



ARTIFICIAL INTELLIGENCE AND THE LIMITS OF KNOWING: REASSESSING A.J. AYER'S VERIFICATION PRINCIPLE IN A MACHINE- INTELLIGENT AGE

Chiedu Ifeakor Ph.D and Nkwor, Chukwuka Simon

Nnamdi Azikiwe University, Akwa, Anambra State

Abstract

In an age increasingly defined by machine intelligence, the question of what it means to know has acquired a new philosophical urgency. This work probes that urgency by re-examining A.J. Ayer's verification principle as a critical framework for evaluating the epistemic claims of artificial intelligence. Ayer's famous assertion that a proposition is meaningful only if analytically true or empirically verifiable, establishes a demanding threshold, one that appears to clash with the probabilistic, non-transparent, and model-driven procedures that power contemporary AI systems. Yet it is precisely this clash that illuminates the shifting boundaries of epistemic legitimacy. The study argues that AI does not overturn the classical verificationist project; rather, it exposes its hidden resilience. As machine-learning architectures generate outputs that are often persuasive without being strictly verifiable, they force a re-evaluation of long-standing assumptions about meaning, justification, and empirical warrant. By examining the opacity of algorithmic inference, the quasi-linguistic nature of AI-generated propositions, and the emerging gap between pragmatic usefulness and epistemic grounding, the article reveals how AI unsettles but ultimately enriches the discourse on verifiability. Against the optimism that credits machines with quasi-cognitive authority, this paper contends that the verification principle remains an incisive tool for distinguishing between what appears to be knowledge and what can defensibly claim that status. At the end, AI does not expand the limits of knowing; it redefines the terrain on which those limits must be critically negotiated, renewing the urgency of Ayer's insights in a machine-intelligent age.

Keywords: *Artificial Intelligence, Epistemology, Logical Positivism, Meaning, Verification Principle.*

Introduction

The question of what qualifies as knowledge has regained an unexpected urgency in a world increasingly shaped by computational reasoning. While earlier debates in epistemology focused on human cognition, perception, and linguistic competence, contemporary discourse must also account for systems that generate assertions without consciousness, intention, or experiential access to the world. This shift has reopened fundamental inquiries about meaning, justification, and the grounds on which truth-claims stand. It is no longer sufficient to describe how such systems operate; one must ask whether the statements they produce meet the standards long associated with meaningful discourse. The re-emergence of these concerns marks a moment of philosophical significance.

A.J. Ayer's verification principle enters this conversation with a sharpness that belies its historical distance. In *Language, Truth and Logic*, Ayer contends that a proposition is meaningful only if it is analytic or capable of empirical confirmation through observation (Ayer 1936, 16). This principle, rooted in the logic of the Vienna Circle, sought to purge philosophical argument of speculative assertions lacking testable content. Its enduring power lies in its insistence that meaning requires a demonstrable connection between language and experience. Although critics have challenged its strictness, the principle remains a compelling benchmark for evaluating the status of claims produced in contemporary contexts.

The demand for verifiability takes on new salience when applied to systems whose procedures unfold behind layers of mathematical abstraction. These systems produce outputs that can be



influential, even authoritative, yet their justificatory pathways often resist inspection. Here lies the philosophical tension that this article seeks to explore: the appearance of meaningfulness without a corresponding route to verification. Without a clear criterion by which such claims may be tested, the status of these assertions becomes epistemically fragile. Ayer's principle thus provides a lens through which this fragility may be scrutinized more precisely and without sentimentality.

This tension also highlights a broader issue concerning the nature of justification. Philosophers from Carnap to Quine demonstrated how linguistic rules and empirical checks form the twin pillars of epistemic stability in analytic inquiry (Carnap 1950, 67; Quine 1951, 39). Assertions lacking grounding in either pillar risk becoming mere verbal performances rather than genuine claims about the world. When new systems generate statements whose internal workings remain concealed, they may achieve practical success but fail to fulfil the deeper requirements of epistemic legitimacy. Philosophical evaluation, therefore, must distinguish between operational adequacy and genuine meaningfulness.

Scholars in recent years have examined the difficulty of assessing claims when the processes generating them cannot be fully articulated. Humphreys argues that certain computational procedures render their own justification inaccessible, leading to "epistemic opacity" of an unprecedented kind (Humphreys 2009, 617). Burrell similarly notes that opacity can arise from mathematical complexity, proprietary restrictions, or conceptual novelty (Burrell 2016, 3). These concerns reinforce the importance of examining whether meaningfulness can be preserved when the grounds for verification retreat from observable scrutiny. Ayer's framework becomes indispensable at precisely this juncture.

The present study therefore advances the position that earlier epistemological standards have not been outgrown but have become newly urgent. The confrontation between classical demands for verifiability and novel forms of inference illuminates the limits of what can properly be called knowledge. Rather than relaxing epistemic boundaries in the face of technological sophistication, philosophical inquiry must reaffirm the distinction between persuasive output and justified assertion. It is only by doing so that the integrity of the concept of knowledge can be preserved against conceptual dilution or unwarranted inflation.

Accordingly, this article proposes that Ayer's verification principle remains an essential tool for evaluating the epistemic status of contemporary claims. The principle does not merely expose the limitations of certain forms of reasoning; it clarifies the conditions under which meaning itself is secured. By revisiting the principle in light of modern debates, the study aims to show that the limits of knowing are not a constraint to be overcome but a contour to be understood. The enduring task of philosophy, therefore, is not to expand meaning indiscriminately but to defend its coherence through rigorous standards of justification.

Statement of the Problem

The rapid development of machine-intelligent systems has introduced a profound epistemic challenge: outputs generated by such systems are often opaque, complex, and probabilistic, raising questions about their cognitive legitimacy. While A.J. Ayer's verification principle historically demarcated meaningful propositions from pseudo-statements by requiring a clear route to empirical or analytic confirmation, the opacity inherent in contemporary computational processes complicates the application of this criterion. This tension generates a pressing philosophical problem: how can claims produced by systems whose internal mechanisms are inaccessible be meaningfully assessed, and to what extent must verificationist standards be adapted to preserve the distinction between genuine knowledge and mere assertion in an age of machine intelligence?



Methodology

This study employs analytical methodology rooted in classical and contemporary philosophical analysis. It involves critical exegesis of primary sources, including Ayer's *Language, Truth and Logic*, supplemented by canonical texts in epistemology, philosophy of language, and philosophy of mind (Aristotle, Descartes, Quine, Searle). The study further integrates insights from contemporary scholarship on epistemic opacity, machine learning, and information theory (Humphreys 2009; Burrell 2016; Floridi 2011), applying logical, comparative, and evaluative reasoning to assess the relevance, adaptability, and limits of the verification principle in evaluating knowledge claims produced by opaque inferential systems. The approach is strictly qualitative, emphasizing normative philosophical analysis over empirical measurement.

Literature Review

The debate about meaning and verification has its roots in the early analytic tradition, where philosophers sought to clarify the conditions under which propositions could legitimately claim cognitive significance. The logical positivists, particularly Moritz Schlick and Rudolf Carnap, argued that meaningfulness required a demonstrable link between a statement and potential experience, a view that heavily shaped Ayer's later formulation (Schlick 1932, 27; Carnap 1932, 51). Their project aimed to eliminate speculative metaphysics by grounding language in empirically testable criteria. Although later revisions softened the strictness of their early formulations, the central insight—that verification anchors meaning—remained a defining signature of the movement. Ayer's own contribution, articulated most forcefully in *Language, Truth and Logic*, sharpened the distinction between analytic truths, grounded in logical relations, and synthetic truths, grounded in empirical observation (Ayer 1936, 15–20). Yet this framework did not emerge unchallenged. Karl Popper argued that verifiability was too restrictive a standard and proposed falsifiability as a more flexible criterion for scientific significance (Popper 1959, 40). While Popper's critique marked a major turning point, it did not undermine the philosophical core of verificationism: the demand for transparent, testable grounding for knowledge-claims. This shared concern remains instructive for evaluating claims produced by contemporary systems whose justificatory pathways are not readily accessible.

Later critiques from W.V.O. Quine and Hilary Putnam further complicated the verificationist picture. Quine's attack on the analytic–synthetic distinction destabilized one of Ayer's foundational pillars, arguing instead for a holistic conception of meaning where statements face experience not individually but as a collective body (Quine 1951, 38). Putnam advanced similar concerns, contending that meaning cannot be reduced to verification-conditions alone, as linguistic practices draw on broader social, semantic, and pragmatic factors (Putnam 1975, 215). These critiques do not negate the relevance of verificationism; rather, they illuminate the difficulties involved in preserving strict criteria in light of complex linguistic realities.

The question of epistemic legitimacy becomes even more pronounced in the context of computational procedures. Hubert Dreyfus argued early that formal systems lack the embedded understanding required for genuine cognitive claims (Dreyfus 1972, 81). In a similar vein, John Searle maintained that symbolic manipulation, however sophisticated, does not produce intrinsic understanding, a point that underscores the distinction between generating outputs and grounding meaning (Searle 1980, 417). These concerns echo the verificationist suspicion of claims whose justificatory grounding cannot be rendered explicit. They also highlight the philosophical divide between functional performance and epistemic warrant.

Recent literature deepens this divide. Paul Humphreys introduced the influential concept of “epistemic opacity,” describing systems whose internal operations cannot be fully inspected or understood by human observers (Humphreys 2009, 618). Jenna Burrell identified three forms of



opacity—intentional, technical, and conceptual—each of which complicates the project of verifying the basis of system-generated claims (Burrell 2016, 2–4). These insights reinforce the idea that epistemic confidence requires more than effective output; it requires traceable justification. Without this, claims risk falling into the category Ayer sought to exclude.

Other scholars point to structural limits in the epistemic status of such systems. Luciano Floridi argues that computational procedures produce “semantic artefacts” that require external interpretation to achieve genuine meaning (Floridi 2011, 89). José Luis Bermúdez, examining the nature of representation, maintains that meaningful content requires both structural mapping and availability to conscious appraisal (Bermúdez 2003, 72). This requirement aligns closely with verificationist principles, which insist that propositions must be open to confirmatory or disconfirmatory experience to count as cognitively significant.

In parallel, scholars such as Bas van Fraassen and Nancy Cartwright highlight the role of observable phenomena in grounding scientific claims. Van Fraassen’s constructive empiricism holds that theories need only be empirically adequate, not metaphysically complete, to be acceptable (van Fraassen 1980, 12). Cartwright’s analysis of models and laws similarly emphasises the importance of empirical grounding over abstract coherence (Cartwright 1983, 102). These perspectives strengthen the argument that claims lacking a clear empirical foothold cannot be granted full epistemic authority.

Developments in philosophy of language deepen the analysis. Michael Dummett argues that meaning is ultimately tied to conditions of use and recognizability, which require public, verifiable criteria (Dummett 1973, 58). Donald Davidson’s work on truth and interpretation likewise stresses that meaning depends on triangulation—publicly accessible standards that distinguish correct from incorrect assertions (Davidson 1984, 125). J.L. Austin and P.F. Strawson further underscore the inseparability of meaning from context, precision, and verifiable conditions of application (Austin 1962, 51; Strawson 1959, 32). Taken together, these positions sustain the philosophical demand that claims must connect with observable or analytically demonstrable standards to be judged meaningful.

The literature therefore reveals a consistent thread, despite internal disagreements: genuine epistemic authority requires intelligible, inspectable, and confirmable grounding. Whether through the strict criteria of early verificationism, the holistic critiques of mid-century philosophy, or the contemporary concerns about opacity, the central philosophical task remains the same. Meaning, to be meaning, must be anchored in grounds that reason can interrogate. This central insight positions Ayer’s verification principle not as an obsolete relic but as a still-potent standard by which claims, especially those produced in the present technological climate—may be critically evaluated.

The Epistemic Demand for Transparency

A central question emerges from the preceding literature: what conditions must a statement satisfy in order to count as meaningful within a disciplined epistemological framework? For Ayer and the verificationists, the answer is uncompromising—meaning requires a clear criterion through which the truth-value of a proposition can be determined, at least in principle (Ayer 1936, 15). This standard reflects a deeper philosophical insistence that language must maintain an intelligible relation to the world it purports to describe. Without such a relation, a statement may provoke interest, admiration, or even practical utility, but it cannot claim cognitive significance. This distinction between usefulness and meaningfulness lies at the heart of verificationist philosophy.

Yet the demand for transparency in justification is not unique to verificationism. It echoes broader epistemological concerns from Aristotle’s requirement for demonstrable *aitia* (causes) in scientific knowledge to Descartes’ insistence on clarity and distinctness as preconditions for certainty (Aristotle *Posterior Analytics* II.2; Descartes 1641, 24, Okhueigbe, 2025, 19). The thrust of these



traditions is the same: knowledge cannot float freely; it must rest upon grounds that reason can examine. Quine's later critique of the analytic–synthetic distinction did not eliminate this demand but reframed it, arguing that justification arises holistically through the web of belief (Quine 1951, 38). Even within this holistic model, the requirement for some form of accessible grounding remains non-negotiable.

This grounding becomes crucial when evaluating assertions generated by increasingly complex inferential systems. Their outputs often present themselves as definitive, yet the processes that yield these outputs may be opaque, whether through conceptual complexity, mathematical depth, or proprietary restrictions. Humphreys describes this opacity as an “epistemic black box,” wherein even experts cannot fully inspect the steps linking input to output (Humphreys 2009, 618). Such opacity threatens the very criterion of meaningfulness articulated by Ayer: if a proposition lacks a clear route to verification, its epistemic status becomes precarious.

Ayer himself anticipated variants of this problem in his criticism of metaphysical propositions. He argued that statements which could not be tested, even in principle, amounted to expressions of sentiment rather than assertions about the world (Ayer 1936, 33). Although his immediate targets were theological and speculative claims, the logical structure of his critique applies with equal force to any assertion lacking verifiable grounding. The philosophical challenge today is not that new systems produce sentiments but that they produce propositions whose justificatory pathways retreat from public inspection. This retreat places pressure on the very idea of cognitive significance.

To defend the meaningfulness of such outputs, one must therefore articulate criteria that allow their verification to be reclaimed rather than abandoned. Several thinkers attempt to soften the verificationist requirement. Putnam argues that meaning depends on a network of social practices rather than single verification-conditions (Putnam 1975, 227). Davidson similarly contends that meaning arises through triangulation, publicly accessible norms that determine correct application (Davidson 1984, 125). Even these broader accounts, however, presuppose transparency. Meaning requires not only use but *intelligible* use, accessible to rational adjudication. If the path from assertion to justification remains hidden, the assertion risks sliding into the realm of what Ayer would call “pseudo-propositions.”

Thus, the epistemic demand for transparency cannot be dismissed as an outdated positivist relic. Rather, it represents a persistent and unavoidable condition for meaningful discourse. Whether traced through the rigor of the Vienna Circle, the critiques of mid-century analytic philosophy, or the contemporary concerns about opacity, the conclusion is the same: propositions that cannot be traced to confirmable or analyzable grounds cannot claim the status of knowledge. This section therefore establishes the conceptual foundation for the next argument: that the limits of knowing in the present age emerge not merely from technological complexity but from philosophical negligence concerning the standards that make knowledge possible in the first place.

Prospects and Limits in the Context of Epistemic Opacity

Building on the preceding discussion of transparency and epistemic opacity, a pressing question emerges: can Ayer's verification principle, conceived in the early twentieth century, be meaningfully adapted to evaluate claims generated by contemporary systems whose inferential processes are inaccessible? At first glance, the principle appears ill-suited to contexts characterized by complexity, abstraction, and opacity. Its demand for direct empirical verification or analytic certainty seems incompatible with probabilistic models, recursive algorithms, and computational heuristics that defy straightforward inspection (Humphreys 2009, 618; Burrell 2016, 3). Yet a closer examination suggests that the principle's normative insight—linking meaningfulness to verifiable justification—remains relevant, albeit in a revised form.



One avenue for adaptation lies in extending the notion of verification beyond strict observational accessibility to include *indirect epistemic checks*. Bas van Fraassen (1980, 12) and Nancy Cartwright (1983, 104) suggest that models and theories may be deemed empirically adequate if their outputs systematically align with observable phenomena, even when the internal mechanisms remain opaque. In this sense, verification is understood as a probabilistic or inferential confirmation rather than a direct inspection of every operational step. Such an extension retains the verificationist commitment to empirical accountability while acknowledging the epistemic complexity of modern systems.

Philosophers of mind and language provide complementary insights. Searle (1980, 417) and Dreyfus (1972, 81) warn against conflating functional performance with genuine understanding; outputs alone do not constitute knowledge. Similarly, Floridi (2011, 89) and Bermúdez (2003, 72) emphasize the role of interpretation: computational artefacts acquire meaning only through structured engagement with human cognitive capacities. Verification, therefore, must encompass not only the empirical reliability of outputs but also their intelligibility to rational agents capable of interrogating them, even if only through mediated or indirect means.

Several contemporary theorists offer further refinements. Paul Humphreys (2009, 620) proposes that opacity can be mitigated through *epistemic scaffolding*: intermediate representations, model validation, and cross-system comparisons that render outputs sufficiently accountable for verification purposes. Burrell (2016, 4) similarly identifies mechanisms—such as logging procedures, explanatory modules, and transparency protocols—that preserve the possibility of justification without demanding total observability. These strategies demonstrate that the verificationist requirement can be preserved in spirit even when its strict letter is unattainable.

Yet significant limitations remain. As Quine (1951, 38) and Putnam (1975, 227) underscore, meaning and justification are relational and holistic: verification cannot be reduced to isolated confirmation of outputs. If opacity permeates the entire network of inference, no amount of indirect checking can fully restore epistemic authority. Moreover, the probabilistic and heuristic nature of many contemporary systems introduces indeterminacy: an output may be reliably predictive yet remain contingent, provisional, or context-dependent, challenging traditional verificationist criteria (van Fraassen 1980, 15). Philosophical rigor thus demands careful calibration between pragmatism and normative standards.

Ayer's verification principle can be adapted to contemporary opaque systems only by extending its criteria to include indirect, probabilistic, and mediated forms of confirmation. Such adaptation preserves the epistemic demand for meaningful justification while accommodating the novel complexity of modern inferential systems. However, this approach is not without tension: it requires balancing fidelity to philosophical standards with recognition of technological realities. Verification, properly understood, remains a critical lens for assessing the epistemic legitimacy of outputs whose internal processes resist direct scrutiny.

Delineating the Boundaries of Knowledge

The preceding analysis demonstrates that the epistemic demands articulated by Ayer's verification principle retain critical relevance in evaluating knowledge-claims generated by contemporary computational systems. While the strict original formulation appears challenged by opacity and complexity, extending verification to include indirect, probabilistic, and mediated forms of confirmation allows the principle to remain a normative benchmark. In essence, verification continues to define the boundary between statements that qualify as meaningful and those that merely simulate cognitive authority (Ayer 1936, 15–20; Humphreys 2009, 618). This boundary becomes especially significant in an era in which outputs are often persuasive, operationally effective, or rhetorically compelling, yet epistemically underdetermined.



A critical insight emerges from the convergence of classical epistemology and contemporary computational concerns. Philosophers from Aristotle to Descartes and beyond emphasized the necessity of accessible justification, whether through demonstrable causes or clarity and distinctness (Aristotle *Posterior Analytics* II.2; Descartes 1641, 24). Modern debates, articulated by Dreyfus (1972, 81), Searle (1980, 417), and Floridi (2011, 89), suggest that opacity in algorithmic systems does not merely complicate verification but raises profound questions about the very status of knowledge. In these cases, outputs may function as epistemically provisional artifacts rather than fully justified claims, highlighting the distinction between reliability and meaningfulness.

This tension is further sharpened by holistic and pragmatic perspectives. Quine (1951, 38) and Putnam (1975, 227) reveal that verification is relational, depending not on isolated propositions but on their integration within a network of beliefs and practices. Opacity, when pervasive, threatens this network, potentially undermining the epistemic authority of the entire system. Even when indirect checks are applied, probabilistic or interpretive confirmation cannot fully substitute for the transparency that classical verificationism demands. Yet, as van Fraassen (1980, 12) and Cartwright (1983, 104) suggest, such probabilistic adequacy may suffice for empirical accountability, establishing a practical compromise between ideal epistemic standards and technological realities. The synthesis indicates that verificationism, when thoughtfully adapted, provides both a diagnostic and normative tool. It diagnoses the limits of knowledge by identifying outputs whose justification cannot be traced or inspected. It also serves a normative function, prescribing the minimal conditions required for meaningful claims, whether through empirical alignment, interpretive engagement, or mediated checks (Burrell 2016, 4; Humphreys 2009, 620). In this sense, the verificationist framework functions as a bridge between classical epistemic rigor and the novel challenges posed by modern inferential systems.

A further implication is that epistemic responsibility shifts. Philosophers, designers, and users of complex systems must recognize that meaningfulness is not automatically conferred by technological sophistication. Verification, even in its extended form, requires active interrogation and interpretive engagement. Outputs devoid of accessible justification cannot be treated as knowledge, however functionally useful they may appear. This reflects a broader moral and intellectual imperative: epistemic authority entails accountability, not mere efficacy (Floridi 2011, 95; Bermúdez 2003, 72). The analysis suggests a conceptual equilibrium: while verificationist criteria must be adapted to accommodate opacity, they cannot be abandoned without dissolving the very notion of knowledge. Meaningfulness, whether analytic, empirical, or inferential, depends on some accessible path to confirmation. Outputs that evade scrutiny entirely risk being epistemically hollow. The verificationist principle, therefore, continues to serve as a vital touchstone, clarifying where knowledge ends and mere assertion begins in an age dominated by machine intelligence.

Evaluation

The analysis conducted across Sections I–IV demonstrates that Ayer’s verification principle remains a philosophically robust framework for evaluating the meaningfulness of propositions, even in contexts marked by opacity and algorithmic complexity. Its central insistence—that meaning requires a clear route to verification—continues to distinguish genuine knowledge from pseudo-statements (Ayer 1936, 15–20). Contemporary systems, characterized by probabilistic outputs, recursive reasoning, and heuristic approximations, present challenges to the original strictures of verificationism. However, extending the principle to include mediated, inferential, or probabilistic forms of confirmation preserves its epistemic utility (Humphreys 2009, 620; Burrell 2016, 4).

Moreover, the study reveals that epistemic opacity is not merely a technical or procedural problem; it is fundamentally philosophical. Scholars such as Dreyfus (1972, 81), Searle (1980, 417), and Floridi (2011, 89) emphasize that functional performance does not equate to knowledge. Outputs lacking intelligible, inspectable, or interpretable grounding risk eroding the distinction between



effective prediction and epistemically legitimate claims. Therefore, verificationism, even when adapted, provides a necessary benchmark to guard against the epistemic inflation of machine-generated assertions.

The evaluation further confirms the need for a dual lens: verification must assess both empirical adequacy and interpretive accessibility. Holistic perspectives from Quine (1951, 38) and Putnam (1975, 227) reinforce that justification is relational, relying on the integration of outputs into a coherent network of beliefs. Where opacity threatens the integrity of this network, the epistemic status of all dependent claims is compromised. In sum, the principle's enduring relevance lies not in strict literalism but in its normative insistence on accountability and intelligibility.

Conclusion

This study has argued that the verification principle, when thoughtfully extended, remains an indispensable tool for delineating the boundaries of knowledge in a machine-intelligent age. Classical verificationism established the critical link between meaning and verifiability; contemporary epistemic challenges, particularly opacity, do not invalidate this insight but demand its adaptation. By accommodating indirect, probabilistic, and mediated forms of verification, the principle can address the realities of complex inferential systems while preserving philosophical rigor. Ayer's framework clarifies the limits of knowing: outputs may be operationally effective, rhetorically compelling, or empirically predictive, but they attain epistemic legitimacy only when accessible to rational scrutiny. Verificationism thus functions as both a diagnostic and normative standard, distinguishing genuine knowledge from mere assertion or pseudo-knowledge. The principle's application ensures that intellectual authority, rather than technological sophistication alone, governs the status of propositions in contemporary epistemic landscapes.

Recommendations

1. Philosophers and epistemologists should actively engage with computational and algorithmic systems, applying verificationist principles to evaluate the meaningfulness and legitimacy of outputs.
2. System designers and institutions should implement mechanisms such as explanatory modules, logging protocols, and interpretive scaffolds to maintain traceability of outputs (Humphreys 2009, 620; Burrell 2016, 4).
3. Curricula in philosophy, information science, and cognitive studies should include training on verificationist evaluation and epistemic accountability in contexts involving opaque inferential systems.
4. Policymakers in AI and computational ethics should employ verificationist-inspired criteria to assess claims made by algorithmic systems, ensuring that epistemic authority is grounded in transparency and intelligibility rather than mere performance metrics.
5. Further studies should explore the intersection of verificationism, epistemic opacity, and emerging machine-intelligent systems, integrating insights from philosophy, computer science, and cognitive theory to refine normative standards for knowledge in the digital age.

References

- Aristotle. 1952. *Posterior Analytics*. Translated by G.R.G. Mure. London: Heinemann.
- Austin, J. L. 1962. *How to Do Things with Words*. Oxford: Clarendon Press.
- Ayer, A. J. 1936. *Language, Truth and Logic*. London: Gollancz.
- Bermúdez, José Luis. 2003. *Thinking Without Words*. Oxford: Oxford University Press.



- Burrell, Jenna. 2016. "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms." *Big Data & Society* 3 (1): 1–12.
- Carnap, Rudolf. 1932. "The Elimination of Metaphysics." *Erkenntnis* 2: 219–241.
- Carnap, Rudolf. 1950. *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- Davidson, Donald. 1984. *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press.
- Descartes, René. 1641. *Meditations on First Philosophy*. Translated by John Cottingham. Cambridge: Cambridge University Press, 1996.
- Dreyfus, Hubert. 1972. *What Computers Can't Do*. New York: Harper & Row.
- Dummett, Michael. 1973. *Frege: Philosophy of Language*. London: Duckworth.
- Floridi, Luciano. 2011. *The Philosophy of Information*. Oxford: Oxford University Press.
- Humphreys, Paul. 2009. "The Philosophical Novelty of Computer Simulation Methods." *Synthese* 169 (3): 615–626.
- Popper, Karl. 1959. *The Logic of Scientific Discovery*. London: Hutchinson.
- Putnam, Hilary. 1975. *Mind, Language, and Reality*. Cambridge: Cambridge University Press.
- Quine, W.V.O. 1951. "Two Dogmas of Empiricism." *The Philosophical Review* 60 (1): 20–43.
- Schlick, Moritz. 1932. "Positivism and Realism." *Erkenntnis* 3: 1–30.
- Searle, John. 1980. "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3: 417–457.
- Strawson, P. F. 1959. *Individuals*. London: Methuen.
- van Fraassen, Bas. 1980. *The Scientific Image*. Oxford: Clarendon Press.