject to nominal human oversight constitute legally binding decisions, challenging assumptions about "solely automated" processing. In response, Malgieri's proposal of a "right to legibility" emphasized the need for substantive interpretability in automated decision-making, aligning with SCOT's view that legal meaning is not intrinsic but constructed through sociotechnical negotiation [30].

This dynamic is echoed in the Italian DPA's bans on Replika and ChatGPT, which exemplify the dilemma of ex post regulation. While motivated by valid concerns such as privacy violations and lack

[7] Court of Justice of the European Union, Case C-634/21, SCHUFA Holding AG v. CJEU Judgment of 7 December 2023.

of safeguards the sudden withdrawal of AI systems already embedded in users' emotional routines created new harms, particularly for vulnerable individuals. These cases highlight that regulation acts not on neutral tools but on systems deeply entangled in affective, economic, and institutional dependencies. Concepts such as "harm," "protection," or "appropriate use" are not self-evident, but emerge from competing claims and contextual realities a core SCOT insight. The interpretive instability of Article 22 becomes even more visible when applied to semi-automated systems. Scholars such as Wachter and Mittelstadt (2019, 2024) argue that terms like "meaningful oversight" function more as aspirational goals than enforceable duties, blurring the boundary between human and machine agency [31,32]. Malgieri (2021, 2024) builds on this by distinguishing algorithmic functions profiling, prediction, and decision-making and calls for legal norms calibrated to degrees of agency and harm. His functionalist reading of "decision" shifts attention from procedural form to substantive impact, revealing how presumed consent and minimalist safeguards obscure deeper structural asymmetries [33,34]. Early critiques, such as Mendoza and Bygrave (2017), warned that vague thresholds like "significant effect" could obscure systemic bias under a veneer of legal formalism [35]. Building on this, Binns and Veale (2021, 2023) demonstrate how algorithmic triaging and multi-stage profiling can maintain the automated character of decisions even with nominal human involvement, reinforcing concerns about the inadequacy of formal safeguards [36–38]. From a SCOT perspective, these contributions exemplify how regulatory meaning is always provisional shaped by institutional constraints and technical design choices. Dominant interpretations often reflect not ethical consensus, but organizational interests and operational feasibility. Rather than treating Article 22 as a fixed or self-evident safeguard, these analyses reveal it as a contested regulatory artifact, with its scope and normative force continuously negotiated [38].

Davis (2023), for example, critiques the strategic framing of AI outputs as "advisory," which allows institutions to avoid liability while retaining decision-shaping power [39]. Van Kolfschooten (2024) calls for a "health-conformant" interpretation that centers patient autonomy in clinical contexts [40], while Lazcoz and de Hert warn that Article 22 risks becoming a "second-class right" absent systemic enforcement mechanisms such as DPIAs [41].These are not mere doctrinal refinements but interventions in an ongoing socio-technical negotiation that seek to embed legal protection in institutional workflows, not abstract text. Ultimately, Article 22 must be understood not as a fixed legal threshold but as a site of regulatory construction its meaning co-produced through legal ambiguity, infrastructural inertia, and the asymmetries of design agency. Reform, in this light, is not a matter of textual clarification alone, but of normative reconstruction, requiring participatory governance and resistance to premature closure. A further line of critique advanced by Netter (2022) proposes distinguishing between "open" and "opaque" AI systems, arguing that legal categories must evolve in light of the varying levels of system transparency and interpretability [42]. This insight reinforces SCOT's claim that regulatory meaning is shaped by technological affordances: legal norms must be responsive not only to abstract principles but also to the operational logics of different AI architectures. Netter's call to differentiate regulatory obligations based on system intelligibility directly challenges the one-size-fits-all approach embedded in Article 22. Davis (2023) similarly interrogates how regulatory ambiguity enables institutional actors to frame AI outputs as merely "advisory," thereby displacing accountability while retaining influence over decision outcomes [39]. This strategic ambiguity what SCOT would describe as the performance of compliance reveals how legal meaning is negotiated through institutional routines rather than determined by textual clarity [39]. Tosoni (2021) highlights another axis of interpretive tension: whether Article 22

should be treated as a general prohibition or an individual right. This doctrinal ambiguity determines not only how the rule functions but also where the burden of resisting automation is placed on the system or the individua. SCOT's concept of legal closure is evident in how conflicting interpretations of Article 22 hinder enforcement [43]. The Dutch childcare benefits scandal shows how algorithmic opacity displaced meaningful human oversight, despite nominal sign-off highlighting that procedural compliance may mask systemic flaws. From a SCOT perspective, Article 22 is not a fixed safeguard, but a socio-technical artifact shaped by institutional routines and design choices. Reform efforts like Malgieri's "right to legibility" and Netter's "open–opaque" distinction aim to restore its normative force but could face resistance from entrenched regulatory structures and evolving AI systems.

Applied legal scholarship on Article 22 particularly by Netter, Davis, Tosoni, Wachter, and Mittelstadt demonstrates that legal meaning is not settled in the text alone, but forged through dynamic, situated negotiations that reflect the interplay of normative claims, design logics, and institutional constraints. Article 22's regulatory potential lies in its co-constructed character: it is not a finished legal boundary, but a continuously evolving space of ethical and technological contestation. While Article 22 governs the outcomes of automated systems, Article 25 GDPR redirects focus to the design stage, marking a new site of socio-technical negotiation. It asks how legal obligations can shape the architecture of data processing from the outset. From a SCOT perspective, this shift reflects a move from post-decision accountability to the co-construction of safeguards during system development. Given the interpretive instability and susceptibility to minimal compliance practices, we argue that Article 22 requires a normative reconstruction grounded in participatory governance and functional accountability. Rather than relying solely on textual refinement, this means embedding ongoing stakeholder engagement particularly from affected communities into how oversight, decision-making, and harm are interpreted in practice. Echoing SCOT's insights, legal protections like Article 22 must be treated not as completed rules, but as evolving artifacts whose ethical legitimacy depends on inclusive, context-sensitive negotiation across the AI lifecycle.

4. Article 25 GDPR as a contested design norm: a SCOT analysis of privacy by design and interpretive closure

Article 25 GDPR shifts regulatory focus upstream by embedding legal obligations into AI system design through the principle of Data Protection by Design and by Default. While it appears to enshrine ethical ideals such as transparency and the principle of Respect for Persons (PRP), a SCOT perspective reveals Article 25 as a contested regulatory artifact its meaning shaped by institutional routines, technical affordances, and socio-technical negotiation. Rather than serving as a fixed safeguard, its interpretation often varies in practice, resulting in performative or minimal compliance. Critics highlight Article 25's definitional vagueness, which enables formalistic implementation and risks undermining its normative aims. Edwards & Veale (2017) were among the first to warn that its ambiguous obligations encourage technical compliance such as privacy toggles, dashboards, or pseudonymization without substantive engagement with autonomy or informed consent [44]. Veale (2018) adds that the language of "appropriate technical and organizational measures" is typically interpreted in institutionally convenient, risk-averse ways, often at the expense of meaningful privacy protections [45]. Rubinstein & Good (2020) similarly contend that while Article 25 formally elevates privacy by design into law, it lacks the concrete technical requirements necessary for robust implementation [46]. Waldman (2020) goes further, arguing

that the provision has strayed from its conceptual roots and requires a teleological reading one grounded in dignity and autonomy to make normative sense [47]. Together, these critiques reinforce SCOT's central insight: that legal meaning is not embedded in legal text alone but constructed through regulatory discretion and socio-technical mediation. In practice, Article 25 often gives rise to performative compliance. Veale & Borgesius (2021) note that formal features like consent mechanisms or anonymization may satisfy legal checklists while deeper power asymmetries such as backend data analytics and cross-platform sharing remain unchallenged [48]. The adtech sector's use of Real-Time Bidding (RTB) offers a compelling parallel to AI governance, particularly in how legal norms are interpreted and operationalized. As Veale *et al.* (2022) argue, RTB exemplifies the structural difficulty of reconciling GDPR's foundational requirements consent, transparency, and security with complex, distributed technological infrastructures. Despite appearing to operate within a regulated environment, RTB systems persistently circumvent the spirit of data protection law through institutionalized opacity and proceduralist consent regimes [49]. From a SCOT perspective, this reflects interpretive flexibility at scale: legal meaning around 'lawful basis,' 'informed consent,' and 'transparency' is not inherent to the GDPR but co-constructed through industry standards, technical affordances, and regulatory inertia. The failure of Article 25's data protection by design in this context reveals a stabilization of minimal compliance logics over normative intent, reinforcing SCOT's claim that closure often reflects institutional expediency rather than ethical consensus. In digital health, for instance, default settings may obscure data flow even as the system adheres to the letter of Article 25, raising concerns about the ethical adequacy of such designs under the data minimization principle. SCOT's concept of interpretive flexibility helps illuminate these dynamics. Legal safeguards are not pre-given, but shaped through socio-technical negotiations among regulators, engineers, compliance teams, and standards bodies. These negotiations often lead to what SCOT calls *closure*[8] a point at which one interpretation stabilizes not because it is ethically robust, but because it is institutionally convenient and technically feasible. Von Grafenstein *et al.* (2024) illustrate this through an interdisciplinary study on privacy icons, which shows that achieving genuine transparency under Article 25 requires coordination between law, UX design, and technical implementation [50]. Similarly, Mike (2022) found that in 49 enforcement cases, Article 25 was never cited as the sole legal basis for sanctions suggesting its limited practical traction and reinforcing the SCOT view that legal force must be constructed in practice, not presumed from textual authority [51]. Kalsi (2024) extends the critique initiated by Stalla-Bourdillon (2020), who warned that Article 25 risks becoming a hollow compliance ritual without enforceable rights and transparency mechanisms. While Stalla-Bourdillon (2020) highlights internal procedural weaknesses, Her call to extend legal responsibility across the full lifecycle of digital systems echoes SCOT's concern that closure can entrench convenient but ethically impoverished interpretations [52]. Kalsi (2024) focuses on an external jurisdictional gap: Article 25 rarely reaches upstream actors engineers, platform architects who play a pivotal role in shaping privacy outcomes [53]. To resist this closure, we argue for a more participatory and pluralistic interpretation of Article 25 one that involves civil society, affected users, and interdisciplinary voices in shaping what data protection means in practice. This aligns with

[8] In the Social Construction of Technology (SCOT) framework, "closure" refers to the point at which interpretive flexibility around a technology or legal norm diminishes, and one dominant meaning or design becomes stabilized. This stabilization is not necessarily driven by ethical superiority or technical excellence, but often by institutional convenience, power asymmetries, or what is seen as pragmatically feasible. This concept is elaborated in foundational SCOT work on sociotechnical change.

SCOT's broader claim: that privacy, consent, and transparency are not self-executing norms, but socio-technical artifacts, built and maintained through institutional choices and cultural values. Whereas Article 25 emphasizes embedding protections into design, Recital 27 of the EU AI Act complicates this logic by carving out exemptions for AI systems developed solely for research. This introduces a different kind of boundary work no longer between design and deployment, but between domains of use. From a SCOT perspective, Recital 27 functions as a regulatory artifact that performs ethical triage, demarcating which AI practices warrant oversight, and which are granted latitude in the name of innovation.

5. Recital 27 EU AI act as ethical boundary work: a SCOT and normative critique of research exemptions

To prevent ethical dilution under Recital 27, we call for tighter exemptions, layered oversight for high-risk AI research, and mandatory transparency tools like ethics impact assessments. Consent conditions must evolve as systems move toward deployment, reflecting SCOT's view that legal categories must be actively reshaped to uphold PRP and justice. This concern is echoed in recent scholarship. Fraser, Belloy, and Villarino (2024), for instance, argue that the research exemption relies on an implicit and inconsistently applied notion of "reasonableness," enabling sector-specific disparities and a form of regulatory trust whereby developers are expected to follow ethical norms without formal oversight [54]. Floridi (2021) critiques this framework as permitting ethically questionable experimentation under the guise of innovation, undermining ethical design from the earliest stages of AI development [55]. Mantelero (2025) offers a rights-based critique, contending that broad research exemptions threaten foundational legal safeguards such as informed consent, privacy, and human dignity core to the principle of Respect for Persons (PRP) [56]. He calls for narrowly defined exemptions, bound by rigorous ethical oversight, particularly in high-risk domains such as genomics and healthcare where harm may be profound and unequally distributed. Colonna (2023) extends this critique by illustrating how Recital 27 enables regulatory arbitrage in hybrid environments that blend public research, private enterprise, and commercial dissemination [57]. This flexibility, she argues, facilitates circumvention of both legal and ethical responsibilities, especially when research outputs cross into quasi-deployment scenarios. Malgieri (2023) emphasizes that exemptions like Recital 27 can exacerbate structural vulnerabilities, particularly in contexts such as healthcare where the burden of harm falls disproportionately on already marginalized populations [58]. His vulnerability aware reading of data protection law reinforces the need for stricter ethical safeguards to avoid embedding inequality in regulatory exceptions. Together, these critiques illuminate how Recital 27 operates as more than a regulatory carve-out: it constructs a hierarchy of values in which innovation and competitiveness are prioritized over moral accountability and ethical scrutiny. By enabling powerful actors such as major tech firms or elite research institutions to frame their work as ethically benign, the exemption narrows the interpretive space for precautionary or participatory responses. From a SCOT perspective, concepts like "informed consent," "meaningful oversight," and "high-risk classification" are not universal or fixed; they are socially negotiated and reflect the closure of alternative, potentially more ethically robust, interpretations. The scholarship of Veale and Borgesius (2021) further critiques the AI Act's vague categorizations of risk, transparency, and human oversight, demonstrating how such ambiguities can be exploited for selective compliance [48]. In this regulatory gray zone, autonomy itself becomes a contested concept ranging from minimal procedural transparency to substantive moral agency. To address this tension between technological 'ought' and

ethical 'ought,' we propose a hybrid SCOT-informed regulatory ethic. While SCOT illuminates how regulatory meaning is constructed through socio-technical negotiation, it must be paired with a normative framework capable of evaluating whether those constructions uphold justice, equity, and autonomy. As Winner (1986) and Hamilton (2022) argue, the social construction of technology must be ethically interrogated particularly in terms of how it impacts structurally marginalized groups [59,60]. This is especially important in AI governance, where those most affected by design decisions often lack the resources or representation to shape regulatory outcomes. Combining SCOT with the principle of Respect for Persons enables a richer critique of Recital 27: one that goes beyond institutional analysis to ask whether exemptions genuinely uphold the rights and agency of research participants and affected publics. Legal texts like Recital 27 must therefore be treated not as static boundaries but as evolving sites of ethical negotiation where law, technology, and morality intersect and where protections must be actively constructed, not presumed. Recital 27, in this light, exemplifies the broader structural tension at the heart of AI governance: the conflict between a technological 'ought' that valorizes innovation and a normative 'ought', grounded in human dignity, participatory accountability, and distributive justice. To navigate this conflict, regulatory instruments must not merely define permissible conduct but also foster environments in which ethical practices are institutionally supported and socially co-produced.

While SCOT helps us understand how consent, oversight, and autonomy are negotiated, it does not inherently assess whether those outcomes are ethically sufficient. A hybrid model that combines SCOT with PRP enables us to evaluate whether regulatory exemptions like Recital 27 truly respect persons not only through formal rights but through practices that enable those rights to be meaningfully exercised. Recital 27 offers a revealing case study in how legal texts perform ethical boundary setting. To govern AI in a way that reflects moral agency and respects human dignity, such texts must not be treated as static or technocratic. Instead, they should be seen as arenas of negotiation, where law, ethics, and technology continuously evolve in response to competing claims and shifting societal values. This dynamic understanding of legal texts as evolving ethical instruments sets the stage for examining a deeper structural tension at the heart of AI governance: the conflict between technological 'ought' and ethical 'ought'. To resist the ethical dilution enabled by Recital 27's broad exemption, we argue for a refined governance approach that reintroduces ethical safeguards without stifling innovation. This includes narrowing the exemption scope, introducing layered oversight for high-risk AI research, and mandating transparency measures such as pre-registration and ethics impact assessments. Additionally, consent and data-use conditions must evolve with AI system functionality, particularly in research outputs as they transition toward deployment. Such a model reflects SCOT's insight that legal categories are socially constructed and therefore must be actively reconstructed to sustain PRP, ensure distributive justice, and uphold human dignity within AI governance.

6. Negotiating the conflict between technological 'ought' and ethical 'ought' in AI development

The push to develop ever more powerful AI systems often prioritizes scalability, efficiency, and precision what AI *ought* to achieve in terms of technical performance. Yet this focus frequently overshadows ethical *ought* such as fairness, transparency, and accountability. When performance metrics dominate, developers and institutions may neglect their moral responsibilities, leading to biased or harmful outcomes. For instance, pulse oximeters have misread oxygen levels in non-white patients, while predictive policing algorithms have reinforced racial biases by replicating discriminatory historical data.

Such tools reduce individuals to probabilities, ignoring ethical imperatives like equity and dignity [61–63]. Similarly, generative AI systems have reproduced harmful stereotypes embedded in their training data. In healthcare, performance-driven AI may compromise explainability and patient autonomy, illustrating how technical benchmarks can dehumanize vulnerable groups when ethical considerations are sidelined. This misalignment between ethical and technological priorities underscores the need to integrate moral reasoning across the AI lifecycle. Waldman's critique of Article 25 of the GDPR reinforces this point: he argues that technological '*ought*' such as efficiency and scalability can eclipse ethical values when legal provisions are implemented procedurally rather than substantively. His emphasis on the discretionary and ambiguous nature of GDPR Article 25 aligns with the SCOT concept of *interpretive flexibility*, illustrating how "data protection by design" may function more as symbolic compliance than meaningful ethical governance [47]. By sidelining ethical reflection, this legal-technological emphasis diminishes the moral agency of stakeholders. Developers under commercial or institutional pressure may pursue benchmarks over values, while end-users often lack insight into AI system design and are disempowered from contesting problematic outcomes. These conditions are further intensified by global regulatory fragmentation, which encourages companies to operate in jurisdictions with weaker safeguards. The result is a climate of ethics washing, where superficial claims to fairness or transparency obscure the absence of genuine accountability mechanisms. Implementation [64–66]. From a SCOT perspective, these patterns illustrate how institutional power and commercial pressures shape not only the design of AI systems but also the weakening of legal norms intended to protect ethical principles.

Moral agency in AI governance involves the ethical responsibility of developers, regulators, and users throughout the AI lifecycle. It requires moral competence the ability to make and be accountable for ethical decisions a normative framework to guide behavior, and awareness of the situational constraints that shape choices. Active human oversight during AI training, deployment, and governance is essential to preserve human values and prevent ethics from being displaced by market logics. Yet power asymmetries in the AI ecosystem often prioritize rapid deployment over ethical development, undermining the moral agency of both individuals and institutions. Limited transparency hampers the ability of regulators and civil society to hold companies accountable, while weak enforcement mechanisms reduce opportunities for ethical deliberation. In this context, legal provisions like Articles 22 and 25 of the GDPR and Recital 27 of the EU AI Act must be understood not as fixed guarantees but as socio-technical artifacts contested and shaped by institutional interests, resource constraints, and competing interpretations.

7. From ethics washing to reconstructing moral agency: a SCOT perspective on human subject research

These dynamics are particularly salient in the context of AI systems used in human subject research, where legal and ethical responsibilities converge. Articles 22 and 25 of the GDPR, and Recital 27 of the EU AI Act, were designed to protect autonomy, privacy, and informed consent. Article 22 restricts solely automated decision-making that produces significant effects; Article 25 mandates data protection by design; Recital 27 exempts research-only AI from regulation. However, the enforcement of these provisions remains inconsistent and institution-dependent, shaped by ongoing tensions between innovation and ethical safeguards. Under a SCOT lens, these legal texts are not neutral rules but evolving socio-technical constructs whose meaning and impact depend on how they are interpreted and

implemented across different institutional settings. The technical opacity of machine learning (ML) and deep learning (DL) systems compounds these challenges. AI is not a singular, fixed tool it is a constellation of algorithmic processes shaped by human design choices and embedded institutional contexts. As complexity increases, so do the risks to transparency, fairness, and accountability. Regulatory efforts seek to address these risks by embedding principles such as explainability, auditability, and human oversight into system design. However, implementation gaps persist, especially when ethical safeguards are treated as compliance checkboxes rather than substantive commitments.

8. SCOT and the fragility of moral agency: legal accountability in AI governance

Applying SCOT to AI in human subject research reveals how legal frameworks shape governance by influencing power relations, assigning responsibility, and defining acceptable risk. High-risk applications such as AI in predictive healthcare demand proactive ethical safeguards long before deployment. Importantly, AI does not possess consciousness, intentionality, or moral reasoning; its ethical implications are determined entirely by the humans who design, deploy, and regulate it. While debates about transparency and accountability in algorithmic systems dominate the discourse, there has been comparatively little attention to how legal and institutional environments affect the moral agency of those involved. Traditional notions of legal and moral agency assume the capacity for intentional action and ethical deliberation. But as Ha (2020) notes, the rise of technological determinism the belief that AI development is inevitable and autonomous challenges these assumptions by diminishing the role of human agency [67].This deterministic framing risks absolving institutions of ethical and legal accountability, particularly in healthcare, where decisions about treatment or diagnosis are increasingly shaped by opaque algorithms. In contrast, SCOT emphasizes that technologies, including AI, are shaped by human actors, social values, and institutional forces and that legal responsibility must be understood within this socio-technical context. In predictive models trained on biased or incomplete historical data, AI may recommend suboptimal care for underrepresented populations. When harm results, it becomes difficult to determine accountability: is it the developers, the deploying institution, or the clinician using the tool? This ambiguity disrupts traditional legal models that rely on clearly attributable human intent. To address such gaps, ethical safeguards must be embedded throughout the AI lifecycle. Legal frameworks like the GDPR and AI Act seek to restore human oversight through requirements for explainability, robustness, and transparency particularly for high-risk systems like healthcare AI. But when interpreted narrowly or implemented minimally, these provisions risk becoming hollow. Ethical responsibility must remain central, even when automation reduces direct human involvement.

9. Conclusion: AI governance as a contested and negotiated socio-technical space

AI governance is not the neutral application of fixed legal rules, but an ongoing socio-technical negotiation shaped by legal interpretation, institutional routines, and technological constraints. Scholars such as Floridi, Veale, Wachter, and Mantelero have emphasized that law must do more than codify protections it must respond to evolving power dynamics and the ethical complexities embedded in AI systems. This article has shown, through the lens of SCOT, that legal provisions such as GDPR Articles 22 and 25 and Recital 27 of the EU AI Act are not static safeguards but contested artifacts. Their meaning is co-produced by designers, regulators, institutions, and affected publics. Article 22's effectiveness hinges not on its

textual clarity, but on how terms like "meaningful oversight" and "significant effect" are interpreted in context. Article 25, while promising privacy by design, often collapses into minimal compliance due to vague standards and implementation gaps. Recital 27 performs a different kind of boundary work valorizing innovation by exempting research AI from oversight, but in doing so, it risks diluting core ethical protections, especially for vulnerable populations. What emerges is a structural tension between a technological 'ought' centered on efficiency, scalability, and performance and an ethical 'ought' grounded in autonomy, transparency, and human dignity. To navigate this tension, we argue for a hybrid approach that combines SCOT's attention to interpretive flexibility with the normative commitments of the principle of PRP. This enables a more reflexive and ethically grounded understanding of how legal protections must be actively constructed and sustained not assumed. Ultimately, meaningful AI governance requires not only anticipatory legal design but also participatory processes that elevate moral agency and embed accountability across the AI lifecycle. By recognizing law as a dynamic, negotiated terrain, we can better ensure that AI systems serve human ends, rather than displacing the very values they purport to uphold. By examining how developers, regulators, and ethicists co-construct the meaning of legal norms, it becomes possible to re-embed ethics not only in the design of AI but in the very structures through which it is governed.

Acknowledgments

This research did not receive any specific funding. The authors would like to express their sincere gratitude to the professors at the University of Bradford for their invaluable support and guidance, particularly in the development of the first author. Their insights into the intersection of law and ethics were instrumental in shaping the theoretical framework of this paper. Special thanks go to the faculty members who taught the author about the ethical dimensions of AI and its regulation, fostering a deeper understanding of the complex relationship between legal structures and ethical considerations. Their contributions have significantly enriched this article.

Authors' contribution

Conceptualization, Nabila Khwaja; formal analysis, Nabila Khwaja; investigation, Nabila Khwaja; resources, Amal Robay; writing—original draft preparation, Nabila Khwaja; writing—review and editing, Amal Robay. All authors have read and agreed to the published version of the manuscript.

Conflicts of interests

The authors declare no conflict of interest.

References

[1] Lyytinen KJ, Klein HK. The critical theory of Jürgen Habermas as a basis for a theory of information systems. In *Research Methods in Information Systems,* 1st ed. Amsterdam: North-Holland Publishing Co., 1985. pp. 1–20.

[2] Sunstein CR. Beyond the precautionary principle. *U. Pa. L. Rev.* 2002, 151:1003.

[3] Brooman, S. Politics, law, and grasping the evidence in fur farming: a tale of three continents. In *The Ethics of Fur: Religious, Cultural, and Legal Perspectives*, 1st ed. Lanham: Rowman & Littlefield Publishing Group, 2023, pp. 37–55.

[4] European Commission. GDPR—the fabric of a success story. 2020, pp. 1–12. Available: https://commission.europa.eu/law/law-topic/data-protection/eu-data-protection-rules/gdpr-fabric-success-story_en (accessed on 16 June 2025).

[5] European Union. Proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain U nionlegislative acts. 2021. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206 (accessed on 16 June 2025).

[6] Djeffal C. The EU AI Act at a crossroads: generative AI as a challenge for regulation. 2023. Available: https://www.pymnts.com/cpi-posts/the-eu-ai-act-at-a-crossroads-generative-ai-as-a-challenge-for-regulation/?utm_source=chatgpt.com (accessed on 30 Januray 2025).

[7] Castro D, McLaughlin M. Ten ways the precautionary principle undermines progress in artificial intelligence. 2019. Available: https://itif.org/publications/2019/02/04/ten-ways-precautionary-principle-undermines-progress-artificial-intelligence/ (accessed on 30 Januray 2025).

[8] Zhang Q, Wang Y. AI in biology: transforming genomic research with machine learning. *Comput. Mol. Biol.* 2024, 14(3):106–114.

[9] Zhang S, Bamakan SMH, Qu Q, Li S. Learning for personalized medicine: a comprehensive review from a deep learning perspective. *IEEE Rev. Biomed. Eng.* 2018, 12:194–208.

[10] Wray NR, Lin T, Austin J, McGrath JJ, Hickie IB, *et al.* From basic science to clinical application of polygenic risk scores: a primer. 2021, pp. 101–109. Available: https://jamanetwork.com/journals/jamapsychiatry/article-abstract/2771079 (accessed on 4 February 2025).

[11] Mittelstadt BD, Allo P, Taddeo M, Wachter S, Floridi L. The ethics of algorithms: mapping the debate. *Big Data Soc.* 2016, 3(2):2053951716679679.

[12] Durán JM, Jongsma KR. Who is afraid of black box algorithms? On the epistemological and ethical basis of trust in medical AI. *J. Med. Ethics* 2021, 47(5):329–335.

[13] Doerr M, Meeder S. Big health data research and group harm: the scope of IRB review. *Ethics & Hum. Res.* 2022, 44(4):34–38.

[14] Doezema T, Frahm N. The new spirit of technoscience: recalibrating symmetrical STS critique. *J. Responsible Innov.* 2023, 10(1):2281112.

[15] Winner L. Upon opening the black box and finding it empty: social constructivism and the philosophy of technology. *Sci., Technol. Hum. Values* 1993, 18(3):362–378.

[16] Bircan T, Özbilgin MF. Unmasking inequalities of the code: disentangling the nexus of AI and inequality. *Technol. Forecasting Social Change* 2025, 211:123925.

[17] Pinch TJ, Bijker WE. The social construction of facts and artefacts: Or how the sociology of science and the sociology of technology might benefit each other. *Social stud. of Sci.* 1984, 14(3): 399–441.

[18] Bijker WE. *Of bicycles, bakelites, and bulbs: toward a theory of sociotechnical change*, 1st ed. Cambridge: MIT press. 1997.

[19] Klein HK, Kleinman DL. The social construction of technology: structural considerations. *Sci., Technol. Hum. Values* 2002, 27(1):28–52.

[20] Neumann M. Did Kant respect persons? *Res Publica* 2000, 6:285–299.

[21] Gostin LO. Informed consent, cultural sensitivity, and respect for persons. *Jama* 1995, 274(10):844–845.
[22] Allen C, Wallach W. Moral machines: contradiction in terms or abdication of human responsibility. In *Robot Ethics: The Ethical and Social Implications of Robotics*, 1st ed. Cambridge: MIT Press, 2012, pp. 55–68.
[23] Ashcroft RE. The declaration of Helsinki. In *The Oxford textbook of clinical research ethics*, 1st ed. New York: Oxford University Press, Inc., 2008. pp. 141–148.
[24] Adashi EY, Walters LB, Menikoff JA. The Belmont report at 40: reckoning with time. *Am. J. Public Health* 2018, 108(10):1345–1348.
[25] Balagurunathan Y, Sethuraman RR. An analysis of ethics-based foundation and regulatory issues for genomic data privacy. *J. Inst. Eng. India Ser. B* 2024, 105(4):1097–1107.
[26] Analytica O. US states step up AI rule-making amid federal lag. 2024. Available: https://www.emerald.com/insight/content/doi/10.1108/oxan-db287615/full/html (accessed on 13 February 2025).
[27] Nashawaty P. Analyst insight the evolving role of developers in the AI revolution. 2024. Available: https://futurumgroup.com/insights/the-evolving-role-of-developers-in-the-ai-revolution/ (accessed on 14 February 2025).
[28] Griffin TA, Green BP, Welie JVM. The ethical agency of AI developers. *AI Ethics* 2024, 4(2):179–188.
[29] Holm S. *Principles of Biomedical Ethics*, 5th ed. Oxford: Oxford University Press, 2001. p. 454.
[30] Basu S. Three decades of social construction of technology: dynamic yet fuzzy? The methodological conundrum. *Social epistemology* 2023, 37(3):259–275.
[31] Malgieri G, Hildebrandt JPDEM, Fuster GG, Stalla-Bourdillon S, Sibony AL, *et al.* Making choices in the digital economy. PhD Thesis, Vrije Universiteit Brussel, 2024.
[32] Wachter S, Mittelstadt B, Russell C. Do large language models have a legal duty to tell the truth? *R. Soc. Open Sci.* 2024, 11(8):240197.
[33] Wachter S, Mittelstadt B. A right to reasonable inferences: re-thinking data protection law in the age of big data and AI. *Colum. Bus. L. Rev.* 2019:494.
[34] Malgieri G, Comandé G. Why a right to legibility of automated decision-making exists in the general data protection regulation. *Int. Data Privacy Law* 2017, 7(4):243–265.
[35] Malgieri G, Pasquale F. Licensing high-risk artificial intelligence: toward ex ante justification for a disruptive technology. *Comput. Law & Secur. Rev.* 2024, 52:105899.
[36] Mendoza I, Bygrave LA. The right not to be subject to automated decisions based on profiling. In *EU Internet Law: Regulation and Enforcement*, 1st ed. Cham: Springer, 2017. pp. 77–98.
[37] Binns R, Veale M. Is that your final decision? Multi-stage profiling, selective effects, and Article 22 of the GDPR. *Int. Data Privacy Law* 2021, 11(4):319–332.
[38] Veale M, Silberman MS, Binns R. Fortifying the algorithmic management provisions in the proposed Platform Work Directive. *Eur. Labour Law J.* 2023, 14(2):308–332.
[39] Binns R, Veale M. Is that your final decision? Multi-stage profiling, selective effects, and Article 22 of the GDPR. *Int. Data Privacy Law* 2021, 11(4):319–332.

[40] Davis P, Schwemer SF. Rethinking decisions under Article 22 of the GDPR: implications for semi-automated legal decision-making. In *Proceedings of the Third International Workshop on Artificial Intelligence and Intelligent Assistance for Legal Professionals in the Digital Workplace (LegalAIIA 2023)*, Braga, Portugal, June 19, 2023, pp. 1–9.

[41] van Kolfschooten HB. A health-conformant reading of the GDPR's right not to be subject to automated decision-making. *Med. law Rev.* 2024. 32(3):373–391.

[42] Lazcoz G, De Hert P. Humans in the GDPR and AIA governance of automated and algorithmic systems. Essential pre-requisites against abdicating responsibilities. *Comput. Law & Secur. Rev.* 2023. 50:105833.

[43] Netter E. The human and machine part in "automated" decisions. Proposals for a rewriting of Article 22 of the GDPR. 2022, pp. 1–21. Available: https://hal.science/hal-03701045/ (accessed on 17 Febuary 2025).

[44] Tosoni L. The right to object to automated individual decisions: resolving the ambiguity of Article 22(1) of the General Data Protection Regulation. *Int. Data Privacy Law* 2021. 11(2):145–162.

[45] Edwards L, Veale M. Slave to the algorithm? Why a 'right to an explanation' is probably not the remedy you are looking for. *Duke L. & Tech. Rev.* 2017. 16:18.

[46] Veale M, Binns R, Ausloos J. When data protection by design and data subject rights clash. *Int. Data Privacy Law* 2018. 8(2):105–123.

[47] Rubinstein IS, Good N. The trouble with Article 25 (and how to fix it): the future of data protection by design and default. *Int. Data Privacy Law* 2020. 10(1):37–56.

[48] Waldman AE. Data protection by design? A critique of Article 25 of the GDPR. *Cornell Int'l LJ* 2020, 53:147.

[49] Veale M, Zuiderveen Borgesius F. Demystifying the draft EU Artificial Intelligence Act—analysing the good, the bad, and the unclear elements of the proposed approach. *Comput. Law Rev. Int.* 2021, 22(4):97–112.

[50] Veale M, Borgesius FZ. Adtech and real-time bidding under European data protection law. *Ger. Law J.* 2022, 23(2):226–256.

[51] von Grafenstein M, Kiefaber I, Heumüller J, *et al.* Privacy icons as a component of effective transparency and controls under the GDPR: effective data protection by design based on art. 25 GDPR. *Comput. Law & Secur. Rev.* 2024, 52:105924.

[52] Mike N. A case study discovering the potential for algorithmic decision-making on setting GDPR fines. *Acta Univ. Sapientiae, Leg. Stud.* 2021, 10(2):215–230.

[53] Stalla-Bourdillon S, Thuermer G, Walker J, Carmichael L, Simperl E. Data protection by design: building the foundations of trustworthy data sharing. *Data Policy* 2020. 2:e4.

[54] Kalsi M. Still losing the race with technology? Understanding the scope of data controllers' responsibility to implement data protection by design and by default. *Int. Rev. of Law, Comput. Technol.* 2024, 38(3):346–368.

[55] Fraser H, y Villarino JMB. Acceptable risks in Europe's proposed AI act: reasonableness and other principles for deciding how much risk management is enough. *Eur. J. Risk Regul.* 2024, 15(2):431–446.

[56] Floridi L. The European legislation on AI: a brief analysis of its philosophical approach. *Philos. Technol.* 2021, 34(2):215–222.

[57] Mantelero A. The AI Act: a realpolitik compromise and the need to look forward. In *Digital Constitutionalism*, 1st ed. Baden-Baden: Nomos Verlagsgesellschaft mbH & Co. KG, 2025. pp. 1–34.

[58] Colonna L. The AI Act's research exemption: a mechanism for regulatory arbitrage? In *YSEC Yearbook of Socio-Economic constitutions 2023: Law and the governance of Artificial Intelligence*, 1st ed. Cham: Springer. 2023. pp. 51–93.

[59] Malgieri G. *Vulnerability and data protection law*, 1st ed. Oxford: Oxford University Press, 2023.

[60] Winner L. *The whale and the reactor: a search for limits in an age of high technology*, 2nd ed. Chicago: University of Chicago Press, 2020.

[61] Hamilton E, Feenberg A. The technical codes of online education. *E-learn. and Digital Media* 2005, 2(2):104–121.

[62] Tobin MJ, Jubran A. Pulse oximetry, racial bias and statistical bias. *Ann. Intensive Care* 2022, 12:1–2.

[63] Shi C, Goodall M, Dumville J, Hill J, Norman G, *et al.* The accuracy of pulse oximetry in measuring oxygen saturation by levels of skin pigmentation: a systematic review and meta-analysis. *BMC Med.* 2022, 20(1):267.

[64] Valbuena VS, Merchant RM, Hough CL. Racial and ethnic bias in pulse oximetry and clinical outcomes. *JAMA Intern. Med.* 2022, 182(7):699–700.

[65] Bietti E. From ethics washing to ethics bashing: a view on tech ethics from within moral philosophy. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, Barcelona, Spain, January 27–30, 2020, pp. 210–219.

[66] Wagner B. Ethics as an escape from regulation. From "ethics-washing" to ethics-shopping? 2018, pp. 1–6. Available: file:///D:/Downloads/10.1515_9789048550180-016.pdf (accessed on 24 Febuary 2025).

[67] Reisman D, Schultz J, Crawford K, Whittaker M. Algorithmic impact assessments: a practical frame work for public agency. 2018, pp. 1–22. Available: https://www.nist.gov/system/files/documents/2021/10/04/aiareport2018.pdf (accessed on 25 February 2025).

[68] Ha YJ. For or against progress?: Institutional agency in a time of technological exceptionalism. In *ETHICOMP* 2020, Logroño, Spain, June 17–19, 2020. pp. 425–427.