

# **Structural Dissolution of the Gettier Problem through Address-Theoretic Epistemology**

## **Abstract**

Edmund Gettier's 1963 paper in *Analysis* demonstrated that justified true belief is insufficient for knowledge: justification and truth can coincide accidentally. The impasse has a structural source: all parties assume justification is specifiable independently of the truth-making state of affairs. This paper introduces an alternative drawn from the Modulign Standard — a Dimensional Address Grammar for Observable Reality — replacing the justification condition with constitutive world-engagement: a classification procedure is epistemically adequate only if each step directly engages the observable features causally produced by the truth-making state of affairs, with no step epistemically isolated from those features. This yields the Address-Theoretic Non-Accidentality Principle (ATNA) and the Address-Theoretic Knowledge (ATK) formulation. Under ATK, Gettier cases cannot arise because the structural precondition — the logical independence of justification from truth — is absent from the architecture. This structural dissolution (§6) is stronger than any fourth-condition solution and differs from knowledge-first dissolution in preserving the analysis project. The argument is restricted to observational knowledge; its scope boundary is architecturally encoded.

**Keywords:** Gettier problem, justified true belief, epistemic luck, address-theoretic epistemology, structural dissolution, constitutive world-engagement, scale-sensitive epistemology.

## **1. The Problem**

Edmund Gettier's three-page paper in *Analysis* (1963) demonstrated that the tripartite analysis of knowledge as justified true belief states insufficient conditions for knowledge. Two

counterexamples share a common structure: the subject holds a justified belief in a false intermediate proposition, deduces by valid inference a second proposition that happens to be true, and satisfies all three JTB conditions. Intuitively, however, the subject does not know the second proposition, because truth and justification coincide by accident.

Case I: Smith has strong evidence that Jones will get a job and that Jones has ten coins in his pocket. He deduces: the person who will get the job has ten coins in their pocket. In fact, Smith will get the job, and Smith happens to have ten coins in his own pocket. The proposition is true, believed, and justified. Smith does not know it. His justification traces to facts about Jones; the truth traces to facts about Smith.

Case II: Smith is justified in believing that Jones owns a Ford. He deduces: either Jones owns a Ford, or Brown is in Barcelona. In fact, Jones owns no Ford, but Brown happens to be in Barcelona. Again: justified, true, believed, not knowledge. His justification traces to the false Ford-belief; the truth traces to Brown's location, of which Smith is wholly ignorant.

The philosophical content of both cases is not their artificiality but their structural revelation. They expose that in the JTB framework, justification is a property of the subject's mental or inferential states, while truth is a property of the world, and these two properties are logically independent. A justification of sufficient strength is compatible with the proposition being false; when the proposition happens to be true for independent reasons, we get a justified true belief that is, structurally, an accident.

## **2. The Landscape of Responses and Their Common Limitation**

The post-Gettier literature has pursued two strategies: amend JTB with a further condition, or abandon the analysis. Neither has succeeded. Fourth-condition proposals include the no-false-lemmas condition (Harman 1973), causal conditions (Goldman 1967), tracking conditions (Nozick 1981), defeasibility conditions (Lehrer and Paxson 1969), safety conditions (Sosa 1999; Pritchard 2005), and virtue conditions (Greco 2010; Zagzebski 1996). Each handles some cases and fails against others. Goldman's fake barns (1976) is the canonical pressure case: Henry drives through a countryside of barn façades, looks at the one genuine barn in the area, and forms the true justified perceptual belief that he sees a barn. No false lemma is involved; the barn causally produces the belief through normal perceptual channels. Yet Henry does not know he is looking at a barn. The fake barns case survives every fourth-condition proposal that permits fallible justification.

The impasse has a structural source the literature has not isolated: all parties assume justification and truth are specified along independent dimensions. Zagzebski's argument shows that any external bridge across this gap can be circumvented — but it does not establish that the gap must remain open. We argue it can be closed from the inside, by changing the architecture of the epistemic act itself.

### **3. The Structural Diagnosis**

Let  $T(P)$  denote the state of affairs that makes  $P$  true. Let  $J(S,P)$  denote the justificatory basis on which  $S$  believes  $P$ . In a Gettier case,  $J(S,P)$  and  $T(P)$  are both present but causally and epistemically unconnected:  $J(S,P)$  traces to facts about Jones;  $T(P)$  traces to facts about Smith.

J(S,P) traces to the false Ford-belief; T(P) traces to Brown's location. The structural requirement for knowledge is not merely that J(S,P) and T(P) co-occur, but that they are non-accidentally connected.

Every post-Gettier proposal attempts to specify this connection externally — by adding a condition C that constrains the relationship between J and T. But Zagzebski's recipe shows that any external C, as long as it fails to constitute T's obtaining, can always be satisfied while the connection between J and T remains accidental. The alternative is to change the architecture of the epistemic act: to make the non-accidental connection between justification and truth constitutive of the classification procedure, rather than imposed upon it externally. This is the move the Modulign architecture makes.

#### **4. The Modulign Architecture**

The Scale axis encodes which tier of physical reality the observation is sampling — from |QNT| (quantum/subatomic) through |MLC| (molecular), |BIO| (biological), |HMN| (human/built environment), |GEO| (geological/planetary), |STL| (stellar/solar system), to |CSM| (cosmological). Scale is not domain. It is an independent axis of observable reality required for the exhaustiveness condition of ATK.

The Decision Protocol is the epistemologically critical feature of the architecture. Every valid address is the output of a mandatory eight-step procedure: (1) scene parsing; (2) dominance resolution; (3) locus resolution; (4) scale assignment; (5) medium assignment; (6) suffix

assembly including Activity State, Observer annotation, and Jurisdiction; (7) conflict elimination; (8) emission with confidence score and full segment-level reasoning record. A code produced outside this protocol is explicitly marked approximate. The procedure is public, auditable, and reproducible. It is not a description of what a classifier happens to do — it is constitutive of what a valid classification is.

## **5. The Address-Theoretic Non-Accidentality Principle**

Let  $J(C,O)$  denote the justificatory basis for Modulign classification  $C$  of observation  $O$ , produced by the Decision Protocol. Let  $T(C)$  denote the state of affairs that makes  $C$  true. Within the Modulign architecture,  $J(C,O)$  is not specified independently of  $T(C)$ . The Protocol's justificatory steps directly engage the observable features causally produced by  $T(C)$ .

For a Domain classification TRN (Transportation) to be justified, the Protocol must have identified transportation infrastructure as the dominant observable context — and the dominant observable context being transportation infrastructure is precisely what constitutes  $T(TRN)$  at that node. In each case, the justification-producing procedure and the truth-making state of affairs are not separate relata connected externally. They are aspects of the same structured observation act.

**Address-Theoretic Non-Accidentality Principle (ATNA):** A Modulign classification  $C$  of observation  $O$  is non-accidentally true if and only if: (a)  $C$  was produced by the full Decision Protocol applied to the observable features of  $O$ ; (b) the Protocol's justificatory steps at each segment directly engage the observable

features causally produced by the state of affairs  $T(C)$  that makes that segment true; and (c) the Protocol contains no inference step epistemically isolated from the observable features of  $O$  — no step whose justificatory basis traces to facts other than those constituting  $T(C)$ .

*Under ATNA, there is no residual independence of  $J(C,O)$  from  $T(C)$ . There is no gap into which Gettier luck can insert itself.*

Constitutive engagement means more than reliable causal connection, and the distinction is the load-bearing difference between ATK and reliabilism. A thermometer reliably causally covaries with temperature but does not constitutively engage it: the thermometer's justificatory basis is its past reliability across a population of cases, not the temperature's current causal manifestation at the moment of measurement. The Protocol's justificatory basis at each step is the observable feature itself — the dominant context that constitutes the domain, the geo-evidence that constitutes the locus, the physical scale that constitutes the scale tier. This is not a claim that the Protocol is infallible. It is a claim about the structure of its justification: each step derives its justification from engaging the observable features causally produced by  $T(C)$  on this occasion, not from the Protocol's historical accuracy across occasions. It is this structural property — not reliability — that blocks the Gettier gap.

Condition (c) is load-bearing. In Gettier's Case I, Smith's inference from the false proposition that Jones will get the job is an inference step whose justificatory basis traces to facts about Jones, while  $T(e)$  is constituted by facts about Smith. Condition (c) is violated: the inferential step is epistemically isolated from the truth-making state of affairs. In Case II, Smith's

disjunction inference is justified by his belief about Jones' Ford ownership, while  $T(h)$  is constituted by Brown's location. Condition (c) is violated again. In each case the violation is identifiable at the level of the reasoning record.

The fake barns case requires separate analysis because it involves no false lemma and no inferential isolation in the standard sense. The real barn causally produces Henry's belief through normal visual channels. The failure is one of exhaustiveness. Condition (b) of ATNA requires that the Protocol engage the observable features causally produced by  $T(C)$  — and in a region of barn façades,  $T('that is a barn')$  is partly constituted by the non-prevalence of fake barns in the immediate environment. The spatial distribution of structures at regional scale is causally relevant to the truth-making conditions for knowledge-grade barn-identification in that environment. Henry's visual glance does not engage this feature. The Protocol's exhaustiveness requirement blocks the fake barns case through condition (b).

## **6. The Structural Dissolution**

**6.1 Dissolution versus solution.** A solution to the Gettier problem shows that Gettier cases, though possible within a framework, are caught by a further condition  $X$ . A dissolution shows that Gettier cases cannot arise within a framework because the structural precondition for their arising is absent. The distinction is not terminological. A solution leaves the JTB architecture intact and adds  $X$ ; Zagzebski's recipe then applies to the  $X$ -augmented analysis. A dissolution changes the architecture such that the recipe cannot be instantiated. The Modulign architecture removes the gap:  $J(C,O)$  is produced by engaging  $T(C)$ 's observable manifestations, so  $J$  and  $T$

are not separate relata requiring a bridge. They are internally related through the structure of the observation procedure.

**6.2 The Zagzebski recipe applied directly.** The recipe requires, as its first step, a classification that satisfies the ATK conditions but is false. For this to hold, the Decision Protocol must have been followed correctly at every step and still produced a false classification. This is possible in exactly two ways. First, performance error: the classifier incorrectly identified the observable features. A performance error is not a Gettier case. It is identifiable in the reasoning record, auditable through independent re-application of the Protocol to the same observation, and flagged by a lower confidence score. The distinction between a Gettier gap (structural: justification present, connected to truth accidentally) and a performance error (procedural: justification incomplete, detectable) is architecturally explicit.

Second, perceptual deception: the observable features themselves were systematically misleading. This is the problem of epistemic access to the world — a genuine epistemological problem, but not Gettier's problem. Gettier cases arise when observable features reliably represent the world but the resulting belief is accidentally true. Perceptual deception cases arise when observable features do not reliably represent the world. No existing epistemological framework solves the problem of radical perceptual deception; ATK is in this respect no weaker than any competitor.

The Zagzebski recipe therefore cannot be instantiated against ATK. Recipe step one requires a justified false classification — where  $J(C,O)$  is satisfied but the conclusion is false and the

connection to T(C) is accidental. Under ATNA, a classification satisfying (a)–(c) is non-accidentally true by definition. A classification that is false but produced by the Protocol in good faith is a performance error, not a justified false classification in the sense the recipe requires. The recipe's entry point does not exist in the architecture.

## 7. The Address-Theoretic Knowledge Formulation

**Address-Theoretic Knowledge (ATK):** S knows that P if and only if: (i) P is true; (ii) S believes that P; and (iii) S's belief that P is the output of a classification procedure that is: (a) constitutively world-engaged — each step directly engages the observable features causally produced by T(P), with no step epistemically isolated from those features; (b) transparent — its steps are explicit and auditable; (c) intersubjective — reproducible by any competent observer with access to the same observable features; (d) exhaustive — covering all segments of the reality-space relevant to P; and (e) scale-sensitive — operating at the scale at which T(P) is constituted.

*Condition (iii) replaces the traditional justification condition. Conditions (a)–(e) specify what it means for a procedure to be epistemically adequate for the observational domain.*

Condition (iii)(a) is the core replacement of traditional justification. It makes the non-accidental connection between J and T constitutive of the epistemic act rather than imposed externally.

Conditions (iii)(b)–(d) operationalise what constitutive world-engagement requires: transparency

ensures the reasoning can be audited; intersubjectivity ensures the procedure is not idiosyncratic; exhaustiveness ensures no epistemically relevant feature is overlooked. These are not additional conditions bolted onto JTB; they are specifications of what it means for a procedure to satisfy (iii)(a).

Condition (iii)(e), scale-sensitivity, generalises ATK beyond human-scale perceptual cases to the full range of scientific observation. The fake barns case is handled directly: Henry's visual procedure fails (iii)(d) and (iii)(e). Gettier's original cases fail (iii)(a): Smith's inferential steps are epistemically isolated from T(P).

## 8. Objections

**8.1 The protocol is still applied by a subject.** Even if the Protocol is constitutively world-engaged, it is applied by a subject who can apply it incorrectly or face deceptive perceptual input. Does this restore the Gettier gap? No, and the reason is precise. Incorrect application is a performance error, not a Gettier gap. A performance error is detectable in the reasoning record, auditable through independent re-application, and flagged by the confidence score and the ^CONT mechanism. It is procedural failure with a defined remediation path. Gettier gaps are structural: justification is present, correctly formed, and connected to truth accidentally. These are categorially different failure modes. Systematic perceptual deception — the problem of radical epistemic access — is a genuine problem no framework has solved. ATK neither solves it nor pretends to.

A further objection holds that we have merely relocated the Gettier problem one level up: the Protocol is itself an epistemic procedure, and one might ask whether the Protocol's design was arrived at by a Gettier-susceptible epistemic act. This objection has force against externalist reliability accounts, since the regress question — what makes the mechanism reliable? — is always available. ATK is not a reliability account. The Protocol's adequacy is assessed by whether its design satisfies the ATK conditions themselves — conditions on the structure of each procedure application, not on output reliability across a population of cases. The evaluation is self-applicable: one asks whether the philosophical evaluation of the Protocol's design constitutively engages the observable features of that design, transparently, intersubjectively, and exhaustively. This produces no vicious regress; it produces the ordinary structure of philosophical argument under public scrutiny. The regress that threatens externalist reliabilism does not arise here.

**8.2 Scope: a priori knowledge.** ATK does not apply to a priori knowledge. By observational knowledge we mean knowledge whose truth-making state of affairs produces detectable physical signals accessible to an observer — this includes all instrument-mediated and recorded evidence but excludes mathematical truths, logical relations, and first-person conscious experience. Mathematical truths and logical relations have no observable manifestations in the sense required by condition (iii)(a). This is an acknowledged scope limitation, not a disqualifying one. The correct scope claim is: the Gettier problem for observational, empirical knowledge — perceptual knowledge, scientific knowledge, historical and evidentiary knowledge derived from records and observation — is structurally dissolved. This domain is precisely where Gettier cases have been

most troubling and most practically significant. The Modulign architecture makes this boundary architecturally explicit.

**8.3 Jurisdiction.** ATK's evidentiary implications are argued at the structural level. The claim that the Decision Protocol satisfies authentication, chain of custody, best-evidence, and expert-foundation requirements is advanced as a formal claim about the architecture's properties, not as a positive assertion that any particular jurisdiction has adopted Modulign certification. The formal argument applies wherever those structural concerns apply. Mapping the architecture to specific jurisdictional requirements is the work of the legal implementation layer, which this paper does not prescribe.

## **9. Situating ATK**

The most important theoretical comparison is with process reliabilism (Goldman 1979). ATK is not a form of reliabilism. Process reliabilism evaluates epistemic procedures by their output reliability across a population of cases. ATK's constitutive world-engagement condition makes no reference to output frequency or population statistics. It is a structural requirement on each individual procedure application: the justificatory basis of each step must trace constitutively to the observable features causally produced by T(P) for that specific classification of that specific observation. A procedure that reliably produces true beliefs because it typically engages the right features satisfies reliabilism but need not satisfy ATK. This is why ATK escapes the Gettier problem while reliabilism does not: a reliable process can still produce a belief whose truth is accidental relative to that particular application — exactly what fake barns demonstrates. ATK's constitutive engagement condition blocks accidental truth at the level of the individual procedure application, not by appeal to aggregate output statistics.

ATK entails safety as a downstream consequence: in nearby possible worlds where the same Protocol is applied to the same observable features, the classification tracks the same truth-making conditions. Safety follows from the architecture; it is not an additional constraint. Conditions (iii)(b)–(e) formalise what epistemic virtue requires of an observation procedure: transparency, intersubjectivity, exhaustiveness, and scale-sensitivity. Unlike knowledge-first epistemology, ATK completes rather than forecloses the analysis project.

## **10. Conclusion**

Gettier's paper is correct. JTB is insufficient for knowledge. The gap is real. Sixty years of subsequent literature has not closed it because all proposed closures have been external bridges across a gap whose existence is presupposed by the architecture. Zagzebski establishes that external bridges, given fallibilism, can always be circumvented.

The Modulign architecture closes the gap from inside: by replacing the traditional justification condition with a requirement of constitutive world-engagement, it eliminates the logical independence of  $J(C,O)$  from  $T(C)$  that Gettier exploits. Under the Address-Theoretic Knowledge formulation, there is no residual independence to exploit. Gettier cases cannot arise not because they are caught by a further condition but because the structural precondition for their arising is absent from the architecture.

This is structural dissolution: stronger than any fourth-condition solution, and different from knowledge-first dissolution in that it preserves and completes the analysis project. The result extends to the full range of observational knowledge from human-scale perception through

cosmological observation. Its scope boundary at a priori knowledge is architecturally encoded rather than asserted in prose.

The Modulign Standard was developed as a practical classification system for observational data. Its epistemological properties emerged in the course of formalising its theoretical foundations rather than being constructed to address the Gettier problem — an accidental discovery in the course of technical work rather than a constructed response to the problem.

## References

- Gettier, E. L. 1963. Is justified true belief knowledge? *Analysis* 23(6): 121–23.
- Goldman, A. I. 1967. A causal theory of knowing. *Journal of Philosophy* 64(12): 357–72.
- Goldman, A. I. 1976. Discrimination and perceptual knowledge. *Journal of Philosophy* 73(20): 771–91.
- Goldman, A. I. 1979. What is justified belief? In G. Pappas (ed.), *Justification and Knowledge*. D. Reidel.
- [Author]. 2026. Modulign Standard v3.0: A dimensional address grammar for observable reality. Open technical specification. Available on request.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge University Press.
- Harman, G. 1973. *Thought*. Princeton University Press.
- Lehrer, K. and Paxson, T. 1969. Knowledge: undefeated justified true belief. *Journal of Philosophy* 66(8): 225–37.
- Nozick, R. 1981. *Philosophical Explanations*. Harvard University Press.

- Pritchard, D. 2005. *Epistemic Luck*. Oxford University Press.
- Sosa, E. 1999. How to defeat opposition to Moore. *Philosophical Perspectives* 13: 141–53.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford University Press.
- Zagzebski, L. 1994. The inescapability of Gettier problems. *Philosophical Quarterly* 44(174): 65–73.
- Zagzebski, L. 1996. *Virtues of the Mind*. Cambridge University Press.

**[Author note — title page only, not in manuscript file]**

The Modulign Standard referenced in this paper is an open specification developed for the practical classification of observational data. Its philosophical implications for the Gettier problem emerged in the course of formalising its epistemological foundations — an accidental discovery in the course of technical work rather than a constructed response to the problem. The author is the originator of the Modulign Standard and holds the associated trademarks. The Standard is published as an open specification free to implement without licence or fee; the trademark covers the name and prefix only. No institutional affiliation. Correspondence: [withheld for review].