

Trabajo de Fin de Máster
Autora: Susana Rodríguez Verdugo
Tutor: Manuel de Pinedo García

The Body as an Epistemic Condition: Reframing Proof of Personhood in Postdigital Sociotechnical Systems



Image: Frame from *Morphe* (2014), a short film commissioned by Aesop and created by Lucy McRae. Reproduced for academic and non-commercial purposes.

*To Manuel, for helping transform a technical question into a philosophical problem,
and for accompanying the inquiry with the rare generosity of following it wherever it led*

Table of Contents

Abstract.....	4
Resumen.....	4
Introduction.....	5
I. From Identity to Agency.....	10
I.1 Speculative design as conceptual approach.....	10
I.2 Personhood and the attribution of agency.....	11
I.2.1 From personal identity to agency attribution.....	11
I.2.2 Agency, embodiment, and second nature.....	12
I.3 Embodiment and individuation.....	14
I.3.1 Embodiment and minimal self.....	14
I.3.2 Individuation as process.....	16
I.4 Identity, individuation, and technical instance.....	18
I.5 Artificial systems and the limits of agency attribution.....	18
II. Reframing Proof of Personhood.....	20
II.1 From digital scarcity to presence scarcity.....	20
II.1.1 Scarcity, replication and participation.....	20
II.1.2 Proof of Personhood as a minimal ontological question.....	21
II.2 Limits of identity-based approaches.....	22
II.3 Structural decoupling between identity, individuation, and instance.....	24
II.4 Identification vs instantiation.....	25
II.5 The body as epistemic infrastructure.....	26
II.5.1 The body as constraint rather than identity.....	26
II.5.2 <i>Körper</i> , <i>Leib</i> and embodied presence.....	27
II.5.3 Embodied agency.....	28
II.5.4 The body as a minimal epistemic condition.....	29
II.6 The symbiont as speculative exploration.....	30
II.6.1 An embodied interface for participation.....	30
II.6.2 Epistemic function and limitations.....	31
III. Beyond Identification: Implications and Limits.....	32
III.1 Privacy and governance implications.....	32
III.2 Inclusion and epistemic participation.....	34
III.3 Technical limits and theoretical implications.....	36
IV. Conclusions and Future Perspectives.....	37
References.....	41

Abstract

Contemporary sociotechnical systems are increasingly shaped by digital infrastructures and by the near-costless replication of informational entities. This development weakens the traditional connection between participation and agency, since a single entity may generate multiple operational instances without corresponding to a plurality of agents. As a result, the attribution of participation becomes an epistemological problem with implications for representation, deliberation, and governance. Against this background, this dissertation examines Proof of Personhood (PoP), an approach aimed at ensuring that each unit of participation corresponds to a single agent and cannot be trivially multiplied within a system. The central claim of this work is that PoP should not be understood simply as a problem of identification or uniqueness, but as a question of how agency can be non-trivially attributed. To address this issue, the dissertation brings together insights from philosophy of technology, social epistemology, and distributed systems research. Drawing on the distinction between identity, individuation, and technical instance, it argues that the body could be considered a minimal epistemic condition for the attribution of agency. The dissertation further introduces a speculative conceptual artefact designed to explore the implications of this approach.

Keywords: Individuation, presence scarcity, speculative research, distributed systems.

Resumen

Los sistemas sociotécnicos contemporáneos se caracterizan por la integración de infraestructuras digitales y la posibilidad de replicar entidades informacionales casi sin coste. Esta transformación menoscaba la relación tradicional entre participación y agencia, ya que una misma entidad puede generar múltiples instancias operativas que no corresponden a una pluralidad de agentes. En consecuencia, la atribución de participación se convierte en un problema epistemológico que afecta a procesos de representación, deliberación y gobernanza. En este contexto, el presente trabajo examina la denominada *Proof of Personhood* (PoP), cuyo objetivo es garantizar que cada unidad de participación corresponda a un solo agente y no pueda multiplicarse trivialmente dentro de un sistema. La tesis central de esta investigación sostiene que la PoP no debe entenderse simplemente en términos de identificación o unicidad, sino como un problema de atribución no trivial de agencia. Para desarrollar esta propuesta, el trabajo articula aportaciones procedentes de la filosofía de la tecnología, la epistemología social y los estudios sobre sistemas distribuidos. A partir de la distinción entre identidad, individuación e instancia técnica, se propone considerar el cuerpo como condición epistémica mínima para la atribución de agencia y se introduce un artefacto conceptual especulativo para explorar las implicaciones de este enfoque.

Palabras clave: Individuación, escasez de presencia, investigación especulativa, sistemas distribuidos.

Introduction

Contemporary sociotechnical systems are characterised by a postdigital configuration in which digital infrastructures are fully embedded within social, economic, and political life, becoming largely indistinguishable from the environments they mediate, as described by Cramer (2015) and, in related terms, by Floridi (2016) through the concept of the infosphere. Distributed networks, artificial intelligence, and platform-based systems now form the background of action, interaction, and representation.

A defining feature of this configuration is informational replicability at negligible cost. Digital entities—accounts, agents, or signals—can be reproduced indefinitely, unlike in the physical world, where presence is constrained by embodiment, location, and temporality. As a result, multiple operational instances can be generated without a corresponding plurality of agents.

This new configuration destabilises the structure of participation under postdigital conditions. The capacity to act and exert influence is no longer tied to materially constrained units of agency, but to replicable informational structures, as discussed by Ohlhaber et al. (2022). Consequently, the distinction between presence and representation becomes uncertain: apparent plurality may in fact derive from a single source of agency.

A single actor may therefore appear as many participants, whether through multiple accounts, automated agents, or other digitally replicated instances. In such cases, the observable plurality of participation does not guarantee a plurality of underlying human agents, undermining the reliability of collective signals, enabling the manipulation of consensus, and distorting processes of representation, governance, and resource allocation.

The result is a structural asymmetry between physical and digital environments. While physical participation remains constrained by embodiment, digital participation can be multiplied at virtually no cost, weakening the traditional connection between participation and agency.

This contrast suggests that the problem may not concern participation alone, but also the conditions that constrain its multiplication. In physical environments, embodiment introduces limits of location, simultaneity, and continuity that make the proliferation of participation non-trivial. Under postdigital conditions, many of these constraints are weakened or disappear altogether. The question therefore arises whether the human body, understood as a source of such material constraints, plays a more fundamental role in the attribution of participation than is usually assumed.

Within this context, the central problem of this dissertation emerges: how participation in postdigital systems can be non-trivially attributed to distinct agents when observable participation can be replicated or coordinated at negligible cost.

This problem has been addressed in recent work under the label of Proof of Personhood (PoP), which seeks to ensure that each unit of participation corresponds to a distinct agent and cannot be trivially multiplied within a system. Attempts to achieve this objective are frequently framed in terms of identification or uniqueness. Yet neither approach fully captures the underlying problem. Identification is not sufficient because a verified identity may still generate multiple accounts, operational instances, or artificial agents. Nor is the challenge resolved simply by establishing uniqueness. Even if participation could be shown to be unique, the question would remain of how that uniqueness relates to the attribution of agency. As Ohlhaber et al. (2024) observe, systems may verify unique human participants while leaving open the possibility that control over participation becomes concentrated in the hands of “puppeteers” who coordinate multiple human-verified accounts.

These conceptual limitations are reflected in the practical trade-offs involved in the design of participation frameworks. Sociotechnical systems thus face a dilemma: identity-based solutions do not fully resolve the problem and may introduce surveillance and exclude legitimate participants, while anonymity and pseudonymity protect privacy and support more inclusive forms of participation but enable artificial amplification, coordinated influence and epistemic opacity. What is lacking is not merely a mechanism capable of establishing unicity in participation, but a way of doing so while ensuring that plurality is not merely the result of technical replication or concentrated control, without sacrificing privacy. The challenge is therefore to establish the conditions under which participation can be reliably linked to distinct agents.

The significance of this dilemma extends beyond questions of technical architecture. Systems of participation, representation, and governance presuppose that observable acts can be attributed to distinct agents in a sufficiently reliable manner. Voting, deliberation, resource allocation, and collective decision-making all depend on the assumption that the apparent plurality of participants corresponds to a plurality of underlying agents. Informational replicability and coordinated control weaken this assumption by allowing apparent participation to be multiplied or concentrated independently of the material conditions that ordinarily constrain presence. Under these circumstances, the problem is no longer simply how participation is organised, but how agency itself can be publicly attributed. The challenge is therefore not only technological or institutional, but also philosophical, transforming a technical question into a problem concerning agency, participation, and individuation.

As argued by De Filippi and Wright (2018) and Floridi (2016), the structural features of postdigital systems give rise to ethical and political tensions that cannot be adequately addressed within existing frameworks of identity and participation. As digital infrastructures increasingly mediate governance, resource allocation, and public deliberation, the conditions under which agency is recognised acquire normative significance.

At the core of this problem lies a fundamental tension. On the one hand, systems require fairness, non-manipulability, and proportionality, presupposing that each unit of participation corresponds to a distinct agent. On the other, they must preserve privacy, anonymity, and freedom of expression, particularly where identification entails risks of surveillance or repression. This tension is exacerbated by the proliferation of bots, multi-account strategies, and AI-driven agents, which enable the simulation of consensus and the artificial amplification of influence. As a result, it becomes increasingly difficult to distinguish between genuine collective support and manipulated signals, undermining the epistemic reliability of participation.

Within the literature on distributed systems, this issue is commonly framed as a Sybil attack problem. As originally formulated by Douceur (2002), a Sybil attack occurs when a single entity creates multiple apparently independent identities in order to obtain disproportionate influence within a network. Responses to this problem generally follow two broad strategies. One seeks to detect, mitigate, or increase the costs of such behaviour (Bonneau et al., 2015). The other seeks to restore accountability through identification mechanisms such as Know Your Customer (KYC) procedures and regulated intermediaries. De Filippi and Wright (2018) discuss how such mechanisms link blockchain activity to legally identifiable actors in contexts of compliance and enforcement. In the present analysis, these mechanisms can be understood as attempts to stabilise participation by connecting accounts or transactions to verifiable individuals.

However, linking participation to verifiable individuals does not by itself resolve the problem of agency attribution. The Proof of Personhood literature has therefore proposed a range of approaches to address this challenge, including pseudonymous verification gatherings linking physical persons to anonymous cryptographic identities (Borge et al., 2017); subjective social-verification approaches that derive legitimacy from distributed peer judgement rather than central authorities (Siddarth et al., 2020); reputation and trust-based frameworks built around non-transferable social credentials (Ohlhaver et al., 2022); uniqueness-based Proof of Personhood systems exemplified by the Idena protocol and critically examined by Ohlhaver et al. (2024); or hybrid personhood-verification systems combining physical-world verification with blockchain-based activity analysis (Cao et al., 2024). Together, these approaches highlight

the growing importance of Proof of Personhood while revealing significant differences in how participation, multiplicity, and agency are conceptualised.

Yet the diversity of these approaches also reveals a persistent difficulty. As Ohlhaber et al. (2024) argue, achieving *de jure* Sybil resistance by distinguishing humans from bots does not necessarily resolve the deeper problem of *de facto* Sybil resistance: distinguishing humans from humans acting like bots. Even when systems successfully verify human participants and establish uniqueness, participation may still become concentrated through collusion, coordination, or other forms of collective control. The issue therefore extends beyond verifying identity or uniqueness. It concerns the conditions under which participation can be attributed to distinct agents rather than merely to apparently distinct participants.

The tensions outlined above point to a more fundamental issue: the absence of reliable epistemic conditions under which participation can be non-trivially attributed to distinct agents, a problem that connects with Lackey's (2021) analysis of epistemic agency in collective contexts. The issue is not only that participation can be manipulated, but that the very basis for recognising participation has become unstable.

Under these conditions, the central challenge becomes epistemological: how can the presence of distinct agents be reliably attributed when participation can be replicated independently of the constraints that ordinarily restrict it? What remains missing is a criterion capable of establishing minimal, non-trivial grounds for attributing participation to distinct agents. Without such a criterion, responsibility, legitimacy, and epistemic trust become increasingly difficult to sustain.

To address this gap, this dissertation combines conceptual analysis with insights from phenomenology, philosophy of technology, social epistemology, and research on distributed systems. The argument proceeds by distinguishing identity, individuation, and technical instance, examining the role of embodiment in the attribution of agency, and exploring the implications of these distinctions through speculative research.

The significance of this problem extends beyond the technical design of participation systems. Once participation acquires social, political, or economic consequences, questions concerning the attribution of agency become inseparable from questions of legitimacy, accountability, and trust.

Accordingly, this work is situated at the intersection of social and political epistemology. The question of whether a distinct human agent can be reliably attributed is not only epistemic, but also fundamental to fairness, responsibility, and legitimacy in systems where participation has normative consequences.

Against this background, the core hypothesis of this dissertation emerges. The problem of Proof of Personhood in postdigital sociotechnical systems is not primarily one of identity or uniqueness, but of establishing a minimal, non-trivial criterion for the attribution of agency under conditions of informational replicability. Such a criterion must satisfy two requirements: it must constrain the multiplication of participation, and it must do so without introducing the surveillance, exclusion, and loss of anonymity often associated with identity-based approaches.

The hypothesis advanced here is that the human body, understood not as a bearer of identity but as a locus of material friction, can fulfil this role. By introducing constraints on simultaneity, location, and continuity, the body enables a distinction between the multiplication of instances and the presence of a non-trivially individuated unit of agency. Accordingly, the body functions as an epistemic condition: not by identifying who a subject is, but by making it possible to attribute the presence of a distinct individual on non-arbitrary, non-trivial grounds. From this perspective, embodiment functions as an epistemic infrastructure that shifts the problem from representational criteria centred on identity towards structural conditions of participation.

To address this question, the dissertation introduces a distinction between identity, individuation, and technical instance. Identity concerns the question of who a subject is across contexts and over time. Individuation refers to the process through which a unit of agency emerges and achieves relative coherence in relation to its *milieu*. A technical instance, by contrast, denotes an operational entity capable of acting within a system and, in digital environments, of being replicated indefinitely. This distinction provides the conceptual framework through which the problem of multiplicity can be analysed beyond the language of identity and identification.

To explore the practical implications of this framework, this work develops a speculative conceptual artefact inspired by the notion of wearable *symbionts* proposed by Ayala-García et al. (2026). In this dissertation, the symbiont is developed as an embodied interface for participation and conceived as a wearable printed-electronics patch that establishes a temporary relation between a body and a digital system without relying on persistent identity data. The symbiont functions as a speculative epistemic device rather than a technological solution. Its purpose is to explore the implications of grounding participation in embodied constraints rather than in external verification procedures, thereby bringing to light assumptions that often remain unexamined in contemporary approaches to Proof of Personhood.

The remainder of the dissertation is structured as follows: section I develops the conceptual and methodological framework, drawing on analytic philosophy, phenomenology, and philosophy of technology to articulate the shift from personal identity to agency attribution through the concepts of embodiment, individuation, and technical instance. Section II presents the central argument of the dissertation, reframing Proof of Personhood as an epistemological problem through the distinction between identity, individuation, and technical instance, introducing the body as a minimal epistemic condition for the non-trivial attribution of agency, and concluding with a speculative exploration of these ideas through the concept of the symbiont. Section III examines the implications and limitations of this framework for privacy, governance, inclusion, and epistemic participation, while assessing its broader significance for postdigital sociotechnical systems. Finally, section IV summarises the main findings, discusses the dissertation's contributions, and identifies limitations and directions for future research.

I. From Identity to Agency

This section establishes both the conceptual framework and the methodology that guide the dissertation. It introduces the theoretical resources required to analyse Proof of Personhood as a problem of agency attribution, drawing on philosophy of technology, social epistemology, and distributed systems research. It also outlines the role of speculative design in exploring alternative configurations of participation and agency through a conceptual artefact. Together, these elements provide the analytical tools through which the problem will be reformulated and examined in the following sections.

I.1 Speculative design as conceptual approach

This dissertation does not propose a fully developed technical solution, but adopts a speculative research perspective to explore alternative configurations of agency in postdigital sociotechnical systems. Rather than projecting hypothetical futures, it examines how different assumptions about technology, embodiment, and social organisation shape modes of participation.

Instead of treating the future as a predictable extension of existing technological trajectories, speculative forms of experimentation explore alternative possibilities (Hui, 2020). Speculative research addresses this question by interrogating existing assumptions and articulating alternative configurations. This approach aligns with traditions of speculative and critical design that treat conceptual artefacts as tools for exploring the epistemic and normative implications of emerging technologies, as proposed by Dunne and Raby (2013).

Within the framework proposed in this work, the notion of an embodied interface for Proof of Personhood is conceptualised as a wearable device that mediates between body and digital systems. Specifically, the dissertation draws on the concept of a wearable symbiont, a skin-mounted electronic patch proposed by Ayala-García et al. (2026), whose technical characteristics make it a useful speculative interface for exploring embodied participation.

Speculative design is particularly appropriate for the study of postdigital sociotechnical systems because such systems are not merely infrastructures but assemblages in which technical, social, political, and epistemic dimensions are deeply intertwined. As Cramer (2015) argues, the postdigital condition is characterised precisely by the disappearance of clear boundaries between the digital and the non-digital, making it increasingly difficult to isolate technological artefacts from the social worlds they help constitute. Similarly, Allen et al. (2020) emphasise that contemporary institutional technologies shape the conditions under which participation, coordination, and governance become possible, while De Filippi and Wright (2018) show how blockchain-based systems simultaneously encode technical procedures and normative assumptions. In this context, speculative design provides a useful methodology because it allows alternative configurations of these relations to be explored without requiring their immediate implementation. Rather than asking how existing systems can be optimised, it asks how they might be otherwise.

The epistemic value of speculative artefacts lies precisely in this capacity to render assumptions visible. The symbiont is not presented as a prediction, prototype, or engineering proposal. Instead, it functions as a conceptual device that materialises questions that remain implicit within existing Proof of Personhood frameworks. By translating abstract assumptions into a concrete, though speculative, configuration, the artefact makes it possible to examine the consequences of grounding participation in embodied constraints rather than in external verification procedures.

The symbiont is therefore employed here as a methodological instrument for examining alternative forms of participation and agency.

I.2 Personhood and the attribution of agency

I.2.1 From personal identity to agency attribution

The concept of person has traditionally been approached in analytic philosophy through the problem of personal identity. In Locke's account, personhood is grounded not in the persistence of a particular substance, whether material or immaterial, but in the continuity of consciousness. A person is understood as a thinking being capable of self-awareness and of recognising itself as the same subject across different times and places (Locke, 1690/1975).

Personal identity therefore provides the basis for linking actions, responsibility, and accountability to a specific individual over time. The central question becomes how an individual can be recognised as the same person despite changes in bodily or psychological states.

Contemporary discussions have expanded this framework in response to the growing importance of digital and informational environments. Within the philosophy of information, Floridi and Sanders (2004) shift the discussion from traditional metaphysical questions of personal identity towards the conditions under which agency can be attributed within informational environments. From this perspective, persons are increasingly understood as agents operating within the infosphere. Rather than focusing exclusively on continuity of consciousness, the discussion moves towards the conditions that enable participation and agency in complex sociotechnical systems.

This shift is particularly relevant under postdigital conditions, where many forms of social, economic, and political participation are mediated by digital infrastructures. Informational approaches offer a useful framework for understanding how agency may be exercised across distributed environments.

The increasing importance of artificial systems nevertheless raises questions about the relationship between observable capacities and the attribution of agency, a problem that will be examined in greater detail in section I.5.

These developments reveal a limitation shared by both identity-based and capacity-based approaches. Identity addresses the question of who a person is, while capacity-based accounts focus on what a person can do. Yet neither framework fully resolves the problem of agency attribution in postdigital systems. Thus, the attribution of agency remains problematic whenever participation can be replicated, delegated, or artificially generated.

For this reason, the discussion must move beyond personhood understood solely in terms of identity or capacities and examine the conditions under which the presence of a distinct individual can be attributed. The question shifts from “who is this person?” to “under what conditions can the presence of a distinct individual be attributed?”. Addressing this issue requires a closer examination of agency itself, particularly its relation to normativity, embodiment, and the space of reasons. It is to this question that the next subsection turns.

I.2.2 Agency, embodiment, and second nature

The limitations of both identity-based and capacity-based approaches suggest that agency cannot be reduced either to personal identity or to the possession of observable abilities. The

question is not only whether an entity can perform certain tasks, but whether it can participate in the normative practices through which actions are justified, criticised, and evaluated.

McDowell (1994) argues that human action cannot be adequately understood solely in terms of causal explanation. Human agents operate within what he calls the “space of reasons”, a normative domain in which beliefs, actions, and judgements are assessed not simply according to their causal origins but according to the reasons that support them. To act as an agent is therefore not merely to produce behaviour, but to participate in practices of justification, commitment, and responsibility.

For McDowell, this normative dimension is acquired through processes of enculturation and learning rather than being opposed to nature. The concept of second nature addresses this issue. Through *Bildung*, individuals acquire the capacities necessary to participate in shared normative practices, interpret reasons, and respond appropriately to social expectations. Rationality is therefore neither a purely biological mechanism nor an abstract faculty detached from experience, but a capacity developed through participation in historically and socially situated forms of life.

This is where embodiment becomes relevant. If participation in the space of reasons is acquired through situated forms of life rather than through abstract cognitive mechanisms, then agency cannot be detached from the embodied conditions under which such participation takes place. The question is no longer only how agents reason, but how they become capable of inhabiting normative practices in the first place.

Such a perspective challenges approaches that understand cognition primarily in representational terms. Anti-representationalist traditions argue that cognition should not be conceived as the manipulation of internal representations standing for an external world. Instead, cognition emerges through active engagement with the environment. This view is developed across ecological, phenomenological, and enactive traditions. Gibson (2014) and Chemero (2009) emphasise the relational character of perception and action, while Merleau-Ponty (1945) and Gallagher (2023) highlight the embodied and situated nature of cognition.

Agency is thus understood as situated, embodied, and relational rather than as the product of an internal representational system. Within ecological and enactive approaches, opportunities for action emerge through the relation between organism and environment. Affordances are understood as possibilities for action relative to the capacities of embodied agents, such that human action takes place within a field of constraints and possibilities that cannot be understood independently of embodiment.

The phenomenological tradition provides a complementary account of embodiment by emphasising the lived character of experience. Rather than treating the body as a neutral object, phenomenology understands it as the condition through which the world becomes accessible to an embodied agent. This distinction is often expressed through the contrast between *Körper*, the body as an organism available to measurement and technical intervention, and *Leib*, the living body as the locus of perception and situated action, a distinction developed within the phenomenological tradition and reflected in Merleau-Ponty's *Phenomenology of Perception* (1945).

This distinction helps explain why agency cannot be reduced to externally observable behaviour alone. Human action is always situated within a bodily perspective that provides continuity, orientation, and practical engagement with the world. Participation in normative practices presupposes the embodied conditions through which agents become responsive to reasons and sustain commitments over time.

The result is a conception of agency that cannot be reduced to identity, behavioural regularities, or informational processing alone. Agency emerges through the intersection of normativity, embodiment, and situated participation in shared practices. This conception of agency provides the basis for examining embodiment and minimal selfhood before turning to the problem of individuation.

I.3 Embodiment and individuation

I.3.1 Embodiment and minimal self

The relation between agency and embodiment raises a further question concerning the nature of the self. If agency cannot be reduced to identity, behavioural performance, or informational processing alone, it becomes necessary to examine the more basic conditions under which a subject can experience itself as an agent.

The concept of the minimal self provides a useful starting point for understanding the connection between embodiment and agency. For Zahavi (2005, 2014), every conscious experience is characterised by a pre-reflective first-personal dimension. Before individuals develop narrative identities, social roles, or explicit self-conceptions, experience is already given as someone's experience. This minimal self is not an object of reflection, nor a stable psychological substance, but the immediate first-personal character of consciousness itself. As such, it constitutes a more fundamental level of subjectivity than personal identity and provides a basis for understanding agency independently of institutional or representational forms of identification.

Gallagher (2000) develops this perspective through an explicitly embodied account of selfhood. More recent developments within embodied and enactive approaches further reinforce this view by emphasising the role of action, interaction, and situated engagement in the constitution of agency (Gallagher, 2020, 2023). Rather than treating the self as a detached cognitive entity, he argues that self-awareness emerges through bodily engagement with the world. Agency is inseparable from sensorimotor capacities, bodily orientation, and situated interaction. The self is therefore not located exclusively within the mind but emerges through the ongoing interaction between embodied subjects and their environments. Embodiment is not simply a vehicle through which agency is expressed; it is one of the conditions through which agency becomes possible.

The embodied perspective is further developed within ecological approaches to perception. Gibson (2014) introduces the concept of affordances to describe the possibilities for action that environments offer to organisms. Affordances are neither purely objective properties of the environment nor subjective projections of agents. Rather, they emerge from the relation between environmental features and the capacities of embodied organisms. The significance of affordances therefore depends upon the bodily capacities through which agents engage with their surroundings. Building on this perspective, Chemero (2009) argues that agency emerges from the dynamic interaction between organism and environment. Cognition is therefore realised through embodied engagement with ecological contexts rather than through internal representations, making agency relational rather than internally generated.

Clark's (2012) theory of extended cognition broadens this perspective by showing how cognitive processes frequently extend beyond the biological organism into tools, technologies, and environmental structures. However, even within extended cognition, embodiment remains a grounding condition. Cognitive systems may be distributed, but they remain anchored in embodied agents whose capacities enable such extensions. Similarly, Heras-Escribano and de Pinedo (2016, 2018) argue that affordances should be understood as genuinely relational structures emerging from the interaction between agents and environments, reinforcing the anti-representationalist view that action is grounded in situated engagement rather than detached internal representations.

Collectively, these approaches converge on a common conclusion. The self is not primarily a matter of identity, representation, or informational description, but emerges through embodied relations with the world. Embodiment is therefore a constitutive dimension of agency rather than a mere support for cognitive processes.

These considerations lead directly to the problem of individuation, understood as the process through which units of agency emerge and stabilise within a given *milieu*.

I.3.2 Individuation as process

The embodied conception of agency developed above can now be extended through the concept of individuation. If agency is situated, embodied, and relational, how do distinct units of agency emerge and maintain coherence over time? Addressing this question requires moving beyond identity and examining individuation as a process.

For Simondon (1958/2005), individuation should not be understood as the expression of a pre-existing identity but as an ongoing dynamic through which a relatively stable individual emerges from a metastable field of potentials. The individual is therefore not primary; individuation always remains relational and incomplete, preserving a preindividual dimension that exceeds any fixed identity.

A central implication of this framework is that individuals are never fully self-contained entities. Every individuated being retains a connection to the preindividual field from which it emerged, meaning that individuation remains an ongoing process rather than a completed state. The environment is not an external backdrop against which individuation takes place, but one of its constitutive conditions.

This perspective contrasts sharply with approaches that treat identity as a fixed attribute that can be simply recognised, represented, or verified. Whereas identity-based frameworks focus on determining whether a subject remains the same over time, Simondon's approach focuses on the processes through which coherence and consistency are produced in the first place. The question shifts from what an individual is to how individuality emerges, stabilises, and transforms.

The relevance of Simondon's framework for postdigital systems will be examined in greater detail in section II, where individuation is contrasted with identity and technical replication.

If individuation depends upon the relation between individuals and their *milieu*, it becomes necessary to examine how technical systems participate in shaping those conditions. This is precisely the question addressed by Bernard Stiegler through his reinterpretation of Simondon's theory of individuation.

Stiegler (1994/2018, 2015) extends Simondon's account by emphasising the constitutive role of technics in processes of individuation. Technical systems are not external instruments acting upon already formed individuals; they participate in the production of psychic and collective individuation by shaping memory, attention, desire, and social organisation. Digital infrastructures therefore participate in reorganising the *milieu* through which subjects become individuated.

For Stiegler, human individuation has always depended upon processes of technical exteriorisation. Memory, knowledge, and cultural practices are not contained exclusively within individual minds but are partially inscribed in technical supports that preserve and transmit them across time. Technical objects thus mediate connections between individual and collective processes of individuation.

This perspective becomes particularly important in contemporary digital systems. Whereas earlier forms of technical exteriorisation expanded the possibilities for psychic and collective individuation, processes of automation increasingly transfer cognitive and social functions to technical systems. Digital infrastructures now participate extensively in organising attention, memory, and decision-making, thereby shaping contemporary forms of subjectivity.

Building on Simondon and Stiegler, Hui (2020) argues that contemporary digital systems should be understood as particular technical *milieux* rather than neutral infrastructures. Different technological arrangements give rise to different forms of social organisation, participation, and subject formation. The concept of techno-diversity highlights the possibility that alternative technical configurations may support different modes of individuation. From this perspective, digital environments are not merely contexts within which agency operates, but active forces shaping how agency emerges and stabilises.

Rather than assuming a single and universal trajectory of technological development, Hui argues that different technical systems reflect different cosmological and cultural assumptions. Such an approach reinforces the idea that digital infrastructures should be analysed not simply in terms of functionality or efficiency, but in terms of the forms of agency, participation, and collective life that they enable or constrain.

The combined insights of Simondon, Stiegler, and Hui suggest that individuation should be understood as a dynamic and relational process rather than as the manifestation of a fixed identity. Individuals emerge and maintain coherence through ongoing relations with a *milieu*, while technical systems actively participate in shaping the conditions under which such processes unfold. This distinction will prove crucial for the subsequent analysis, where identity, individuation, and technical instance are examined as distinct but interconnected dimensions of participation in postdigital sociotechnical systems.

I.4 Identity, individuation, and technical instance

Building on the processual account of individuation outlined in the previous subsection, it becomes possible to distinguish between three dimensions that are frequently conflated in

discussions of digital participation: identity, individuation, and technical instance. Each operates according to a different logic in the relation between agents and systems.

Identity concerns the conditions under which a subject can be uniquely identified as a specific individual over time, answering the question “who is this person?” and grounding most identification systems, as formulated by Locke (1690/1975).

Individuation, in the sense developed by Simondon (1958/2005), refers to the process through which a unit of agency emerges and maintains relative coherence in relation to its *milieu*, without presupposing a fixed or fully determined identity.

A technical instance, by contrast, denotes a distinct operational entity—such as an account, profile, or software process—capable of acting within a system and, in digital environments, of being replicated indefinitely. This category becomes particularly salient in the context of the Sybil attack (Douceur, 2002) and subsequent discussions of Proof of Personhood (Ohlhaber et al., 2022).

In postdigital systems, these levels are no longer aligned. A single identity may correspond to multiple instances, while instances may operate without corresponding to distinct processes of individuation. As a result, the correspondence traditionally assumed between identity, individuation, and participation can no longer be taken for granted.

The implications of this distinction for Proof of Personhood will be examined in section II. Before turning to that discussion, however, a further question must be addressed. While the distinction between identity, individuation, and technical instance clarifies the structure of the problem, it does not by itself explain why contemporary artificial systems are capable of generating forms of operational multiplicity that challenge traditional assumptions about agency. The next subsection therefore examines the epistemic limits of artificial systems and the extent to which functional performance can be associated with genuinely individuated agents.

I.5 Artificial systems and the limits of agency attribution

Artificial systems further intensify the problem by enabling a single actor to scale their presence through multiple technical instances, thereby generating forms of operational multiplicity that challenge traditional assumptions about agency. This development raises an important question for the present dissertation: under what conditions can the behaviour of an artificial entity justify attributing participation to a distinct agent? The issue is whether

operational performance alone is enough to justify the attribution of agency in a non-trivial sense.

Recent work in machine learning suggests that caution is warranted. Schölkopf et al. (2021) argue that contemporary AI systems excel at identifying statistical regularities in data, yet often lack access to the causal structures that generate those regularities. A system may produce highly effective outputs while remaining dependent upon correlations that fail to capture the underlying organisation of the environment in which it operates.

This limitation becomes particularly relevant when observable behaviour is used as a proxy for agency. However, behavioural patterns do not necessarily provide a stable basis for identity, as they can be replicated, imitated, or generated by artificial systems. The problem is analogous to that faced by autonomous poker-playing programs, whose behavioural signatures can be designed to resemble those of human players despite the absence of a corresponding human agent (Noble, 2002). Observable conduct alone therefore cannot establish the presence of a distinct agent.

This concern is reinforced by research on shortcut learning. Geirhos et al. (2020) show that machine learning systems frequently solve problems by exploiting superficial statistical cues rather than by learning the underlying structure of a task. Models may therefore appear competent while relying on accidental or context-dependent correlations. Functional success can mask a significant gap between performance and understanding.

Related concerns emerge in studies of distribution shift. Koh et al. (2021) demonstrate that machine learning systems often perform reliably only within the environments in which they were trained, while changes in the statistical distribution of data across contexts can significantly reduce their performance. Such systems are therefore highly dependent upon the statistical characteristics of particular environments and frequently lack the robustness associated with agents capable of adapting across changing conditions.

Taken together, these findings have significant epistemic implications. Postdigital systems offer no reliable means of distinguishing between genuine plurality and technical multiplication. The problem is not merely that artificial systems can act, but that increasingly sophisticated forms of action can be generated without providing clear criteria for determining whether the presence of a distinct individual can be attributed.

These limitations have led recent discussions to emphasise the importance of causal grounding. Veloz (2025) argues that the attribution of agency cannot be reduced to successful optimisation within a task environment. What matters is whether behaviour is grounded in meaningful and stable causal relations with the environment. From this perspective, the

question is not only whether a system can act effectively, but whether its activity reflects the emergence of a distinct individual. Causal grounding therefore points beyond performance metrics and towards the broader conditions under which agency can be attributed.

Viewed as a whole, these limitations do not imply that artificial systems are incapable of sophisticated behaviour. Rather, they suggest that operational performance should not be conflated with individuation. Optimisation can generate effective action within a particular environment without thereby establishing the conditions under which a distinct individual emerges, maintains coherence, and persists over time. The analysis therefore suggests that behavioural performance alone provides insufficient grounds for attributing agency.

II. Reframing Proof of Personhood

Building on the conceptual and methodological foundations developed in the previous section, this part presents the central analytical contribution of the dissertation. It develops a reinterpretation of Proof of Personhood as an epistemological problem concerning the non-trivial attribution of agency under conditions of informational replicability. Drawing on the distinction between identity, individuation, and technical instance, it examines the limitations of existing approaches and introduces embodiment as a minimal epistemic condition for the attribution of participation to a distinct individual. The section concludes with a speculative exploration of these ideas through the concept of the symbiont.

II.1 From digital scarcity to presence scarcity

II.1.1 Scarcity, replication and participation

A useful comparison can be drawn with the introduction of scarcity in digital monetary systems. In cryptocurrencies, the problem of double spending arises because digital tokens can be copied without loss. Bitcoin addresses this issue by introducing a verifiable cost—Proof of Work or PoW—making the creation of new units non-trivial, as proposed by Nakamoto (2008) and analysed by Bonneau et al. (2015). Scarcity is thus engineered through constraints linking digital operations to physical expenditure.

A similar issue arises in the domain of participation. In postdigital systems, the problem extends from the duplication of value to the duplication of presence. Without constraints, participation can be multiplied indefinitely, leading to a form of “presence inflation”, where apparent plurality does not reflect a plurality of agents.

To address this problem, Proof of Personhood seeks to ensure that each instance of participation requires a non-trivial constraint, preventing the multiplication of agency while preserving genuine plurality.

This parallel is particularly significant because both cryptocurrency systems and Proof of Personhood respond to challenges generated by digital replicability. In both cases, scarcity is not a natural property of the digital environment but a sociotechnical achievement produced through the introduction of non-trivial constraints.

This insight has broader implications for the organisation of postdigital systems. As Allen et al. (2020) argue, many digital ecosystems rely upon rule-based mechanisms that enable coordination without requiring centralised authority. Blockchain systems achieve trust not through identification, but through verifiable constraints embedded in sociotechnical protocols. Similarly, De Filippi, Mannan, and Reijers (2020) emphasise that decentralised governance depends upon mechanisms capable of establishing reliable conditions for participation while preserving openness and autonomy.

From this perspective, Proof of Personhood can be understood as an attempt to extend the logic of scarcity beyond economic value and into the domain of participation. The focus is therefore on preserving the integrity of participation under conditions of digital replicability. Presence, like value, must become resistant to arbitrary replication if collective signals are to retain epistemic and political significance.

II.1.2 Proof of Personhood as a minimal ontological question

This understanding of presence scarcity also illuminates the logic of Proof of Personhood. Its aim is not to identify participants, but to ensure that each instance of participation is grounded in conditions that make the multiplication of agency non-trivial, thereby addressing the epistemological problem of plurality without incurring the ethical and political costs of identity-based systems.

This shift reflects a deeper transformation in the way participation is conceptualised. Rather than asking who a participant is, Proof of Personhood asks whether participation can be attributed to a distinct individual within a given system. For this reason, it can be interpreted as a minimal ontological question concerned not with personal identity, but with the presence of a distinct participant.

This perspective is consistent with the work of Ohlhaver et al. (2022, 2024), who argue that the primary objective of Proof of Personhood is to establish uniqueness rather than identity. The problem is not to determine the personal characteristics of participants but to ensure that one participant cannot trivially appear as multiple participants.

In philosophical terms, this orientation resonates with Zahavi's (2005) account of minimal selfhood, understood as a basic form of first-personal presence that precedes narrative identity, social roles, and institutional recognition.

The Sybil attack provides the structural formulation of this problem. The attack succeeds because digital environments allow the multiplication of operational instances at negligible cost. What is threatened is not identity as such, but the assumption that each instance of participation corresponds to a distinct agent.

From this perspective, the significance of the Sybil problem extends beyond computer security. It reveals a structural instability within postdigital systems: the possibility that operational multiplicity can become detached from ontological plurality, as apparent participation no longer guarantees the presence of multiple agents.

Proof of Personhood therefore emerges as an attempt to restore a non-trivial relation between participation and agency by establishing the minimal grounds for recognising the presence of a distinct participant.

II.2 Limits of identity-based approaches

A common response to the problems discussed above is to restore the link between participation and agency through identification mechanisms such as Know Your Customer (KYC) procedures and regulated intermediaries. As discussed by De Filippi and Wright (2018), such mechanisms connect digital activity to legally identifiable actors through processes of verification and accountability. Within this framework, the attribution of agency is presumed to follow from the verification of identity.

However, such approaches are insufficient. First, identity does not prevent the multiplication of operational instances. A single identified individual may control multiple accounts or deploy several artificial agents, maintaining disproportionate influence within a system, a dynamic that Ohlhaber et al. (2024) analyse in their discussion of digital personhood. Identification secures correspondence between an account and an individual, but it does not constrain multiplicity.

On the other hand, the failure to constrain multiplicity becomes particularly visible in systems that seek to preserve anonymity while maintaining meaningful participation. Within such systems, anonymity plays an ambivalent role: it protects legitimate political expression, especially in high-risk contexts, yet also enables manipulation and the fabrication of artificial majorities. Pseudonymity offers partial mitigation by allowing persistent participation without

full identification, but it does not prevent a single agent from generating multiple identities or deploying several artificial agents.

From this perspective, the Sybil problem reveals a structural weakness in identity-based approaches. As Douceur (2002) highlighted, the central challenge is not simply the absence of identity, but the possibility that one entity can appear as many.

The limitations of identification become even more apparent when its ethical and political implications are considered. Identity-based systems typically rely on centralised infrastructures that collect, store, and manage personal information. Such systems introduce risks of surveillance, data breaches, and institutional or corporate control, while potentially excluding participants who require anonymity for safety or political reasons, as analysed by Zuboff (2019) and discussed by Sunstein (2017).

According to Zuboff (2019), contemporary digital infrastructures increasingly operate through the extraction and accumulation of behavioural data, transforming participation into a source of information subject to economic and political exploitation. From this perspective, identity verification does not merely solve a technical problem; it also reconfigures the relationship between participants and institutions by expanding the capacity for monitoring and control.

These concerns become particularly acute where anonymity performs an important political function. As Sunstein (2017) observes, anonymity can protect freedom of expression and dissent, while De Filippi and Hassan (2016) emphasise that decentralised systems often emerge in response to concerns about concentration of power. The challenge is therefore to secure inclusion without undermining the conditions that make open participation possible.

More fundamentally, identity-based systems frequently displace rather than resolve the problem. By focusing on administrative verification, they leave intact the structural conditions that enable the replication of participation. What matters, therefore, is not whether identities can be verified, but whether participation can be prevented from being trivially multiplied.

This difficulty becomes even more apparent when participation is delegated to artificial systems. Contemporary digital environments increasingly involve software agents, automated decision-making systems, and algorithmic infrastructures. Floridi and Sanders (2004) and Varian (2019) both highlight how technical systems can perform actions, generate effects, and mediate social activity independently of the characteristics traditionally associated with personhood. Identifying the individual associated with a system therefore does not necessarily clarify how agency is distributed, mediated, or amplified through technical infrastructures.

The issue therefore extends beyond identity itself. Participation may be delegated, distributed, automated, or scaled through technical systems while remaining formally linked to a single verified individual.

For these reasons, identity-based mechanisms cannot adequately address Proof of Personhood. What is required is not stronger identification, but a mechanism capable of constraining the multiplication of participation without relying on identity disclosure. This shift requires a more fundamental diagnosis of the structural transformations introduced by postdigital systems, a line of inquiry developed in the following subsection.

II.3 Structural decoupling between identity, individuation, and instance

The limitations of identity-based approaches point towards a deeper transformation affecting contemporary sociotechnical systems. The deeper issue is not simply that identity verification fails to prevent manipulation, but that the relation between participation and agency has itself become unstable.

In materially constrained environments, coordination among agents remains grounded in the plurality of bodies. Even when strategic or deceptive, coordination remains constrained by embodiment, location, and temporality, dimensions central to phenomenological accounts of embodied experience (Merleau-Ponty, 1945). By contrast, postdigital systems allow a single source of agency to generate potentially unbounded operational instances—accounts, bots, or AI agents—producing not coordination, but the simulation of a vast plurality, a dynamic central to the analysis of digital personhood by Ohlaver et al. (2022).

This transformation can be understood as a decoupling between individuation and operation, as outlined in subsection I.3.2. Operational instances may be replicated indefinitely without requiring the emergence of new individuated agents.

This decoupling is reinforced by the postdigital condition described by Cramer (2015), in which digital infrastructures are no longer experienced as separate technological domains but as pervasive environments embedded within everyday life. Under such conditions, the distinction between representation and presence becomes increasingly difficult to maintain. Systems register activity, interaction, and participation, yet they frequently lack reliable criteria for determining whether these activities originate from distinct agents or from the multiplication of operational instances.

Contemporary artificial systems further intensify this dynamic by enabling the replication of operational capacities without corresponding processes of individuation, as argued in subsection I.5.

The result is a structural asymmetry between operational presence and ontological grounding. Systems are capable of registering participation, counting instances, and coordinating interactions, yet they possess no intrinsic means of determining whether apparent plurality corresponds to a plurality of agents. The central challenge of postdigital systems therefore lies not in identifying participants more accurately, but in determining whether apparent multiplicity corresponds to a plurality of individuated agents. This structural diagnosis prepares the transition from identification to instantiation developed in the next subsection.

II.4 Identification vs instantiation

The diagnosis developed above suggests that the challenge posed by digital replicability cannot be resolved through stronger forms of identification alone. The question becomes how participation can be attributed to a distinct individual under conditions in which apparent plurality no longer guarantees the presence of multiple individuated agents.

This distinction reflects the difference between two underlying assumptions. Identification systems assume that the problem of ensuring that participation corresponds to distinct individuals can be solved through increasingly reliable representations of those individuals. Instantiation-based systems, by contrast, focus on how participation comes to be realised within a system rather than on the identification of participants. From this perspective, participation should be grounded in conditions that make the multiplication of agency non-trivial, regardless of whether the identity of participants is known.

Within this framework, the relevant question is no longer whether an individual can be accurately recognised, but whether participation corresponds to a distinct agent. What matters is not the persistence of a representation, but the existence of conditions that prevent one agent from appearing as many.

This distinction helps clarify why identity-based solutions remain insufficient even when technically effective. A perfectly reliable identification system may establish correspondence between an individual and an account, yet still fail to prevent the multiplication of participation through networks of automated agents, technical infrastructures, or other forms of mediated action. Identification secures representation; it does not necessarily secure individuation.

Instantiation-based approaches seek to address this limitation by introducing constraints on the conditions under which participation occurs. Rather than relying on the disclosure of

personal information, they aim to ensure that each instance of participation is grounded in conditions that cannot be trivially replicated.

The distinction between identification and instantiation therefore marks a broader epistemic shift. Traditional systems rely on representational criteria, treating participation as something that can be secured through increasingly reliable representations of individuals. Instantiation-based approaches move towards structural constraints, seeking to establish conditions that make participation resistant to arbitrary replication.

This shift from representational to structural criteria provides the conceptual bridge to the central contribution of this dissertation. If the problem is no longer one of identification but of instantiation, the question becomes what kinds of constraints can support the non-trivial attribution of participation to distinct agents in systems characterised by informational replicability. The following subsection develops the hypothesis that embodiment may provide such a condition, allowing the body to be understood not as an identity marker, but as an epistemic infrastructure for participation.

II.5 The body as epistemic infrastructure

II.5.1 The body as constraint rather than identity

The preceding subsections have argued that the central challenge of Proof of Personhood is not the identification of individuals but the non-trivial attribution of participation. If identity-based approaches fail because they do not constrain the multiplication of operational instances, a different criterion is required. The proposal developed in this dissertation is that embodiment may provide such a criterion.

The central claim of this dissertation is that the human body can function as a minimal epistemic infrastructure capable of introducing constraints on the multiplication of presence, drawing on Merleau-Ponty's account of embodiment (Merleau-Ponty, 1945) and Gallagher's analysis of the embodied self (Gallagher, 2000). This does not involve treating the body as an identity marker or a source of biometric data—an approach present in several existing frameworks that rely on the body for identification—but rather as a material condition that limits simultaneous replication.

In digital environments, where replication is near-costless, such constraints are absent. Introducing the body as a condition of participation reintroduces material friction, making the multiplication of participation non-trivial. This does not identify the agent, but anchors each instance of participation in a condition that cannot be indefinitely replicated.

The significance of this proposal lies in the distinction between embodiment as a source of information and embodiment as a source of constraint. Existing identification systems frequently rely on the body as a repository of data, treating biometric characteristics as markers through which subjects can be recognised and tracked. The proposal developed here departs from this logic. Rather than functioning as a source of information, the body is understood as a limit on the proliferation of instances.

II.5.2 *Körper*, *Leib* and embodied presence

The proposal to use embodiment as a condition of participation immediately raises an important objection. If the body is introduced as a constraint within digital systems, does this not reduce it to a purely mechanical substrate, stripping it of its phenomenological and normative significance?

This concern can be addressed by distinguishing between two dimensions of embodiment. On the one hand, the body can function as a physical substrate and can therefore be understood in terms of its objectifiable properties—such as biometric signals, physiological parameters, or spatial localisation.

At this level, the body corresponds to what phenomenological traditions describe as *Körper*: the body considered as an object among objects, measurable, describable, and available to technical intervention. By contrast, *Leib* refers to the body as lived and experienced from the first-person perspective. Most contemporary identification systems operate primarily at the level of *Körper*, treating bodily characteristics as data that can be collected, stored, and compared.

This reliance on *Körper* raises an important issue. Once bodily characteristics are transformed into data, they become detachable from the lived presence of the individual. Biometric systems operate through representations of the body that can be stored, processed, and circulated independently of the person to whom they originally belonged. As a result, they raise concerns not only regarding privacy, surveillance, and control, but also regarding the relation between bodily data and participation itself. The existence of a bodily representation does not necessarily establish the presence of the individual whose body is represented. For this reason, the proposal developed here does not appeal to the body primarily as a source of identifying information, but as a form of embodied presence that cannot be reduced to its informational representation.

In this sense, the body functions as a condition of non-trivial presence. Its epistemic relevance lies not in the information it provides, but in the constraints it imposes on replication: a body cannot be simultaneously present in multiple places nor sustain unlimited instances of action.

This distinction allows a response to the objection of reduction. The body is not reduced to a mechanism, but reconfigured in its role: from a semantic function—as bearer of identity—to an ontological one—as a condition that makes the distinction between one and many non-trivial. The approach proposed here seeks to avoid the reduction by reintroducing, in a minimal sense, the *Leib* level—the living body—not as a basis for identity, but as a condition of situated and first-personal presence, in line with Zahavi's (2005) account of minimal selfhood and Gallagher's (2000) account of embodied selfhood.

At first sight, this may appear paradoxical. The dimension of embodiment that most closely relates to subjectivity—the *Leib*—is not invoked here to establish personal identity. Rather, the distinction highlights that the problem addressed by Proof of Personhood concerns agency rather than subjectivity. The aim is not to determine who a subject is or to provide access to first-person experience, but to establish the minimal grounds for recognising the presence of an individual. In this respect, the constraints that limit replication belong to the body as *Körper*, while the appeal to *Leib* serves to prevent those constraints from being understood merely as properties of an objectifiable substrate detached from embodied existence. Their significance for participation derives from the fact that they are constraints of embodied existence rather than properties of abstract informational instances.

Thus, the distinction between *Körper* and *Leib* is rearticulated: the body is engaged as an objectifiable substrate at the level of *Körper*, while its epistemic significance derives from its status as *Leib*. This avoids both the reduction of the body to data and the need to ground participation in a comprehensive phenomenological theory of subjectivity, highlighting embodiment instead as a source of constraints that ground non-trivial agency.

II.5.3 Embodied agency

The argument developed thus far has focused on embodiment as a source of constraint. However, embodiment is not only relevant because it limits replication. It is also constitutive of agency itself. If participation is to be attributed to distinct agents, it becomes necessary to clarify the relation between embodiment and action.

As discussed in section I, embodied and ecological approaches reject the view that agency can be understood solely in terms of internal representations or informational processing. Instead, agency emerges through the situated engagement of organisms with their environments. Perception, action, and cognition are mutually constitutive dimensions of embodied activity, a view developed across phenomenological, ecological, and enactive approaches.

This perspective becomes particularly important in the context of postdigital systems. The multiplication of operational instances may reproduce behavioural outputs, decision procedures, or patterns of interaction, but it does not necessarily reproduce the embodied conditions through which agency emerges. Functional performance alone therefore cannot establish the existence of a distinct agent.

McDowell's (1994) concept of the space of reasons provides a useful way of understanding this limitation. Human agency is not exhausted by causal responsiveness to environmental stimuli. Actions are situated within normative practices in which agents are capable of giving, evaluating, and responding to reasons. Participation therefore presupposes more than behavioural competence; it requires a form of situated engagement within normative contexts.

From this perspective, embodiment is not merely a practical constraint on participation but one of the conditions that make agency possible in the first place. The role of the body therefore extends beyond its capacity to limit replication. Embodiment provides one of the foundations for the situated, relational, and normative dimensions of agency that cannot be reduced to informational descriptions or behavioural outputs alone.

The distinction developed in this dissertation does not imply that embodiment guarantees agency, nor that every embodied entity necessarily constitutes an individuated participant. Rather, it suggests that agency cannot be adequately understood independently of the embodied conditions through which action becomes situated within a shared world of embodied and first-personal presence. This insight supports treating embodiment not as a marker of identity but as a condition relevant to the attribution of participation.

II.5.4 The body as a minimal epistemic condition

The preceding analysis allows the central claim of this dissertation to be formulated more precisely. The proposal is not that the body provides proof of identity, authenticity, or personhood in any strong metaphysical sense.

Instead, the body is understood as a minimal epistemic condition for the non-trivial attribution of agency under conditions of informational replicability. The significance of embodiment lies not in the information it provides but in the constraints it imposes. Bodies occupy specific locations, unfold within finite temporal horizons, and remain subject to material limitations that restrict the simultaneous multiplication of presence.

Under postdigital conditions, where operational instances can be replicated at negligible cost, such constraints acquire renewed importance. Informational systems can reproduce accounts, profiles, software agents, and behavioural outputs indefinitely. The argument therefore does

not rest upon any essentialist conception of the body. The body is not treated as the source of a unique identity, nor as a guarantee of authenticity. Its role is structural rather than semantic. What matters is not who a participant is, but whether participation remains connected to conditions that make the distinction between one and many non-arbitrary.

This shift has important implications for the design of participation systems. If embodiment functions as a condition of participation rather than as a source of identity, there is no need to rely on the persistent storage of biometric information or on mechanisms that enable long-term identification and traceability. The relevant conditions need not be continuously recorded or accumulated over time, but may instead be linked to transient forms of activation associated with particular acts of participation. What matters is not the accumulation of information about participants, but the existence of constraints that make participation non-trivially attributable.

From this perspective, embodiment provides a way of reframing Proof of Personhood. Rather than grounding participation in identity, representation, or behavioural verification, it becomes possible to ground participation in constraints that support the non-trivial attribution of agency. The body therefore functions not as a credential or repository of information, but as a condition through which participation can be distinguished from arbitrary replication.

The value of this proposal lies in the shift it introduces from representational to structural criteria. The question is no longer who a participant is, but on what grounds the presence of a distinct individual can be attributed. It is this shift that ultimately motivates the speculative exploration developed in the following part of this section, where the concept of the symbiont is introduced as a conceptual artefact for exploring how embodied constraints might be articulated within postdigital sociotechnical systems.

II.6 The symbiont as speculative exploration

II.6.1 An embodied interface for participation

The argument developed throughout this dissertation has sought to reframe Proof of Personhood as an epistemological problem concerned with the non-trivial attribution of agency under conditions of informational replicability. If identity-based approaches prove insufficient and embodiment can be understood as a source of constraints rather than identification, a further question arises: how might such constraints be articulated within postdigital sociotechnical systems?

To examine this question, this dissertation presents the symbiont as a speculative conceptual artefact. The symbiont is not introduced as a fully developed technological proposal nor as a

prediction regarding future systems. Rather, it functions as an exploratory device that enables the implications of the theoretical framework developed thus far to be examined. In this respect, the proposal follows the tradition of speculative design described by Dunne and Raby (2013), in which conceptual artefacts are used to interrogate assumptions embedded within existing technological paradigms and to investigate alternative possibilities.

The symbiont can be understood, in basic technical terms, as a thin wearable patch attached to the body and designed to establish temporary connections with digital infrastructures. Recent developments in flexible and printed electronics suggest the technical plausibility of such interfaces, including skin-mounted systems capable of sensing, signalling, and harvesting energy from surrounding environments (Ayala-García et al., 2026). Participation is instantiated through an event rather than through the disclosure of personal information, linking it to embodied conditions without relying on persistent identity.

In this respect, the symbiont differs significantly from conventional identification technologies. It is not conceived as a biometric system, nor as a device for storing personal information. Rather than generating persistent representations of participants, it operates through temporary activations that establish the conditions under which participation can occur. This orientation aligns with recent developments in privacy-preserving approaches to digital credentials and personhood verification (Adler et al., 2024), while extending these approaches by grounding them in embodied constraints.

The value of the symbiont therefore lies not in the device itself, but in the conceptual shift it articulates. Participation becomes linked to constraints rather than identities, to instantiation rather than representation, and to embodied conditions rather than informational descriptions.

As such, the symbiont provides a way of exploring how embodied constraints might support the non-trivial attribution of agency without relying on identity-based verification.

II.6.2 Epistemic function and limitations

The epistemic role of the symbiont can now be clarified. Its function is not to establish identity or authenticate a person. Rather, it serves as a conceptual artefact for exploring how participation might be grounded in conditions that make the multiplication of agency non-trivial. At the same time, the proposal raises important objections that must be acknowledged.

From a privacy perspective, any use of the body as a basis for participation risks being interpreted as surveillance or biometric control, given the history of systems that extract and store bodily data. This framework addresses the concern by distinguishing between persistent data capture and ephemeral activation: it explicitly rejects the storage of biometric

information, relying instead on transient signals that do not enable identification or long-term traceability.

Concerns of exclusion arise in relation to access and bodily diversity. Systems requiring specific devices or interfaces may exclude individuals lacking resources or whose conditions do not fit design assumptions. This highlights the need for contextual deployment and makes clear that the proposal is not a universal solution, but a response to specific problems arising from the multiplication of agency.

These limitations reveal the ethical and political tensions that accompany any attempt to introduce embodied constraints into systems of participation. The objective is therefore not to advance a universally applicable technological solution, but to investigate the consequences of treating embodiment as an epistemic condition within postdigital environments.

The symbiont should therefore be understood as a speculative extension of the framework developed in this dissertation. Its contribution lies less in technical implementation than in bringing to light assumptions that remain implicit within contemporary Proof of Personhood approaches.

III. Beyond Identification: Implications and Limits

This section examines the implications and limits of the framework proposed in the dissertation. Rather than introducing new concepts, it evaluates the consequences of understanding Proof of Personhood in terms of agency attribution and embodied constraints. Particular attention is given to questions of privacy, governance, inclusion, and epistemic participation, as well as to the technical and theoretical limitations of the proposal. The aim is to assess the broader significance of the framework and to situate its contribution within ongoing debates concerning postdigital sociotechnical systems.

III.1 Privacy and governance implications

Building on the distinction developed in subsection II.4, participation may be attributed without requiring the continuous disclosure or storage of identity information. This shift has significant implications for privacy and governance.

From a privacy perspective, the proposal offers a potential alternative to models that depend on the accumulation of personal data. As Zuboff (2019) argues, contemporary digital infrastructures increasingly operate through the extraction, storage, and analysis of

behavioural information, transforming participation into a source of economic and political value. Identity-based systems frequently reinforce this tendency by requiring persistent records that enable the tracking and verification of individuals over time. By contrast, grounding participation in instantiation rather than identification minimises the amount of information required for participation. This orientation is consistent with privacy-preserving personhood credentials, which aim to attest the presence of a distinct human participant without exposing personal identity or generating extensive behavioural records (Adler et al., 2024).

The introduction of embodied constraints also affects how trust is established. Rather than eliminating the need for trust, it shifts the basis upon which trust is grounded. O'Neill (2018) argues that trust should be grounded not in assumptions about actors but in conditions that make trustworthiness assessable. From this perspective, the value of embodied constraints lies not in proving who a participant is, but in providing verifiable grounds for recognising the presence of an individual. The emphasis shifts from identity verification to the design of systems capable of supporting trustworthy participation. Similar concerns appear in discussions of blockchain governance, where technical infrastructures seek to replace reliance on personal trust with verifiable procedural conditions (De Filippi, Mannan, & Reijers, 2020).

At the same time, the proposal should not be interpreted as a complete solution to questions of governance. As De Filippi and Hassan (2016) observe, technical systems increasingly perform regulatory functions by embedding rules and constraints directly within infrastructures. Yet governance cannot be reduced to technical design alone. Decisions concerning participation, access, legitimacy, and accountability remain normative and political questions. De Filippi and Mannan (2025) similarly argue that the legitimacy of emerging technological systems depends not only on technical effectiveness but also on their alignment with public values and broader social expectations.

These considerations suggest that the framework developed here should be understood neither as a replacement for governance nor as a purely technical solution to problems of participation. Rather, it offers a different way of conceptualising the relationship between participation, privacy, and trust under postdigital conditions. By shifting attention from persistent identity to the grounds on which participation can be attributed to distinct agents, the proposal opens new possibilities for designing systems that seek to balance legitimacy, privacy, and coordination without relying exclusively on identification.

III.2 Inclusion and epistemic participation

The framework developed in this dissertation seeks to address the problem of multiplicity without relying on persistent identity. However, introducing embodied constraints into participatory systems raises important questions concerning inclusion and epistemic participation. If engagement becomes linked to embodied presence, it is necessary to consider who may be excluded, disadvantaged, or rendered less visible within such systems.

One potential concern is that any mechanism designed to constrain participation may inadvertently reproduce existing social and technological asymmetries. Access to participation is never distributed evenly. Differences in material resources, technological infrastructure, physical ability, and social context can affect the capacity of individuals to engage with a given system. While the proposal advanced here does not advocate a specific implementation, the introduction of embodied conditions necessarily raises questions regarding accessibility and accommodation.

On the other hand, the framework developed here seeks to avoid some of the exclusionary effects associated with existing approaches. Identity-based systems often depend on access to official documentation, civil registries, or forms of institutional recognition that are not equally available to all participants. Biometric approaches may reduce reliance on documents, but they also raise concerns regarding privacy and surveillance and can create barriers for individuals whose bodily characteristics cannot be easily captured, verified, or accommodated within standardised identification procedures. Other approaches rely on accumulated reputation, social graphs, or community verification, making participation dependent on prior integration into a network and potentially disadvantaging newcomers or individuals occupying more marginal positions. Participation mechanisms linked to economic resources may introduce further asymmetries by granting greater influence to those who control larger amounts of capital.

By contrast, the proposal developed here does not ground participation in identity verification, institutional membership, social reputation, or economic stake. The relevant criterion is instead linked to embodied presence and the constraints associated with it. Although any practical implementation would still need to address important questions of accessibility and accommodation, the criterion itself appeals to a condition that is universal rather than contingent upon nationality, citizenship, social status, institutional affiliation, or prior digital history. In this respect, the framework seeks to introduce non-trivial conditions for participation while minimising dependence on forms of recognition that may systematically exclude certain groups.

For this reason, inclusion is not treated here as an external limitation imposed on Proof of Personhood systems, but as one of the considerations that motivates the search for alternatives to identity-based, reputation-based, or capital-based approaches. The objective is not simply to constrain multiplicity, but to do so in a manner that preserves the possibility of broad and diverse participation in shared epistemic and political processes.

These concerns can be understood more broadly as questions of epistemic participation. As Lackey (2021) argues, the production and circulation of knowledge frequently depend on collective processes involving multiple agents and institutions. Participation is therefore not merely a matter of access to a platform or system, but also relates to the conditions under which individuals can contribute to shared epistemic practices. This observation is particularly relevant for Proof of Personhood systems, since any mechanism that regulates access simultaneously shapes who can contribute to collective epistemic processes and how those contributions are recognised. From this perspective, the challenge is not simply to prevent the multiplication of agency, but to do so without undermining the plurality of perspectives necessary for collective deliberation and knowledge production.

Related concerns emerge in contemporary discussions of informational environments. Nguyen (2020) and Sunstein (2017) highlight the ways in which participation can be shaped by the structure of the informational spaces in which individuals operate. Digital systems do not merely facilitate interaction; they influence visibility, attention, and the distribution of influence. Moreover, both authors show that distortions in informational environments may affect not only what participants know, but also which voices become salient, credible, or influential within a collective process. Consequently, a system that successfully constrains multiplicity may still reproduce forms of epistemic inequality if participation remains unevenly distributed or if certain groups face greater barriers to engagement than others.

These concerns become particularly relevant in the context of artificial amplification. The proliferation of automated accounts and coordinated networks can distort public discourse by creating the appearance of support, consensus, or opposition that does not correspond to an underlying plurality of distinct agents (Ferrara et al., 2016). The framework proposed in this dissertation seeks to address this problem by introducing non-trivial conditions for participation. Yet reducing artificial amplification is not sufficient on its own. A desirable participation framework must also preserve openness, diversity, and the possibility of legitimate collective action.

The challenge, therefore, is to balance two competing demands. On the one hand, participation must remain resistant to trivial replication and artificial inflation. On the other, it must remain sufficiently inclusive to accommodate diverse forms of agency and engagement.

This tension cannot be reduced to a simple technical solution. Rather, it highlights the need to evaluate Proof of Personhood systems not only in terms of security or Sybil resistance, but also in terms of their broader epistemic and social consequences. Accordingly, inclusion should be understood not as an external constraint on system design, but as one of the requirements that participation frameworks must take into account.

III.3 Technical limits and theoretical implications

The proposal developed in this dissertation should not be understood as a technical solution to the problem of Proof of Personhood, but as a conceptual approach for rethinking the conditions under which agency can be attributed in postdigital systems. Consequently, its limitations are not only technical but also conceptual. Recognising these limitations is essential for clarifying both the scope and the contribution of the argument.

One important limitation concerns implementation. The symbiont introduced in section II is not presented as a deployable technology but as a speculative artefact designed to explore the implications of embodied constraints. Although recent developments in wearable and flexible electronics suggest that interfaces of this kind are technically plausible (Ayala-García et al., 2026), many practical questions remain unresolved. Issues of reliability, accessibility, interoperability, security, and large-scale deployment fall outside the scope of this dissertation. The significance of the symbiont therefore derives less from its immediate feasibility than from its capacity to open a space for reflecting on alternative conditions of participation. Rather than providing a technical answer, it functions as a conceptual intervention that allows alternative configurations of participation and agency to be explored. In this respect, the proposal is closer to the speculative design tradition described by Dunne and Raby (2013) than to an engineering blueprint.

A related limitation emerges from the relation between embodiment and agency. Throughout this dissertation, embodiment has been presented as a condition capable of introducing non-trivial constraints on participation. However, while such constraints may support the non-trivial attribution of agency, they do not by themselves establish a full account of agency, intentionality, or responsibility. The proposal advanced in this dissertation therefore identifies minimal conditions under which agency may be attributed, but it does not eliminate the need for further normative and epistemological accounts of agency.

Additional challenges arise from increasingly sophisticated artificial systems. As discussed in subsection I.5, operational performance alone does not provide sufficient grounds for attributing agency. The analysis developed here proposes embodied constraints as a basis for supporting the non-trivial attribution of agency in postdigital systems.

From this perspective, current artificial systems do not exhibit the kind of embodiment discussed in this dissertation. As Veloz (2025) argues, significant differences remain between biological forms of cognition and contemporary machine-learning systems, particularly regarding the causal organisation through which agents relate to their environments. Nevertheless, whether alternative forms of artificial individuation may become possible in the future remains a separate question.

Overall, these limitations indicate that the contribution of this work should be understood primarily as an epistemological clarification rather than a complete solution. The analysis identifies a set of distinctions and conditions that may help reframe Proof of Personhood, but it leaves unresolved many of the technical, normative, and ontological questions that continue to shape debates concerning agency in postdigital sociotechnical systems.

Building on the distinction developed in subsection I.4, future discussions of Proof of Personhood may benefit from focusing on individuation rather than identity. Instead of treating participation primarily as a problem of identification, the argument developed here shifts attention towards the grounds for recognising the presence of distinct individuals. While this perspective does not yield a complete Proof of Personhood architecture, it suggests that further progress may benefit from moving beyond the language of identity and towards a broader investigation of agency, individuation, and participation in postdigital sociotechnical systems.

IV. Conclusions and Future Perspectives

This dissertation argues that Proof of Personhood should be understood not primarily as a problem of identity verification or uniqueness, but as a problem of agency attribution under conditions of informational replicability. By distinguishing identity, individuation, and technical instance, the analysis shows how postdigital environments destabilise traditional correspondences between participation and agency. In response, the dissertation proposes embodiment as a minimal epistemic condition for the non-trivial attribution of agency and explores the implications of this proposal through the speculative figure of the symbiont. The analysis then extends to questions of privacy, governance, inclusion, and epistemic participation, while also considering the conceptual and practical limitations of the proposal.

Beyond its reinterpretation of Proof of Personhood as a problem of agency attribution, the dissertation also brings into dialogue bodies of literature that are rarely examined together within discussions of Proof of Personhood. While existing work is primarily situated within

cryptography, distributed systems, and digital governance, this dissertation draws on philosophy of technology, phenomenology, social epistemology, and distributed systems research to bring questions of embodiment, individuation, technical mediation, and agency attribution into dialogue with debates on Sybil resistance, blockchain governance, and decentralised participation.

The analysis nevertheless remains subject to important conceptual and practical limitations. Acknowledging them helps clarify both the scope of the argument and directions for future research.

A first limitation concerns implementation. The proposal developed in this dissertation is conceptual rather than technical, and the symbiont introduced in section II is presented as a speculative artefact rather than a deployable system. Although developments in wearable technologies and human-computer interfaces suggest that some of its components may become feasible, significant questions concerning security, governance, scalability, and maintenance remain open. The proposal should therefore be understood not as an engineering solution to Proof of Personhood, but as a conceptual intervention for clarifying the problem. Future research should investigate how embodied constraints might be translated into concrete sociotechnical architectures while preserving the privacy, inclusiveness, and non-trivial participation conditions that motivate the framework.

At a deeper conceptual level, the relationship between embodiment and agency remains insufficiently understood. Throughout this dissertation, embodiment has been proposed as a minimal epistemic condition capable of introducing non-trivial constraints on participation. While such constraints may support agency attribution, they do not by themselves provide a complete account of agency, intentionality, responsibility, or subjectivity. The proposal therefore addresses a specific aspect of the problem—the trivial replication of participation—rather than a comprehensive theory of agency. Future research should therefore examine more closely how embodiment relates to processes of individuation, agency, and normative responsibility, and whether these relations can provide a more comprehensive account of participation in postdigital sociotechnical systems.

The status of artificial agency presents a further limitation. While this dissertation argues that operational performance should not be equated with individuation, whether future artificial systems might develop forms of agency comparable to those associated with embodied human agents remains uncertain. The extent to which agency attribution, individuation, and embodiment may apply to increasingly sophisticated artificial systems therefore remains an open question.

Substantive issues also arise regarding privacy and governance. Although the framework seeks to avoid persistent identity systems and minimise data retention, introducing embodied conditions into participation raises questions concerning surveillance, control, and institutional power. Recent developments in privacy-preserving personhood credentials illustrate the possibility of moving beyond conventional identity systems, yet questions of governance and legitimacy remain unresolved. How embodied constraints might be integrated with privacy-preserving infrastructures and legitimate governance mechanisms therefore remains an important question for further investigation.

Questions of inclusion also require further research. While the framework emphasises non-trivial conditions for agency attribution, access to participation remains unevenly distributed across technological, bodily, economic, and social conditions. Because participation is shaped by epistemic environments, an important challenge is determining how systems grounded in embodied constraints might accommodate diverse forms of participation without reproducing existing inequalities.

A broader limitation concerns the sociotechnical implications of the argument. If Proof of Personhood is understood as a problem of agency attribution rather than as a question of identification or participant verification, its significance extends beyond personhood protocols to issues of governance, collective decision-making, digital institutions, decentralised coordination, and human-AI interaction. Such questions cannot be fully addressed within the scope of the present dissertation and open new avenues of research concerning participation in postdigital societies. Exploring them may contribute both to the development of Proof of Personhood systems and to a wider understanding of participation under postdigital conditions.

Collectively, these limitations suggest that the framework proposed here should be understood as an initial contribution rather than a definitive solution. Its contribution lies less in the viability of its specific proposals than in the shift of perspective it enables. If the problem is not simply determining who participates, but understanding how participation remains connected to distinct individuals, then Proof of Personhood becomes part of a wider inquiry into the conditions under which human presence and agency can be recognised and sustained alongside other human and artificial agents within increasingly complex sociotechnical environments.

Ultimately, the fundamental question is no longer “who are you?” but rather “is there a distinct individual here?” The answer to that question cannot be settled through identity or uniqueness alone. It is precisely this difficulty that reveals Proof of Personhood as an epistemological problem for the postdigital age.

References

- Adler, S., Hitzig, Z., Jain, S., Brewer, C., Chang, W., DiResta, R., Lazzarin, E., McGregor, S., Seltzer, W., Siddarth, D., Soliman, N., South, T., Spellicy, C., Sporny, M., Vasudevan, V., Bailey, J., Critch, A., Falcon, R., Flanagan, H., ... Zick, T. (2024). *Personhood credentials: Artificial intelligence and the value of privacy-preserving tools to distinguish who is real online*. arXiv preprint. <https://arxiv.org/abs/2408.07892>
- Allen, D. W. E., Berg, C., Markey-Towler, B., Novak, M., & Potts, J. (2020). Blockchain and the evolution of institutional technologies: Implications for innovation policy. *Research Policy*, 49(1), 103865. <https://doi.org/10.1016/j.respol.2019.103865>
- Ayala-Garcia, C., Merino Sanchez-Fayos, I., Petti, L., & Cohen, N. (2026). Sustainable wearable symbionts. *The Design Journal*, 29(1), 229-246. <https://doi.org/10.1080/14606925.2025.2590074>
- Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., & Felten, E. W. (2015). SoK: Research perspectives and challenges for Bitcoin and cryptocurrencies. In *2015 IEEE Symposium on Security and Privacy (SP)* (pp. 104-121). IEEE. <https://doi.org/10.1109/SP.2015.14>
- Borge, M., Kokoris-Kogias, E., Jovanovic, P., Gasser, L., Gailly, N., & Ford, B. (2017). Proof-of-personhood: Redemocratizing permissionless cryptocurrencies. In *2017 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)* (pp. 23-26). IEEE. <https://doi.org/10.1109/EuroSPW.2017.46>
- Cao, Y., Cao, J., Liu, H., Cui, Z., & Wen, L. (2024). Eden: An edge computing empowered proof-of-personhood protocol for anti-Sybil in blockchain-based metaverse. In *2024 2nd International Conference on Intelligent Metaverse Technologies & Applications (iMETA)* (pp. 182-189). IEEE. <https://doi.org/10.1109/iMETA62882.2024.10808133>
- Chemero, A. (2009). *Radical embodied cognitive science*. MIT Press. <https://doi.org/10.7551/mitpress/8367.001.0001>
- Clark, A. (2012). Embodied, embedded, and extended cognition. In K. Frankish & W. M. Ramsey (Eds.), *The Cambridge handbook of cognitive science* (pp. 275-292). Cambridge University Press.
- Cramer, F. (2015). What is 'post-digital'? In D. M. Berry & M. Dieter (Eds.), *Postdigital aesthetics: Art, computation and design* (pp. 12-26). Palgrave Macmillan. https://doi.org/10.1057/9781137437204_2

De Filippi, P., & Hassan, S. (2016). Blockchain technology as a regulatory technology: From code is law to law is code. *First Monday*, 21(12). <https://doi.org/10.5210/fm.v21i12.7113>

De Filippi, P., & Mannan, M. (2025). Regulatory equivalence in blockchain systems: The role of public values and legitimacy. In J. Goossens, E. Keymolen, & A. Stanojević (Eds.), *Public governance and emerging technologies*. Springer. https://doi.org/10.1007/978-3-031-84748-6_8

De Filippi, P., Mannan, M., & Reijers, W. (2020). Blockchain as a confidence machine: The problem of trust & challenges of governance. *Technology in Society*, 62, 101284. <https://doi.org/10.1016/j.techsoc.2020.101284>

De Filippi, P., & Wright, A. (2018). *Blockchain and the Law: The Rule of Code*. Harvard University Press. <https://doi.org/10.2307/j.ctv2867sp>

Douceur, J. R. (2002). The Sybil attack. In P. Druschel, F. Kaashoek, & A. Rowstron (Eds.), *Peer-to-Peer Systems (Lecture Notes in Computer Science, Vol. 2429)*, pp. 251–260. Springer. https://doi.org/10.1007/3-540-45748-8_24

Dunne, A., & Raby, F. (2013). *Speculative everything: design, fiction, and social dreaming*. MIT Press.

Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The rise of social bots. *Communications of the ACM*, 59(7), 96–104. <https://doi.org/10.1145/2818717>

Floridi, L., & Sanders, J. W. (2004). On the morality of artificial agents. *Minds and Machines*, 14(3), 349–379. <https://doi.org/10.1023/B:MIND.0000035461.63578.9d>

Floridi, L. (2016). *The fourth revolution: How the infosphere is reshaping human reality*. Oxford University Press.

Gallagher, S. (2000). Philosophical conceptions of the self: implications for cognitive science. *Trends in Cognitive Sciences*, 4(1), 14–21. [https://doi.org/10.1016/S1364-6613\(99\)01417-5](https://doi.org/10.1016/S1364-6613(99)01417-5)

Gallagher, S. (2020). Action and interaction. In *The Oxford handbook of 4E cognition*. Oxford University Press. <https://doi.org/10.1093/oso/9780198846345.001.0001>

Gallagher, S. (2023). *Embodied and enactive approaches to cognition*. Cambridge University Press. <https://doi.org/10.1017/9781009209793>

Geirhos, R., Jacobsen, J.-H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., & Wichmann, F. A. (2020). Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11), 665–673. <https://doi.org/10.1038/s42256-020-00257-z>

Gibson, J. J. (2014). *The ecological approach to visual perception: Classic edition*. Psychology Press. <https://doi.org/10.4324/9781315740218>

Heras-Escribano, M., & de Pinedo, M. (2016). Are affordances normative? *Phenomenology and the Cognitive Sciences*, 15(4), 565–589. <https://doi.org/10.1007/s11097-015-9440-0>

Heras-Escribano, M., & de Pinedo-García, M. (2018). Affordances and landscapes: Overcoming the nature–culture dichotomy through niche construction theory. *Frontiers in Psychology*, 8, Article 2294. <https://doi.org/10.3389/fpsyg.2017.02294>

Hui, Y. (2020). Machine and ecology. *Angelaki*, 25(4), 54–66. <https://doi.org/10.1080/0969725X.2020.1790835>

Koh, P. W., Sagawa, S., Marklund, H., Xie, S. M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R. L., Beery, S., Leskovec, J., Kundaje, A., Pierson, E., Levine, S., Finn, C., & Liang, P. (2021). WILDS: A benchmark of in-the-wild distribution shifts. In *Proceedings of the 38th International Conference on Machine Learning (ICML)* (pp. 5637–5664). PMLR. <https://doi.org/10.48550/arXiv.2012.07421>

Lackey, J. (2021). *The epistemology of groups*. Oxford University Press. <https://doi.org/10.1093/oso/9780199656608.001.0001>

Locke, J. (1690/1975). *An essay concerning human understanding* (P. H. Nidditch, Ed.). Clarendon Press.

McDowell, J. (1994). *Mind and world*. Harvard University Press.

Merleau-Ponty, M. (1945). *Phénoménologie de la perception*. Gallimard.

Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system*. <https://doi.org/10.2139/ssrn.3440802>

Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme*, 17(2), 141–161. <https://doi.org/10.1017/epi.2018.32>

Noble, J. (2002). Finding robust Texas Hold'em poker strategies using Pareto coevolution and deterministic crowding. In M. A. Wani (Ed.), *Proceedings of the 2002 International Conference on Machine Learning and Applications (ICMLA'02)* (pp. 233–239). CSREA Press.

Ohlhaber, P., Weyl, E. G., & Buterin, V. (2022, May 10). *Decentralized society: Finding Web3's soul* [Working paper]. SSRN. <https://doi.org/10.2139/ssrn.4105763>

Ohlhaber, P., Nikulin, M., & Berman, P. (2024, March 6). *Compressed to 0: The silent strings of Proof of Personhood* [Working paper]. SSRN. <https://doi.org/10.2139/ssrn.4749892>

- O'Neill, O. (2018). Linking trust to trustworthiness. *International Journal of Philosophical Studies*, 26(2), 293–300. <https://doi.org/10.1080/09672559.2018.1454637>
- Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward causal representation learning. *Proceedings of the IEEE*, 109(5), 612–634. <https://doi.org/10.1109/JPROC.2021.3058954>
- Siddarth, D., Ivliev, S., Siri, S., & Berman, P. (2020). Who watches the watchmen? A review of subjective approaches for Sybil-resistance in Proof of Personhood protocols. *Frontiers in Blockchain*, 3, 590171. <https://doi.org/10.3389/fbloc.2020.590171>
- Simondon, G. (1958/2005). *L'individuation à la lumière des notions de forme et d'information*. Éditions Jérôme Millon.
- Stiegler, B. (1994/2018). *La technique et le temps 1: La faute d'Épiméthée*. Fayard.
- Stiegler, B. (2015). *La société automatique 1: L'avenir du travail*. Fayard.
- Sunstein, C. R. (2017). *#Republic: Divided democracy in the age of social media*. Princeton University Press.
- Varian, H. R. (2019). Artificial intelligence, economics, and industrial organization. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 399–419). University of Chicago Press. <https://doi.org/10.7208/chicago/9780226613475.001.0001>
- Veloz, T. (2025). Toward aitiopoietic cognition: Bridging the evolutionary divide between biological and machine-learned causal systems. *Frontiers in Cognition*, 4, 1618381. <https://doi.org/10.3389/fcogn.2025.1618381>
- Zahavi, D. (2005). *Subjectivity and selfhood: Investigating the first-person perspective*. MIT Press. <https://doi.org/10.7551/mitpress/6541.001.0001>
- Zahavi, D. (2014). *Self and other: Exploring subjectivity, empathy, and shame*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199590681.001.0001>
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.