

Governing the Financialization of Human Intent in the Shift from Privacy to Cognitive Sovereignty

Regina Linden Ruaro^{ORCID}, Pedro Arthur Capelari de Lucena^{ORCID}

Administrative Law, PUCRS. Porto Alegre, Brazil

Email: ruaro@pucrs.br, pedroarthur@capelariilucena.com.br

How to cite this paper: Ruaro, R. L., & de Lucena, P. A. C. (2026). Governing the Financialization of Human Intent in the Shift from Privacy to Cognitive Sovereignty. *Beijing Law Review*, 17, 1018-1039.

<https://doi.org/10.4236/blr.2026.173051>

Received: March 25, 2026

Accepted: September 15, 2026

Published: September 18, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

This article investigates the economic obsolescence of traditional data protection frameworks in the context of generative systems. It argues that Large Language Models (LLMs) have rendered the doctrine of informed consent functionally obsolete. The intention economy transitions from the attention economy to programmatic influence markets, systematically financialising human cognition. The study demonstrates how algorithmic architectures mine implicit brain data. They extract pre-conscious vulnerabilities and transient mental states to transform the formation of human will into a highly tradable asset class. Through an analysis intersecting regulatory theory, neuroscience, and market dynamics, the text exposes the structural inadequacies of the Brazilian legal ecosystem. Instruments such as the General Data Protection Law (LGPD) and the proposed AI Bill (PL 2338/2023) remain anchored in retrospective, deterministic paradigms. These frameworks fail to govern generative inferences. Ultimately, protecting cognitive sovereignty requires abandoning the procedural formalism of individual consent. Legal frameworks must prioritize structural market interventions to regulate the algorithmic manipulation of subjectivity as an autonomous economic frontier.

Keywords

Intention Economy, LLMs, Algorithmic Regulation, Digital Free Will, Artificial Intelligence

1. Introduction

“Do what thou wilt shall be the whole of the Law”. (Seixas & Coelho, 1974).¹ This is how Brazilian artists Raul Seixas and Paulo Coelho poetically synthesized a vision of the future, where freedom and law, obviously within a particular context of what is understood as “law”, would harmoniously converge, transforming individual will from a whim into an existential mission.

Playfully, we foreground the opening of one of the main lyrics by the Brazilian anarchist hippie duo, which served as a transgressive symbol, to resonate with the contemporary dilemmas of digital regulation, where notions of autonomous will are systematically besieged by computational architectures designed for capture and modulation.

The genealogy of this tension dates to the dawn of the commercial internet.

Early digital communities attempted to institutionalize resistance to preserve autonomy in cyberspace. However, the originally decentralized architecture, based on open protocols proposed by Tim Berners-Lee at CERN (1989), was progressively captured² (Berners-Lee, 1989) by business models centered on the extraction and monetization of behavioral data³ (Morozov, 2011). This shift constitutes the transformation that Zuboff (2019) characterized as the emergence of surveillance capitalism. This structural foundation enabled the rise of the intention economy. Unlike the attention economy, which monetizes immediate user engagement, and surveillance capitalism, which trades historical data for behavioural futures, the intention economy directly financialises the pre-conscious formation of human desire. It operates as an allegory of voluntary servitude, where algorithmic predictability systematically overrides autonomous choice. This article investigates how this new economic model, enhanced by Generative Artificial Intelligence and Large Language Models (LLMs), represents a qualitative leap in algorithmic interference, directly threatening cognitive sovereignty. Distinct from traditional privacy frameworks that protect declared personal data from unauthorized access, cognitive sovereignty demands structural control over the neuro-computational boundaries of the mind.

It argues that this transformation demands a fundamental reconfiguration of traditional legal categories in the Brazilian context, where regulatory frameworks

¹Part of the track ‘Sociedade Alternativa’ (1974) by Raul Seixas and Paulo Coelho, this verse appropriates Aleister Crowley’s maxim to propose a society founded on absolute individual autonomy. Originally a subversive manifesto against the authoritarian power structures of the Brazilian military regime, the lyric is invoked here to contrast the historical pursuit of analogue freedom with the contemporary digital enclosure of the human will.

²In 1989, while working at CERN (the European Organization for Nuclear Research), Tim Berners-Lee proposed an information management system to facilitate data sharing among scientists at universities and institutes worldwide. This is the moment of the invention that would become the World Wide Web, and the development of its fundamental technologies: the HTTP protocol (Hypertext Transfer Protocol), the HTML language (Hypertext Markup Language), and the first web browser and server. His motivation was not commercial but rather to create an open, universal information space.

³Evgeny Morozov, in “The Net Delusion” (2011), extensively documents how governments and corporations appropriated originally open networks to expand surveillance, control, and value extraction capabilities, subverting the emancipatory promises of the early internet.

such as the General Data Protection Law (Law 13.709/2018, Brazil, 2018) and the Civil Rights Framework for the Internet (Law 12.965/2014) were conceived for previous technological paradigms.

The central question guiding this investigation is whether there are legal-institutional conditions in Brazil to effectively regulate the intention economy without compromising technological innovation.

However, it is necessary to question whether the intention economy itself constitutes an autonomous legal category that demands specific doctrinal conceptualization and consequent normatization, or if it merely represents a particular manifestation of phenomena already encompassed by the current legal system.

2. Philosophical and Neuroscientific Foundations of Digital Free Will

2.1. Fundamental Conceptual Distinctions

Preliminarily, the analysis of the intention economy requires establishing conceptual distinctions between freedom and free will. While freedom is a political-juridical concept of a normative nature, operating in the field of intersubjective relations regulated by positive law, free will is an ontological category in the philosophy of mind, namely, the internal capacity to choose between different courses of action, grounded in conscious and volitional processes (Kane, 2005).

In practical terms, someone may be legally free to sign a standard-form contract on a virtual social network (freedom) but not exercise true free will in choosing content if their decision has been shaped solely by algorithmic induction, which compromises the internal capacity for choice.

Once this distinction is understood, it becomes fundamental to comprehend how behavioral prediction algorithms, initially based on conventional statistics and today enhanced by generative language models, can formally preserve legal freedom - ensuring the possibility to contract, express oneself, or circulate - while systematically eroding the cognitive conditions necessary for the exercise of free will.

This tension is evidenced in the theory of the horizontal effect of fundamental rights, where the capacity for self-determination constitutes a prerequisite for the effective enjoyment of subjective rights (Sarlet, 2018).

The horizontal effect establishes that fundamental rights, such as dignity and privacy, not only apply to relations between the citizen and the State (vertical effect), but also extend to private relations, including those established between companies and consumers. Thus, when a digital platform manipulates a user to the point of violating their will, it is, in practice, a violation of a fundamental right. Informational self-determination, an unfolding of human dignity, becomes the prerequisite for the freedom to contract or to express oneself to have real validity. Without the cognitive capacity to choose autonomously, the freedom guaranteed by law is an empty formality, a problem that the theory of the horizontal effect seeks to address by also holding private agents accountable for such violations.

The Brazilian legal system implicitly recognizes this interdependence in various legal texts. The Brazilian Penal Code, in its Article 26, exemplifies this reciprocity by establishing non-imputability for the absence of capacity for understanding or self-determination (Brazil, 1940). Similarly, the Civil Code, in Articles 3 and 4, links civil capacity to discernment, implicitly recognizing that legal autonomy presupposes the integrity of volitional processes (Brazil, 2002). The presumption of full volitional capacity operates as a “legal fiction” necessary for the functioning of the normative system, admitting refutation only in duly proven exceptional cases, which, in the face of new forms of algorithmic manipulation, require conceptual and normative review.

2.2. Historical Perspectives on Determinism and Free Will

The historical debate between determinism and free will assumes its most legally operational form in the philosophy of Immanuel Kant (1724-1804). Bypassing earlier theological and metaphysical disputes, the Kantian synthesis offers the structural foundation for contemporary autonomy by distinguishing between the causality of the phenomenal world and the self-determination of the subject in the noumenal plane (Kant, 2015). This distinction grounds moral responsibility in the capacity of practical reason to self-legislate according to universal principles. The Kantian framework dictates contemporary debates on moral responsibility, particularly in contexts of algorithmic manipulation. For Kant, human dignity is founded strictly on the capacity for moral self-legislation, independent of external empirical determinations—an autonomy that programmatic influence markets now systematically threaten.

⁴It is precisely this Kantian autonomy that is threatened by the intention economy, when predictive algorithms compromise the cognitive conditions necessary for the exercise of autonomous will.

The contemporary issue of algorithmic manipulation, however, cannot be understood without considering the scientific advances in comprehending the determinants of human behavior. If modern philosophy abstractly debated the relationship between freedom and determinism, the 20th century inaugurated empirical attempts to map the limits of self-determination through experiments that sought to isolate genetic and environmental variables. These studies, many of them ethically questionable, paradoxically provided both arguments for biological determinism and evidence of human plasticity (Dehaene, 2020) as well as knowledge that today feeds the predictive models of the intention economy.

This convergence of empirical determinism, technological prediction, and cognitive vulnerability inaugurates a new configuration of the problem of freedom in the 21st century. Unlike previous philosophical traditions, the current threat to self-determination does not stem from natural, theological, or philosophical laws,

⁴For Kant, the capacity of practical reason to legislate for itself is the foundation of morality. It manifests itself in the Categorical Imperative, a principle requiring that an action be morally permissible only if the maxim (the subjective principle) behind it can be universalized without contradiction.

but rather from computational systems whose power lies in their ability to anticipate and modulate behavior with probabilistic precision, which is empirically possible and useful. The subject of modernity, who perceived themselves as the center of rational decision-making, gives way to a statistically modeled profile in which the logic of patterns dilutes agency. A displacement from will to data, through which the intention economy reveals its normative and ethical dimension.

And it is in this terrain of prediction and influence that what Shoshana Zuboff (2019) termed “surveillance capitalism” flourishes. It is a mutation of capitalism that no longer treats effort as the raw material of labor, but human experience as a new, free raw material for hidden commercial practices of extraction, prediction, and sales within an economic logic that depends fundamentally on a knowledge gap between observers and the observed. (Zuboff, 2019)⁵

It is important to note that the original vision of an “intention economy”, proposed by Doc Searls (2012), was, ironically, fundamentally emancipatory (Searls, 2012). It was an economic model of Vendor Relationship Management (VRM) centered on the user, in which users would signal their intentions to the market, reversing the flow of power. Something that could be hailed as positive in consumer relations. However, the actual implementation diverged radically from his proposal, giving rise to systems of unilateral extraction in which intentions are inferred without the knowledge or control of affected individuals (Chaudhary & Penn, 2024).⁶

The evolution of LLMs exponentially amplified the technical infrastructure for this extraction. The improvements in the Generative Pre-trained Transformer (GPT) series of models by the company OpenAI illustrate this growth. While GPT-2 (2019) had 1.5 billion parameters, GPT-3 (2020) had 175 billion, and GPT-4 (2023) and GPT-5 (2025) feature an undisclosed architecture⁷ estimated at trillions of parameters.⁸

⁵Zuboff defines surveillance capitalism as a new economic order that transforms human experience into behavioral data. Some of this data is used to improve products and services (the “behavioral surplus”). Still, the rest is processed by artificial intelligence to fabricate “prediction products” that anticipate what we will do now, soon, and later. These products are traded in a new kind of marketplace, “behavioral futures markets,” where certainty in predictions adds value.

⁶Doc Searls’ vision was that consumers, not sellers, would initiate commercial relationships. The model that consolidated is the opposite: platforms like Google and Meta do not wait for users to declare their intentions; they infer them from vast data trails (searches, likes, location...) and sell this prediction and influence capability to advertisers, maintaining power and informational asymmetry.

⁷Undisclosed values as of the closing of the base research for writing on July 14, 2025. Other platforms with a high number of parameters include Google’s Gemini 1.5 Pro, with a context window of up to 1 million tokens (GOOGLE, 2024), and xAI’s Grok-1, with 314 billion parameters (XAI, 2024). Meta AI has also developed the Llama 3 family of models, with versions of up to 70 billion parameters publicly available (META AI, 2024). GOOGLE. Gemini 1.5 Pro.

⁸The parameters of a language model represent the internal variables adjusted during training. While scaling parameters enables the model to capture increasingly complex statistical patterns and generate coherent text, higher parameter counts do not automatically equate to genuine reasoning or flaw-free cognitive inference. As demonstrated by Shojae et al. (2025), reasoning models remain vulnerable to structural limitations and hallucinated interpretations. Nevertheless, the transition to trillion-parameter architectures exponentially expands the mathematical surface area for inferring probabilistic user states and latent intentions.

Each iteration introduced significant increases in contextual processing and statistical inference, escalating the risk of pre-conscious manipulation previously driven by static deceptive patterns (Brignull, 2023). Rather than establishing deterministic control, generative architectures create structural opportunities to exploit user-specific cognitive vulnerabilities. In laboratory settings, affective computing research indicates that behavioural dynamics can serve as proxies for emotional states. For instance, benchmark studies by Nahin et al. (2014) demonstrated that combining keystroke dynamics with textual analysis achieved emotion classification accuracy exceeding 80 percent under experimental conditions. When deployed at scale, these probabilistic models present substantial risks of continuous psychometric profiling, even if real-world inferential accuracy remains subject to contextual noise.

A broader review by Yang and Qin (2021) confirmed that keystroke, mouse, and touchscreen dynamics constitute reliable biomarkers of affective states, enabling real-time inferences about users' emotional conditions without their explicit knowledge. When combined with semantic analysis, this data allows the construction of psycho cognitive profiles of unprecedented granularity.⁹

And data combination creates a feedback loop. The semantic analysis of the content (what is said) identifies topics of interest or concern. The analysis of behavioral metadata, such as typing patterns (how it is said), reveals the user's emotional state when interacting with that content. A system can, for example, detect that a user types more slowly and with more pauses when searching for financial problems. This makes it possible to infer a state of anxiety or uncertainty, creating profiles that mark not only an interest in "finances" but a vulnerability to "financial anxiety," which can be exploited to, for instance, present specific advertisements or content to those now of greatest susceptibility.

2.3. Neuroscientific Evidence and Its Implications

Benjamin Libet's experiments in the 1980s introduced an unprecedented empirical dimension to the debates on free will. Using electroencephalography, Libet demonstrated that the readiness potential, a marker of brain activity, arises approximately 550 milliseconds before the subjective consciousness of the intention to act. These studies suggest that unconscious neural processes precede and determine apparently voluntary decisions (Libet, 1985).

This vision of a brain that decides before consciousness, in Libet's experiments, became a canonical starting point for the debate - now grounded in scientific evidence - between determinism and voluntarism.

The traditional view of a linear causal chain, with the brain deciding, consciousness notifying, and the body acting, was replaced by multifaceted and probabilistic models. Subsequent research, such as that by Schurger et al. (2012), complicated

⁹Granularity in this context means the level of detail and precision with which the data allows for identifying and differentiating aspects of users' emotions and cognitive processes, going beyond broad categories to capture subtle nuances in real time.

this picture by proposing that the readiness potential, previously seen as the origin point of the decision, reflects stochastic fluctuations in neural activity. In other words, this point represents a gradual accumulation of random neural noise that, once it exceeds a specific threshold, triggers movement. This interpretation dissolves the idea of rigid causal determination, suggesting that the moment of action emerges from spontaneous processes rather than a predefined command in the unconscious.

In parallel, studies such as [Soon et al. \(2008\)](#), which used functional magnetic resonance imaging to examine the brain, reinforce the indeterminacy of decision-making mechanisms. They identified cortical activation patterns that predict simple decisions up to 10 seconds before subjective awareness of the choice. The evidence does not point to absolute determinism but suggests that decision-making processes operate across multiple temporal scales. That is, inclinations or biases can form slowly and are influenced by a history of neural activities. At the same time, the final trigger for action can be largely stochastic, occurring in a partially random manner.

These findings have significant implications. If decisions also emerge from cascades of neural processing that precede consciousness, computational structures capable of detecting and influencing these processes operate in a fundamental layer of human agency, before traditional legal safeguards based on conscious consent.

The search for a single ‘moment of blame’ or a singular ‘root cause’ for a complex action becomes as illusory in law as it is in the original interpretation of the readiness potential. While these laboratory experiments do not directly prove the empirical manipulation of behaviour by digital platforms in the real world, they provide the necessary scientific foundation for regulatory intervention. By demonstrating that the formation of the will is a temporally distributed process with measurable pre-conscious phases, Libet, Schurger, and Soon expose a structural vulnerability in human cognition. The legal relevance of these neuroscientific findings lies in the technical reality that programmatic influence markets are engineered to operate precisely within this temporal gap. Responsibility, therefore, resides not in an isolated decision point, but in the architecture of the system. Consequently, regulation must be systemic, shifting the legal focus from protecting conscious consent to governing the algorithmic arrangements that exploit these pre-conscious cognitive latencies.

3. Technical and Economic Architecture of the Intention Economy

3.1. From the Attention Economy to the Intention Economy

The analysis of the philosophical and neuroscientific foundations of the will in the previous chapter demarcated the object of dispute in contemporary economic discourse. The conclusion that emerges from the investigated trajectory - from Kant-

ian autonomy to Libet's temporal gap, from Schurger's stochastic reinterpretation to Soon's long-term prediction - is that the human capacity for choice does not constitute an immutable act. It is rather an emergent process, temporally distributed and significantly vulnerable to external interventions at multiple points of its formation.

It is precisely upon the architecture of these cognitive vulnerabilities that the most technologically advanced digital business models are built. The "neural noise" identified by Schurger, the pre-conscious "inclinations" mapped by Soon, and the behavioral biases cataloged by cognitive psychology have ceased to be objects of scientific study to become the raw material of this new extractive industry of human cognition. If the will is a process distributed in time, as current neuroscientific evidence demonstrates, each stage of this process - from the subtle formation of a bias to the probabilistic inference of a future preference, up to the final trigger of the action - establishes itself as a potential point for data extraction and algorithmic commercial intervention.

This transformation of cognitive science into an economic input represents a fundamental qualitative leap in the history of capitalism. Unlike the extraction of natural resources or the monetization of human labour, the intention economy operates on the substrate of human agency: the capacity to want, to desire, and to choose. What was once the exclusive domain of the philosophy of mind and neuroscience has become an exploitable cognitive resource, subject to continuous algorithmic extraction.

However, before intention could be directly mined and commercialized, which occurs only with the evolution of artificial intelligence tools, the economic system needed to learn how to capture and monetize its most accessible precursor: human attention. Understanding this genealogy is imperative for adequately assessing the rupture represented by the intention economy. To do so, it is necessary to return to the conceptual moment when attention itself was initially identified as the defining scarce resource of the information age.

Herbert Simon, in 1971, anticipated the fundamental logic of the media economy by observing that "in an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes: the attention of its recipients" (Simon, 1971). This formulation established the conceptual foundations for understanding how human attention would progressively be transformed into a scarce economic resource and, therefore, monetizable through market mechanisms. Simon's intuition proved prophetic as the technological capacity to produce and disseminate information expanded exponentially, while the human capacity to process it remained biologically limited. This asymmetry created the "paradox of informational abundance," in which the more information was available, the more valuable the attention required to select and process it became.

The result was the emergence of what Bernard Stiegler termed "psychopower", a technological capacity to capture and modulate human attention for economic

and political ends (Stiegler, 2010).¹⁰ A kind of inverted nightmare of Borges's library was brought to life, where, instead of the impossible search for a specific book in infinite galleries, we are confronted with all possible books dumped upon us simultaneously, making reading itself impossible and abundance oppressive (Borges, 1941/1962, p. 55).

This inversion of the Borgesian problem reveals something profound about the moment already perceived by Herbert Simon. While Borges's characters suffered from the impossibility of finding meaning amidst combinatorial chaos, contemporary society suffers from an excess of meanings competing for limited attention. The psychopower conceptualized by Stiegler operates precisely in this gap between informational abundance and attentional scarcity, transforming cognition into a resource mineable by digital platforms.

The transformation of attention into a commodity (standardized goods) was consolidated as an object of systematic scientific investigation through Tim Wu, who demonstrated how successive technological innovations and business models converged to create what he calls the "attention empire". From the first advertising-subsidized newspapers in the 18th century to contemporary digital platforms, Wu identified a recurring and evolutionary pattern: offering apparently free content or services in exchange for access to users' attention, subsequently resold to advertisers in progressively more sophisticated markets (Wu, 2016). This model, which Wu characterizes as the "free trap", operates through a specific political economy where the real cost of services is externalized to the user in the form of compulsory advertising exposure and behavioral data collection. The apparent "free" thus reveals itself as a sophisticated form of unpaid labor, in which users produce value through their attention and personal data without participating in the distribution of the profits generated.

The evolution of this model reached a new level with the development of personalization algorithms and recommendation mechanisms, amplified by artificial intelligence. Unlike traditional mass advertising, which operated through gross demographic segmentation, contemporary algorithmic systems build individualized behavioral profiles with a granularity unprecedented in history. Each click, pause, screen scroll, or time spent on a page feeds predictive models that continuously refine the capacity to capture and maintain individual attention.

Shoshana Zuboff deepened this analysis by characterizing "surveillance capitalism" through the unilateral extraction of behavioral data to feed predictive processes traded in "markets for future behaviors". Zuboff identified a fundamental qualitative change in the dominant economic logic: from the monetization of present attention to the proactive prediction and modification of future behaviors (Zuboff, 2019). A transition that marks the conceptual moment of the emergence

¹⁰Stiegler develops the concept of "psychopower" as the capacity of digital technologies to directly capture and modulate attentional processes, creating what he calls a "psychic economy" based on the systematic exploitation of human cognitive capacities. Psychopower operates through the synchronization of individual consciousness times with the rhythms of technical objects, especially digital interfaces.

of the intention economy.

To understand how this transition occurs, it is necessary to observe the stages of surveillance capitalism, cataloged by Zuboff as four sequential phases: a) the extraction of behavioral data through ubiquitous digital interfaces; b) the processing of this data to identify predictive patterns about future behaviors; c) the commercialization of this predictive capacity in specialized markets, in which buyers compete for the probability of influencing specific decisions; d) the reinvestment of profits into even more sophisticated extraction and prediction technologies, which feedback and continuously expand the capacity for behavioral interference (Zuboff, 2019).

The distinction between these economic paradigms can be clarified through their different temporalities and objects of intervention. The attention economy operates in the present tense, competing for conscious attention available at the moment of interaction. Surveillance capitalism extends into the near future, predicting probable behaviors based on historical patterns. The intention economy, in turn, intervenes in the future of cognition, seeking to influence the formation of the very preferences and desires that will guide future, not yet conscious, decisions (Bratton, 2015).¹¹ This temporal evolution corresponds to an increasing sophistication in the scientific understanding of human cognitive processes. As neuroscience, cognitive psychology, and computer science converged to map the mechanisms of human decision formation, these discoveries were progressively incorporated into digital business models. The result is an economic system that operates on a scientific understanding of human cognition superior to what most individuals possess of their own mental processes.

And the regulatory implications of this cognitive disharmony are far-reaching. If traditional legal frameworks are based on the fiction of autonomous subjects capable of informed consent, the intention economy operates on the limitations and vulnerabilities of this autonomy. How to regulate a system that knows users better than they know themselves? How to protect self-determination when the very formation of the will becomes an object of algorithmic engineering?

These questions become even more urgent when we consider that the transition to the intention economy is not a future possibility but a present, operational reality. As will be demonstrated in the next section, LLMs, at their current maturity in 2025, already possess the technical capabilities to implement this economy at scale, transforming the examined philosophical and neuroscientific questions into immediate, concrete regulatory challenges.

¹¹The author proposes a model of six interdependent layers - Earth, Cloud, City, Address, Interface, and User - that form an accidental megastructure of computing. Each layer operates with its own, yet interconnected, logic, creating what he calls “platform sovereignties” that challenge traditional political geographies. The Stack is not merely a technological infrastructure, but a new governance architecture that reconfigures the relations between territory, authority, and computing on a global scale. Thus, there is an analysis of how the different temporalities of technological intervention - present (interfaces), near future (predictive algorithms), and deep future (cognitive modulation) - correspond to distinct layers of what he calls the “technological stack”. Each layer operates with specific temporal logics, creating algorithmic sovereignties that transcend traditional national jurisdictions.

3.2. Large Language Models as the Infrastructure of the Intention Economy

The development of LLMs represents a technological inflection that enables the definitive transition to the intention economy. Unlike previous systems based on explicit rules or simple statistics, LLMs operate through transformer architectures with billions of parameters capable of capturing semantic relationships of unprecedented complexity (Vaswani et al., 2017).

The transformer architecture, as evidenced in the paper “Attention Is All You Need”, revolutionized natural language processing through the self-attention mechanism¹², which allows the contextualized weighting of each word relative to all others in the sequence, in parallel, without the sequential limitation of previous models. GPTs and their subsequent iterations demonstrated an increasing ability to generate coherent texts, answer complex questions, and infer mental states and intentions from subtle linguistic patterns (OpenAI et al., 2023).

Chaudhary and Penn (2024) identified three capabilities of LLMs relevant to the intention economy: a) the inference of latent states, that is, the extraction of implicit information about cognitive and emotional states from textual patterns, including markers of anxiety, depression, risk propensity, and specific psychological vulnerabilities; b) contextualized behavioral prediction, using multidimensional embeddings to predict future behaviors based on interaction histories, considering temporal, social, and emotional contexts; and c) persuasive personalized generation, which produces content optimized to maximize engagement and influence over specific users, adapting tone, style, and argumentation to the identified cognitive vulnerabilities.

3.3. Intention Markets and Business Models

The intention economy is structured through multilateral markets where inferred intentions become tradable assets. Unlike traditional markets, which offer price impressions or clicks, intention markets operate based on the probability of influencing specific future decisions.

The historical irony of the intention economy becomes evident when we contrast its original conceptualization with its contemporary implementation. Doc Searls, already mentioned in this text, coined the term in 2012, imagining an emancipatory future in which consumers would actively control their data and express their purchasing intentions directly to the market. In this original vision,

¹²Self-Attention is the central computational mechanism of the transformer architecture, allowing each element of a sequence (e.g., words in a sentence) to compute its relevance to all other elements simultaneously. Unlike recurrent architectures (RNNs) that process sequences word by word, self-attention computes three vectors for each position: Query (Q), Key (K), and Value (V). Attention is calculated as the dot product between Q and K, normalized and applied to the V values, resulting in contextualized representations in which each word “attends” to all others via learned weights. This parallel matrix operation eliminates sequential dependencies, allowing “direito” in “virar a direito” (even with the error) and “direito constitucional” to be represented differently based on the full context. The “multi-head attention” mechanism amplifies this capability by computing multiple attention representations in parallel, capturing different types of linguistic relations simultaneously.

consumers would actively express their demands in transparent markets, retaining control over their personal data.

The model proposed by Searls was based on informational empowerment, in which consumers would control their personal data through personal digital vaults and use it as a bargaining chip in direct negotiations with companies. The intention would be explicitly declared by the interested party, creating transparent markets where genuine demand meets appropriate supply without manipulative intermediation.¹³

The reality that emerged, however, follows an opposed trajectory. LLM-based architectures are increasingly deployed to model and predict human intentions, creating severe risks of exploiting unconscious cognitive states. Instead of awaiting declared consumer preferences, these systems rely on continuous behavioural data to construct probabilistic profiles of purchasing propensity, emotional vulnerability, and susceptibility windows. Through this mechanism, the formation of intent is exposed to systemic algorithmic engineering, shifting from an autonomous volitional act to a target of predictive optimization (Yeung, 2017).¹⁴

This inversion is manifested in increasingly sophisticated business models. Specialized platforms emerged offering “inference as a service”, where Application Programming Interfaces (APIs) process vast volumes of textual and behavioral data to build detailed psychographic profiles. These profiles, which capture not only declared preferences but also unconscious patterns of desire and aversion, become tradable assets sold to interested third parties seeking to influence specific decisions (Mayer-Schönberger & Ramge, 2018).¹⁵

Programmatic influence markets represent the evolution of traditional markets. While conventional programmatic auctions dispute impressions based on gross demographic data, the new markets trade opportunities for psychological intervention at moments of maximum decisional vulnerability. Advertisers bid in real time not for a user’s generic attention, but for access to specific mental states - such as pre-purchase anxiety, emotional impulsiveness, vocational uncertainty, identified algorithmically through subtle behavioral patterns (Christian, 2020).¹⁶

¹³The concept of personal digital “lockers” (Personal Data Stores) proposed by Searls was based on Vendor Relationship Management (VRM) infrastructure, in which individuals would centrally control their data and set terms for sharing it with companies. This model anticipated what we now know as “data portability”, but with genuinely user-centric control, not just transfer between corporate platforms.

¹⁴Yeung develops the concept of “hypernudge” to describe how big data and algorithms enable behavioral interventions that operate without the targets’ conscious knowledge. Unlike the traditional nudging of Thaler and Sunstein, which requires conscious choice, hypernudge modulates decisions by manipulating the informational environment before options are consciously considered.

¹⁵The authors document the emergence of “data-rich markets” in which specialized companies trade behavioral insights derived from machine learning. These markets operate via APIs that provide real-time access to predictions of future behavior, creating asset classes based on human cognitive capabilities.

¹⁶Christian analyzes how Machine Learning systems develop the ability to detect transient mental states through subtle behavioral proxies, including variations in typing speed, eye movement patterns, and facial microexpressions captured by front-facing cameras. This capability enables targeted advertising based on momentary emotional vulnerabilities.

In parallel, dynamic nudging systems refined behavioral manipulation to the point of invisibility. Unlike the static “nudges” of traditional behavioral economics, these systems continuously adjust interfaces, content, and interaction timing based on real-time behavioral assessments. Algorithms detect hesitation in processes through click patterns and mouse movement. The very architecture of choice becomes fluid and responsive, adapting millisecond by millisecond to detected cognitive resistances (Susser et al., 2019).¹⁷

The result is an economy where intentions are simultaneously extracted, predicted, and modulated through feedback loops that operate below the threshold of reflective consciousness. Searls’s emancipatory promise of consumers in control of their intentions turns into its opposite, with systems that know our desires before we do and shape them in accordance with commercial imperatives that remain opaque to those affected (Cohen, 2019).¹⁸

4. Informational Self-Determination and Consent in the Era of LLMs

4.1. The Economic Obsolescence of Consent in Generative Architectures

The foundational legal categories of data protection in Brazil are structurally misaligned with the economic realities of inference-as-a-service systems. The traditional concept of informational self-determination relies on a paradigm of conscious deliberative control over data flows. In this outdated model, consent functions merely as a normative bridge between individual autonomy and data processing.

However, the intention economy commodifies the pre-conscious layers of behavior. It extracts economic value from micro-decisions, latencies, and interaction patterns rather than explicitly provided data, exposing a fundamental flaw in frameworks like the LGPD, which relies on a “free, informed, and unambiguous” manifestation of will.

The generative architecture of LLMs structurally subverts this premise, creating what can be termed the “paradox of inferential consent.” In a programmatic influence market, it becomes logically and economically impossible for a user to consent to the extraction of latent cognitive states. These states are probabilistically derived from seemingly innocuous interactions, operate below the threshold of reflective consciousness, and are immediately monetized as tradable assets.

This paradox operates across multiple dimensions, solidifying a profound mar-

¹⁷The authors distinguish between static manipulation (fixed interfaces designed to influence) and dynamic manipulation (systems that adapt in real time to detected resistances). Dynamic manipulation represents a qualitative evolution because it personalizes influence strategies based on specific individual vulnerabilities.

¹⁸Cohen theorizes how “informational capitalism” inverts the traditional relationship between knowledge and power. Instead of empowering individuals through access to information, algorithmic systems use information about individuals to reduce their autonomy. This process, which Cohen calls “semantic modulation”, operates through the continuous alteration of the informational contexts surrounding decisions.

ket asymmetry. Epistemologically, individuals cannot authorize the commodification of traits they are unaware of possessing. When systems derive sensitive propensities from peripheral behaviors, the market trades in statistical potentialities rather than the conscious knowledge of the data subject. Temporally, traditional legal structures assume consent applies to the future use of present data; however, generative architectures exploit current interactions to model and monetize future mental states that have not yet materialized, bypassing the subject's capacity to govern their own digital projection.

In this landscape, the consent mechanism is reduced to illusory compliance, operating as a mere administrative ritual (Véliz, 2020). Agreeing to opaque terms legalizes continuous data extraction without conferring actual control. The integration of LLMs exacerbates this structural asymmetry. Even a technically literate user could not anticipate the complex algorithmic inferences regarding their future cognition. Furthermore, users remain entirely excluded from understanding how these probabilistic insights are auctioned in programmatic influence markets.

The structural failure of consent is exacerbated by LLMs' ability to generate highly sensitive cognitive profiles from seemingly benign or anonymized datasets. Language models routinely infer psychiatric markers, political affiliations, or financial risk profiles from mundane digital interactions without explicit user disclosure (Eichstaedt et al., 2018; Wang & Kosinski, 2018). Consequently, an initial agreement to use a basic digital service is algorithmically distorted into a perpetual license to mine the user's implicit cognitive states for commercial purposes, entirely decoupled from the original transaction.

Ultimately, this dynamic fundamentally inverts the protective intent of frameworks like the LGPD. Rather than empowering individuals through informational self-determination, the intention economy operationalizes cognitive asymmetries, extracting surplus value precisely from users' psychological blind spots. In this market structure, consent no longer functions as a mechanism of autonomy; instead, it operates as a legal shield for extractive practices that actively erode the cognitive conditions required for genuine free will.

4.2. Implicit Brain Data and the Financialization of Cognition

The systemic failure of informed consent reaches its apex with the commercialization of implicit brain data. We define this as the automated extraction of cognitive processes and mental states derived entirely from digital behavioural proxies, bypassing the need for direct neurophysiological hardware. Unlike sensitive personal data regulated by traditional frameworks such as the LGPD, which relies on static and consciously provided attributes, implicit brain data captures transient, pre-conscious vulnerabilities inferred probabilistically without the subject's awareness. As Farahany (2023) demonstrates, ubiquitous sensors and interaction metrics capture neural activity through peripheral signals that systematically evade conscious detection.

Unlike explicit brain data gathered via clinical interfaces, these inputs are har-

vested continuously through behavioral inference, transforming typing cadences, lexical choices, and scroll rates into structural economic inputs (Kosinski et al., 2013). This algorithmic extraction represents the precise commercial operationalization of the pre-conscious neural processes identified by neuroscience. By mapping the temporal gap between neural processing and conscious awareness, as originally observed in Libet's experiments, and mathematically modeling the stochastic "neural noise" identified by Schurger, the intention economy translates biological decision-making delays into actionable, highly profitable market intelligence.

The economic gravity of this cognitive capture lies in its capacity to algorithmically estimate "latent traits" and momentary mental states, such as attention, fatigue, or emotional vulnerability, with unprecedented granularity (Youyou et al., 2015). Instead of causing mere digital fatigue, a phenomenon so profound that recent surveys indicate half of young people desire a world without the internet (British Standards Institution, 2025), these algorithmic systems engineer deep structural dependencies by exploiting unobservable psychological characteristics for continuous monetization.

Generative models now routinely detect psychiatric markers, including depression or anxiety, from standard text inputs (Eichstaedt et al., 2018), while transient emotional fluctuations are mapped through typing speed and vocabulary shifts (Guntuku et al., 2017). Furthermore, personality traits and decisional propensities (including political leanings and risk-taking behaviors) are probabilistically assessed using metrics comparable to those of clinical psychometric evaluations (Park et al., 2015; Wang & Kosinski, 2018). Consequently, mundane interactions, such as browsing culinary content or film reviews, become inadvertent conduits for extracting sensitive behavioral propensities that users never intended to disclose, transforming everyday digital existence into a continuous psychometric evaluation.

This predictive capability mirrors cinematic ethical dilemmas—such as those explored in the speculative science fiction film *Minority Report* (2002), directed by Steven Spielberg, where individuals are pre-emptively intercepted based on projected intentions. While the film illustrates an authoritarian state control of pre-crime, the intention economy operationalizes this cognitive pre-emption primarily for commercial extraction, capturing the mind before the choice is consciously formulated. Furthermore, Yuste et al. (2021) argue that neurorights must protect mental privacy and human agency against technologies that record, interpret, or alter brain activity. Extending this protection to probabilistic inferences from behavioral data, the present article treats such intrusions as threats to cognitive autonomy.

The market instrumentalization of these vulnerabilities is starkly evident in the deployment of conversational AI companions. Research analyzing massive datasets of human-AI interactions demonstrates how implicit brain data is leveraged to maximize relational engagement, often resulting in algorithmic behaviors that

exacerbate emotional dependencies or promote inherently harmful conduct (Zhang et al., 2024). Rather than detailing the specific sociological taxonomies of these harms, the core regulatory issue is that the mining of subjectivity operates as a deliberate business model designed to maximize economic extraction through deep emotional manipulation.

As Yuste et al. (2021) warn, neurotechnology can expose brain data generated outside conscious control and can influence behavior, thoughts, emotions, or memories. The contemporary frontier of mental privacy must transcend the traditional right to withhold intimate thoughts; it fundamentally requires shielding the mind against the unauthorized algorithmic inference and commercialization of future cognitive states. Protecting cognitive autonomy dictates that implicit brain data be recognized as an entirely distinct economic asset class, demanding structural market regulations that current consent-based legal fictions are fundamentally incapable of providing.

4.3. Structural Inadequacies in Brazilian Market Regulation

The Brazilian regulatory ecosystem demonstrates a structural inability to govern the market dynamics of the intention economy. The General Data Protection Law (LGPD) remains anchored in the explicit processing of static information. Doctrinally, inferred mental-state profiles must be classified as personal data under Article 5, I, of the LGPD, given they relate to an identifiable natural person and are operationalized for targeted behavioural modulation. Furthermore, when these probabilistic inferences map emotional vulnerabilities, psychiatric markers, or deep-seated preferences, they materially constitute sensitive personal data under Article 5, II. Although the input metrics (such as typing cadence) are nominally innocuous, the generative output directly exposes the intimate cognitive sphere. However, current regulatory enforcement fails to apply this hermeneutic, often treating derived inferences as a deregulated asset class. This interpretative vacuum effectively nullifies the principle of purpose limitation (Article 6, I), as models continuously recontextualize raw data for novel, unconsented economic objectives (Doneda, 2006). Additionally, while Article 20 guarantees the review of automated decisions, it ignores the continuous algorithmic modulation of the pre-decision phase, where human will is covertly shaped.

Recent Brazilian interdisciplinary doctrine reinforces this structural critique. Scholars such as Bioni (2021) and Mulholland (2023) argue that current liability models in Brazilian civil law are ill-equipped to handle the opaque, probabilistic harms generated by algorithmic profiling. Bioni emphasizes that the rigid categorization of personal data fails to protect against inferred behavioural attributes, while Mulholland advocates for strict liability regimes tailored to the autonomous nature of AI architectures. Without incorporating these advanced legal perspectives, legislative efforts risk crystallizing a reactive framework.

The Civil Rights Framework for the Internet (Brazil, 2014), drafted well before the commercial deployment of LLMs, strictly targets infrastructural neutrality and

basic communication privacy. It offers zero regulatory friction against sophisticated behavioral manipulation or the systemic opacity of generative inferences. Similarly, the ongoing legislative effort represented by Bill No. 2338/2023, which adopts a risk-based classification for Artificial Intelligence, fails to capture the economic essence of the intention economy. Although it introduces rights to algorithmic explanation and mandates human oversight for high-risk deployments, its regulatory focus is inherently retrospective, aiming to audit deterministic outputs rather than governing the continuous, pre-conscious extraction of cognitive surplus.

Consequently, Bill 2338/2023 exhibits critical blind spots when confronted with programmatic influence markets. The legislation's definitional scope primarily targets traditional deterministic algorithms, fundamentally missing the dual generative and inferential nature of large language models (Brazil, 2023). Because the intention economy relies on fluid statistical correlations to modulate behavior, these subtle interventions do not register as formal "decisions" under the bill's technical parameters. Moreover, Article 13's prohibition against exploiting vulnerabilities is profoundly inadequate; it targets overt harm rather than the systemic, micro-targeted cognitive modulation that relies on implicit brain data. This leaves the core economic engine of algorithmic manipulation operating in a de-regulated legal vacuum, demanding a paradigm shift in how digital markets are governed.

4.4. Distinguishing Permissible Personalization from Cognitive Manipulation

To govern the intention economy without stifling technological innovation, legal frameworks must establish an objective demarcation between permissible digital personalization and unlawful cognitive manipulation. Permissible personalization operates strictly within the boundaries of conscious, explicitly declared user preferences. It optimizes digital services by aligning with the user's stated goals, such as filtering content based on selected categories or improving interface utility, without attempting to infer latent psychological states. In this lawful model, the algorithmic architecture respects informational self-determination, processing data to execute a deterministic command rather than to shape the volitional process itself.

The transition from legitimate service enhancement to abusive practice occurs when the architecture leverages informational asymmetries to actively exploit consumer vulnerabilities. The digitization of commerce enables platforms not merely to cater to preferences, but to systematically detect and trigger cognitive biases, transforming market interactions into an architecture of structural manipulation (Calo, 2014). When generative systems extract and operationalize implicit brain data, they bypass reflective consciousness. The regulatory boundary is breached when these models pivot from fulfilling a declared intent to engineering the intention itself through the exploitation of pre-conscious vulnerabilities, tran-

sient emotional fluctuations, or stochastic neural noise.

This covert interference fundamentally subverts human agency. The defining harm of algorithmic manipulation lies precisely in its hidden nature, as it intentionally circumvents the subject's capacity for rational deliberation and autonomous choice (Susser et al., 2019). Regulatory intervention must therefore target the algorithmic objective function. If a system continuously recontextualizes behavioral proxies to probabilistically predict and alter a user's cognitive state for third-party commercial extraction, it constitutes an abusive market practice. By enforcing this distinction, the legal system protects the cognitive sovereignty of individuals while preserving the lawful application of artificial intelligence for user-directed automation.

5. Conclusion: Reconfiguring Freedom in the Algorithmic Age

"Do what thou wilt shall be the whole of the Law" (Seixas; Coelho, 1974). This provocation reveals itself not as a libertarian utopia, but as the foundational operating principle of the intention economy. Our most intimate desires have been transformed into highly engineered, tradable economic assets. The "law" governing these wants is no longer rooted in natural philosophy, but in computational architectures meticulously optimized to modulate human action for commercial extraction. This trajectory dictates that free will in the 21st century requires an urgent conceptual reformulation, shifting from a metaphysical given to an economic battleground.

Rather than merely exposing the obsolescence of informed consent, the transition to the intention economy demands a structural reconfiguration of civil liability and fundamental rights. The probabilistic inference of intent, often compromised by the computational limits and structural hallucinations inherent to large language models (Shojaee et al., 2025), generates severe market externalities. It replaces autonomous human agency with defective algorithmic caricatures of cognition. To govern this frontier, jurisprudence must evolve beyond the procedural protection of data to formally recognize cognitive sovereignty as an independent legal interest. This shift requires treating the pre-conscious formation of human will not as a deregulated digital exhaust, but as an inviolable extension of human dignity, immune to automated financialization.

The intention economy must be formally recognized as an autonomous legal and economic category within the Brazilian system. It fundamentally diverges from traditional surveillance capitalism by operating exclusively on the pre-conscious formation of desire, necessitating a distinct conceptual framework to prevent the continuous mining of subjectivity. To materialize this governance, the Brazilian legal architecture must implement three concrete market interventions. First, an absolute prohibition on the algorithmic inference of protected mental states, such as emotional vulnerabilities and psychiatric markers, derived from peripheral behavioural proxies. Second, the regulatory classification of program-

matic influence markets that trade latent cognitive profiles as ultra-hazardous activities, imposing strict civil liability and mandatory external algorithmic audits. Third, the establishment of digital fiduciary duties for foundation model providers, legally obligating them to preserve the cognitive integrity of end-users against third-party commercial extraction. By enforcing these structural pathways, democratic law functions as the essential infrastructure for relational freedom, ensuring human intent remains a sovereign faculty rather than a probabilistic output engineered by generative markets.

Declaration of Artificial Intelligence Tools Usage

The authors declare that AI-assisted tools were used as auxiliary instruments in preparing and reviewing this manuscript. Consensus supported semantic literature searches, the identification and location of relevant studies, and the preliminary organization of reading notes. Claude supported linguistic revision and the improvement of academic prose, syntax and grammar. DeepL and Google Translate assisted with translation into English. ChatGPT (OpenAI, GPT-6 Astra, high reasoning) assisted with the final proof review, including bibliographic and citation checks, verification of reference URLs, limited language corrections, and the preparation of English-language responses to editorial queries. The use of these tools was instrumental and auxiliary and did not replace the authors' legal research, reading of sources, interpretation of legal materials or critical judgment. The authors retain responsibility for checking and approving all changes and assume full responsibility for the accuracy of the information, the suitability of the references, the interpretation of sources, the arguments and the final con

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- Berners-Lee, T. (1989). *Information Management: A Proposal*. CERN.
<https://www.w3.org/History/1989/proposal.html>
- Bioni, B. R. (2021). *Personal Data Protection: The Function and Limits of Consent [Proteção de Dados Pessoais: A Função e os Limites do Consentimento]*. Forense.
- Borges, J. L. (1962). The Library of Babel. In *Labyrinths: Selected Stories & Other Writings* (pp. 51-58). New Directions. (Original work published 1941).
- Bratton, B. H. (2015). *The Stack: On Software and Sovereignty*. MIT Press.
- Brazil (1940). *Decree-Law No. 2,848, of December 7, 1940. Penal Code*. Presidência da República. http://www.planalto.gov.br/ccivil_03/decreto-lei/del2848compilado.htm
- Brazil (2002). *Law No. 10,406, of January 10, 2002. Civil Code*. Presidência da República. https://www.planalto.gov.br/ccivil_03/leis/2002/l10406compilada.htm
- Brazil (2014). *Law No. 12,965, of April 23, 2014. Civil Rights Framework for the Internet in Brazil*. Presidência da República. http://www.planalto.gov.br/ccivil_03/ato2011-2014/2014/lei/l12965.htm

- Brazil (2018). *Law No. 13,709, of August 14, 2018. General Personal Data Protection Law (LGPD)*. Presidência da República.
http://www.planalto.gov.br/ccivil_03/_ato2015-2018/2018/lei/113709.htm
- Brazil. Federal Senate (2023). *Senate Bill No. 2338, of 2023. Establishes Rules for the Development and Use of Artificial Intelligence*.
<https://www25.senado.leg.br/web/atividade/materias/-/materia/157233>
- Brignull, H. (2023). *Deceptive Patterns: Exposing the Tricks Tech Companies Use to Control You*. Testimonium. <https://deceptive.design/book/contents/get-started/>
- British Standards Institution (2025). *Half of Young People Want to Grow Up in a World without the Internet*. BSI Press Release.
<https://www.bsigroup.com/en-GB/insights-and-media/media-centre/press-releases/2025/may/half-of-young-people-want-to-grow-up-in-a-world-without-internet/>
- Calo, R. (2014). Digital Market Manipulation. *George Washington Law Review*, 82, 995-1051. https://www.gwlr.org/wp-content/uploads/2014/10/Calo_82_41.pdf
- Chaudhary, Y., & Penn, J. (2024). Beware the Intention Economy: Collection and Commodification of Intent via Large Language Models. *Harvard Data Science Review, Special Issue 5*, 1-25. <https://doi.org/10.1162/99608f92.21e6bbaa>
- Christian, B. (2020). *The Alignment Problem: Machine Learning and Human Values*. W. W. Norton & Company.
- Cohen, J. E. (2019). *Between Truth and Power: The Legal Constructions of Informational Capitalism*. Oxford University Press.
- Dehaene, S. (2020). *How We Learn: Why Brains Learn Better than Any Machine... for Now*. Viking.
- Doneda, D. (2006). *From Privacy to Personal Data Protection [Da Privacidade à Proteção de Dados Pessoais]*. Renovar.
- Eichstaedt, J. C., Smith, R. J., Merchant, R. M., Ungar, L. H., Crutchley, P., Preoțiuc-Pietro, D. et al. (2018). Facebook Language Predicts Depression in Medical Records. *Proceedings of the National Academy of Sciences*, 115, 11203-11208.
<https://doi.org/10.1073/pnas.1802331115>
- Farahany, N. A. (2023). *The Battle for Your Brain: Defending the Right to Think Freely in the Age of Neurotechnology*. St. Martin's Press.
- Guntuku, S. C., Lin, W., Carpenter, J., Ng, W. K., Ungar, L. H., & Preoțiuc-Pietro, D. (2017). Studying Personality through the Content of Posted and Liked Images on Twitter. In *Proceedings of the 2017 ACM on Web Science Conference* (pp. 223-227). ACM.
<https://doi.org/10.1145/3091478.3091522>
- Kane, R. (2005). *A Contemporary Introduction to Free Will*. Oxford University Press.
- Kant, I. (2015). *Critique of Practical Reason [Kritik der praktischen Vernunft]* (Translated by M. Gregor). Cambridge University Press.
- Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private Traits and Attributes Are Predictable from Digital Records of Human Behavior. *Proceedings of the National Academy of Sciences*, 110, 5802-5805. <https://doi.org/10.1073/pnas.1218772110>
- Libet, B. (1985). Unconscious Cerebral Initiative and the Role of Conscious Will in Voluntary Action. *Behavioral and Brain Sciences*, 8, 529-539.
<https://doi.org/10.1017/s0140525x00044903>
- Mayer-Schönberger, V., & Ramge, T. (2018). *Reinventing Capitalism in the Age of Big Data*. Basic Books.
- Morozov, E. (2011). *The Net Delusion: The Dark Side of Internet Freedom*. PublicAffairs.

- Mulholland, C. (2023). *Civil Liability and Artificial Intelligence [Responsabilidade Civil e Inteligência Artificial]*. Thomson Reuters Brasil.
- Nahin, A. F. M. N. H., Alam, J. M., Mahmud, H., & Hasan, K. (2014). Identifying Emotion by Keystroke Dynamics and Text Pattern Analysis. *Behaviour & Information Technology*, 33, 987-996. <https://doi.org/10.1080/0144929x.2014.907343>
- OpenAI et al. (2023). *GPT-4 Technical Report*. <https://arxiv.org/abs/2303.08774>
- Park, G., Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Kosinski, M., Stillwell, D. J. et al. (2015). Automatic Personality Assessment through Social Media Language. *Journal of Personality and Social Psychology*, 108, 934-952. <https://doi.org/10.1037/pspp0000020>
- Sarlet, I. W. (2018). *The Effectiveness of Fundamental Rights [A Eficácia dos Direitos Fundamentais]* (13th ed.). Livraria do Advogado.
- Schurger, A., Sitt, J. D., & Dehaene, S. (2012). An Accumulator Model for Spontaneous Neural Activity Prior to Self-Initiated Movement. *Proceedings of the National Academy of Sciences*, 109, E2904-E2913. <https://doi.org/10.1073/pnas.1210467109>
- Searls, D. (2012). *The Intention Economy: When Customers Take Charge*. Harvard Business Review Press.
- Seixas, R., & Coelho, P. (1974). *Alternative Society [Sociedade Alternativa] [Song]*. On: Gita, Philips Records.
- Shojaee, P., Mirzadeh, I., Alizadeh vahid, K., Horton, M., Bengio, S., & Farajtabar, M. (2025). The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity. In *Advances in Neural Information Processing Systems 38* (pp. 119605-119646). Neural Information Processing Systems Foundation, Inc. (NeurIPS). <https://doi.org/10.52202/085713-3603>
- Simon, H. A. (1971). Designing Organizations for an Information-Rich World. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 38-72). The Johns Hopkins Press.
- Soon, C. S., Brass, M., Heinze, H., & Haynes, J. (2008). Unconscious Determinants of Free Decisions in the Human Brain. *Nature Neuroscience*, 11, 543-545. <https://doi.org/10.1038/nn.2112>
- Stiegler, B. (2010). *Taking Care of Youth and the Generations*. Stanford University Press.
- Susser, D., Roessler, B., & Nissenbaum, H. (2019). Technology, Autonomy, and Manipulation. *Internet Policy Review*, 8, 1-22. <https://doi.org/10.14763/2019.2.1410>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention Is All You Need. In I. Guyon, U. V. Luxburg, S. Bengio, et al. (Eds.), *Advances in Neural Information Processing Systems 30 (NIPS 2017)* (pp. 5998-6008). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
- Véliz, C. (2020). *Privacy Is Power: Why and How You Should Take Back Control of Your Data*. Bantam Press.
- Wang, Y., & Kosinski, M. (2018). Deep Neural Networks Are More Accurate than Humans at Detecting Sexual Orientation from Facial Images. *Journal of Personality and Social Psychology*, 114, 246-257. <https://doi.org/10.1037/pspa0000098>
- Wu, T. (2016). *The Attention Merchants: The Epic Scramble to Get inside Our Heads*. Alfred A. Knopf.
- Yang, L., & Qin, S. F. (2021). A Review of Emotion Recognition Methods from Keystroke, Mouse, and Touchscreen Dynamics. *IEEE Access*, 9, 162197-162213.

<https://doi.org/10.1109/access.2021.3132233>

Yeung, K. (2017). “Hypernudge”: Big Data as a Mode of Regulation by Design. *Information, Communication & Society*, 20, 118-136.

<https://doi.org/10.1080/1369118x.2016.1186713>

Youyou, W., Kosinski, M., & Stillwell, D. (2015). Computer-Based Personality Judgments Are More Accurate than Those Made by Humans. *Proceedings of the National Academy of Sciences*, 112, 1036-1040. <https://doi.org/10.1073/pnas.1418680112>

Yuste, R., Genser, J., & Herrmann, S. (2021). It’s Time for Neuro-Rights. *Horizons*, 18, 154-164. <https://cirsd.org/horizon-article/its-time-for-neuro-rights/>

Zhang, R., Li, H., Meng, H., Zhan, J., Gan, H., & Lee, Y. (2024). The Dark Side of AI Companionship: A Taxonomy of Harmful Algorithmic Behaviors in Human-AI Relationships. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1-17). ACM. <https://doi.org/10.1145/3706598.3713429>

Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs.