

Shadows to Referents: An Ontological Interpretation of Plato's Cave for Anomaly Resolution

John H. Bittner^{1,2}, John Beverley^{1,2,3} and Timothy W. Coleman^{1,2}

¹Department of Philosophy, University at Buffalo, NY, USA

²National Center for Ontological Research, NY, USA

³Institute for Artificial Intelligence and Data Science, University at Buffalo, NY, USA

Abstract

This paper develops an ontological interpretation of Plato's Cave to clarify the lifecycle of anomalies and their resolution. We argue that anomalies, as epistemic disruptions, exist only within bounded temporal regions; beginning with an initial observation and ending with resolution that reclassifies them into known categories. To formalize this lifecycle, we introduce the Anomaly Resolution Ontology (ARO) aligned with the Basic Formal Ontology (BFO) and Common Core Ontologies (CCO). These ontologies model anomalies as temporally bounded continuants, together with the acts of observation and resolution that define their existence. The framework is illustrated through the Act of Epistemic Enlightenment, modeled as a resolution act that ends the anomaly and transforms unresolved information into resolved information. By linking the Cave allegory with anomaly resolution, this paper advances a semantic framework for representing epistemic transitions in AI, cognitive modeling, and intelligence analysis.

Keywords

Anomaly, Information Content Entity, Epistemic Uncertainty, Temporal Region, Basic Formal Ontology

1. Introduction

Anomalies represent unresolved deviations from patterns expected by observers, systems, or institutional models that create significant operational and epistemic risks across multiple domains. When observers or systems are unable to assign data to a determinate referent, such unresolved phenomena reflect semantic configurations shaped by context, prior information, and representational constraints.

For example, in military intelligence, an unresolved thermal signature may lead to tactical missteps with life-threatening consequences. In AI systems, misclassifications erode user trust and propagate faulty downstream decisions. In scientific discovery, as Lindley Darden notes, unresolved anomalies may stall theoretical progress by preventing determinate referential assignments to observed data [1]. Modern detection systems flag anomalies as deviations from expected patterns [2], often stemming from representational rather than material mismatches; semantic artifacts produced by interpretation rather than direct observation.

These practical consequences motivate the need for a structured ontological approach capable of formally capturing the lifecycle of anomalies and guiding resolution processes. To address these challenges, we introduce the Anomaly Resolution Ontology (ARO) [3], which offers a formalized ontological approach to anomaly lifecycle modeling.

The ARO provides a formal, interoperable framework for modeling the semantic and temporal dimensions of anomalies. Unlike statistical models that merely detect deviations, or rule-based systems that enforce fixed classifications, ARO captures the evolving epistemic states, referential gaps, and semantic transitions that characterize anomaly resolution [4, 5, 6, 7]. This ontological approach is especially critical in domains where anomalies reflect not just data irregularities but interpretive uncertainty that demands conceptual realignment.

Proceedings of the Joint Ontology Workshops (JOWO) - Episode XI: The Sicilian Summer under the Etna, co-located with the 15th International Conference on Formal Ontology in Information Systems (FOIS 2025), September 8–9, 2025, Catania, Italy

✉ jhbittne@buffalo.edu (J. H. Bittner); johnbeve@buffalo.edu (J. Beverley); tcoleman@buffalo.edu (T. W. Coleman)

🆔 0009-0001-5196-8279 (J. H. Bittner); 0000-0002-1118-1738 (J. Beverley); 0009-0009-3931-4393 (T. W. Coleman)



© 2025 Copyright © 2025 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

The ARO [3] is developed as a domain-specific extension of the Common Core Ontologies (CCO) [8] and aligned with Basic Formal Ontology (BFO) [9], which serves as the ISO/IEC 21838-1 standardized upper-level ontology for scientific and applied domains. BFO provides formal distinctions between universals (general types) and particulars, which are divided into continuants (entities persisting through time) and occurrents (entities unfolding in time), enabling consistent categorization of entities across diverse contexts [9].

CCO builds on BFO by introducing mid-level ontologies that define common types such as material artifacts, information content entities, and acts of processing [8]. This layered foundation allows ARO to formally model the lifecycle of anomalies while ensuring cross-domain interoperability and ontological realism as instantiated in the commitments of BFO [4]. As Barry Smith argues in his critique of fantology [10], many logical frameworks obscure the temporal and indeterminate nature of unresolved phenomena, treating referential absences as formal artifacts rather than structurally meaningful states.

To orient the reader, we briefly preview several key concepts introduced in this framework. **Anomaly:** an Information Content Entity (ICE), understood as a generically dependent continuant that is about some entity and depends on an information-bearing entity for its existence, here reflecting unresolved deviation. Ontologically, an Anomaly participates in an Act of Anomaly Detection, exists during a Bounded Temporality Interval, and is assigned an Epistemic State that reflects its unresolved interpretive status. **Bounded Temporality:** a Prescriptive Information Content Entity (PICE) defining the temporal span between anomaly detection and referential resolution. **Act of Epistemic Enlightenment:** a Planned Act of Information Processing in which an agent achieves referential clarity by resolving an Anomaly and terminating its interpretive ambiguity. We use ‘Epistemic Enlightenment’ to denote a formal process of referential clarification, not a metaphorical or cognitive state.

An Information Content Entity (ICE) is a type of generically dependent continuant that depends on a bearer for its existence, such as a database record, sensor observation, or digital document, and is about some portion of reality to which it refers. The notion of aboutness is central to its ontological characterization [11, 12]. As Dretske observes, informational structures can carry semantic content even in the absence of verified referents [13].

An Act of Epistemic Enlightenment is not limited to cognitive agents. It may occur when a sensor fusion system classifies a thermal signature as narcotics rather than mechanical heat loss [2], or when a machine learning anomaly is resolved through human or rule-based reasoning [14]. In each case, the anomaly, sustained by unresolved content, no longer functions as an anomaly once its referent is identified. This epistemic and ontological transition is formally captured by the the Act of Epistemic Enlightenment.

Our argument proceeds in four stages. We first reinterpret Plato’s Cave [15] as a parable of semantic mismatch, framing shadows as information content entities linked to type-based intersections. We then define unresolved anomalies using the ontological distinction between generically dependent continuants and information content entities, introduce the Act of Epistemic Enlightenment as a subclass of Act of Information Processing, and finally illustrate the framework with examples from ISR and AI misclassification.

Anomaly resolution is the transition from shadow (i.e., an unresolved Information Content Entity representing an unknown or indeterminate referent) to referent, modeled as a temporally bounded Act of Anomaly Resolution and Act of Information Processing, where unresolved content is formally processed and resolved by Agents or Systems. The framework relies on Bounded Temporality, a subclass of PICE which specifies a temporal interval determined by epistemic or contextual conditions that constrain when entities such as anomalies exist. An Anomaly exists from the moment of its initial Act of Anomaly Detection, the observation introducing unresolved content, until its resolution through an Act of Epistemic Enlightenment. Once referential clarity is achieved, the conditions sustaining the Anomaly collapse.

In this paper, we focus specifically on anomalies that involve referential indeterminacy, where an Information Content Entity exists but its aboutness relation cannot yet be determinately assigned to a portion of reality. Such anomalies differ from routine deviations or noise: their lifecycle depends

on resolving the gap between unresolved informational content and the entity it purports to describe. While other forms of anomalies (e.g., measurement error within known tolerances) can also be modeled, our concern is with those whose ontological status is sustained only until determinate aboutness is achieved.

Table 1: Ontology classes referenced in the ARO, BFO, and the CCO.

Label	Subclass Of	Definition/Elucidation
Anomaly	Information Content Entity	An Information Content Entity that describes an unresolved deviation from expected patterns, entities, or configurations within a given context.
Unresolved Anomaly Information	Information Content Entity	An Information Content Entity that describes ambiguous or unverified information about an Anomaly.
Resolved Anomaly Information	Information Content Entity	An Information Content Entity that describes confirmed or interpreted explanation about an Anomaly.
Anomaly Lifecycle	Descriptive Information Content Entity	A Descriptive Information Content Entity that describes the lifecycle of an Anomaly, beginning with its detection, continuing through its unresolved state, and ending with its resolution via an Act of Epistemic Enlightenment that terminates the Bounded Temporality Interval.
Bounded Temporality	Prescriptive Information Content Entity	A Prescriptive Information Content Entity that specifies conditions under which some entity is considered to exist, apply, or hold; where these conditions include the initiation and termination of that status based on epistemic, institutional, biological, or social criteria.
Epistemic State	Prescriptive Information Content Entity	A Prescriptive Information Content Entity that encodes the interpretive status of an ICE's referent, such as whether it is known, unknown, or partially understood.
Known Known	Epistemic State	An Epistemic State that prescribes that an Anomaly has a referent that is identified and acknowledged, even if its cause remains unresolved, during a Bounded Temporality Interval.
Known Unknown	Epistemic State	An Epistemic State that prescribes that an Anomaly lacks an identified or understood referent during a Bounded Temporality Interval.
Act of Resolution	Act of Information Processing	An Act of Information Processing in which an Agent clarifies, revises, or replaces an ICE in order to resolve ambiguity or uncertainty in what the content is about.
Act of Anomaly Resolution	Act of Resolution	An Act of Resolution in which an Agent interprets or clarifies Unresolved Anomaly Information to produce Resolved Anomaly Information.
Act of Detection	Act of Observation	An Act of Observation in which an Agent uses a sensor, system, or analytic method to determine whether an entity of interest is present or identifiable within a physical, informational, or computational environment.
Act of Anomaly Detection	Act of Detection	An Act of Information Processing in which an Agent analyzes input data to identify the presence of one or more Anomalies.
Bounded Temporality Interval	temporal interval	A temporal interval that represents the continuous span of time during which a temporally bounded entity, such as an Anomaly or role or process, exists.
Act of Pseudo-Closure	Act of Invalid Anomaly Resolution	An Act of Resolution in which an Anomaly is dismissed, reclassified, or collapsed without adequate justification, resulting in a premature termination of the anomaly's unresolved status.
Act of Invalid Anomaly Resolution	Act of Resolution	An Act of Resolution that replaces an Unresolved Anomaly Information Content Entity without achieving referential stability.

Label	Subclass Of	Definition/Elucidation
Act of Epistemic Enlightenment	Act of Information Processing	An Act of Information Processing in which an Agent achieves referential clarity by identifying the actual entity or cause underlying the Anomaly, thereby eliminating its unresolved status.
generically dependent continuant	continuant	A generically dependent continuant is an entity that exists in virtue of there being at least one (possibly many) copies which share the same content or pattern.
temporal interval	one-dimensional temporal region	A one-dimensional temporal region that is continuous, without gaps or breaks.
immaterial entity	independent continuant	An independent continuant which is such that there is no time at which it has a material entity as continuant part.
Planned Act	Act	A process in which at least one Agent plays a causative role.
Information Content Entity	generically dependent continuant	A generically dependent continuant that generically depends on some Information Bearing Entity and stands in relation of aboutness to some Entity.
Prescriptive Information Content Entity	Information Content Entity	An Information Content Entity that consists of a set of propositions or images (e.g., a blueprint) that prescribe some entity.
Act of Information Processing	Planned Act	A Planned Act in which one or more input ICEs are received, manipulated, transferred, or stored by an Agent.
Act of Observation	Planned Act	A Planned Act of acquiring information from a source via one or more senses.
Agent	material entity	A material entity that bears an Agent Capability.

The definition of Anomaly should be understood to include both expected and unexpected cases. Some anomalies, such as diagnostic alerts designed into medical monitoring systems, are anticipated possibilities; others arise as unanticipated deviations. In both situations, their ontological status as anomalies persists only while referential indeterminacy remains.

In this light, anomalies are not to be dismissed as mere errors or noise, but recognized as epistemic events that reveal the limits of current understanding. Each anomaly marks a transition in knowledge, shifting from an unresolved shadow to a clarified referent, and from indeterminate information to structured content that can be integrated into existing ontologies. This framing establishes anomalies as critical moments in the movement from ignorance to recognition, setting the stage for a formal account of their lifecycle.

2. Ontological Structures and Shadows

Plato’s Allegory of the Cave [15] casts shadows as ambiguous projections lacking reference to their true sources. Ontologically, shadows can be understood as referentially indeterminate information structures—semantic placeholders that arise when phenomena fail to align with determinately referenced entities. These are not errors or hallucinations, but genuine Information Content Entities (ICEs) marked by epistemic uncertainty. Such entities correspond to anomalies: ICEs whose aboutness is incomplete or unresolved at the time of detection.

We adopt insights from the Unreal Patterns framework [16], which conceptualizes unresolved reference using anonymous class-level intersections of types. While our ontology does not formally implement anonymous class expressions, it inherits the semantic intuition that anomalies represent type-level uncertainty without asserting placeholder individuals. For example, an anomaly may concern a thermal signature exhibiting elevated heat, anomalous cargo patterns, and insulation characteristics; yet until resolved, no determinate referent is assigned. This avoids ontological inflation by capturing uncertainty without positing fictitious entities like “unknown vehicle 47”. In this framework, what Plato metaphorically described as a “shadow” is represented as an ICE that encodes unresolved semantics, echoing the anonymous class strategy in Unreal Patterns.

This approach aligns with ontological realism, which permits modeling incomplete knowledge structures without reifying nonexistent particulars [4]. In ARO, these commitments allow unresolved intervals to be treated as ontologically significant, temporally bounded entities within an anomaly lifecycle [3].

In this framework, the “shadow” is not a material entity, but an ICE representing unresolved semantic content as class-level intersections. It corresponds to a referential gap that persists until referential assignment resolves the aboutness relation. This approach avoids ontological overreach by capturing representational uncertainty without requiring commitment to fictitious individuals. As Smith and Ceusters note, ontological realism permits

representations of incomplete knowledge without inflating the domain [4]. Building on Smith’s critique of fantology [10], which highlights how standard logical frameworks can obscure change and indeterminacy, we extend these insights to the modeling of anomalies as temporally bounded entities. Building on this foundation, the ARO explicitly models these periods of unresolved reference and interpretive uncertainty as ontologically significant intervals within an entity’s lifecycle, capturing both their temporal boundaries and their semantic status [3].

3. Act of Epistemic Enlightenment

The Act of Epistemic Enlightenment, subclassed under Act of Information Processing [8], is a Planned Act that transforms unresolved semantic content into resolved form through referential clarification. This act terminates the anomaly’s status as a temporally bounded, unresolved representational gap [5]. The Unresolved Anomaly Information that sustains it, an Information Content Entity (ICE) carrying semantic content about an as-yet undetermined referent, is either resolved or clarified through referential assignment.

In Plato’s Allegory of the Cave [15], this corresponds to the prisoner turning from shadows toward the source, gradually acquiring referential clarity. Ontologically, anomaly resolution is modeled as a transition between distinct ICEs. Initially, the anomaly is sustained as an ICE encoding Unresolved Anomaly Information: semantic content that reflects a referential gap when observed data cannot yet be fully determined or dismissed. The Act of Epistemic Enlightenment brings referential clarity and marks the end of the Bounded Temporality Interval. Only then does the aboutness of the anomaly become referentially determinate [1].

The transition from unresolved to referentially determinate status alters the unity conditions of the anomaly, as unresolved fragments coalesce into an individuated entity [7]. For example, in medical diagnosis, an initial psychiatric evaluation may yield Unresolved Anomaly Information when symptoms resist classification, corresponding to a Known Unknown. As diagnostic inquiry proceeds, resolution may assign referential clarity, for instance, as bipolar II disorder under DSM-5 criteria, thereby generating Resolved Anomaly Information and completing the Act of Epistemic Enlightenment [17]. This transformation completes the anomaly’s role as a semantic placeholder. Once determinate aboutness is assigned, the anomaly no longer functions as a marker of uncertainty. Its existence is temporally bounded between its detection and resolution, and the Act of Epistemic Enlightenment defines that boundary. In Plato’s allegory, the shadows lose their interpretive force once their referents become known [15]. Similarly, anomalies are not merely eliminated; rather, their ambiguity is ontologically extinguished through assignment of determinate aboutness. The information structures that sustained them are no longer active [11, 13].

Table 2
OWL subclass axioms (fragment) formalizing selected ARO classes.

Class	Subclass Of (Axioms)
Anomaly	Information Content Entity; has_interpretive_status some Epistemic State; participates_in some Act of Anomaly Detection; is_about some Entity; exists_at some Bounded Temporality Interval; generally_depends_on some Information Bearing Entity.
Unresolved Anomaly Information	Information Content Entity; is_about some Anomaly; exists_at some Bounded Temporality Interval.
Resolved Anomaly Information	Information Content Entity; output_of some Act of Anomaly Resolution.
Act of Anomaly Detection	Act of Information Processing; has_participant some Anomaly; marks_start_of some Bounded Temporality Interval.
Act of Epistemic Enlightenment	Act of Information Processing; participates_in some Act of Anomaly Resolution; marks_end_of some Bounded Temporality Interval.

While anomalies are detected under conditions of referential uncertainty, their bounded temporal existence

is modeled ontologically: each Anomaly is treated as an ICE with real temporal boundaries defined by its participation in the detection and resolution process, not merely by our epistemic access to it.

The Act of Epistemic Enlightenment may be carried out by agents, AI systems, or hybrid workflows. It does not require introspection or conscious deliberation; what matters is that the event results in the creation of Resolved Anomaly Information, an ICE that is about a confirmed or interpreted explanation of the anomaly. This formal act operationalizes the referential closure of unresolved semantic content, transforming ambiguity into stable representation and terminating the anomaly's existence. This act transforms unresolved content into a grounded representation, thereby terminating the anomaly. The Bounded Temporality Interval, a constrained temporal interval, is delimited by this transformation; it begins with the initial Act of Anomaly Detection and ends with the occurrence of the Epistemic Enlightenment Event, which serves as the semantic and temporal closure of the anomaly's existence. The core ontological structure of this transition is illustrated in Figure 1.

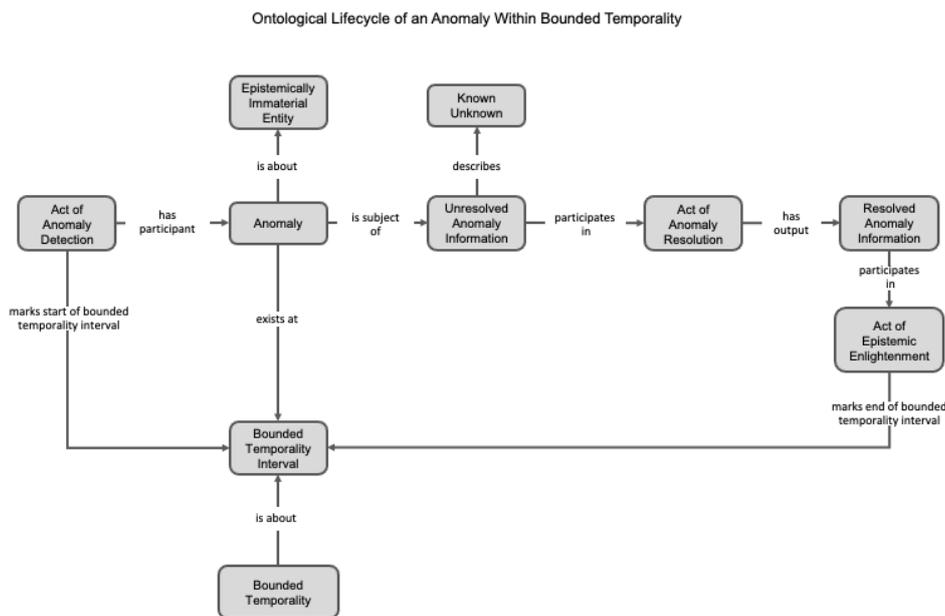


Figure 1: Ontological lifecycle of an anomaly within Bounded Temporality.

4. The Reach and Limits of Epistemic Enlightenment

This section demonstrates how anomaly resolution is represented ontologically through the Act of Epistemic Enlightenment, across domains including intelligence, AI, and human cognition. A simple case illustrates the pattern. A heart monitor flags an irregular signal that is later traced to sensor movement rather than pathology. The anomaly exists only from detection until clarification, bounded by an unresolved interval that terminates with referential resolution.

4.1. Intelligence, Surveillance, and Reconnaissance

In satellite-based ISR, sensors capture thermal, spectral, or radar signatures that deviate from expected environmental patterns. A multispectral satellite might detect an infrared signature inconsistent with the chlorophyll index of an agricultural zone [2]. The anomaly is not the land feature itself, but the system's failure to interpret the observed configuration; this generates Unresolved Anomaly Information.

An analyst or automated process may later fuse the data with historical overlays and identify the signature as a decoy installation [2]. This interpretive clarification links the previously unresolved content to a determinate entity through assignment of aboutness, producing Resolved Anomaly Information and terminating the anomaly's ontological lifecycle. More than classification, this is an Act of Epistemic Enlightenment, closing the interpretive gap.

4.2. Artificial Intelligence and Language Models

Large Language Models (LLM) often produce confident-sounding responses that are semantically unstable. As Bender et al. explain, “an LM is a system for haphazardly stitching together sequences of linguistic forms it has observed in its vast training data, according to probabilistic information...but without any reference to meaning: a stochastic parrot” [18]. A fabricated citation or conflated historical account may seem plausible but fails under scrutiny. Such cases represent mismatches between the user’s expectations and the model’s output. When the user intervenes to detect and correct the representational error, this constitutes an Act of Epistemic Enlightenment.

These errors are often described as hallucinations in the language model literature, with distinctions drawn between intrinsic and extrinsic types depending on whether the output contradicts the input or introduces unsupported content [19]. What is often described as “chain-of-thought reasoning” in large language models can instead be interpreted as a sequence of generated outputs that function as provisional anomaly resolutions, where each step represents unresolved content that may or may not converge on determinate aboutness. This framing avoids treating such outputs as reasoning while still highlighting how ARO can capture their status as unresolved or resolved informational artifacts.

4.3. Behavioral Psychology and Misperception

In eyewitness testimony [20], a person may confidently identify a suspect based on distorted perception shaped by stress, lighting, or bias. As Elizabeth Loftus observes, “a robust impairment of memory [can be] produced by exposure to misinformation,” causing false details to be recalled as real [21]. The claim persists as unresolved content about a type-level configuration, such as clothing or proximity. When subsequent evidence disproves the claim, referential clarification eliminates the anomaly. The original certainty was not about a person. It was about a semantic shadow. Anomaly resolution through referential clarification constitutes an Act of Anomaly Resolution, which leads to Resolved Anomaly Information that participates in an Act of Epistemic Enlightenment.

Social anomalies can arise from collective expectation. Consider someone claiming to see a snake in a tree. There is no snake, but the claim propagates uncertainty. Others begin to perceive something coiled among the leaves, not through observation but shared suggestion. This echoes Sherif’s 1935 study on conformity in ambiguous conditions [22] and Asch’s 1951 experiment on group pressure and perceptual override [23]. Each participant sustains the anomaly based on unresolved visual content. Resolution occurs when the collective inspects the tree and determines that no corresponding entity exists, thereby completing a shared Act of Epistemic Enlightenment. Building on Smith’s critique of fantology [10], which highlights how standard logical frameworks can obscure change and indeterminacy, we extend these insights to the modeling of anomalies as temporally bounded entities. Figure 1 illustrates this ontological transition, showing how the anomaly exists within a Bounded Temporality Interval and is resolved through the Act of Epistemic Enlightenment.

In many application domains, multiple agents may hold divergent epistemic states regarding the same anomaly. For example, one analyst may classify a signal as resolved while another maintains it as unresolved. ARO accommodates such plurality by allowing different Epistemic States to be assigned in parallel, reflecting the coexistence of conflicting interpretations until referential closure is achieved.

5. The Ontology

Many conventional systems, including logic-based AI classifiers, operate under assumptions of ontological completeness that obscure representational ambiguity. Such systems prematurely assume referential stability, collapsing the distinction between resolved and unresolved content. This erases the very boundary conditions epistemic transitions depend on. As Barry Smith observes, “When it comes to truths about things marked by change...[first-order logic] needs to be extended by some sort of new machinery” [10].

The ARO [3] introduces that machinery, modeling the transition from unresolved to resolved semantic content. This approach responds to fantological limitations in representing context-sensitive transitions, such as those Smith identifies in logic-based systems that reduce temporal processes to static tuples [10]. ARO defines a set of formal ontology classes: Anomaly, Unresolved Anomaly Information, Resolved Anomaly Information, Act of Epistemic Enlightenment, and Bounded Temporality Interval, that collectively model the full semantic lifecycle of anomalies. These classes and their formal definitions are summarized in Table 1. ARO is implemented in OWL2 and validated using SHACL (Shapes Constraint Language) constraints, ensuring compatibility with BFO 2020 [9] and the Common Core Ontologies [8]. The OWL2 and SHACL artifacts implementing ARO are publicly available at [3] for direct inspection and reuse. The Act of Epistemic Enlightenment, subclassed under Act of Information Processing, marks the terminus of this transition.

The ontology also models failed or misleading forms of resolution, defined as Acts of Pseudo-Closure: Acts of Anomaly Resolution in which an anomaly is dismissed or misassigned without referential clarity, resulting in semantically invalid closure. Such pseudo-closures create superficial coherence without achieving true epistemic grounding. Some anomalies are explained away through overfitting, correlation-based reasoning, or premature classification. As Guarino and Welty argue, misapplied class or property structures can yield referential assignments that appear valid but are ontologically unsound [6, 5]. These acts may create apparent resolution without eliminating ambiguity. As Yablo notes, aboutness can persist even without a definite target, so long as structural reference remains [24]. The anomaly appears resolved, but its semantic gap remains open.

The model presented here does not introduce new upper ontology terms, but it deepens the interpretation of existing structures. The Anomaly Lifecycle is positioned as an epistemically sensitive temporal construct: beginning with representational misalignment, proceeding through semantic re-evaluation, and terminating with assignment of determinate aboutness. This mirrors Darden's account of theory change, where the repair of conceptual misalignment enables more accurate model-world correspondence [1].

Anomalies are modeled as generically dependent continuants sustained by ICE that represent unresolved semantic content. Their existence reflects an unresolved semantic configuration and is temporally bounded by the absence of a determinate referential assignment. This allows anomalies to be treated formally as ontologically significant semantic gaps, situated within the generically dependent continuant framework while preserving coherence across BFO categories.

Not all resolutions achieve referential clarity. Some events terminate Anomalies without resolving the underlying representational mismatch, producing what this paper terms pseudo-closure. These pseudo-resolutions mimic epistemic closure but leave unresolved semantic gaps vulnerable to reactivation. This distinction draws on Guarino and Welty's critique of ontological misuse, where taxonomic inconsistencies yield apparent but invalid resolutions [6, 5].

Related work on anomaly handling often arises in engineering maintenance, risk management, and failure analysis. For example, production line failures are modeled as anomalies relative to tolerance thresholds, while risk management frameworks describe anomalies as deviations from safety norms. ARO complements these approaches by providing a formal account of referentially indeterminate anomalies, cases where uncertainty stems from semantic rather than purely material mismatch. This distinguishes our contribution from statistical anomaly detection surveys [14] or engineering risk ontologies, while ensuring interoperability through BFO and CCO alignment.

6. Conclusion

This paper reinterprets Plato's Cave not just as a metaphor for ignorance, but as an ontological model of anomaly resolution [15]. The shadows represent anomalies: generically dependent continuants that persist only while their referents remain unresolved. An Act of Epistemic Enlightenment terminates the anomaly by enabling referential resolution, marking the point at which Unresolved Anomaly Information gives way to Resolved Anomaly Information.

By introducing the Act of Epistemic Enlightenment as a subclass of Act of Information Processing, we refine the formal structure of the ARO [3]. Rather than adding new primitives, the model clarifies resolution as the grounding of uncertainty in referential stability. This transition unfolds in time and is operationalized by the Act of Epistemic Enlightenment, which marks the replacement of the ICE that sustains the anomaly, from Unresolved Anomaly Information to Resolved Anomaly Information. The ARO formalizes this structure to support semantic alignment, anomaly lifecycle modeling, and epistemic transitions across domains [3].

Across domains, from ISR to medical diagnosis to large language model inference, anomalies exist ontologically only while uncertainty remains. Their resolution, whether gradual or abrupt, is marked by an Act of Epistemic Enlightenment, which ends the Bounded Temporality Interval. An anomaly does not end simply with consensus; rather, it terminates with a shift in ontological status enabled by the resolution of representational mismatch [4].

Resolution is also constrained by Bounded Temporality. An anomaly begins with the detection of a representational mismatch and ends with referential closure. Yet this boundary can be extended by diagnostic lag, institutional resistance, or social refusal to accept clarification. Temporality reveals where epistemic inertia delays resolution [9].

While ARO provides a structured ontological framework for modeling anomaly resolution, it is not without limitations. Acts of Pseudo-Closure may terminate anomalies prematurely without true referential grounding, introducing epistemic risk when apparent resolutions mask persistent uncertainty. The Bounded Temporality model can also be sensitive to diagnostic lag, institutional biases, or conflicting social interpretations, delaying or distorting resolution. Furthermore, while alignment with BFO and CCO ensures interoperability, it may constrain adoption in domains where these foundational ontologies are not yet widely integrated. Future work may extend

ARO's formalism to better detect invalid closures and incorporate adaptive mechanisms for epistemic validation across dynamic contexts.

Not all resolutions are valid. Acts of Pseudo-Closure or Invalid Anomaly Resolution may appear to close anomalies but fail to establish determinate referential assignments. These do not qualify as Acts of Epistemic Enlightenment. Anomalies resolved through convenience remain open to reactivation. True resolution assigns determinate aboutness, replacing what is unresolved with what is determinate.

The Bounded Temporality pattern introduced here generalizes beyond anomalies to roles, procedures, legal identities, and natural phenomena, including the brief existence of a flame or the lifecycle of a star; all unfolding within bounded temporal intervals. While fully formalizing this as a reusable design pattern is beyond the scope of this paper, the structure offers a reusable model for any entity with temporally bounded existence.

Temporal boundaries alone do not suffice to make something meaningful in an ontology. Both anomalies and their resolutions require unity conditions [7], enabling entities to be treated as coherent and bounded, not just as loose fragments. This draws on Guarino and Welty's framework for individuation [7] and aligns with BFO-based modeling, where structural coherence is necessary for realist representation [10]. As the prisoner emerges from the Cave [15], what matters is not just that the shadow disappears, but that what lacked identity can now be seen, counted, and known.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT to paraphrase and reword. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] L. Darden, Reasoning in Biological Discoveries: Mechanisms, Interfield Relations, and Anomaly Resolution, Cambridge University Press, Cambridge, 2006.
- [2] C.-I. Chang, S.-S. Chiang, Anomaly detection and classification for hyperspectral imagery, IEEE Transactions on Geoscience and Remote Sensing 40 (2002) 1314–1325. doi:10.1109/TGRS.2002.800280.
- [3] J. Bittner, T. W. Coleman, J. Ravenel, Anomaly resolution ontology: Ontology accompanying the jowo 2025 paper "from shadows to resolution", <https://github.com/JJBittner/Anomaly-Resolution-Ontology>, 2025.
- [4] B. Smith, W. Ceusters, Ontological realism: A methodology for coordinated evolution of scientific ontologies, Applied Ontology 5 (2010) 139–188. doi:10.3233/AO-2010-0079.
- [5] N. Guarino, C. A. Welty, Evaluating ontological decisions with OntoClean, Communications of the ACM 45 (2002) 61–65. doi:10.1145/503124.503150.
- [6] N. Guarino, C. A. Welty, A formal ontology of properties and taxonomic relations, in: Proceedings of KR-2000 Workshop on Applications of Ontologies and Problem-Solving Methods, 2000.
- [7] N. Guarino, C. A. Welty, Identity, unity, and individuation: Towards a formal toolkit for ontological analysis, in: Proceedings of EKAW 2000, 2000, pp. 97–112.
- [8] M. Jensen, G. De Colle, S. Kindya, C. More, A. P. Cox, J. Beverley, The common core ontologies, in: Formal Ontology in Information Systems (FOIS 2024), volume 394 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2024, pp. 47–58. doi:10.3233/FAIA241292.
- [9] Iso/iec 21838-2: Information technology – top-level ontologies (tlo) – part 2: Basic formal ontology, 2021. URL: <https://www.iso.org/standard/74572.html>.
- [10] B. Smith, Against fantology, in: J. Jadacki, J. Pasniczek (Eds.), *Philosophy and Logic in Search of the Polish Tradition*, Rodopi, Amsterdam, 2005, pp. 153–170.
- [11] F. Dretske, *Knowledge and the Flow of Information*, MIT Press, Cambridge, MA, 1981.
- [12] W. Ceusters, B. Smith, Aboutness: Towards foundations for the information artifact ontology, in: W. Ceusters, B. Smith (Eds.), *Proceedings of the Sixth International Conference on Biomedical Ontology (ICBO)*, CEUR vol. 1515, 2015, pp. 1–5.
- [13] F. Dretske, Précis of Knowledge and the Flow of Information, Behavioral and Brain Sciences 6 (1983) 55–63.
- [14] V. Chandola, A. Banerjee, V. Kumar, Anomaly detection: A survey, ACM Computing Surveys 41 (2009) 15:1–15:58. doi:10.1145/1541880.1541882.
- [15] Plato, Republic, Book VII, Hackett, Indianapolis, 1992.
- [16] J. Beverley, Unreal patterns, 2025. Manuscript in preparation, University at Buffalo.
- [17] American Psychiatric Association, Diagnostic and Statistical Manual of Mental Disorders, 5 ed., American Psychiatric Publishing, Arlington, VA, 2013.

- [18] E. M. Bender, T. Gebru, A. McMillan-Major, S. Shmitchell, On the dangers of stochastic parrots: Can language models be too big?, *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT)* (2021) 610–623. doi:10.1145/3442188.3445922.
- [19] Z. Ji, Z. Lee, R. Frieske, et al., Survey of hallucination in natural language generation, *ACM Computing Surveys* (2023). doi:10.1145/3571730.
- [20] E. F. Loftus, *Eyewitness Testimony*, Harvard University Press, Cambridge, MA, 1979.
- [21] E. F. Loftus, Planting misinformation in the human mind: A 30-year investigation of the malleability of memory, *Learning & Memory* 12 (2005) 361–366. doi:10.1101/lm.94705.
- [22] M. Sherif, A study of some social factors in perception, *Archives of Psychology* 27 (1935) 1–60.
- [23] S. E. Asch, Effects of group pressure upon the modification and distortion of judgments, in: H. Guetzkow (Ed.), *Groups, Leadership and Men*, Carnegie Press, 1951, pp. 177–190.
- [24] S. Yablo, *Aboutness*, Princeton University Press, Princeton, NJ, 2014.