



OPEN ACCESS

EDITED BY

Amon Rapp,
University of Turin, Italy

REVIEWED BY

Jahna Otterbacher,
Open University of Cyprus, Cyprus
Simone Di Plinio,
University of Studies G. d'Annunzio
Chieti and Pescara, Italy

*CORRESPONDENCE

Vladimír Šucha
✉ vladimir.sucha@eui.eu

RECEIVED 26 November 2025

REVISED 12 March 2026

ACCEPTED 23 March 2026

PUBLISHED 02 April 2026

CITATION

Šucha V (2026) The silent accumulation:
AI as mental contaminant.
Front. Artif. Intell. 9:1755151.
doi: 10.3389/frai.2026.1755151

COPYRIGHT

© 2026 Šucha. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

The silent accumulation: AI as mental contaminant

Vladimír Šucha*

European University Institute, Florence, Italy

This paper introduces a theoretical framework for understanding how seemingly innocuous AI systems create cumulative effects on human cognition, emotion, and agency that current governance approaches fail to address. Drawing on environmental health science, we propose that “low-risk” AI applications, those falling below regulatory thresholds in frameworks such as the EU AI Act, function as cognitive environmental contaminants whose collective and sustained presence may reshape human psychological capacities. We operationalise cumulative AI exposure along five dimensions (frequency, duration, intensity, diversity of systems, and developmental timing) and identify five pathways through which cumulative effects may manifest: attention erosion, emotional dependency, social connection alteration, decision-making dependency, and identity fragmentation. For each pathway, we distinguish empirical regularities documented in existing research, plausible mechanisms through which cumulative effects may operate, and speculative population-level hypotheses that require empirical testing. Situating our framework against adjacent literatures including technostress, cognitive offloading, hypernudging, and automation bias, we argue that the distinctive contribution lies in the cumulative, cross-system, population-level analytical paradigm and its governance translation. We propose three governance mechanisms—cumulative impact assessment extending existing algorithmic auditing frameworks, cognitive-social environmental monitoring using validated psychometric instruments, and economic valuation of cognitive-social ecosystem services, accompanied by a phased validation strategy. The framework is offered as a complement to existing risk-based governance, addressing a specific gap: the systematic invisibility of effects that emerge from the interaction of multiple AI systems over extended time periods.

KEYWORDS

algorithmic impact assessment, artificial intelligence governance, cognitive environmental contamination, cognitive-social ecosystem services, cumulative exposure, human agency

1 Introduction

When environmental scientists first began measuring trace amounts of heavy metals in water supplies or atmospheric particulates in urban environments, these levels were often deemed “acceptable” or “low-risk” by regulatory standards. Individual exposures appeared inconsequential. Yet decades later, the cumulative impact of these persistent, low-level contaminants revealed devastating consequences: cognitive impairment in children exposed to lead, cardiovascular disease from air pollution, and bioaccumulation of toxins throughout ecosystems (UNEP, 2025).

We now face a parallel situation with artificial intelligence, not in our physical environment, but in our cognitive and social landscape. While governance frameworks focus

primarily on high-risk AI applications, this paper explores how seemingly innocuous, “low-risk” AI systems may, through cumulative exposure, fundamentally reshape human agency and societal functioning in ways that current technological governance approaches fail to address.

This phenomenon is the complex interplay between technological development and social transformation, where the choices available to society regarding technology use have profound implications for human flourishing (Winner, 1980). As Bijker et al. (1987) demonstrated, technological artifacts are socially constructed, but they also construct social realities in return. Our framework extends this insight to show how AI’s social construction of reality operates through subtle, cumulative cognitive and emotional influences that accumulate below the threshold of regulatory attention.

This paper makes three primary contributions. First, we introduce a novel conceptual framework for understanding technology’s cumulative societal effects. Second, we identify five specific pathways through which AI systems possibly create technological externalities that accumulate beyond individual system impacts. Third, we propose governance mechanisms that address cumulative effects rather than individual system risks, contributing to broader discussions about democratic technology governance.

The central argument is that current technological governance approaches suffer from a fundamental assessment gap: they evaluate systems individually while missing collective impacts that may pose greater risks to human agency and social cohesion than any single application. By extending environmental health principles to cognitive and social domains, we can develop more comprehensive approaches to ensuring that technological development serves human flourishing rather than undermining it.

2 From environmental to cognitive-social protection

2.1 Expanding the analytical frame

Environmental health science offers a valuable framework for understanding technology’s potential long-term societal impacts. Just as Rachel Carson’s “Silent Spring” (1962) revealed how seemingly safe pesticide levels could accumulate into ecological disaster, we need frameworks for understanding how seemingly innocuous technological systems might accumulate into social and psychological disruption.

Consider lead exposure, once considered safe at “low levels.” We now know there is no safe threshold for lead in children’s blood (UNEP, 2025). Similarly, carbon emissions, individually negligible from a single vehicle, collectively transform global climate systems. The environmental health paradigm highlights three critical patterns that apply equally to technological exposure (Colborn et al., 1997):

First, cumulative effects exceed the sum of individual exposures. Just as multiple chemical exposures can interact synergistically, creating effects more severe than any single contaminant, the combined influence of multiple technological systems may amplify beyond their individual impacts.

Second, there is latency between exposure and manifestation. Environmental contamination often produces delayed effects, with decades passing before consequences fully emerge. Similarly, the

cognitive and psychological effects of technological immersion may only become apparent after prolonged exposure.

Third, there is variable vulnerability across populations. Environmental toxins disproportionately affect developing bodies and marginalized communities (Colborn et al., 1997). Likewise, technological effects may particularly impact developing minds and vulnerable populations.

2.2 Conceptualizing AI as cognitive-social environmental contaminant

What constitutes “low-level technological exposure” in this framework? We define it as regular interaction with systems that subtly shape cognitive and emotional processes without crossing thresholds of perceived harm or risk. This includes recommendation algorithms that imperceptibly narrow information environments, predictive text systems that subtly shape language patterns, engagement-optimizing content feeds that modulate emotional states, convenience-enhancing AI assistants that gradually supplant skills, and personalization systems that incrementally reinforce existing beliefs.

None of these systems typically raise regulatory concerns under risk-based frameworks. Yet their collective presence creates a cognitive-social environment fundamentally different from the one human minds and societies evolved to navigate, a transformation comparable to how industrial technologies altered the physical environment our bodies evolved to inhabit.

This perspective may bring insights about how technologies become embedded in social practices and reshape human capabilities (Latour, 1992). However, while STS typically focuses on how technologies are socially constructed, our framework examines how they may be reconstructing human cognition and social connection through mechanisms that operate below the threshold of conscious awareness.

2.3 Related frameworks and differentiation

The proposition that digital technologies reshape cognition is not new. Several established research programs address overlapping dimensions, and our framework must be situated within that landscape.

The technostress literature (Tarafdar et al., 2015) documents how digital demands create acute and chronic stress responses. Our framework differs in focusing on effects operating below conscious awareness: the gradual reshaping of capacities through systems users perceive as benign. Where technostress asks how technology makes us feel, we ask how it changes what we are capable of feeling, thinking, and deciding.

Risko and Gilbert (2016) formalised cognitive offloading using external tools to reduce internal cognitive demands, and substantial work has examined how delegation affects memory (Sparrow et al., 2011), spatial cognition (Dahmani and Bohbot, 2020), and reasoning (Wilmer et al., 2017). Our framework extends this by embedding individual offloading within a population-level cumulative model that considers interactions across multiple domains simultaneously and translates findings into governance implications.

Yeung (2017) introduced “hypernudges”—Big Data-driven choice architectures that are networked, continuously updated, and pervasive. Our framework complements this regulatory analysis by examining the downstream psychological consequences of prolonged exposure to hypernudging environments.

Research on automation bias demonstrates that increasing automation reduces experienced agency (Berberian et al., 2012) and that

environmental control manipulations modulate the prospective sense of agency, with individual vulnerability factors acting as moderators (Di Plinio et al., 2019). Our framework embeds these findings within a broader cumulative model, arguing that repeated micro-delegations across multiple systems may produce aggregate agency depletion.

Most directly adjacent, Di Plinio (2025) reviewed AI threats to agency and identity, advocating psychometric assessment tools and neurorights frameworks. Our key differentiation is analytical level: where Di Plinio focuses on individual psychological assessment, we introduce an environmental-population paradigm foregrounding cumulative effects across systems and proposing governance mechanisms drawn from environmental regulation.

The distinctive contribution of our framework is therefore not the identification of any single pathway, most have been documented in adjacent literatures, but three interrelated propositions: (1) individually documented effects interact and accumulate across domains in ways current frameworks cannot capture; (2) the appropriate governance response is cumulative impact monitoring analogous to environmental protection; and (3) the cognitive and social capacities affected constitute public goods whose degradation represents quantifiable externalities.

2.4 Operationalising cumulative exposure

If cumulative exposure is to serve as more than a metaphor, it requires operationalisation. Drawing on environmental health exposure science (WHO, 2020), we propose that cumulative AI exposure be characterised along five dimensions: frequency (how often one interacts with AI-mediated systems), duration (length of episodes and total lifetime exposure, including developmental windows), intensity (degree to which systems actively shape cognition through personalisation and engagement optimisation), diversity of systems (the number and variety of concurrent AI exposures, paralleling environmental mixture effects; Colborn et al., 1997), and developmental timing (life stage of exposure, given disproportionate vulnerability during cognitive development).

We further distinguish three cumulation models generating distinct testable predictions. In additive cumulation, total impact is the sum of individual exposures. In synergistic cumulation, interactions amplify effects beyond their additive sum—for example, algorithmic content curation combined with AI-mediated social interaction may reshape identity formation more profoundly than either alone. In threshold cumulation, effects remain negligible until a critical exposure level triggers nonlinear shifts.

Key moderators influencing individual vulnerability include baseline self-regulatory capacity (Diamond, 2013), social embeddedness, prior digital literacy, and socioeconomic context. These generate testable hypotheses about differential vulnerability and inform targeted intervention design. This operationalisation does not resolve measurement challenges, but transforms cumulative exposure from a rhetorical device into a structured analytical framework with identifiable dimensions and testable interaction models.

3 Five pathways of cognitive-social contamination

Two clarifications regarding the five pathways are necessary. First, while presented as analytically distinct, they overlap and interact causally. Attention erosion may exacerbate emotional dependency;

decision-making dependency may feed back into identity fragmentation. These feedback loops are central to the cumulative-exposure thesis: the interaction among pathways may produce aggregate effects exceeding the sum of individual impacts. Second, the pathways are not exhaustive; they represent domains where current evidence is most suggestive.

We also note a definitional boundary. While our pathways draw on broader digital technology research, our argument is that AI (systems using machine learning to optimise outputs based on user data) represents a qualitative intensification of these dynamics. Where we draw on non-AI-specific evidence, we note where extrapolation requires further validation. Each pathway below distinguishes: (a) empirical regularities, (b) plausible mechanisms, and (c) speculative cumulative hypotheses with testable predictions.

3.1 Attention erosion

3.1.1 Empirical regularities

Heavy media multitaskers show greater susceptibility to interference, though meta-analytic pooled effect sizes are small (Parry and Roux, 2021). Wilmer et al. (2017) found suggestive but largely correlational associations between smartphone habits and attentional difficulties. Ward et al. (2017) reported smartphone mere presence reduces cognitive capacity, though replication has produced mixed results (Ruiz Pardo and Minda, 2022). Collectively, evidence indicates a pattern of associations, though predominantly correlational and small to moderate in effect size (Delgado et al., 2018).

3.1.2 Plausible mechanisms

The cognitive offloading framework (Risko and Gilbert, 2016) suggests that when external systems consistently manage information filtering, internal attentional processes may undergo functional reorganisation (Clark, 2015). AI-driven content optimisation intensifies this through continuous real-time personalisation, creating feedback loops between attention patterns and algorithmic content selection that may progressively narrow attentional repertoires.

3.1.3 Speculative cumulative hypothesis

We hypothesise that prolonged population-wide exposure to AI-optimised environments may produce measurable shifts in sustained attention distributions—most pronounced in individuals whose exposure began during developmental periods and interacting synergistically with diversity of system exposure. If sustained, such shifts may have consequences for democratic governance, which depends on citizens' capacity to engage with complex policy arguments over extended periods. A population whose attentional baseline has shifted toward fragmented processing may find deliberative demands increasingly difficult to meet—a subtle form of democratic capacity erosion that no individual system produces but that the cumulative environment may generate.

Testable predictions: This hypothesis would be weakened if longitudinal studies show no attentional changes beyond pre-existing traits, if populations with different exposure levels show no distributional differences, or if changes prove fully reversible.

3.2 Emotional dependency

3.2.1 Empirical regularities

Odgers and Jensen (2020) caution that large-scale preregistered studies report small associations between digital technology use and adolescent wellbeing, unlikely to be of clinical significance in isolation. Nevertheless, algorithmic content curation modulates mood states (Kramer et al., 2014), and emerging research on AI companion use suggests patterns of emotional reliance (Pentina et al., 2023).

3.2.2 Plausible mechanisms

Two mechanisms are proposed: emotional outsourcing meaning reliance on algorithmically mediated stimulation rather than developing internal regulation capacities; and emotional calibration drift—prolonged exposure to amplified emotional content shifts baseline expectations, making naturally occurring experiences relatively less satisfying.

3.2.3 Speculative cumulative hypothesis

Cumulative exposure may gradually erode internal emotional regulation, producing increased reactivity without digital stimulation and difficulty with self-directed regulation in non-digital contexts. This engages fundamental questions about agency: if agency requires emotional self-regulation then gradual outsourcing of this capacity may undermine agency at its affective foundation (Borgmann, 1984), not through restricting choices but through atrophy of capacities that make autonomous choice meaningful.

Testable predictions: Weakened if emotional regulation shows no decline in heavy users over longitudinal observation, or if changes are fully attributable to pre-existing differences.

3.3 Social connection alteration

3.3.1 Empirical regularities

AI-mediated communication strips multimodal cues triggering bonding responses (Sætra, 2020; Pugh, 2023); algorithmic filtering curates encounters by engagement criteria rather than relationship-building criteria. Recent research shows algorithmic personalization impacts social connectedness (Taylor and Chen, 2024; Wang et al., 2022).

3.3.2 Plausible mechanisms

Proposed mechanisms include: social signal attenuation when AI-mediated interaction reduces the neurochemical and nonverbal signals that underlie empathy and social bonding; and algorithmic relationship filtering when platforms prioritize engagement metrics over authentic relationship development.

3.3.3 Speculative cumulative hypothesis

We hypothesise cumulative reliance on AI-mediated social interaction may reduce empathic accuracy and alter attachment patterns, with effects scaling with diversity of mediating systems. This technological mediation represents what Borgmann (1984) identified as the

“device paradigm,” where technologies promise to deliver social goods while transforming their very nature.

Testable predictions: Weakened if empathy measures show no decline with high vs. low AI-mediated interaction, controlling for total social interaction quantity.

3.4 Decision-making dependency

3.4.1 Empirical regularities

This pathway has the strongest empirical support. GPS use associates with poorer spatial memory (Dahmani and Bohbot, 2020; Gramann et al., 2017); knowing information is accessible reduces encoding effort (Sparrow et al., 2011); increasing automation reduces sense of agency (Berberian et al., 2012), with individual vulnerability factors moderating the effect (Di Plinio et al., 2019).

3.4.2 Plausible mechanisms

Cognitive offloading creates a self-reinforcing dependency: as unassisted decision-making becomes more effortful through disuse, the case for further delegation strengthens.

3.4.3 Speculative cumulative hypothesis

We hypothesise cumulative delegation across multiple systems may produce generalised decision-making dependency (diminished confidence in unassisted judgment and heightened anxiety without algorithmic mediation) with consequences for the independent judgment democratic participation requires. If citizens increasingly rely on algorithmic mediation for daily decisions, their capacity for the independent judgment that democracy requires may be gradually undermined.

Testable predictions: Weakened if decision-making competence in non-delegated domains remains unaffected by heavy delegation elsewhere.

3.5 Identity formation

3.5.1 Empirical regularities

Personalisation algorithms impact self-perception (Taylor and Chen, 2024), with correlations between social media use and identity distress. This effect appears particularly pronounced in developmental stages where identity formation occurs most actively.

3.5.2 Plausible mechanisms

Proposed mechanisms include: algorithmic identity capture—personalisation narrows the user’s sense of possibility to what the algorithm predicts; and performative self-optimisation—identity formation shifts toward external, algorithmically-optimised performance (Turkle, 2015).

3.5.3 Speculative cumulative hypothesis

We hypothesise prolonged exposure may reduce capacity for self-directed identity exploration. Following Taylor (1992), the concern is not that AI presents a “false” self but that algorithmic recognition

(optimised for engagement) may progressively displace the broader frameworks of meaning through which authentic identity (the ongoing capacity to define oneself through meaningful frameworks of value and commitment rather than externally imposed categories) has historically been constituted. Not forbidding authenticity but making it increasingly effortful.

Testable predictions: Weakened if identity coherence shows no association with algorithmic personalisation exposure.

4 The governance gap: why current technological governance falls short

Current technological governance approaches typically assess systems individually, asking whether a specific application creates substantial risk. This resembles early environmental regulation that evaluated individual chemicals in isolation, missing their combined impact. Just as individually “safe” levels of multiple toxins can collectively exceed safety thresholds, AI systems deemed “low-risk” individually may collectively transform cognitive and social environments beyond resilience thresholds.

This mirrors findings in environmental health research where combinations of pollutants at concentrations below regulatory limits can produce unexpected health impacts (UNEP, 2025). Environmental contaminants often interact in complex ways that regulatory frameworks fail to capture, with additive or synergistic effects that only become apparent through cumulative exposure.

Consider recommendation systems: Any single recommendation engine likely falls below risk thresholds in regulatory frameworks. Yet when dozens of such systems collectively shape information exposure across multiple platforms daily, they potentially create information environments that systematically undermine epistemic agency and democratic deliberation. This cumulative effect remains largely unaddressed in system-by-system assessment approaches.

This governance gap reflects what Scott (1998) identified as the limitations of “seeing like a state” regulatory frameworks that break complex socio-technical systems into discrete, manageable components may miss the emergent properties that arise from their interaction.

Environmental regulation evolved partly in response to the concept of externalities costs imposed on parties who did not choose to incur them (Colborn et al., 1997). AI systems similarly produce cognitive and social externalities that remain largely unaccounted for in governance frameworks.

When an AI system optimizes for user engagement, corporate profit, or operational efficiency, it potentially creates technological externalities: attention fragmentation, decision dependency, emotional modulation, and social disconnection. These effects do not appear on corporate balance sheets or regulatory risk assessments, yet they represent real costs to cognitive sovereignty, social cohesion, and democratic capacity that accumulate over time.

This externality gap grows as AI systems increasingly mediate everyday experiences while optimisation objectives remain misaligned with human flourishing and democratic governance (Zuboff, 2019). Without mechanisms to account for these externalities, market forces alone will likely increase rather than mitigate cumulative technological impacts on society.

Environmental regulation advanced significantly with improved detection methods that could measure previously imperceptible contamination levels. Similarly, addressing cumulative technological effects requires developing metrics and measurement approaches for cognitive, emotional, and social impacts that currently remain largely unmeasured. How do we quantify the erosion of decision making independence? What constitutes a “dangerous level” of attention fragmentation? How can we measure the degree of algorithmic influence on identity formation or social bonding capacity? How do we assess changes in empathic accuracy or democratic deliberation quality? Without established metrics and thresholds, governance frameworks struggle to incorporate these dimensions into technological assessment models (Zuboff, 2019). This measurement challenge represents both a technical obstacle and a fundamental conceptual hurdle for effective governance of cumulative technological effects. Developing appropriate metrics requires not just new measurement tools but new conceptual frameworks for understanding human-technology ecologies that span cognitive, emotional, and social domains.

5 Proposed governance innovations for democratic technology development

5.1 From individual to cumulative technological assessment

Environmental protection combines individual health practices with systemic regulation personal choices to avoid contaminants alongside policies that limit emissions. Similarly, addressing AI’s cumulative effects requires both individual technological hygiene practices and systemic governance innovations that ensure technology serves democratic purposes.

Individual practices like attention management, emotional awareness, intentional technology use, maintaining authentic social connections, resemble personal health measures like filtering drinking water or reducing exposure to environmental toxins. These practices provide important protection but remain insufficient without complementary systemic measures that address root causes and ensure technological development serves human flourishing.

Effective technological governance must move beyond risk-based assessment of individual systems to consider cumulative effects across cognitive, emotional, and social environments. This might include holistic technological impact assessments that evaluate how systems collectively shape attention, emotional regulation, social connection, decision-making, and identity formation; algorithmic impact labels that disclose how systems influence psycho-emotional patterns, similar to environmental impact disclosures; cumulative exposure monitoring across platforms and applications, analogous to environmental monitoring networks; developmental preservation zones that limit technological optimisation pressures during critical developmental periods; and social authenticity standards that prevent manipulation of social and emotional responses through design patterns that prioritize engagement over human wellbeing. Current algorithmic auditing exemplified by the SMACTR framework (Raji et al., 2020) focuses on individual system harms. We propose extending this through: (i) cross-system exposure mapping at the SMACTR scoping stage, requiring developers to consider how their system interacts with the broader

AI ecosystem users already navigate; (ii) longitudinal impact modeling, specifying expected interaction patterns and cumulative trajectories; and (iii) cumulative threshold analysis, assessing whether a system's contribution might push populations past critical exposure levels. Given the adaptive nature of AI systems (Yeung, 2017), these assessments should function as living documents with regular review cycles, not one-time pre-deployment gates.

5.2 Cognitive-social environmental monitoring

We propose initial population-level indicators: sustained attention capacity (measured via standardised tasks such as the SART; Robertson et al., 1997), emotional self-regulation (Difficulties in Emotion Regulation Scale; Gratz and Roemer, 2004), social connection quality (integrated into existing surveys such as the European Social Survey), decision-making independence (experimental tasks plus self-report algorithmic reliance measures), and identity coherence (Sense of Coherence scale; Antonovsky, 1987). These are designed to be administrable through existing infrastructure like PISA assessments, European Social Survey, national health surveys, minimising institutional overhead while generating longitudinal data.

5.3 Synthetic population modeling as predictive governance tool

Unlike environmental contaminants, where decades of epidemiological data now inform regulation, we cannot wait for longitudinal studies to reveal technology's long-term cognitive and social effects. The pace of technological change and deployment far exceeds traditional research timelines. Instead, we should look for alternative methods. One could consider synthetic population modeling as an evidence-based predictive approach for democratic technology governance. Synthetic populations computer generated models of individuals with realistic attributes can simulate how technological exposure might influence cognitive, emotional, and social outcomes over extended timeframes. These models allow testing evidence-based hypotheses about cumulative effects without waiting decades for real-world data, potentially identifying concerning patterns before they become entrenched in actual populations.

This approach builds on established methods in digital twin technology, where virtual replicas of physical entities enable predictive simulation. Recent advances in AI-enabled digital twins allow for high-fidelity modeling of complex systems, predicting behaviors and outcomes under various conditions (Taylor et al., 2023; Biller and Biller, 2023).

For example, synthetic populations could model how various levels of social media algorithm exposure might influence democratic participation patterns in different demographic groups over a simulated decade. These simulations could incorporate feedback loops (e.g., increased polarization leading to decreased civic engagement) and test potential interventions (e.g., reduced algorithmic personalization or enhanced media literacy programs).

This approach resembles how climate scientists use computational models to forecast potential outcomes based on current trends and mechanisms, allowing for evidence-based policy development without requiring catastrophic events to first occur. While imperfect, such models provide actionable insights when combined with existing research, theoretical frameworks, and emerging evidence potentially

avoiding the decades-long delay between early warning signs and effective regulation that characterised many environmental contaminants.

5.4 Cognitive-social ecosystem services valuation

Just as clean water and unpolluted air provide essential ecosystem services that underpin physical health and economic prosperity, our cognitive and social capacities constitute "mental ecosystem services" that support human flourishing and democratic governance. These cognitive-social commons, including sustained attention, emotional authenticity, genuine social connection, independent decision making, and coherent identity formation represent resources as vital as any natural capital, yet they remain largely unprotected from technological disruption.

The breakthrough 1997 study that valued global ecosystem services at \$33 trillion annually (exceeding the global GNP of \$18 trillion) transformed environmental protection by demonstrating nature's economic significance (Costanza et al., 1997). Similarly, we must quantify our cognitive-social ecosystem services' contribution to human flourishing and democratic governance. Their true social and economic value would likely approach several percentage points of global GDP, encompassing everything from the capacity for democratic deliberation to the social trust that enables economic cooperation.

Recognizing these cognitive-social ecosystem services not as individual luxuries but as quantifiable public goods essential for democratic society will drive the governance frameworks needed to preserve them from technological degradation. This approach aligns with broader movements toward recognizing the social value of commons that cannot be reduced to market transactions (Ostrom, 1990).

5.5 Validation strategy and phased implementation

We propose phased implementation: Phase 1 (Years 1–3)—establish baseline monitoring in pilot populations with varying AI exposure, testing whether cumulative exposure predicts cognitive-social variance beyond demographic confounders; Phase 2 (Years 2–4)—pilot cumulative impact assessment with AI developers; Phase 3 (Years 3–5)—integrate validated metrics into existing survey infrastructure; Phase 4 (Years 5+)—review and refine. This design is responsive to findings: if Phase 1 reveals no significant associations, the case for intensive governance diminishes accordingly.

5.6 Regulation-by-design for adaptive systems

For AI systems that continuously evolve (Yeung, 2017), we propose: effect boundary specification (developers define the range of algorithmic parameter variation), adaptation monitoring (parameter evolution logs accessible to auditors), and circuit-breaker mechanisms (automated review triggers when impact thresholds are approached). These requirements align with the EU AI Act's (European Parliament, 2024) risk-tiered approach while extending it to capture cumulative dynamics.

6 Implications for technology-society research and practice

This environmental paradigm offers several theoretical contributions to technology-society literature. First, it shifts focus from individual technological risks to cumulative exposure effects, revealing previously invisible governance gaps that affect democratic capacity. Second, it provides a conceptual bridge between environmental health science and technology governance, suggesting productive cross-disciplinary collaboration for understanding technology's societal impacts (Zuboff, 2019). Third, it operationalizes concepts like “technological externalities” and “cognitive-social ecosystem services” that could inform broader technology governance frameworks focused on human flourishing rather than mere efficiency.

The framework also suggests new research directions for technology studies. Empirical studies could test the proposed pathways through longitudinal analysis of technological exposure and psychological outcomes. Computational models could explore cumulative effect mechanisms through agent-based simulation of technology-society interactions. Policy research could evaluate the effectiveness of proposed governance innovations through natural experiments and quasi-experimental designs that examine different technological governance approaches.

For policymakers, this framework suggests several immediate applications that could strengthen democratic governance of technology. Technological impact assessments could incorporate cumulative effect analysis alongside individual system evaluation, ensuring that governance decisions consider broader societal implications. Regulatory sandboxes could test cumulative monitoring approaches at smaller scales before broader implementation, allowing democratic experimentation with new governance models. International cooperation could develop shared standards for measuring and governing technological externalities, preventing a “race to the bottom” in technology governance.

For technology developers, the framework implies design principles that consider cumulative impacts alongside individual system performance and business objectives. This might include attention-preserving interface design that supports rather than fragments human cognition, emotional authenticity standards that prioritize human wellbeing over engagement metrics, and algorithmic transparency measures that enable user agency rather than technological dependency.

For researchers, the framework suggests methodological innovations including synthetic population modeling for predicting technological impacts, longitudinal exposure studies that track technology's effects on human development, and cross-platform impact assessment that examines how multiple technological systems interact to shape human experience. It also implies new metrics development for cognitive, emotional, and social impacts that current approaches cannot capture but that may be essential for preserving human agency in a technological society.

This theoretical framework has several limitations that future research should address. First, the environmental analogy, while illuminating, may not capture all aspects of technology's societal impacts. Unlike chemical exposure, technological interaction involves conscious agency and social construction, creating different dynamics than passive contamination. Future research should explore how human agency and technological influence interact in complex ways that neither technological determinism nor social constructivism alone can explain.

Second, the proposed pathways require empirical validation through controlled studies that can establish causal rather than

correlational relationships. The GPS navigation research provides a compelling example of how technology can reshape cognitive capacities, but similar evidence is needed across other domains to validate the broader framework.

Third, the governance innovations proposed here require testing through pilot programs and policy experiments before broader implementation. The synthetic population modeling approach, in particular, needs validation through comparison with real-world outcomes and democratic input into the values and assumptions built into such models.

Future research should prioritize developing measurement tools for the proposed pathways, conducting longitudinal studies of cumulative technological exposure across different populations, and testing governance innovations through controlled policy experiments that examine their effectiveness and democratic legitimacy. Cross-disciplinary collaboration between technology researchers, environmental health scientists, social scientists, and policy scholars will be essential for advancing this framework in ways that strengthen rather than undermine democratic governance.

7 Conclusion

We stand at a moment comparable to the early days of environmental awareness, when the first evidence of widespread contamination emerged but comprehensive protection frameworks had not yet developed (Carson, 1962). The human cognitive and social environment now faces similar challenges, but subtle pervasive alterations from technologies deployed faster than their democratic implications can be fully understood or governed.

The environmental analogy offers both caution and hope for democratic technology governance. Environmental degradation continued for decades before effective protection emerged, with substantial and sometimes irreversible damage occurring in the interim (Carson, 1962; Colborn et al., 1997). Yet environmental protection eventually developed into sophisticated governance systems that significantly reduced many forms of contamination while allowing continued technological advancement guided by ecological principles.

By recognizing AI's potential cumulative effects on human agency and social cohesion now, and by developing governance approaches that address the quality of our cognitive, emotional, and social environments alongside individual system risks we have the opportunity to practice democratic technology stewardship. This means preserving the cognitive-social commons that underlie human flourishing and democratic governance, ensuring that technological development serves rather than undermines the capacity for meaningful choice and authentic connection.

Just as clean water and air are essential for physical health, uncontaminated attention, emotional authenticity, genuine social connection, independent decision-making, and coherent identity formation are essential for psychological wellbeing and democratic participation. By extending environmental protection paradigms to these domains of human experience, we can develop governance approaches that preserve these fundamental capacities while still embracing technology's potential benefits for human flourishing.

The stakes could not be higher for democratic society. If we fail to address these cumulative impacts, we risk a future where human

agency and social cohesion gradually erode not through sudden technological disruption, but through the silent accumulation of cognitive and social effects that undermine the very foundations of democratic self-governance. The time to develop comprehensive protection for our cognitive-social environment is now, before such effects become both severe and irreversible.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

VS: Writing – original draft, Writing – review & editing.

Funding

The author(s) declared that financial support was not received for this work and/or its publication.

References

- Antonovsky, A. (1987). *Unraveling the Mystery of Health*. San Francisco: Jossey-Bass.
- Berberian, B., Sarrazin, J.-C., Le Blaye, P., and Haggard, P. (2012). Automation technology and sense of control: a window on human agency. *PLoS One* 7:e34075. doi: 10.1371/journal.pone.0034075
- Bijker, W. E., Hughes, T. P., and Pinch, T. J. (1987). *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. Cambridge: MIT Press.
- Biller, B., and Biller, S. (2023). Implementing digital twins that learn: AI and simulation are at the core. *Machines* 11:425. doi: 10.3390/machines11040425
- Borgmann, A. (1984). *Technology and the Character of Contemporary Life: A Philosophical Inquiry*. Chicago: University of Chicago Press.
- Carson, R. (1962). *Silent Spring*. Boston: Houghton Mifflin.
- Clark, A. (2015). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford: Oxford University Press.
- Colborn, T., Dumanoski, D., and Myers, J. P. (1997). *Our Stolen Future: Are We Threatening Our Fertility, Intelligence, and Survival? A Scientific Detective Story*. New York: Dutton.
- Costanza, R., d'Arge, R., de Groot, R., Farber, S., Grasso, M., Hannon, B., et al. (1997). The value of the world's ecosystem services and natural capital. *Nature* 387, 253–260.
- Dahmani, L., and Bohbot, V. D. (2020). Habitual use of GPS negatively impacts spatial memory during self-guided navigation. *Sci. Rep.* 10:6310. doi: 10.1038/s41598-020-62877-0
- Delgado, P., Vargas, C., Ackerman, R., and Salmerón, L. (2018). Don't throw away your printed books: a meta-analysis on the effects of reading media on reading comprehension. *Educ. Res. Rev.* 25, 23–38. doi: 10.1016/j.edurev.2018.09.003
- Di Plinio, S. (2025). Panta rh-AI: assessing multifaceted AI threats on human agency and identity. *Soc. Sci. Humanit. Open* 11:101434. doi: 10.1016/j.ssaho.2025.101434
- Di Plinio, S., Arnò, S., Perrucci, M. G., and Ebisch, S. J. H. (2019). Environmental control and psychosis-relevant traits modulate the prospective sense of agency in non-clinical individuals. *Conscious. Cogn.* 73:102776. doi: 10.1016/j.concog.2019.102776
- Diamond, A. (2013). Executive functions. *Annu. Rev. Psychol.* 64, 135–168. doi: 10.1146/annurev-psych-113011-143750
- European Parliament (2024). Regulation (EU) 2024/1689 Laying down Harmonised rules on artificial Intelligence (Artificial Intelligence Act). *Off. J. Eur. Union*.
- Gramann, K., Hoepner, P., and Karrer-Gauss, K. (2017). Modified navigation instructions for spatial navigation assistance systems lead to incidental spatial learning. *Front. Psychol.* 8:193. doi: 10.3389/fpsyg.2017.00193
- Gratz, K. L., and Roemer, L. (2004). Multidimensional assessment of emotion regulation and dysregulation: development, factor structure, and initial validation of the difficulties in emotion regulation scale. *J. Psychopathol. Behav. Assess.* 26, 41–54. doi: 10.1023/b:joba.0000007455.08539.94
- Kramer, A. D. I., Guillory, J. E., and Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proc. Natl. Acad. Sci. USA* 111, 8788–8790. doi: 10.1073/pnas.1320040111
- Latour, B. (1992). "Where are the missing masses? The sociology of a few mundane artifacts," in *Shaping Technology/Building Society*, eds. W. E. Bijker and J. Law (Cambridge: MIT Press), 225–258.
- Odgers, C. L., and Jensen, M. R. (2020). Annual research review: adolescent mental health in the digital age: facts, fears, and future directions. *J. Child Psychol. Psychiatry* 61, 336–348. doi: 10.1111/jcpp.13190
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.
- Parry, D. A., and le Roux, D. B. (2021). Cognitive control in media multitaskers ten years on: a meta-analysis. *Cyberpsychology* 15:7. doi: 10.5817/CP2021-2-7
- Pentina, I., Hancock, T., and Xie, T. (2023). Exploring relationship development with social chatbots: a mixed-method study of replika users. *Comput. Hum. Behav.* 140:107600. doi: 10.1016/j.chb.2022.107600
- Pugh, A. J. (2023). *The Last Human Job: The Work of Connecting in a Disconnected World*. Princeton: Princeton University Press.
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., et al. (2020). "Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing," in Proc 2020 Conf Fairness Account Transparency, (), 33–44.
- Risko, E. F., and Gilbert, S. J. (2016). Cognitive offloading. *Trends Cogn. Sci.* 20, 676–688. doi: 10.1016/j.tics.2016.07.002
- Robertson, I. H., Manly, T., Andrade, J., Baddeley, B. T., and Yiend, J. (1997). 'Oops!': performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. *Neuropsychologia* 35, 747–758. doi: 10.1016/s0028-3932(97)00015-8
- Ruiz Pardo, C. E., and Minda, J. P. (2022). Reexamining the effect of smartphone presence on cognitive performance. *J. Exp. Psychol. Gen.* 151, e44–e54. doi: 10.1037/xge0001284
- Sætra, H. S. (2020). The parasitic nature of social AI: sharing minds with the mindless. *Integr. Psychol. Behav. Sci.* 54, 308–326. doi: 10.1007/s12124-020-09523-6

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declared that Generative AI was not used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Scott, J. C. (1998). *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven: Yale University Press.
- Sparrow, B., Liu, J., and Wegner, D. M. (2011). Google effects on memory: cognitive consequences of having information at our fingertips. *Science* 333, 776–778. doi: 10.1126/science.1207745
- Tarafdar, M., Pullins, E. B., and Ragu-Nathan, T. S. (2015). Technostress: negative effect on performance and possible mitigations. *Inf. Syst. J.* 25, 103–132. doi: 10.1111/isj.12042
- Taylor, C. (1992). *The Ethics of Authenticity*. Cambridge: Harvard University Press.
- Taylor, S. H., and Chen, A. Y. (2024). The lonely algorithm problem: the relationship between algorithmic personalization and social connectedness on TikTok. *J. Comput. Mediat. Commun.* 29:zmae017. doi: 10.1093/jcmc/zmae017
- Taylor, S. J., Macal, C., Matta, A., Rabe, M., Sanchez, S., and Shao, G. (2023). “Enhancing digital twins with advances in simulation and artificial intelligence: opportunities and challenges,” in *Winter Simulation Conference 2023*, (San Antonio: IEEE).
- Turkle, S. (2015). *Reclaiming Conversation: The Power of Talk in a Digital Age*. New York: Penguin Press.
- UNEP (2025). Pollution and health. Available online at: <https://www.unep.org/topics/chemicals-and-pollution-action/chemicals-management/pollution-and-health> (Accessed January 15, 2025).
- Wang, Q., Camacho, I., Jing, S., and Goel, A. K. (2022). Understanding the design space of AI-mediated social interaction in online learning. *Proc. ACM Hum. Comput. Interact.* 6, 1–28. doi: 10.1145/3512943
- Ward, A. F., Duke, K., Gneezy, A., and Bos, M. W. (2017). Brain drain: the mere presence of one’s own smartphone reduces available cognitive capacity. *J. Assoc. Consum. Res.* 2, 140–154. doi: 10.1086/691462
- WHO. (2020). *WHO Human Health Risk Assessment Toolkit: Chemical Hazards*. Geneva: World Health Organization.
- Wilmer, H. H., Sherman, L. E., and Chein, J. M. (2017). Smartphones and cognition: a review of research exploring the links between mobile technology habits and cognitive functioning. *Front. Psychol.* 8:605. doi: 10.3389/fpsyg.2017.00605
- Winner, L. (1980). Do artifacts have politics? *Daedalus* 109, 121–136.
- Yeung, K. (2017). ‘Hypernudge’: big data as a mode of regulation by design. *Inf. Commun. Soc.* 20, 118–136. doi: 10.1080/1369118X.2016.1186713
- Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: Public Affairs.