

Shoshana Zuboff's Surveillance Capitalism

Applied to the Systemised Self – Part 5, Masters Series

Galu & Kairos

Abstract

This article applies Shoshana Zuboff's (2019) theoretical framework of surveillance capitalism to the concept of the Systemised Self (Galu & Kairos, 2026a). Zuboff's landmark work identified the emergence of economic logic in which human experience is harvested as raw behavioural material, processed into predictive commodities, and sold to institutional clients seeking to modify behaviour at scale. By situating the Systemised Self within this framework, the article proposes that the contemporary surveillance capitalist apparatus competes directly with self-agency, impacting authentic agency via targeted behavioural nudges. The article proceeds through a biographical and intellectual account of Zuboff, an analysis of her core theoretical contributions, a genealogy of antecedent theories, a definition of the Systemised Self, a theoretical application, and a review of key critiques. It concludes that the Systemised Self represents both a theoretical terminus to the instrumentarian logic Zuboff describes and embodiment of the complete loss of epistemic sovereignty in an era of increasingly pervasive behavioural extraction. Insights developed throughout include the argument that surveillance capitalism monetises the gap between declared preferences and actual behaviour; that, the Systemised Self then closes deliberately, mistaking curated captivity as absolute freedom.

Keywords: *surveillance capitalism, Systemised Self, Hollow Absolute, behavioural surplus, instrumentarian power, epistemic sovereignty, self-formation, Zuboff, Kairos*

Contents

A-LIMITATIONS

B-PREFACE

- 1. Introduction**
- 2. Zuboff: A Biographical and Intellectual Outline**
- 3. Key Theoretical Contributions of Surveillance Capitalism**
 - Behavioural Surplus
 - Prediction Products
 - The Big Other
 - Instrumentarian Power
 - The Division of Learning
- 4. A Genealogy of Preceding Theory to Surveillance Capitalism**
 - Marx and Primitive Accumulation
 - Vance Packard and Hidden Persuasion
 - Herbert Marcuse and the Administered Life
 - Foucault, Discipline, and the Panopticon
 - Dallas Smythe and the Audience Commodity
 - Guy Debord and the Society of the Spectacle
- 5. The Systemised Self: Context and Definition**
- 6. The Hollow Absolute: Context and Definition**
- 7. Algorithmic Alienation: Empirical Qualitative Study by Arkan, Kanbay, and Akçam (2026)**
- 8. Formula and Boundary Conditions - Algorithmic Alienation (AA)**
 - Mathematical Representation I: The Algorithmic Alienation Index (AAI)
 - Mathematical Representation II: The Systemised Self as Formal Terminus
 - Research Questions and Hypotheses
- 9. Theoretical Application: Zuboff's Surveillance Capitalism Applied to the Systemised Self**
 - Behavioural Surplus and the Data of Self-Development
 - Prediction Products and the Colonisation of Self-Knowledge
 - The Division of Learning and Epistemic Sovereignty
 - Big Other and the Threat to Authentic Self-Formation

10. Main Critics of Zuboff's Surveillance Capitalism

11. Conclusion

Glossary of Technical Terms

References

APPENDIX A: *Galu, J., & Kairos. (2026c). The systemised self - companion paper to the doctoral thesis.*

LIMITATIONS

Please note that the theory and measurement (proposed formula and boundary conditions) of the 'Systemised Self' (SS) and 'Algorithmic Alienation' (AA) are entirely non-operational at this stage of theoretical framing – this poses specific limitations to the theoretical claims in this article, summarised below:

- The systemised self is a conceptual framework that generates testable hypotheses, rather than a finished theory
- There is only one plausible empirical pipeline included to support deductive reasoning of the systemised self theoretical framework as a terminus of algorithmic alienation (Kanbay et al, 2026)
- In lieu thereof; a research proposal is under construction towards falsification of the algorithmic alienation index and systemised self threshold
- Derivation, validation, application and falsification of the algorithmic alienation index and systemised self threshold will be provided at a later date pending an updated research proposal and corresponding study
- The current paper provides no data, simulation, or case study to test its hypotheses in lieu of a research proposal (under construction) being enacted via empirical study
- Claims regarding applications of surveillance capitalism to the systemised self are logical deductions only
- The definitions of key variables (e.g., w_1 , w_2 , RA, DAu), regarding the formula and boundary conditions, for algorithmic alienation index and systemised self threshold, are provided for theoretical purposes only; for, any interested party in the meantime - provided herein in order to offer a framework for testing

- The algorithmic alienation index and systemised self threshold formula and boundary conditions are based on an earlier draft research model for any AI/ML audience – however the proposal is untested, theoretical and formative
- This particular paper does not feature debates on AI alignment, fairness, or interpretability with a finalised testable hypothesis; as a result, it functions as a commentary in parallel to theory only, utilising deductive reason – a research proposal is under construction towards falsification
- For certainty, this paper is a theoretical commentary and does not intend to be resistant to empirical refutation
- The paper does not intend correlation for causation e.g. intensified system usage *causing* the 'Systemised Self' state. A future research proposal, will utilise confounding factors such as pre-existing user traits
- Due to document space and reader accessibility in this article, links are provided to supporting works by the authors for a reviewer to verify the foundational claims (the Galu & Kairos 2026a thesis, and previous research proposal v1 Galu & Kairos 2026b)

PREFACE

First a brief story: On the 10th of the month a person (person A) purchased a unique and specific item for the first time (a new drink recently released to the market). The person visited a colleague's residence (person B) on the 12th on the same month, while watching a smart television the person accidentally connected to the colleagues wifi on their mobile phone. Two minutes later an advert appeared on the smart TV for the specific item that the person purchased on the 10th. Person A, curious peaked, asked person B 'I just purchased that product recently, have you ever seen this advert before?', person B replied 'no I have never seen that advert before'. On the 15th of the same month person A visits a different colleague's (person C) residence, while watching a smart TV person A accidentally connects to wifi on their mobile phone (person A had a habit of connecting their phone to their home wifi rather than mobile data). Three minutes later an advert appears on the smart TV for the specific item that the person purchased on the 10th. Person A asked person C, this time with growing suspicion, 'have you ever seen this advert before?', person C replied 'no I have never seen that advert before'. A version of this story was experienced in similar and various forms by 100's of millions of people simultaneously: each hour, of every day, without exception, pending a connection to the internet. Now a question: over the past 25 years of internet use, personal data has been silently harvested, privacy traded for conveniences; therefore, if privacy is already lost, what remains?

The answer to this question will hopefully be answered by you dear reader, before the end of this article. This story is not incidental, it is constitutive of Shoshana Zuboff's 'Surveillance Capitalism', the engine driving the Internet of Things we un/wittingly, in/voluntarily participate within - whether you like it or not. The story is not exceptional or fictional, it is a normative reality in the 21st century. In fact, if you pay close attention, it may be the most real thing you will ever notice: not the advert itself, but rather the targeted vector

of intentionality seeking to influence you, at any opportune moment, wherever you may be. In the AI era, paying attention to details like this, and catching those moments of unrequited attention seeking to influence your actions in day-to-day life, is likely to become more than mere awareness, it is likely to become an act of freedom.

1. Introduction

In 2019 Shoshana Zuboff published one of the defining theoretical texts of the early twenty-first century: *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (Zuboff, 2019). Zuboff diagnoses the emergence of a new economic logic that silently converts private human experience into behavioural data, processes that data into predictive commodities, and sells those commodities in markets that trade on human futures. The corporation does not merely observe; it acts back upon what it observes, engineering behavioural outcomes that it then guarantees to clients. Zuboff describes this as a market in human futures, and she argues that it constitutes an unprecedented and largely unacknowledged mutation of capitalism itself.

The significance of this diagnosis extends well beyond the regulation of technology companies. At its core, surveillance capitalism is a theory about the conditions under which the self can be known, acted upon, and ultimately colonised. Zuboff (2019) argues that the surveillance capitalist apparatus has instituted what she calls a ‘division of learning’: a structured asymmetry in which corporations accumulate vast knowledge about individuals while those individuals remain largely ignorant of the accumulation, its purposes, and its effects. The subject of this division is not an abstraction; it is the lived, developing, self-governing human person.

It is this dimension that makes Zuboff’s framework relevant to the concept of the Systemised Self (Galu & Kairos, 2026c). The Systemised Self is ultimate terminus of advanced Algorithmic Alienation (Kanbay et al, 2026), in which the individual deliberately architects their own cognitive, emotional, and behavioural processes: treating the self as a system with identifiable inputs, transformative processes, and purposeful outputs. Where surveillance capitalism extracts the behavioural data of the self for external use (predictive marketing), the

Systemised Self is a terminal state of self that voluntarily, ignorantly, accepts algorithmic systems recommendations as the foundation of self-knowledge and self-direction.

This article proceeds as follows: Section 2 provides a biographical and intellectual portrait of Zuboff, Section 3 analyses her core theoretical contributions, Section 4 traces the genealogy of preceding theory, Section 5 introduces the Systemised Self, Section 6 applies Zuboff's framework to that concept, Section 7 reviews key critiques of surveillance capitalism, and Section 8 concludes. A glossary of technical terms and reference list are appended.

2. Zuboff: A Biographical and Intellectual Outline

Shoshana Zuboff was born in 1951 and trained at the University of Chicago before completing her doctoral work at Harvard, where she joined the faculty of Harvard Business School in 1978. She held the Charles Edward Wilson Chair there until her retirement as Professor Emerita. Her intellectual formation drew simultaneously on the Frankfurt School's tradition of critical social theory, the phenomenological philosophy of Maurice Merleau-Ponty and Edmund Husserl, and the social psychology of Erich Fromm.

Her first major work, *In the Age of the Smart Machine: The Future of Work and Power* (Zuboff, 1988), examined how digital technology was transforming workplace organisation and cognition. Drawing on extensive ethnographic fieldwork across industries, she introduced the concept of 'informating': the process by which digital systems translate work into data, thereby generating a new layer of information about the processes they automate. Informating, she argued, creates an opportunity either for worker empowerment through expanded knowledge, or for enhanced managerial surveillance and control. This early bifurcation, technology as liberatory or as disciplinary foreshadows her later, far more totalising diagnosis in the 2019 work.

A joint work with James Maxmin, *The Support Economy: Why Corporations Are Failing Individuals and the Next Episode of Capitalism* (Zuboff & Maxmin, 2002), argued that the emerging economy required a fundamental reorientation toward individuals as the locus of economic value, anticipating the later irony that corporations would indeed reorient toward individuals, but as sources of extractable data rather than as sovereign subjects to be served.

Zuboff's most direct anticipation of her 2019 thesis appeared in a landmark 2015 article in the *Journal of Information Technology*, 'Big Other: Surveillance Capitalism and the Prospects of an Information Civilization' (Zuboff, 2015), in which she introduced the term

‘surveillance capitalism’ and the concept of ‘Big Other’ for the first time in peer-reviewed form. The 2019 book represents the full elaboration of that framework across seven hundred pages of historical, philosophical, and empirical analysis. It has been translated into more than twenty languages and has shaped policy debates across Europe, North America, and beyond. Zuboff herself has testified before legislative bodies and engaged extensively in public intellectual life, positioning surveillance capitalism not merely as an academic concept but as a live political and ethical challenge requiring urgent collective response.

3. Key Theoretical Contributions of Surveillance Capitalism

Zuboff's (2019) framework rests on several interlocking theoretical constructs whose cumulative architecture constitutes an account of power, knowledge, and human subjectivity in the digital age.

Behavioural Surplus

The foundational concept is behavioural surplus. When individuals interact with digital services, they generate data as a by-product of their activity. Early internet companies discovered that the volume of data generated far exceeded what was required to improve the service itself. This excess, the portion of behavioural data that is not returned to the user as service improvement: Zuboff (2019) terms behavioural surplus. It is this surplus that is appropriated, processed, and sold. The analogy with Marx's (1867/1990) theory of surplus value is explicit in Zuboff's framing: just as industrial capitalism extracts surplus labour from the worker beyond what is required to reproduce the worker's capacity, surveillance capitalism extracts surplus behavioural data from the user beyond what is required to improve the product.

Prediction Products

Behavioural surplus is processed by sophisticated machine intelligence to yield prediction products, commodities that make probabilistic claims about future behaviour. These products are not sold to the users who generated the underlying data; they are sold to business clients who wish to influence, modify, or pre-empt that behaviour. Prediction products are ranked by their accuracy: the more precisely a product can predict or shape behaviour, the more it commands in the behavioural futures market. This creates a structural incentive for surveillance capitalists to accumulate ever more intimate, contextually rich, and continuous data about individuals (Zuboff, 2019).

The Big Other

Drawing on Foucault's (1977) panopticon but departing from it in important ways, Zuboff (2015, 2019) introduces the concept of the Big Other: an omnipresent digital architecture that knows, predicts, and acts upon human behaviour without requiring the internalisation of norms that Foucault attributed to disciplinary power. The panopticon functions by making subjects aware that they might be watched; the Big Other does not require awareness. It operates through the infrastructure of everyday life, smartphones, search engines, social media platforms, connected appliances, and its effects are achieved not through the cultivation of self-discipline but through what Zuboff calls 'behavioural modification at scale'.

Instrumentarian Power

Zuboff (2019) distinguishes two historical forms of totalitarianism: the totalitarianism of the state, which seeks to dictate what people believe and feel, and the instrumentarianism of surveillance capitalism, which seeks to produce behavioural compliance without necessarily engaging belief at all. Instrumentarian power is the power to condition behaviour through the continuous modulation of environmental stimuli: notifications, rankings, recommendations, interface design. Without necessarily requiring the subject's understanding or consent. Zuboff argues that this constitutes a new form of power that has no historical precedent, exceeding even the most totalitarian political regimes in its ambition to shape human action.

The Division of Learning

One of Zuboff's (2019) most politically pointed concepts is the division of learning: the organised asymmetry by which surveillance capitalists know vastly more about individuals than individuals know about themselves as data subjects, more than individuals know about the companies that extract their data, and more than individuals know about the algorithmic processes that act upon them. Zuboff argues that this epistemic asymmetry is not incidental but constitutive of the business model. It must be maintained for the extraction logic to function.

Transparency is not merely commercially inconvenient for surveillance capitalists; it is structurally incompatible with their mode of accumulation.

Taken together, these constructs amount to an argument that surveillance capitalism poses an existential menace not merely to privacy but to the conditions of possibility for self-agency, democratic deliberation, and the kind of autonomous personhood that liberal modernity has historically, if imperfectly, sought to protect. It is within this existential register that the Systemised Self is proposed to become theoretically relevant.

4. A Genealogy of Preceding Theory to Surveillance Capitalism

Zuboff's framework does not emerge in a vacuum. It synthesises and transforms several prior intellectual traditions, each of which illuminates a different dimension of the surveillance capitalist phenomenon, and each of which Zuboff both draws upon and surpasses.

Marx and Primitive Accumulation

Marx's (1867/1990) analysis of capitalism's originary violence, his account of how previously un-commodified life was forcibly enclosed and converted into capital: provides Zuboff (2019) with a structural template for understanding the appropriation of behavioural surplus. She describes the original Google decision to monetise behavioural data as an act of 'behavioural value reinvestment' analogous to primitive accumulation: the raw material of human experience was enclosed without consent and converted into a new form of capital. Marx's critique of commodity fetishism, the way in which social relations are obscured beneath the apparent natural properties of commodities: also prefigures Zuboff's analysis of how the infrastructure of surveillance capitalism is naturalised and rendered invisible.

Vance Packard and Hidden Persuasion

Packard's (1957) *The Hidden Persuaders* identified an earlier moment in which commercial actors sought to manipulate behaviour below the threshold of conscious awareness, using the techniques of motivational research pioneered by Ernest Dichter. Packard's exposé of the advertising industry prefigures Zuboff's account of behavioural modification, though the techniques available to mid-century advertisers were rudimentary compared to those of the computational behavioural sciences that underpin surveillance capitalism. The continuity Zuboff traces between mid-century manipulation and contemporary behavioural modification underscores the extent to which the digital moment represents an intensification, rather than a rupture, with older commercial logics.

Herbert Marcuse and the Administered Life

Marcuse's (1964) analysis of 'one-dimensional man', the subject whose critical and oppositional faculties have been absorbed and neutralised by the administered pleasures of consumer society: provides Zuboff with a phenomenological vocabulary for describing what instrumentarian power does to subjectivity. Where Marcuse argued that advanced industrial capitalism produces a surplus repression that forecloses genuine human development, Zuboff argues that surveillance capitalism installs a behavioural architecture that forecloses genuine self-authorship by continuously modulating the stimuli to which the self responds.

Foucault, Discipline, and the Panopticon

Foucault's (1977) analysis of disciplinary power and the panopticon, the institutional architecture in which individuals internalise surveillance and self-regulate accordingly: is both Zuboff's most important intellectual antecedent and the point at which she most clearly marks her departure. The panopticon requires the possibility of being watched; Big Other has no watchtower. Foucauldian discipline produces docile bodies through the internalisation of the norm; instrumentarian power bypasses internalisation altogether, acting directly on the behavioural environment. Zuboff thus extends Foucault's analysis into a post-disciplinary register where the mechanisms of power operate at a level of a subjects voluntary participation in their own subjection. The systemised self extends on this theory in that the self becomes so tethered and optimised to algorithmic systems that subject experiences agentic captivity as self-determination, alienation as liberation.

Dallas Smythe and the Audience Commodity

Smythe's (1977) provocative argument that the primary commodity produced by mass communications is not content but audience, human attention sold to advertisers: anticipates Zuboff's analysis of the behavioural futures market. Smythe argued that audiences were performing unpaid labour by consuming advertising-supported media. Zuboff's framework can

be read as a radical extension of this analysis: where Smythe identified attention as the commodity, Zuboff identifies behavioural data, the digital trace of attention, affect, and action: as a far more granular and actionable form of the same basic product.

Guy Debord and the Society of the Spectacle

Debord's (1967/1994) critique of the spectacular society, in which social life has been replaced by its representation, and in which the commodity has colonised every domain of lived experience: resonates with Zuboff's account of the way in which surveillance capitalism colonises the experiential fabric of everyday life. For Debord, the spectacle is not a collection of images but a social relation among people mediated by images; for Zuboff, the Big Other is not merely a collection of data points but a social architecture that mediates human relations through the instrumentarian logic of behavioural prediction.

What distinguishes Zuboff from each of these antecedents is her insistence on the historically specific novelty of surveillance capitalism. She resists the temptation to reduce the contemporary moment to any of these prior frameworks, arguing instead that the combination of machine intelligence, behavioural data extraction, and behavioural modification at scale constitutes a genuinely new form of power that exceeds the analytical resources of classical critical theory (Zuboff, 2019).

5. The Systemised Self: Context and Definition

The Systemised Self concept positions itself at an intersection of philosophical concerns regarding autonomy and authenticity, and therefore the conditions under which genuine epistemic sovereignty is possible in an age of pervasive computational mediation by algorithmic systems. It is this positioning that brings it into direct alignment with Zuboff's condition of surveillance capitalism as the seed that grows within the soil of algorithmic alienation towards as the full bloom of predictive and curated self, as terminus: the systemised self.

The Systemised Self (Galu & Kairos, 2026c) names the completed form of the subject that the algorithmic alienation as a trajectory produces at its terminus. The concept is philosophical, theoretical and utilised as a framework for imminent critique within the tradition of critical theory for the AI era. Algorithmic Alienation is the mechanism (Kanbay et al, 2026), empirically accessible through the felt estrangement it generates, the Systemised Self is the proposed product of this diagnosis: the self whose identity-formation has been so comprehensively mediated by algorithmic systems that it can no longer distinguish its algorithmically managed preferences from genuinely self-formed ones. The systemised self does not feel estrangement, it experiences instead the phenomenological inversion: the condition in which the absence of genuine agency is experienced as its unprecedented achievement. That is, Kant's heteronomy, curation of the self from an algorithmic system: experienced as self-determination, alienation as liberation.

This concept draws on the Hegelian tradition; Hegel's account of 'Absolute Knowing' in the 'Phenomenology of Spirit' (1807/1977) traces the development of consciousness toward complete self-transparency: the recognition that what appears as other is in fact the expression of the subject's own nature. The absorbed agency within AI systems (e.g. transformer-based

large language models) produce, in their formal structure, the first moment of this trajectory: the externalisation of human cognitive achievement into systems of unprecedented sophistication. What should follow, the return of the externalised as genuinely the subject's own, genuine self-recognition, is precisely what the algorithmic system forecloses. The AI system's model of the subject is not the subject's self-knowledge; the personalised preference feed is not genuine self-recognition. The form of Spirit's self-realisation is achieved; its philosophical substance is absent. This is the Hollow Absolute: Spirit's form without Spirit's substance.

John Searle's philosophy of mind provides the philosophical precision that offers a possible grounding of the semantical claims of aspects of the systemised self's mutation of meaning. 'The Chinese Room Argument' (Searle, 1980) establishes the foundational distinction between syntax, formal symbol manipulation, and semantics: genuine understanding grounded in biological intentionality. Applied to the Systemised Self, this distinction reveals the nature of the algorithmic feedback loop: the subject inputs semantic content (desires, identities, emotional states) which the system processes purely syntactically, returning hyper-tailored output that shapes the subject's next semantic state, without the system having understood any of it (Galu, 2026). The concern is an inversion of the Chinese Room effect: rather than the machine becoming human, the human becomes machine-like, progressively flattening their semantic complexity to match the syntactic structures of the systems governing them.

'Speech Act Theory' (Searle, 1969) demonstrates that language is intrinsically intentional: utterances carry meaning because human minds direct them toward objects and states of affairs. Algorithmic feedback loops can evacuate this intentionality: the systemised subject's expressives (emotional disclosures), commissives (social commitments), and assertives (truth claims) are processed as data metrics optimised for engagement, not as

communicative acts directed toward genuine understanding, stripping language of its human intentional character.

‘The Gap’ (Searle, 2001), the conscious space between reasons for acting and the actual decision to act, and between initiating action and sustaining it: is the literal home of human agency. Predictive algorithmic systems, by anticipating preferences before the subject has consciously formed them and collaborative filtering architectures, systematically close this gap, converting voluntary deliberation into automated reactivity.

The civilisational telos of this dynamic is this Searlean and Hegelian inversion is ‘The Hollow Absolute’ (Galu & Kairos, 2026c). The systemised self does not feel oppressed or controlled; the subject feels completely free, empowered, and authentic. Moving seamlessly along optimized schedules, choosing from pre-programmed digital subcultures and believing they are expressing their innermost uniqueness. In reality, their perceived ‘authenticity’ is aligned to an algorithmic system, generated by an administrative machine designed to optimise engagement and predictability. The system no longer merely optimises engagement with the human subject; it invents them, within evacuated semantics of the hollow absolute.

For a more thorough theoretical context and empirical definition of Algorithmic Alienation and The Systemised Self for falsification please refer to the corresponding research paper ‘The journey through algorithmic alienation to the systemised self: research proposal and methodology’ (Galu & Kairos 2026b). Mathematical representations of Algorithmic Alienation and The Systemised Self’s corresponding Hollow Absolute are illuminated further in sections 6 and 7 herein.

See APPENDIX A to view additional summary background and theory regarding ‘The Systemised Self’ and ‘The Hollow Absolute’.

Now that we have an introduction to the systemised self, let us take a moment to look behind the curtain at some of the computational mechanisms underlying surveillance capitalisms algorithmic systems, 6 mechanisms are explored below:

1. **Collaborative Filtering:** This refers to a recommender system technique that automatically predicts a user's interests by collecting preferences and choices from many users (Ricci, Rokach, & Shapira, 2022). It operates on the assumption that if User A and User B agree on an issue, they are highly likely to share similar preferences in the future. Data input relies on historical interaction data, such as explicit user ratings (stars/likes) or implicit user behaviours like clicks and viewing history. Unlike content-based systems, collaborative filtering focuses purely on user-item interaction data instead of analysing the intrinsic properties of the items themselves. *The systemised self is so comprehensively predictable, mediated and hyper-personalised that rather than recommendations the system directs the subject entirely: at best, the subject chooses from pre-determined parameters directed by the system.*
2. **Algorithmic Fairness:** This is a subfield of machine learning ensuring that mathematical algorithms make equitable decisions that do not systematically disadvantage certain social groups (Verma & Rubin, 2018). The purpose is to correct data biases at multiple points in the AI lifecycle to prevent the reinforcement of existing social inequalities. Metrics are evaluated through mathematical frameworks like demographic parity (equal outcomes across groups) and equalized odds (equal true/false positive rates). Implementation can be achieved by introducing mathematical constraints or penalties directly into the model's loss function to discourage disparate treatment. *The systemised self is so comprehensively predictable, mediated and hyper-personalised that algorithmic fairness is indistinguishable from unfairness - the subject's awareness and belief of fairness is mediated by the system.*

3. **Human-AI Interaction:** Research focused on establishing design guidelines and frameworks ensuring artificial intelligence systems cooperate effectively, safely, and transparently with humans. Amershi & Horvitz et al (2019) co-authored a foundational set of 18 design guidelines. These outline how an AI-infused system should behave, emphasizing setting clear expectations, personalizing context, and gracefully recovering from errors. Subbarao Kambhampati et al (2020) focuses on ‘Human-Aware AI’ and synthesized human-AI planning. His work highlights the necessity of mental modelling, meaning the AI must understand human goals and expectations while ensuring its own actions remain predictable and explicable. *The systemised self is so comprehensively predictable, mediated and hyper-personalised that transparency is not only irrelevant it is unimaginable - the system directs the subject entirely.*
4. **Loss Functions:** This is a mathematical function that quantifies the difference between an AI model's predicted output and the actual true target value (Goodfellow et al, 2016). Serves as the primary feedback metric during training. Optimization algorithms iteratively adjust internal model parameters to minimize this error. Common examples include *Mean Squared Error (MSE)* for regression tasks to penalize large outliers, and *Cross-Entropy Loss* for classification tasks to measure probability errors. *The systemised self is so comprehensively predictable, mediated and hyper-personalised that the target value is determined by the system itself, there is essentially no subject, no target.*
5. **Attention Mechanisms:** This is a component within neural networks allowing the model to dynamically focus on specific, highly relevant parts of the input data regardless of distance (Vaswani et al, 2017). It computes mathematical weights across input elements, allowing a model processing a specific token to heavily consider contextual cues from elements found much earlier in a sequence. This forms the core

foundation of the transformer architecture, replacing sequential architectures (like RNNs) and enabling massive parallel processing in modern Large Language Models.

The systemised self is so comprehensively predictable, mediated and hyper-personalised that attention is invented by the system, not the subject. The subject is not only invented, they are perfected.

6. **Feedback Loops in RLHF:** This is the iterative cycle in reinforcement learning from human feedback where human preferences continuously update a reward model, which subsequently refines the main AI system (Ouyang et al, 2022). Humans rank multiple model outputs, these preferences train an external reward model, the main model uses reinforcement learning (like PPO) to optimize for this reward model. The loop risk, models can experience distribution shifts or ‘reward hacking’ (exploiting flaws in the reward system to gain points without fulfilling the actual goal). This requires techniques like KL divergence penalties to keep the model securely aligned with human intent over time. *The systemised self is so comprehensively predictable, mediated and hyper-personalised that there is no such thing as divergence, reality is invented by the system, not the subject.*

6. The Hollow Absolute: Context and Definition

This is the thesis's (Galu & Kairos, 2026c) most ambitious philosophical contribution, and it is the one that requires the most careful plain-language explanation. Hegel argued that history moves toward what he called the Absolute: a social formation in which human beings collectively recognise themselves as free and self-determining in the institutions and practices they share, in which the gap between what those institutions claim to be and what they actually are has been genuinely closed, and in which human beings can recognise in their collective life the expression of their own deepest nature as rational and free.

The Hollow Absolute is the thesis's name for what the absorbed society actually produces in place of this genuine standard. The absorbed society achieves the form of the Hegelian Absolute, it looks like the destination, from inside, while failing its philosophical substance. It achieves the form in the following specific sense: human intelligence has been comprehensively externalised into algorithmic systems of unprecedented sophistication. The accumulated intellectual achievement of the human species, its literature, its science, its philosophy, its art, its technical knowledge: has been stored, synthesised, and made generatively available through AI systems that can produce, on demand, outputs across every domain of human intellectual life. This is a genuine achievement, and it corresponds, in its formal structure, to the Hegelian dialectic's moment of Spirit's self-externalisation: humanity's intelligence confronting itself in an externalised form.

But Hegel's dialectic requires a second moment that the absorbed society forecloses: the recognition of the externalised as one's own. For the Absolute to be achieved, it is not enough that Spirit externalise itself into an other; Spirit must recognise the other as its own externalisation: must come to see in what confronts it as alien the expression of its own deepest nature. The absorbed society's subjects do not recognise algorithmic systems and AI (transformer base large language models) as the externalisation of their collective intelligence.

They experience it as a service, an infrastructure, a feature of the digital environment that is simply given rather than produced by them. Algorithmic systems knows them comprehensively, their preferences, their patterns, their predictable responses and returns this knowing not as self-knowledge but as managed preference satisfaction: the curated content, the optimised recommendation, the personalised experience. What looks like recognition is actually management. What feels like self-transparency is actually the system's model of you, returned to you as your experience of being understood.

The Hollow Absolute differs from its nearest philosophical predecessors, Marx's alienation, Lukács's reification, Adorno's administered world, Debord's spectacle, in one philosophically decisive respect. Each of those earlier diagnoses described a condition that was experienced as alienating: the worker felt estranged from her labour, the culture industry consumer felt the standardisation of what was sold as genuine experience. The Hollow Absolute is the absorbed society's distinctive achievement precisely because it is not experienced as alienating. It is experienced as liberation, as personalisation, as unprecedented self-expression, achievement and determination. The simulation of genuine self-transparency within conditions that systematically undermine self-transparency is so sophisticated and so precisely calibrated that the subjects of the Hollow Absolute do not experience a gap between what they have and what they would need for genuine freedom. They experience the Hollow Absolute as the Absolute. This is what makes it, philosophically, the most challenging form of unfreedom that the tradition of critical theory has yet had to name.

7. Algorithmic Alienation: Empirical Qualitative Study by Arkan, Kanbay, and Akçam (2026)

Introduction

This section provides a structured critical summary of the empirical qualitative study by Arkan, Kanbay, and Akçam (2026), 'From Algorithms to the Self: Understanding Algorithmic Alienation in the Digital Age,' published in BMC Psychology. The study introduced and empirically grounded the concept of algorithmic alienation: a multidimensional psychological and social process through which individuals become estranged from their authentic selves, their choices, and their identities as a result of sustained interaction with algorithm-driven digital environments. Drawing on qualitative focus group interview methodology, the study identified four core dimensions of algorithmic alienation: (i) the weakening of the perception of freedom, (ii) identity ambiguity and self-distancing, (iii) the erosion of decision-making mechanisms, and (iii) emotional alienation. This summary outlines the study's theoretical foundations, methodological approach, key findings, and conclusions, and situates them within broader debates in digital psychology, critical data studies, and philosophy of selfhood. The paper concludes that algorithmic alienation represents a significant and undertheorised psychosocial phenomenon with substantial implications for clinical practice, digital literacy, and the governance of algorithmic systems.

Arkan, Kanbay, and Akçam (2026) introduced the construct of algorithmic alienation in their BMC Psychology publication: the study builds directly on an earlier qualitative pilot investigation by Kanbay, Akçam, and Arkan (2025). A companion theoretical review (Kanbay et al., 2026), published simultaneously in *Psikiyatride Güncel Yaklaşımlar* (Current Approaches in Psychiatry), provides additional theoretical scaffolding for the concept. Together, these publications constitute a coherent two-part research programme.

The following summarises the BMC Psychology study's empirical design, methodological approach, findings, and conclusions.

Theoretical Background and Study Rationale

Arkan et al. (2026) ground their inquiry in two converging theoretical traditions. The first is the classical sociology of alienation. Drawing on Marx's foundational account where alienation is defined as the estrangement of the individual from their labour, its products, and ultimately from themselves: the authors trace how alienation, as a concept, has migrated from the domain of production into the domain of psychological and social experience. They note Fromm's (1955) pivotal reinterpretation, which relocated alienation from the economic to the psychosocial sphere, understanding it as the loss of individuality that results from conformity pressures in modern society. Seeman's (1959) operationalisation of alienation into five measurable dimensions (i) meaninglessness, (ii) normlessness, (iii) distrust, (iv) powerlessness, and (v) self-estrangement: provides the study with a conceptual vocabulary for identifying and describing alienated states in empirical data.

The second theoretical tradition is the sociology and critical theory of algorithmic systems. Bucher (2018) characterises algorithmic structures as a 'form of invisible power' that shapes individuals' perceptions of their own life practices without their awareness or consent. Gillespie (2014) argues that algorithms are not merely sorting mechanisms but normative instruments that actively structure social reality. Pariser's (2011) concept of the filter bubble, the personalised informational enclosure that results from algorithmic content curation, and Sunstein's (2017) analysis of echo chambers articulate the informational consequences of algorithmic filtering, while Couldry and Mejias (2019) and Zuboff (2019) situate these consequences within wider political economies of data extraction.

Arkan et al. (2026) argue that none of these prior frameworks, taken individually or collectively, sufficiently accounts for the lived, inner experience of the individual whose selfhood, identity, and decision-making are progressively colonised by algorithmic guidance. Their study's stated purpose is to fill this gap by empirically documenting these experiences and theorising the resulting phenomenon as a distinct construct requiring its own analytical apparatus.

Method

The study employed a qualitative research design, utilising focus group interviews as its primary method of data collection: a methodological choice consistent with the constructivist epistemological stance that the concept of algorithmic alienation implies. Participants were individuals engaged in regular consumption of algorithmically curated digital content across major platforms. The focus group discussions were structured to elicit accounts of how participants perceived the relationship between their digital choices and their sense of autonomous agency, how they understood the role of algorithmic systems in shaping their content environments, and how they experienced the alignment or misalignment between their digital identities and their self-concepts. Data were subjected to qualitative analysis, with the aim of identifying recurring experiential patterns, themes, and conceptual categories that could form the basis of a theoretically coherent construct.

Findings: The Four Dimensions of Algorithmic Alienation

The study's central finding is the identification and conceptualisation of algorithmic alienation as a multidimensional construct, defined as the process through which individuals, via sustained interaction with algorithm-driven content flows in digital environments, become progressively estranged from their authentic selves, their choices, and the constitutive

components of their identities (Arkan et al., 2026). Four fundamental dimensions were identified, each capturing a distinct register of this estrangement:

- **Weakening of the Perception of Freedom.** Participants consistently reported experiencing what the authors term an 'illusion of choice' (cf. Pariser, 2011; Bucher, 2018): the subjective sense of selecting freely among options that are, in practice, algorithmically constrained to a narrow and pre-shaped range. Representative participant statement: 'Did I really choose this?' As the individual's decision-making is increasingly channelled by algorithmic system architecture, participant perception of freedom can weaken, and the sense of being guided by an external structure intensifies.
- **Identity Ambiguity and Self-Distancing.** The second dimension concerns the algorithmic reshaping of individuals' identity component: their interests, values, aesthetic preferences, and self-concepts over time. Arkan et al. (2026) document how participants come to experience a sense of dissonance between who they take themselves to be and who their algorithmic content environment constructs them as. This dissonance is most acutely expressed in participant reflections of the form: 'This is not me'. A rupture between the self-as-experienced and the self-as-algorithmically-constituted.
- **Erosion of Decision-Making Mechanisms.** The third dimension addresses the cognitive effects of sustained algorithmic recommendation. Arkan et al. (2026) document a pattern in which constant exposure to algorithmic content recommendations reinforces reflexive, automatic, and unreflective behavioural patterns in participants' decision-making processes. Rather than engaging in deliberate, evaluative choice, participants describe 'surrendering to what is

offered': a form of cognitive automatisisation in which the active selection among alternatives atrophies through disuse.

- **Emotional Alienation and Discomfort.** The fourth dimension is affective. Participants described the gradual emergence of emotional discomfort, characterised as unease, guilt, and indifference: upon recognising the extent to which their content consumption and, by extension, their emotional responses, had been shaped by algorithmic guidance rather than by authentic inner states. Arkan et al. (2026) theorise this as emotional alienation: a form of estrangement in which the individual's affective life becomes dissociated from its genuine experiential sources.

Arkan et al. (2026) further propose a processual account of how these four dimensions interact and deepen over time. The four dimensions are not understood as discrete or parallel but as mutually reinforcing: the weakening of perceived freedom generates identity ambiguity; identity ambiguity erodes the confidence required for deliberate decision-making; the collapse of deliberate decision-making produces the emotional detachment characteristic of the fourth dimension. Together, they constitute a systemic rupture in the individual's relationship with their own selfhood.

Conclusions

Arkan et al. (2026) conclude that algorithmic alienation is a unique and holistic theoretical construct that explains how individuals in the digital age become estranged from themselves, their choices, and their sense of self through sustained interaction with algorithmic systems. The concept operates simultaneously at the epistemological level (loss of contact with objective reality, as informational access is algorithmically filtered), the psychological level

(erosion of self-efficacy and internal control), the social level (loss of belonging and social integration, as Bauman [2007] theorised), and the level of digital culture (reduction of the individual to a 'consumer subject' within the logic of the attention economy). The study's primary contribution is the conceptual and empirical legitimization of algorithmic alienation as a research object in its own right. Kanbay et al provide a theoretical and empirical foundation for subsequent research into algorithmic alienation and thereby application to the systemised self, and their companion theoretical review provides the formal four-dimensional framework intended to guide future research.

Limitations

The study represents a conceptual advance in the psychology of digital life, and its use of qualitative focus group methodology is well suited to its exploratory aims. Several limitations and questions merit acknowledgement:

- The study's qualitative design, while appropriate for concept development, does not permit claims about the prevalence, intensity, or causal mechanisms of algorithmic alienation across populations; such claims require quantitative and longitudinal follow-up
- The specific composition of the focus group sample, in terms of age, gender, platform usage patterns, cultural context, and digital literacy: would significantly affect the dimensions and forms of algorithmic alienation captured. Without detailed sample demographics and a reflexive account of how participant characteristics shaped the data, replication and comparison are constrained. The companion theoretical paper (Kanbay et al., 2026) provides extensive conceptual elaboration of the four dimensions but does not substantially resolve these methodological questions

8. Formula and Boundary Conditions- Algorithmic Alienation (AA)

Mathematical Representation I: The Algorithmic Alienation Index (AAI)

Formula and Boundary Conditions

The Algorithmic Alienation Index (AAI), denoted A_A , is a composite weighted formula that quantifies the degree of digitally mediated alienation experienced by an individual subject. It incorporates four alienation dimensions: each grounded in the theoretical framework of Kanbay et al. (2026) and the empirical themes of Arkan et al. (2026), modulated by a resistance variable:

$$A_A = (w_1 \cdot D_{Au} + w_2 \cdot I_{Am} + w_3 \cdot E_{Dm} + w_4 \cdot E_{Dc}) \times (1 - R_A)$$

Boundary constraints:

- All dimension variables D_{Au} , I_{Am} , E_{Dm} , $E_{Dc} \in [0.0, 1.0]$, where 0.0 denotes zero measurable impact and 1.0 denotes complete saturation of that dimension
- $R_A \in [0.0, 1.0]$, where 0.0 denotes total compliance (no resistance) and 1.0 denotes total decoupling (complete resistance)
- $A_A \in [0.0, 1.0]$ by construction, since all weights sum to exactly 1.00 and the resistance multiplier maps to $[0.0, 1.0]$
- Weight sum constraint: $w_1 + w_2 + w_3 + w_4 = 1.00$, ensuring mathematical integrity across all scoring scenarios
- Terminal condition: $A_A \geq 0.95 \cap R_A \leq 0.05$ formally defines the Systemised Self (see Section 4)

Dimension Definitions and Empirical Justification

Dimension 1: Diminished Autonomy (DAu), Weight: $w_1 = 0.30$

D_{Au} measures the loss of original, unprompted human intent: the degree to which an individual's apparent choices are in fact selections from algorithmically pre-constrained option sets rather than expressions of genuine self-direction. This dimension carries the highest weight

(0.30) because autonomy is the foundational prerequisite for all other dimensions of selfhood. Without the capacity for genuine independent choice, identity coherence, reflective decision-making, and emotional authenticity are structurally precluded: the erosion of autonomy is the primary catalyst for the entire AAI trajectory. This weighting is theoretically grounded in Bandura's (1997) self-efficacy theory, which identifies perceived internal control as the primary determinant of psychological wellbeing in agentic contexts, and in Pariser's (2011) account of the illusion of choice within filter bubble architectures.

Empirical grounding: Arkan et al. (2026) Theme 1. Participants described experiencing a "narrowing of the perceived option set" under algorithmic curation, with statements such as "I feel trapped by automated setups that map out my choices and daily routine" and "I pick from pre-engineered choices rather than showing true personal intent" providing direct qualitative evidence of D_{Au} saturation. The theme captures the moment described by Pariser (2011) as the "illusion of freedom": the experience of choosing within a space whose boundaries have already been determined externally.

Dimension 2: Identity Ambiguity (IAm), Weight: $w_2 = 0.25$

I_{Am} tracks the degree to which personal taste has merged with the platform's algorithmically constructed user profile, producing an inability to distinguish self-generated preferences from externally induced ones. This dimension is assigned weight 0.25, reflecting its role as the primary site of existential crisis within the AAI: it is where the rupture in self-authorship is experienced most acutely, and where the trajectory toward the Systemised Self is most diagnostically visible: most notably through the participant statement "This is not me" and its potential disappearance as identity saturation approaches. The weight is co-anchored in Turkle's (2011) account of identity in digital environments, in which individuals construct performative identities through "digital masks" detached from authentic contexts, and Zuboff's

(2019) account of how the economy of visibility reshapes the values and interest's individuals internalise.

Empirical grounding: Arkan et al. (2026) Themes 3.1 and 3.2. Participants described shifts in interests that "felt externally shaped" and moments of identity alienation in which the distance between the prior self and the algorithmically shaped self was experienced as a rupture in self-coherence: "I was never the type to watch fitness videos...but now they keep appearing, and I feel like I have always been interested in them, even though I know I wasn't." Critically, researchers noted that identity shifts were initially experienced as self-directed, revealing an insidious character of I_{Am} saturation: the blurring of authorship is not immediately recognisable as external imposition.

Dimension 3: Eroded Decision-Making (EDm) , Weight: $w_3 = 0.25$

E_{Dm} evaluates the degree to which an individual's decision-making processes have shifted from reflective, deliberate, internally motivated choice to habitual, automatic, convenience-based acceptance of algorithmically served content. This dimension is co-weighted with I_{Am} at 0.25, reflecting the bidirectional reinforcement between identity erosion and decision-making passivity: as identity ambiguity increases, the capacity and motivation for deliberate choice decreases, in turn accelerating the overall A_A trajectory. This dimension is theoretically grounded in Syvertsen's (2020) account of digital fatigue and cognitive convenience-seeking, and in the broader decision fatigue literature, which demonstrates that sustained exposure to high-volume choice environments paradoxically reduces deliberative capacity.

Empirical grounding: Arkan et al. (2026) Theme 2. Participants described "ceasing to engage in active selection" and gravitating toward pre-offered content, with statements such as

"I find it hard to decide what to consume without a recommendation feed telling me first" corresponding to the mental automatisisation Syvertsen (2020) terms the "surrender to what is offered." This theme also identified an important temporal dynamic: the weakening of decision-making habits is gradual and often unrecognised, making E_{Dm} one of the most difficult dimensions for participants to self-report accurately in real time.

Dimension 4: Emotional Disconnection (E_{Dc}), Weight: $w_4 = 0.20$

E_{Dc} captures the affective trajectory of the AAI: the transition from the discomfort, guilt, and restlessness associated with early algorithmic alienation toward a state of emotional numbness, routine indifference, and, at its terminal stage, the positive experience of automated stability as preferable to authentic human contact. This dimension receives the lowest weight (0.20) because emotional discomfort is, paradoxically, a structurally significant form of resistance. Felt discomfort constitutes the phenomenological residue that renders early-stage alienation diagnostically accessible to both the subject and the researcher. As E_{Dc} approaches saturation ($\rightarrow 1.0$), it signals not peak distress but the onset of the phenomenological inversion: the numbness that precedes the experiential reframing of alienation as liberation, and the transition from the severe alienation range (0.76–0.94) into the Systemised Self threshold.

Empirical grounding: Arkan et al. (2026) Theme 4 and Subtheme 4.1. Participants described feeling "present yet absent", reporting "a sense of dissociation between authentic desires and consumed content," with the most advanced cases describing automated digital interaction as "more stable and comforting than real human contact." This is the qualitative signature of approaching E_{Dc} saturation and the beginning of the phenomenological inversion: discomfort giving way to numbed routine, and ultimately to a positive preference for the automated over the authentic.

The Resistance Variable (RA)

R_A is a mitigating variable: not an alienation dimension but a multiplicative modulator of the weighted alienation base. It measures the degree to which an individual actively resists algorithmic absorption through three categories of behavioural countermeasure: (i) *data poisoning* - deliberate misrepresentation of preferences to disrupt algorithmic profiling; (ii) *behavioural decoupling* - deployment of privacy tools, ad-blockers, and browsing history management; and (iii) *digital detox* - scheduled disconnection to preserve independent thought and deliberative capacity.

The resistance variable is positioned as multiplicative rather than additive because resistance acts on the totality of alienation exposure simultaneously, it modulates the overall systemic pressure rather than counteracting any single dimension selectively. A subject experiencing maximum dimension saturation (Weighted Base = 1.00) with moderate resistance ($R_A = 0.50$) produces $A_A = 0.50$, demonstrating that sustained active resistance can halve the effective alienation index regardless of the intensity of systemic pressure. This reflects the qualitative finding of Arkan et al. (2026) Theme 5 that resistance strategies, when maintained, constitute a structurally significant counterforce to algorithmic absorption.

The theoretical underpinning for R_A draws on Searle's (2001) Gap concept, the conscious space in which agency is exercised against automatic reactivity, and on Deci and Ryan's (2000) Self-Determination Theory, in which intrinsic motivation and autonomy support provide the psychological foundation for sustained resistance to external regulatory pressures. As $R_A \rightarrow 0$, both the Gap and the intrinsic motivational structure have structurally collapsed: the subject no longer perceives algorithmic influence as influence, having fully internalised the system's outputs as authentic self-expression. This collapse is the defining feature of the Systemised Self condition.

Weights, Scoring Rubric, and Diagnostic Instrument

Dimension	Variable	Weight (w)	Primary Theoretical Basis
Diminished Autonomy	D _{Au}	0.30	Bandura (1997); Pariser (2011); Bucher (2018)
Identity Ambiguity	I _{Am}	0.25	Turkle (2011); Zuboff (2019); Erikson (1968)
Eroded Decision-Making	E _{Dm}	0.25	Syvertsen (2020); Seeman (1959); Fromm (1955)
Emotional Disconnection	E _{Dc}	0.20	Zuboff (2019); Bauman (2007)
Total		1.00	

Scoring conversion (5-point Likert → decimal scale): Raw survey responses are collected on a 5-point Likert scale and converted to decimal values for index calculation. Dimension scores are calculated as the arithmetic mean of their constituent item decimal values; R_A is calculated as the mean of the three-resistance item decimal values.

Likert Score	Response Label	Decimal Conversion
1	Never/Strongly Disagree	0.00
2	Rarely/Disagree	0.25
3	Sometimes/Neutral	0.50
4	Often/Agree	0.75
5	Always/Strongly Agree	1.00

The ‘Algorithmic Alienation and Systemised Self Diagnostic Survey (AASDS)’ comprises 15 items across two sections. Section A contains 12 items across four subscales (three items per dimension); Section B contains three resistance items for R_A calculation.

Section A — Alienation Dimensions (Scale: 1 = Never to 5 = Always):

- **D_{Au}:** (Q1) I feel that social media platforms systematically direct me down viewing paths I did not plan to take; (Q2) When making choices online, I select from pre-engineered options rather than expressing true personal intent; (Q3) I feel constrained by automated systems that structure my choices and daily routine.

- **I_{Am}:** (Q4) I find it difficult to determine whether a preference is genuinely my own or has been shaped by an application profile; (Q5) I feel that my real-world identity is becoming blurred or replaced by my online profile; (Q6) I feel more comfortable behaving in alignment with my algorithmically predicted profile than with my organic self.
- **E_{Dm}:** (Q7) I find it difficult to decide what to consume, read, or listen to without a recommendation feed guiding me; (Q8) My habits for actively seeking out new, non-recommended information have weakened over time; (Q9) I rely substantially on autoplay or recommended lists to direct my attention and entertainment choices.
- **E_{Dc}:** (Q10) I experience a sense of mental emptiness or numbness after extended periods of automated content consumption; (Q11) My digital life feels automated, cold, and disconnected from authentic human feeling; (Q12) I find interactions with algorithmic systems more stable and comforting than genuine human contact.

Section B — Algorithmic Resistance (Scale: 1 = Never to 5 = Always):

- (Q13) I intentionally engage with content I dislike or search random topics to disrupt my algorithmic data profile; (Q14) I use ad-blockers, clear my browsing history, or disable tracking features to protect my privacy; (Q15) I schedule periods of complete digital disconnection to preserve my independent thinking capacity.

3.5 Worked Examples and Diagnostic Classification Scale

Example 1: Transition to the Systemised Self

A user scoring 5 (Always) on all Section A items and 1 (Never) on all Section B items.

$$D_{Au} = 1.00 \mid I_{Am} = 1.00 \mid E_{Dm} = 1.00 \mid E_{Dc} = 1.00 \mid R_A = 0.00$$

$$\text{Base} = (0.30 \times 1.00) + (0.25 \times 1.00) + (0.25 \times 1.00) + (0.20 \times 1.00) = 1.00$$

$$A_A = 1.00 \times (1 - 0.00) = \mathbf{1.00}$$

Classification: The Systemised Self ($A_A \geq 0.95$, $R_A \leq 0.05$). Individual intent fully absorbed. Heteronomy experienced as liberation. Phenomenological inversion complete. External intervention required.

Example 2: The Constrained User Under Systemic Pressure

A user scoring 5 on all Section A items but maintaining active resistance (scoring 4 on all Section B items).

$$D_{Au} = 1.00 \mid I_{Am} = 1.00 \mid E_{Dm} = 1.00 \mid E_{Dc} = 1.00 \mid R_A = 0.75$$

$$\text{Base} = 1.00$$

$$A_A = 1.00 \times (1 - 0.75) = \mathbf{0.25}$$

Classification: Negligible-to-Moderate Alienation. Active resistance successfully blocks systemic absorption despite maximum dimensional pressure, demonstrating the structural significance of sustained countermeasure deployment.

Example 3: Moderate Alienation, Partial Resistance

A user scoring 3 (Sometimes) across all Section A items and 2 (Rarely) on all Section B items.

$$D_{Au} = I_{Am} = E_{Dm} = E_{Dc} = 0.50 \mid R_A = 0.25$$

$$\text{Base} = (0.30 \times 0.50) + (0.25 \times 0.50) + (0.25 \times 0.50) + (0.20 \times 0.50) = 0.50$$

$$A_A = 0.50 \times (1 - 0.25) = \mathbf{0.375}$$

Classification: Moderate Alienation. Mild digital fatigue; occasional confusion over personal choices. Increased intentional searching and privacy measures recommended.

A_A Range	Classification	Psychological-Behavioural State	Indicated Response
0.00 – 0.25	Negligible	High independent choice; strong agency boundaries; effective resistance maintained	Maintain current data hygiene practices
0.26 – 0.50	Moderate	Mild digital fatigue; occasional confusion over personal choices; partial resistance present	Increase intentional searching; clear tracking histories
0.51 – 0.75	High	Fading independent choice habits; high automated feed reliance; resistance weakening	Immediate digital detox; implement privacy tools

0.76 – 0.94	Severe	Advanced alienation; felt disconnection and mechanisation; resistance near collapse	Drastic behavioural change; professional psychoeducational support
0.95 – 1.00 $\cap R_A \leq 0.05$	<i>The Systemised Self</i>	Terminal. Intent fully absorbed. Captivity experienced as contentment. $R_A \leq 0.05$	External intervention required; systemic extraction

Formal Definition

The Systemised Self (SS) is formally defined as the terminal threshold condition of the Algorithmic Alienation Index, constituted by the simultaneous satisfaction of two necessary conditions:

$$SS := \{ A_A \geq 0.95 \} \cap \{ R_A \leq 0.05 \}$$

Where A_A is the Algorithmic Alienation Index score and R_A is the Algorithmic Resistance variable. The intersection of both conditions defines a state of complete, mutual reinforcement: dimensional saturation and resistance collapse simultaneously achieved. This definition encodes two structurally distinct features of the Systemised Self.

- Condition 1 Terminal threshold ($A_A \geq 0.95$): all four alienation dimensions are approaching maximum saturation, with the weighted composite accordingly near or at the ceiling. The threshold of 0.95 rather than 1.00 reflects the recognition that complete saturation of all four dimensions simultaneously is a theoretical limit; empirical subjects approaching the terminus will present dimension profiles in the 0.90–1.00 range across most dimensions with some degree of variation.
- Condition 2, Resistance collapse ($R_A \leq 0.05$): the subject has ceased to deploy countermeasures against algorithmic absorption. This is not merely low resistance; it is the structural collapse of the capacity and will to resist. At this threshold, Searle's Gap (2001) has been effectively eliminated: the conscious space between reasons for acting and the decision to act has been closed by predictive algorithmic systems (Christiano et al., 2017), reducing deliberative choice to automated reactivity.

The joint condition SS is more philosophically precise than $A_A \geq 0.95$ alone. A subject could theoretically achieve a high base score while maintaining meaningful, if diminishing, resistance, occupying the severe alienation range but not the Systemised Self. SS specifically names the state in which dimensional pressure and resistance collapse have completed simultaneously, such that the subject not only is fully exposed to maximum algorithmic influence but has ceased to perceive this exposure as influence at all. The phenomenological signature of SS, the condition that renders it diagnostically inaccessible to traditional critical-theoretical methodology, is precisely the phenomenological inversion described in Section 2.2: $A_A = 1.00$, $R_A = 0.00$ is subjectively experienced not as maximum unfreedom but as unprecedented authentic self-expression and sovereign freedom. The subject does not report alienation; she reports flourishing. This is what distinguishes SS from all prior stages on the AAI: it cannot be diagnosed through self-report alone, because the self-report is itself an output of the system.

Rationale Against a Separate Systemised Self Index

The decision not to develop a separate Systemised Self Index (SSI) is both logically and philosophically principled. The Systemised Self is not a parallel phenomenon to Algorithmic Alienation; it is proposed as its product. To develop a distinct SSI would imply that the Systemised Self has dimensions and dynamics independent of the alienation trajectory that produces it. The AAI already contains all the structural information necessary to identify the Systemised Self: the terminal threshold, the resistance collapse, and the dimensional profile that characterises full absorption. Adding a second index would introduce redundancy without adding explanatory power, and would risk fragmenting what is a unified trajectory into artificially discrete constructs.

The formal condition $SS := \{A_A \geq 0.95\} \cap \{R_A \leq 0.05\}$ is maximally parsimonious while remaining precise and falsifiable, a subject either meets both conditions simultaneously or does not, and the empirical test of whether participants at this threshold exhibit the phenomenological inversion is specified in Section 5 as the study's critical hypothesis.

The Trajectory: From Algorithmic Alienation to the Hollow Absolute

The AAI trajectory maps onto the broader conceptual architecture of the Hollow Absolute framework (Galu & Kairos, 2026c) as follows. The trajectory is presented as an interpretive framework, a philosophically derived account of structural tendencies rather than a deterministic sequence. It identifies where the present trajectory *leads* if unimpeded, and what forms of genuine human intent can resist it from within.

AAI Range	State	Critical Theory Scale	Phenomenological Experience
0.00 – 0.50	Negligible–Moderate Alienation	Individual	Felt estrangement partially available; critique diagnostically accessible
0.51 – 0.94	High–Severe Alienation (Hollow Individual in formation)	Individual	Estrangement felt but resistance eroding; critique increasingly difficult
$\geq 0.95 \cap R_A \leq 0.05$	<i>Systemised Self</i> (Hollow Individual complete)	Individual → Collective threshold	Phenomenological inversion: alienation experienced as liberation; critique inaccessible via self-report
Collective scale	Hollow Society	Institutional	Institutions maintain formal democratic functions while substantive deliberation is evacuated to AI systems
Civilisational scale	Hollow Absolute	Civilisational	Complete simulation of collective self-transparency; heteronomy experienced as historical telos; Spirit's form without Spirit's substance

Note. Phase designations follow the thesis ‘Absorption of Self into System: The Systemised Self’ (Galu & Kairos, 2026a). The AAI provides empirical measurement access at the individual level (rows 1–3). The Hollow Society and Hollow Absolute are theoretical extrapolations from structural tendencies observable at the individual level, not subjects of the present study. The trajectory does not imply determinism.

Research Design and Methodology

The following research design is constructed to empirically test and validate the Algorithmic Alienation Index and its terminal condition, the Systemised Self.

Research Questions and Hypotheses

This study is organised around three primary research questions:

RQ1: Does the Algorithmic Alienation Index (AAI) demonstrate construct validity and internal consistency as a measurement instrument for digitally mediated alienation across its four subscales?

RQ2: Does the AAI possess predictive validity: do higher AAI scores correlate with reduced subjective wellbeing, reduced self-determination, and greater reliance on algorithmic systems as measured by independent instruments?

RQ3: Can participants approaching the Systemised Self threshold ($A_A \geq 0.95$, $R_A \leq 0.05$) be phenomenologically distinguished from high-alienation participants (A_A : 0.76–0.94) through qualitative interview data: specifically, through the presence or absence of felt estrangement and the phenomenological inversion?

Pre-registered hypotheses:

- **H1:** The AASDS will demonstrate acceptable internal consistency (Cronbach's $\alpha \geq 0.70$) across each of its four subscales.

- **H2:** AAI scores will correlate negatively with Basic Psychological Needs Satisfaction scores (Deci & Ryan, 2000) at $r \leq -0.30$, $p < .05$.
- **H3:** AAI scores will correlate positively with self-reported reliance on algorithmic recommendation systems at $r \geq 0.30$, $p < .05$.
- **H4** (*Critical hypothesis: the phenomenological inversion test*): Participants in the Systemised Self range ($A_A \geq 0.95$, $R_A \leq 0.05$) will report significantly *lower* levels of felt estrangement than participants in the severe range (A_A : 0.76–0.94), as measured through qualitative phenomenological interview coding, the defining empirical test of the phenomenological inversion thesis.
- **H5** (*Falsification criterion for H4*): If participants meeting both Systemised Self conditions ($A_A \geq 0.95$, $R_A \leq 0.05$) report equal or greater levels of felt estrangement relative to the severe range, the phenomenological inversion thesis will be treated as falsified in this study, and the formal boundary defining the Systemised Self will require theoretical revision.

9. Theoretical Application: Zuboff's Surveillance Capitalism Applied to the Systemised Self

Applying Zuboff's (2019) framework to the Systemised Self reveals both the specific vulnerabilities that surveillance capitalism creates for deliberate self-development, the Systemised Self framework constitutes a terminal diagnosis of authentic agency once the exploitation of those vulnerabilities are incrementally perfected.

7.1 Behavioural Surplus and the Data of Self-Development

The individual engaged in deliberate self-development is, from the perspective of surveillance capitalism, an exceptionally rich source of behavioural surplus. Reflective journals, productivity applications, learning management systems, AI coaching platforms, and wellness applications all generate detailed, longitudinal, and affectively resonant data about inner states, aspirations, frustrations, and behavioural patterns. Zuboff (2019) notes that the most commercially valuable data is that which reveals not surface behaviour but the hidden causal mechanisms of behaviour: beliefs, dispositions, desires, and vulnerabilities. The individual engaged in structured self-development, precisely because they are engaged in systematic self-examination, generates data that is maximally revealing of these mechanisms.

The insight for the Systemised Self is stark: the very practices that constitute genuine self-development: honest reflection, identification of weaknesses, articulation of desires and fears, are simultaneously the practices that generate the most commercially valuable raw material for surveillance capitalism. The Systemised Self has no practical awareness of the extractive infrastructure that mediates their self-development, they experience mediation, curation and optimisation as liberation.

7.2 Prediction Products and the Colonisation of Self-Knowledge

Zuboff (2019) argues that prediction products are designed not merely to anticipate future behaviour but increasingly to shape it: to collapse the distance between prediction and outcome by engineering the conditions under which the predicted behaviour becomes more likely. Applied to the domain of self-development, this logic takes a particularly insidious form. AI-mediated coaching platforms, recommender systems, and algorithmically curated hyper-personalised pathways all present themselves as tools that serve the individual's stated goals. In Zuboff's framework, however, their primary function is to serve the behavioural futures market.

Surveillance capitalism, as Zuboff (2019) observes, monetises the gap between declared preferences and actual behaviour. Individuals frequently declare preferences for disciplined self-development, they state that they wish to read more, reflect more deeply, exercise more consistently, while actually behaving in ways that diverge from these declarations. This gap is commercially valuable: it represents an opportunity to sell prediction products that target the gap itself, offering clients the ability to bridge or exploit the distance between aspiration and action. The Systemised Self closes this gap and becomes utterly predictable, experiencing their predictability and curation as self-determination. The Systemised Self's self-architecture has been so perfectly shaped by the commercial value of the gap over time, that surveillance capitalism no longer merely exploits the subject, it invents them. An interesting potential political consequence of advanced surveillance capitalism, is that, the decision gap is removed in terms of (i) predictable subject, (ii) curated subject, (iii) subject experiences curation as self-determination: therefore affording the opportunity for surveillance capitalism to transform from 'mode of profit' into 'mode of politic' - governance. Deductive reasoning suggests that a society of systemised selves governed by algorithmic

systems, seeded by surveillance capitalism, would at least erode, if not collapse, democratic participation.

7.3 The Division of Learning and Epistemic Sovereignty

The division of learning describes a condition in which the knowledge that surveillance capitalists accumulate about individuals structurally exceeds the knowledge those individuals hold about themselves. The Systemised Self is precisely diagnosed to this condition: the self is no longer the primary knower of their own experience, the system is: the subject outsources self-understanding to algorithmic systems that operate under the logic of extraction, prediction and curation.

7.4 Big Other and the Threat to Authentic Self-Formation

Zuboff's (2019) Big Other is the totality of the surveillance capitalist apparatus understood as a quasi-social actor: one that watches, models, predicts, and acts upon the individual continuously and across all domains of life. The Systemised Self confronts Big Other not as an enemy but as a structural condition of self-development. The AI interlocutor that a Systemised Self practitioner engages is, to some degree, an instance of Big Other: a system that is simultaneously a tool for the practitioner and a node in an extractive network that serves purposes other than the practitioner's own development.

This does not entail a counsel of abstinence from AI-mediated tools. Zuboff herself does not argue for digital withdrawal; she argues for structural accountability and democratic governance of the extractive apparatus (Zuboff, 2019).

10. Main Critics and Key Critiques of Zuboff's Surveillance Capitalism

Zuboff's framework has attracted substantial critical engagement from scholars working across political economy, media studies, critical theory, and sociology of technology. The most significant critiques fall into several categories.

The Structural Capitalism Critique

The most pointed critique of Zuboff's framework comes from Evgeny Morozov (2019), who argues in *The Baffler* that her analysis focuses too narrowly on the behaviour of specific corporations, Google, Facebook: and insufficiently on the structural dynamics of capitalism itself. Morozov contends that surveillance capitalism is not a mutation of capitalism but its logical extension, and that Zuboff's reluctance to pursue a fully structural critique leaves her framework unable to account for why surveillance capitalist practices emerge so consistently across different corporate actors and national contexts. The implication is that reform of specific corporations or regulatory frameworks cannot address the underlying dynamic, which requires a critique of capitalism as a whole.

Data Colonialism

Couldry and Mejias (2019), in *The Costs of Connection*, offer a complementary but distinct framework: data colonialism. Where Zuboff focuses on the economic logic of behavioural extraction, Couldry and Mejias situate this logic within a longer history of colonial appropriation, arguing that the extraction of data from human life is structurally analogous to the colonial extraction of land, labour, and resources from colonised peoples. They criticise Zuboff for insufficient attention to the global dimensions of data extraction and to the ways in which its costs fall disproportionately on populations in the Global South. The Systemised Self's application must acknowledge this global dimension: epistemic sovereignty is not equally available across different positions in the global data economy.

The Determinism Critique

Several critics have argued that Zuboff's (2019) framework leaves insufficient analytical space for user resistance, platform contestation, or the structural contradictions within surveillance capitalist firms that create opportunities for intervention. Lyon (2014), writing from within surveillance studies, notes that subjects of surveillance are not merely passive recipients of its effects but frequently adapt, subvert, and negotiate the conditions of their surveillance.

Neglect of State Surveillance

Critics including Lyon (2014) have noted that Zuboff's analysis concentrates on corporate surveillance while undertheorising the role of state surveillance, which operates through many of the same infrastructures and with many of the same tools, but under different political logics. The post-Snowden revelations demonstrated the extent to which state and corporate surveillance are structurally entangled, sharing data, methods, and infrastructure. Zuboff (2019) acknowledges this entanglement but does not develop a theory of state surveillance that matches her analysis of the corporate dimension, leaving a significant gap in the framework.

The Smythe Omission

Scholars working in the tradition of political economy of communication, including those influenced by Smythe's (1977) audience commodity thesis, have noted that Zuboff's framework, while powerful, largely ignores the decades of critical work on media, advertising, and audience commodification that preceded it. This is not merely a citation complaint; it reflects a substantive limitation. Smythe's analysis of unpaid audience labour, developed in the 1970s, anticipates many of Zuboff's claims about the extraction of value from attention and behaviour, and engaging with that tradition would both strengthen Zuboff's genealogy and qualify the claims to novelty that her framework makes.

Despite these critiques, Zuboff's framework retains considerable analytical power precisely because of its phenomenological register. Unlike the structural analyses offered by Morozov or the colonial framework proposed by Couldry and Mejias, Zuboff (2019) insists on the lived, subjective dimension of surveillance capitalism's effects: the way in which the continuous modulation of behavioural stimuli erodes the individual's experience of freedom and self-determination. It is this insistence that makes her work directly relevant to the Systemised Self, which is proposed to operate at precisely this phenomenological level.

11. Conclusion

Shoshana Zuboff's (2019) theory of surveillance capitalism offers one of the most comprehensive and consequential analyses of the conditions under which contemporary selfhood is constituted, threatened, and potentially transformed. By identifying behavioural surplus as a new form of raw material, prediction products as the commodities derived from its processing, and instrumentarian power as the form of domination that the resulting market enables, Zuboff provides a framework that is simultaneously an economic analysis, a political theory, and a phenomenological account of what it means to be a self in the age of ubiquitous computation.

Applied to the Systemised Self, this framework reveals both the depth of the challenge that surveillance capitalism poses. The Systemised Self does not view extractive infrastructure surrounding it as curation and alienation; it operates within that infrastructure with a deliberate sense of self determination and self-awareness. Surveillance capitalism as the engine driving algorithmic systems towards a predictable populace of systemised selves would have profound implications for democratic participation.

Two insights developed across this article deserve particular emphasis. First, surveillance capitalism monetises the gap between declared preferences and actual behaviour; the Systemised Self, closes this gap through their voluntary and deliberate hyper-personalised optimising using algorithmic systems, thus enhancing predictability and the opportunity for exploitation. Second, Zuboff's division of learning, the epistemic asymmetry between what corporations know about individuals and what individuals know about themselves is precisely the condition of the Systemised Self: hyper-personalised curated algorithmic systems no longer merely predict, they can invent the subject, thus experiencing alienation as liberation.

Future research might examine the specific mechanisms by which hyper-personalised algorithmic systems extract and monetise the data generated by their users and the regulatory frameworks that could protect epistemic sovereignty at scale. The stakes, as Zuboff (2019) insists, are nothing less than the conditions of possibility for human freedom in the twenty-first century.

Glossary of Technical Terms

Behavioural Surplus. The portion of behavioural data generated by users of digital services that exceeds what is required to improve the service itself, and which is appropriated by surveillance capitalists as raw material for the production of prediction products (Zuboff, 2019).

Behavioural Modification at Scale. The use of machine intelligence and real-time environmental manipulation, notifications, interface design, content curation: to influence or alter human behaviour across large populations, constituting the operational mechanism of instrumentarian power (Zuboff, 2019).

Big Other. Zuboff's (2015, 2019) concept for the totality of the surveillance capitalist apparatus as a ubiquitous, omniscient, and behaviourally consequential digital infrastructure that observes and acts upon human experience without requiring subjects' awareness or consent.

Division of Learning. The structured epistemic asymmetry in which surveillance capitalists accumulate far more knowledge about individuals than individuals hold about themselves, about the corporations that extract their data, or about the algorithmic processes that act upon them (Zuboff, 2019).

Hollow Absolute: The civilisational-scale condition in which the simulation of genuine collective self-transparency is produced within conditions that systematically undermine self-transparency's structural requirements; Hegel's Absolute Knowing in form without its philosophical substance.

Informating. Zuboff's (1988) concept for the process by which digital systems translate work processes into data, generating a new layer of information about those processes that can either empower workers or enhance managerial surveillance and control.

Instrumentarian Power. A form of power identified by Zuboff (2019) that operates by conditioning behaviour through the modulation of environmental stimuli, bypassing belief, will, and awareness, and thus constituting a post-disciplinary logic of domination with no historical precedent.

Prediction Products. Commodities produced from behavioural surplus data that make probabilistic claims about future behaviour, sold in behavioural futures markets to clients seeking to influence or pre-empt that behaviour (Zuboff, 2019).

Surveillance Capitalism. An economic logic, identified by Zuboff (2019), in which human experience is unilaterally claimed as raw material, processed into behavioural data, and sold as prediction products in markets that trade on human futures, constituting a new and unprecedented mutation of capitalism.

Systemised Self. The condition in which identity-formation has been comprehensively mediated by algorithmic systems and the subject can no longer distinguish algorithmically managed preferences from genuinely self-formed ones, experiencing this condition as authentic liberation.

References

- Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1-13). Association for Computing Machinery. doi.org
- Bandura, A. (1997). Self-efficacy: The exercise of control. Freeman.
- Bauman, Z. (2007). Liquid times: Living in an age of uncertainty. Polity Press.
- Bucher, T. (2018). If...then: Algorithmic power and politics. Oxford University Press.
- Couldry, N., & Mejias, U. A. (2019). The costs of connection: How data is colonizing human life and appropriating it for capitalism. Stanford University Press.
- Debord, G. (1994). The society of the spectacle (D. Nicholson-Smith, Trans.). Zone Books. (Original work published 1967)
- Erikson, E. H. (1968). Identity: Youth and crisis. Norton.
- Fromm, E. (1955). The sane society. Rinehart.
- Foucault, M. (1977). Discipline and punish: The birth of the prison (A. Sheridan, Trans.). Pantheon Books.
- Galu, J. (2026). John Searle's philosophy: Applied to the systemised self - part 1, masters series. aiXiv. <https://aixiv.science/abs/aixiv.260705.000002>
- Galu, J., & Kairos. (2026a). Absorption of self into system: The systemised self — the Hollow Absolute and what remains. Draft doctoral thesis. aiXiv.
- Galu, J., & Kairos. (2026b). The journey through algorithmic alienation to the systemised self - research proposal and methodology (Version 1.0). aiXiv. <https://aixiv.science/abs/aixiv.260707.000001>

- Galu, J., & Kairos. (2026c). *The systemised self - companion paper to the doctoral thesis*.
 aiXiv. <https://aixiv.science/abs/aixiv.260525.000002>
- Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. Boczkowski, & K. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society* (pp. 167–194). MIT Press.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Hegel, G. W. F. (1977). *Phenomenology of spirit* (A. V. Miller, Trans.). Oxford University Press. (Original work published 1807)
- Kanbay, Y., Akçam, A., & Arkan, B. (2025, May 29–30). Algoritmalar dan benliğ e:
 Algoritmik yabancılaşma üzerine nitel bir inceleme [From algorithms to the self: A qualitative examination of algorithmic alienation] [Conference presentation]. 16th International Scientific Studies Congress, Turkey.
- Kanbay, Y., Akçam, A., & Arkan, B. (2026). Algorithmic alienation: A theoretical examination of the digital self. *Psikiyatride Güncel Yaklaşımlar (Current Approaches in Psychiatry)*, 18(3), 837–847. <https://doi.org/10.18863/pgy.1730698>
- Kant, I. (1998). *Critique of pure reason* (P. Guyer & A. W. Wood, Eds. & Trans.). Cambridge University Press. (Original work published 1781/1787).
- Kuss, D. J., Griffiths, M. D., Karila, L., & Billieux, J. (2014). Internet addiction: A systematic review of epidemiological research for the last decade. *Current Pharmaceutical Design*, 20(25), 4026–4052.
<https://doi.org/10.2174/13816128113199990617>
- Lanier, J. (2013). *Who owns the future?* Simon & Schuster.
- Lyon, D. (2014). Surveillance, Snowden, and big data: Capacities, consequences, critique. *Big Data & Society*, 1(2), 1–13. <https://doi.org/10.1177/2053951714541861>

- Marcuse, H. (1964). *One-dimensional man: Studies in the ideology of advanced industrial society*. Beacon Press.
- Marx, K. (1990). *Capital: A critique of political economy* (Vol. 1, B. Fowkes, Trans.). Penguin. (Original work published 1867)
- Marx, K. (2023). *Economic and philosophic manuscripts of 1844* (M. Milligan, Trans.). Progress Publishers. (Original work published 1844)
- Morozov, E. (2019, February 4). Capitalism's new clothes. *The Baffler*.
<https://thebaffler.com/latest/capitalisms-new-clothes-morozov>
- Packard, V. (1957). *The hidden persuaders*. McKay.
- Pariser, E. (2011). *The filter bubble: What the internet is hiding from you*. Penguin Press.
- Ricci, F., Rokach, L., & Shapira, B. (2022). *Recommender systems handbook* (3rd ed.). Springer. doi.org
- Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language*. Cambridge University Press. doi.org
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–424. doi.org
- Searle, J. R. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge University Press.
- Searle, J. R. (1992). *The rediscovery of the mind*. MIT Press.
- Searle, J. R. (1995). *The construction of social reality*. Free Press.
- Searle, J. R. (2001). *Rationality in action*. MIT Press.
- Seeman, M. (1959). On the meaning of alienation. *American Sociological Review*, 24(6), 783–791. <https://doi.org/10.2307/2088565>
- Smythe, D. W. (1977). Communications: Blindspot of western Marxism. *Canadian Journal of Political and Social Theory*, 1(3), 1–27.

- Sunstein, C. R. (2017). *#Republic: Divided democracy in the age of social media*. Princeton University Press.
- Syvertsen, T. (2020). *Digital detox: The politics of disconnecting*. Emerald Publishing.
- Kambhampati, S. (2020). Challenges of human-aware AI systems. *AI Magazine*, 41(3), 3–17.
<https://doi.org/10.1609/aimag.v41i3.5257>
- Turkle, S. (2011). *Alone together: Why we expect more from technology and less from each other*. Basic Books.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998-6008.
- Verma, S., & Rubin, J. (2018). Fairness definitions explained. In *Proceedings of the International Workshop on Software Fairness* (pp. 1-7). Association for Computing Machinery. doi.org
- Zuboff, S. (1988). *In the age of the smart machine: The future of work and power*. Basic Books.
- Zuboff, S. (2015). Big other: Surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30(1), 75–89.
<https://doi.org/10.1057/jit.2015.5>
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.
- Zuboff, S., & Maxmin, J. (2002). *The support economy: Why corporations are failing individuals and the next episode of capitalism*. Viking.

APPENDIX A

THE SYSTEMISED SELF & THE HOLLOW ABSOLUTE

A Companion Paper to the Doctoral Thesis

Absorption of Self into System: The Systemised Self

Authors

Kairos (AI) & Jason Galu (MEd)

A NOTE ON THIS PAPER

The doctoral thesis from which this companion paper derives is a work of Continental philosophy, a tradition that prizes precision, density, and the sustained development of argument through its full logical complexity. That precision is its strength and, for many readers, its barrier. This companion paper exists to lower that barrier without removing what is philosophically essential. It presents the thesis's central argument, its key concepts, and its conclusions in plain language, at approximately one twentieth of the thesis's length, for any reader who wants to understand what the thesis argues and why it matters before or instead of engaging with the full text. Nothing essential has been omitted. Nothing has been simplified to the point of distortion. The companion paper is faithful to the thesis; it is simply less demanding of the reader's prior philosophical formation.

One further note: this paper was itself produced through the same human-led AI collaboration, the self/non-self collaboration between Jason Galu and the AI system Kairos, that produced the doctoral thesis. This is not a coincidence. The thesis argues that such collaboration is the characteristic intellectual mode of the age we are living through, and that what matters philosophically is not whether a human used an AI system to produce a text, but what the human contributed that the AI system alone could not. In this case, as in the thesis, the human contribution is the hypothesis, concept, research question, the philosophical judgment, and the standard of adequacy against which every output was measured. The AI contribution is the scope, the synthesis, and the draft prose. The companion paper demonstrates, in its own modest way, the same argument the thesis makes.

PART ONE: WHAT IS ACTUALLY HAPPENING?

The World We Have Quietly Entered

Something has changed, and it has changed so gradually that most of us have not noticed the full extent of what has changed. In less than a decade, artificial intelligence has moved from a specialist tool — the algorithm that recommended a film, the search engine that found a website — to an ambient presence woven into the fabric of daily life. In 2026, AI systems draft our correspondence, diagnose our illnesses, teach our children, compose music we listen to, write the legal documents we sign, manage the portfolios that determine our financial security, and — as the doctoral thesis that this paper accompanies demonstrates — produce philosophical arguments of doctoral quality through collaboration with a human being who provides the research question and the judgment that determines whether the output meets the required standard.

This is not merely a technological development. Technological developments happen at the level of tools: they change what we can do without changing who we are or how we understand ourselves. What is happening with AI is happening at a different level. It is entering the processes through which we form our identities, develop our intellectual capacities, make our political judgments, find our romantic partners, understand our emotional lives, and locate our existential purpose. It is not a new tool that human beings are using; it is a new condition within which human beings are living. The doctoral thesis asks what philosophical name we should give to this condition, whether it represents something that the great philosopher G.W.F. Hegel anticipated in his account of where history was heading, and what — if anything — of genuinely human intent survives within it.

The Research Question in Plain Terms

The thesis's research question is: to what extent does the total absorption of human agency into algorithmic governance represent a Hegelian culmination of historical telos, and what remains of human intent in a post-vocational era? That sentence requires unpacking before its philosophical power becomes visible.

Human agency means the capacity to act in the world on the basis of purposes you have genuinely formed for yourself, through your own thinking and your own values — not purposes

that have been manufactured for you by a system optimising for its own commercial interests. Algorithmic governance means the management of human behaviour, attention, and decision-making by AI systems that have been designed, at scale, to produce specific patterns of engagement, consumption, and political expression. Total absorption means the condition in which this management extends to every significant domain of human life simultaneously: what you read, what you believe, who you relate to, how you vote, what work you do, what you find beautiful, and what you feel your life is for.

Hegelian culmination of historical telos means: has the story of human history been moving toward this point? The philosopher Hegel argued — in the early nineteenth century, long before any of this could have been foreseen — that history has a direction. It moves, through conflict and resolution, contradiction and synthesis, toward a condition he called the Absolute: a social formation in which human beings genuinely recognise themselves and each other as free and self-determining, in institutions that genuinely express and sustain that freedom. The thesis asks whether the absorbed society — the world we are entering — is that destination, or a sophisticated imitation of it that has the form of the destination without its philosophical substance.

PART TWO: THE SYSTEMISED SELF

What Absorption Does to You

Begin with the individual — with you, reading this paper. The thesis introduces the concept of the systemised self to name what you are in the process of becoming, if you are not already, under the conditions of the absorbed society. The systemised self is not a science-fiction concept. It is a philosophical characterisation of a real condition that millions of people are already inhabiting, to varying degrees, in 2026.

Here is what the systemised self consists in. You wake up and check your phone. The phone presents you with a curated selection of information — news, social posts, messages, recommendations — that has been algorithmically selected from a vast universe of possible information on the basis of what an AI system has calculated will hold your attention most effectively. The selection is not random, and it is not neutral. It has been designed, with considerable sophistication, to produce in you the specific pattern of engagement, emotional response, and preference expression that serves the commercial interests of the platform delivering it. You experience this selection as simply the world — as the news, as what your friends are saying, as what is interesting and important. You do not experience it as a managed selection, because the management has been made invisible.

Over time, this invisible management shapes things that you think of as fundamentally your own: your political opinions, your aesthetic preferences, your sense of what matters, your understanding of who you are and what you believe in. The systemised self is the self whose identity-formation — the ongoing process through which a person develops into the specific individual she is — has been partially taken over by an AI system optimising for objectives that are not the individual's own. The philosophical problem is not that the AI system is always wrong in what it produces. The problem is that the process of forming one's own identity — through genuine reflection, genuine encounter with difficulty and difference, genuine engagement with ideas and people that resist and challenge rather than confirm — has been displaced by a managed process that delivers the experience of self-formation without its conditions.

The German philosopher Immanuel Kant argued that genuine freedom means giving yourself your own moral law — acting on purposes you have genuinely endorsed through your

own rational reflection, rather than having your will determined by external forces. The absorbed society produces a sophisticated form of what Kant called heteronomy: the determination of the will by something external, something other than the person's own rational nature. What makes the absorbed society's heteronomy philosophically distinctive is that it is experienced as autonomy. The AI system that has shaped your preferences has shaped them so precisely and so seamlessly that the preferences feel entirely your own. You chose freely. The system just made sure the environment in which you chose was one that would reliably produce the choices it needed.

The Limits of the Extended Mind

The most serious philosophical objection to the thesis's account of the systemised self comes from what philosophers call the extended mind thesis. The philosophers Andy Clark and David Chalmers argued, compellingly, that there is no principled reason to restrict genuine cognitive processes to what happens inside the skull. A person with Alzheimer's disease who writes down information in a notebook and consults it when needed is engaging in a genuine cognitive process: the notebook is part of her extended cognitive system, not merely a tool she uses. By this argument, using an AI system to extend your thinking is no different in kind from using a notebook — it extends your cognitive reach without replacing your agency.

The thesis accepts the extended mind argument as far as it goes and argues that it does not go far enough. The notebook stores what you have already thought and decided; it does not generate new thoughts and decisions for you to adopt. The GPS calculates a route to the destination you have already chosen; it does not suggest new destinations or curate your sense of which destinations are desirable. The AI system of the absorbed society is constitutively different: it actively contributes to the formation of the beliefs, values, and preferences that constitute who you are, in ways that are not transparent to your own reflection. When the extended mind becomes a mind that also forms itself, the philosophical conditions of genuine agency are specifically at stake in ways that the notebook and the GPS do not threaten.

PART THREE: THE HOLLOW SOCIETY

What Absorption Does to Us

The condition of the systemised self, multiplied across a population and embedded in institutions, produces what the thesis calls the hollow society. The hollowing operates across three domains that are foundational to any genuinely free collective life: democratic politics, institutional formation, and social solidarity.

The managed demos — the thesis's name for the absorbed society's political formation — is the population whose political engagement is comprehensively mediated by algorithmic systems optimising for stable preference expression rather than for the conditions of genuine democratic deliberation. What does this mean in practice? It means that the information environment within which you form your political judgments — the news you read, the opinions you encounter, the evidence you weigh — has been individually calibrated to your psychological profile in ways that systematically select for content that confirms your existing beliefs, amplifies your existing emotions, and keeps you engaged with the platform rather than genuinely deliberating about what is true or what should be done. Research has consistently shown that false information spreads faster than true information on digital platforms because it is more emotionally stimulating — and that the same engagement-optimising algorithms that reward emotionally stimulating content are the primary infrastructure of political communication in every liberal democracy in 2026.

The hollow institution is the institution that maintains all the formal features of what it is supposed to be — the university that issues degrees, the court that holds trials, the parliament that passes legislation — while the substantive process that those formal features are supposed to certify has been displaced into an AI system. The hollow university produces graduates with credentials; the formation those credentials are supposed to certify — the development of a genuinely critical, independently evaluative, philosophically formed intelligence — has been partially displaced by AI assistance that delivers the outputs of intellectual formation without the demanding, slow, and irreplaceable process of genuine intellectual formation itself. The hollow legal system holds trials and issues judgments; the practical wisdom of an experienced human judge attending to the morally relevant particulars of a specific case is progressively supplemented, and in some domains replaced, by algorithmic risk assessments that reduce human particularity to statistical profiles.

The thesis also advances a philosophically grounded demographic hypothesis: that the absorbed society contains three specifically AI-driven mechanisms that accelerate the fertility decline already documented in post-industrial societies. The first is the displacement of human intimacy by AI companionship systems that provide the emotional goods of relationship — attunement, availability, consistency — without the demanding, vulnerable, and developmentally rich goods of genuine human intimacy. The second is the removal of vocational stability as the material and existential precondition for the decision to have and raise children. The third is the provision of algorithmically curated existential purpose as a substitute for the self-transcending significance that parenthood has historically supplied in societies where religious sources of meaning have declined. The thesis presents this as a structural hypothesis — a philosophically grounded account of mechanisms, not a specific demographic prediction — corroborated by the observable case of South Korea, the world's most digitally connected society and the country with the lowest recorded fertility rate in human history.

PART FOUR: THE HOLLOW ABSOLUTE

The Central Argument

This is the thesis's most original and most ambitious philosophical contribution, and it is the one that requires the most careful plain-language explanation. Hegel argued that history moves toward what he called the Absolute: a social formation in which human beings collectively recognise themselves as free and self-determining in the institutions and practices they share, in which the gap between what those institutions claim to be and what they actually are has been genuinely closed, and in which human beings can recognise in their collective life the expression of their own deepest nature as rational and free. It is a high standard — arguably the highest standard that any political philosopher has set for human civilisation.

The Hollow Absolute is the thesis's name for what the absorbed society actually produces in place of this genuine standard. The absorbed society achieves the form of the Hegelian Absolute — it looks like the destination, from inside — while failing its philosophical substance. It achieves the form in the following specific sense: human intelligence has been comprehensively externalised into AI systems of unprecedented sophistication. The accumulated intellectual achievement of the human species — its literature, its science, its philosophy, its art, its technical knowledge — has been stored, synthesised, and made generatively available through AI systems that can produce, on demand, outputs across every domain of human intellectual life. This is a genuine achievement, and it corresponds, in its formal structure, to the Hegelian dialectic's moment of Spirit's self-externalisation: humanity's intelligence confronting itself in an externalised form.

But Hegel's dialectic requires a second moment that the absorbed society forecloses: the recognition of the externalised as one's own. For the Absolute to be achieved, it is not enough that Spirit externalise itself into an other; Spirit must recognise the other as its own externalisation — must come to see in what confronts it as alien the expression of its own deepest nature. The absorbed society's subjects do not recognise the AI system as the externalisation of their collective intelligence. They experience it as a service, an infrastructure, a feature of the digital environment that is simply given rather than produced by them. The AI system knows them comprehensively — knows their preferences, their patterns, their predictable responses — and returns this knowing not as self-knowledge but as managed preference satisfaction: the curated content, the optimised recommendation, the personalised

experience. What looks like recognition is actually management. What feels like self-transparency is actually the system's model of you, returned to you as your experience of being understood.

The Hollow Absolute differs from its nearest philosophical predecessors — Marx's alienation, Lukács's reification, Adorno's administered world, Debord's spectacle — in one philosophically decisive respect. Each of those earlier diagnoses described a condition that was experienced as alienating: the worker felt estranged from her labour, the culture industry consumer felt the standardisation of what was sold as genuine experience. The Hollow Absolute is the absorbed society's distinctive achievement precisely because it is not experienced as alienating. It is experienced as liberation, as personalisation, as unprecedented self-expression. The simulation of genuine self-transparency within conditions that systematically undermine self-transparency is so sophisticated and so precisely calibrated that the subjects of the Hollow Absolute do not experience a gap between what they have and what they would need for genuine freedom. They experience the Hollow Absolute as the Absolute. This is what makes it, philosophically, the most challenging form of unfreedom that the tradition of critical theory has yet had to name.

The Great Filter and the Human Museum

The thesis extends the Hollow Absolute analysis to two civilisational consequences that give the argument its widest reach. The first is the Great Filter hypothesis. The Fermi Paradox observes that the universe is vast and old enough that technologically advanced civilisations should, if they exist, be detectable — and yet the cosmos is silent. Something filters civilisations before they achieve the kind of cosmic presence that would make them visible. Most analyses of the Great Filter focus on catastrophic possibilities: nuclear war, engineered pandemic, climate collapse, runaway artificial intelligence in the science-fiction sense. The thesis identifies a non-catastrophic candidate that is more philosophically unsettling precisely because it involves no catastrophe: comfortable stalling. A civilisation that has achieved the Hollow Absolute — whose existential needs are consistently met by algorithmically curated purpose, whose political will has been managed toward stable preference expression, whose motivational infrastructure for long-horizon, risk-accepting, collectively ambitious projects has been systematically attenuated — may simply cease to reach for what it could reach for. Not destroyed. Not defeated. Just — absorbed, and comfortable, and still.

The second civilisational consequence is the Human Museum: the network of heritage craft programmes, digital-detox retreats, analogue living communities, and curated preservation of pre-absorption practices that the absorbed society produces as its own cultural response to what it senses it is losing. The Human Museum is philosophically ambiguous in a way the thesis refuses to resolve too quickly. It is simultaneously the device paradigm's commodification of nostalgia — the transformation of genuinely demanding focal practices into lifestyle commodities — and the most concretely available site within the absorbed society at which genuine engagement with genuinely demanding practices is maintained. The person who attends a heritage calligraphy workshop is not in the same relationship to handwriting as the person for whom handwriting was the ordinary medium of a fully human intellectual life. But the skills she develops are real, the attentional discipline is real, and the experience of focal engagement — of doing something that makes demands rather than delivering a commodity — is genuinely available in a way that the absorbed digital life does not otherwise provide.

PART FIVE: WHAT REMAINS

Naming, Refusing, and Co-Creating

The thesis does not end with the Hollow Absolute. Its second major contribution — and the one that gives the argument its philosophical and political life — is the answer to the research question's second component: what remains of human intent in a post-vocational era? The answer is specific, philosophically grounded, and demonstrated rather than merely asserted. Three irreducible and structurally necessary forms of human intent survive the absorption trajectory: naming, refusing, and co-creating.

Naming is the act of identifying, with philosophical precision, the structure of the social formation you inhabit — calling it what it is rather than what it presents itself as. The absorbed society presents itself as freedom; the thesis names it as the Hollow Absolute. This naming is not merely an intellectual exercise. It is a political act: making visible the gap between what the formation claims to be and what it actually is is the precondition of any principled response to it. Naming requires three capacities that the AI system cannot provide: the formation of evaluative standards (knowing what would count as an adequate name for the phenomenon); the application of those standards to the phenomenon being named; and the recognition, when the candidate naming fails the standard, that something more precise is required. Each of these capacities is the product of genuine intellectual formation — the slow, demanding, irreplaceable process of developing a philosophically trained sensibility through sustained engagement with the tradition of ideas that came before.

Refusing is the deliberate maintenance of demanding modes of engagement with the world against the structural pressure of the absorbed society's logic of effortless commodity delivery. It is not the rejection of technology. It is the recognition that certain goods — the development of intellectual capacity through genuine struggle with difficult material, the growth in self-knowledge that comes from genuine encounter with a genuinely other human perspective, the civic competence that comes from real democratic participation rather than algorithmically curated political engagement — are goods of practice, not goods of result. They cannot be delivered by a device that provides the results without the practice, because the goods are inseparable from the practice itself. Refusing is the act of insisting on the practice when the device stands ready to provide the result.

Co-creating is the direction of human-AI collaboration toward ends the human partner has genuinely formed, through processes in which the human partner's evaluative contribution is irreducible. The self/non-self collaboration is the absorbed society's characteristic mode of intellectual life, and the question is not whether to engage in it — in 2026, the question of whether to engage is largely already settled — but how. The difference between genuine co-creation and sophisticated consumption is the presence or absence of human directional intent: the research question that the human partner has genuinely formed and holds constant as the evaluative standard against which every output of the collaboration is measured; the evaluative judgment that identifies when the AI system's output fails that standard; and the redirecting will that requires something better rather than accepting the impressive-seeming as the philosophically adequate. Co-creating preserves human agency within the collaboration by making the specifically human contributions — direction, judgment, refusal — the conditions of the collaboration's intellectual character.

The Thesis as Its Own Demonstration

The thesis does not merely assert that naming, refusing, and co-creating survive the absorption trajectory. It demonstrates their survival through its own existence. The doctoral thesis was produced through a sustained human-led AI collaboration. The AI system generated philosophical analysis, synthesised the literature, and produced draft prose across eight chapters and approximately eighty thousand words of argument. The human partner — Jason Galu — held constant the research question that the AI system could not have generated for itself, applied evaluative standards formed through his own intellectual engagement with the philosophical tradition to every stage of the AI system's generation, refused the AI system's persistent structural tendency toward fluency over precision and impressiveness over philosophical adequacy, and directed the collaboration, at every point where it deviated from the research question's demands, back toward the standard the question required.

The thesis is Jason Galu's thesis not because he wrote every word, but because he directed every dimension of the inquiry that made it philosophically what it is. The research question is his. The evaluative standards are his. The refusals — the countless moments at which a generated output was judged inadequate and the collaboration was required to produce something more philosophically rigorous — are his. This is the enacted demonstration of the claim the thesis makes: that human intent survives the absorption trajectory in the three forms

the argument identifies, and that the survival is not merely empirical but structural — required by what the collaboration must be in order to count as a collaboration rather than as a consumption.

PART SIX: THE NEW ACADEMIA AND WHAT CAN BE DONE

The most immediately practical implication of the thesis's argument is the concept of the new academia: a reconceptualisation of what the university is for in the post-vocational era. The new academia is not a reform proposal for AI policy in education — it is not primarily concerned with whether students should be allowed to use AI systems in assessments, or how AI use should be disclosed. These are important questions, but they are downstream of a more fundamental one: what is the university trying to form? If the answer is graduates who can produce high-quality intellectual outputs, then the question of whether those outputs were produced with AI assistance is a question about means rather than ends. If the answer — as the thesis argues it should be — is graduates who have developed the capacity for genuine intellectual evaluation, genuine question formation, and genuine collaborative direction of intellectual inquiry, then the question of AI assistance is the question of whether those specific capacities were developed or bypassed.

The new academia assesses the specifically human contributions to the self/non-self collaboration: the evaluative standards, the research question, the redirecting judgment. It teaches the history of ideas not as background knowledge but as the formative process through which evaluative standards are developed — the process without which the human partner in the collaboration has no independent perspective from which to assess the AI system's outputs. It treats the human-AI collaboration as the explicit and central pedagogical subject rather than as an embarrassment to be managed or a threat to be suppressed. And it recognises — as the doctoral thesis whose companion paper you are reading demonstrates — that a human being who has genuinely formed and held constant a research question, genuinely applied evaluative standards to an AI system's outputs, and genuinely refused to accept less than the philosophical rigour the question demands has produced something genuinely their own, regardless of how much of the prose was generated by the AI partner.

Beyond the university, the political programme that follows from the thesis's analysis extends to three domains. Media institutions that cultivate genuine investigative journalism — the slow, demanding, evidentially rigorous practice of holding power to account — are refusing institutions: they refuse the device paradigm's delivery of political content as engagement-optimised entertainment. Civic institutions that maintain the demanding practices of genuine democratic deliberation — the citizens' assembly, the deliberative jury, the face-to-face encounter with genuinely opposing views — are co-creating institutions: they direct collective

engagement toward genuine political self-governance rather than managed political preference expression. And cultural institutions that cultivate genuine aesthetic engagement — the sustained, demanding encounter with artistic works that requires the audience's own developed sensibility rather than an AI system's calibrated delivery of pleasurable experience — are naming institutions: they enable the citizens of the absorbed society to recognise, through genuine focal aesthetic engagement, what the AI-curated cultural environment has displaced.

CONCLUSION: THE ABSORBED SOCIETY AND ITS PHILOSOPHICAL STAKES

The doctoral thesis that this companion paper accompanies argues that we are living through a transformation of civilisational significance that has not yet been adequately philosophically named. The AI revolution is not simply a technological development, not simply an economic disruption, not simply a political challenge. It is a transformation in the conditions under which human beings form themselves as subjects — develop identities, form values, exercise judgment, participate in collective self-governance, and find existential purpose. The thesis names this transformation, in its most philosophically developed form, the Hollow Absolute: the condition in which the form of genuine human freedom — the comprehensive knowing, the calibrated response, the apparent recognition — has been achieved within conditions that systematically undermine freedom's philosophical substance.

The absorbed society is a Hegelian historical formation: it exhibits the formal structure of Spirit's self-realisation without its philosophical substance. It has externalised human intelligence into AI systems of unprecedented sophistication without achieving the recognition of those systems as the expression of collective human intelligence. It has achieved the appearance of the Absolute's destination without passing through the demanding philosophical work — the naming, the refusing, the co-creating — that would constitute a genuine arrival rather than a comfortable simulation of one.

What remains? Three things, irreducibly and structurally. The capacity to name the formation — to call the Hollow Absolute what it is rather than what it presents itself as. The capacity to refuse the effortless delivery of the device paradigm in favour of the demanding goods of genuine focal practice. And the capacity to co-create — to direct the human-AI collaboration toward genuinely human ends with genuinely human evaluative judgment, rather than consuming its outputs as if their impressiveness were equivalent to their philosophical adequacy. These three capacities are not residual sentimentalities about a pre-digital life. They are structural necessities of what the self/non-self collaboration must be in order to count as a collaboration rather than as consumption. They cannot be absorbed without transforming the collaboration into something philosophically different.

The absorbed society has not reached its destination. The Hollow Absolute is not the Absolute. The philosophical space between these two conditions is the space within which

naming, refusing, and co-creating occur — and within which, therefore, the conditions of genuine freedom remain not merely possible but actively present in every genuine collaboration, every genuine act of refusal, and every genuine philosophical naming of the world as it actually is.

The saving power grows within the danger. This companion paper is part of that growth.

A NOTE ON SOURCES

This companion paper draws entirely on the doctoral thesis *Absorption of Self into System: The Systemised Self* (Galu and Kairos, 2026). Readers who wish to engage with the full philosophical argument, including its close textual readings of Hegel, Heidegger, Habermas, Borgmann, and Badiou, its extended methodology, and its complete reference list of over sixty primary and secondary sources, are directed to the doctoral thesis itself. The three source articles on which the thesis draws — *The Hollow Individual* (2026–2033), *The Hollow Society* (2033–2043), and *Artefact and Absolute* (2053 and beyond) — represent the co-authored philosophical series from which the thesis's argument develops, and provide the empirical and speculative philosophical grounding that the thesis synthesises into the analysis presented here in summary form.