Negation Before Logic: Inferential Fields, Orientation Reversal, and the Geometry of Admissible Inference

Flyxion

Independent Researcher

June 2026

Abstract

The cognitive cost of negation has been studied for decades, yet the dominant accounts remain unsatisfying. Operator-counting models predict that processing cost should increase monotonically with the number of negative expressions; they do not explain why two negations can be easier than one, why cost tracks monotonicity domains rather than operator occurrence, or why negated representations retain a processing signature long after the sentence has been parsed. Polarity-based accounts handle some scalar phenomena but conflate orientation toward zero with reversal of inferential direction, which are demonstrably distinct costs.

This paper proposes a geometric alternative. We introduce the notion of an *inferential field* $\mathcal{F} = (X, R, \omega, A)$, where X is a space of distinctions, R is an inferential reachability relation, $\omega : X \to [-1, +1]$ is a continuous orientation field, and A is an admissibility structure. Within this framework, negation is not fundamentally a truth-value operator. It is an orientation transformation $N : \mathcal{F} \to \mathcal{F}$ satisfying $N(\omega) = -\omega$. Monotonicity becomes a property of the orientation field rather than of individual lexical items. Upward-entailing environments correspond to regions where $\omega > 0$ and inferential reachability expands; downward-entailing environments correspond to regions where $\omega < 0$ and reachability contracts. Processing cost is not determined by operator count but by two separable quantities: the magnitude of orientation reversal and the curvature of admissibility domain boundaries.

This decomposition yields the *Orientation-Reversal Principle*: the cognitive cost of a semantic operator is a function of the induced distortion of inferential orientation and the local curvature of admissibility domains, not of operator count. Four bodies of experimental evidence follow as consequences: the non-additivity of double negation, the domain-sensitivity of monotonicity processing cost, the sustained effect of negation in delayed verification, and the licensing

conditions of negative polarity items. We further show that negative polarity item licensing is an admissibility probe: an expression such as *ever* or *any* succeeds not because a negative word is present but because the local region of the inferential field carries the required orientation and satisfies the admissibility condition jointly. Neural dissociation evidence from logical negation studies constrains the framework empirically, supporting a four-stage functional architecture in which syntactic field construction, logical field deformation, and scenario verification are partially separable operations.

The framework connects to a broader research program in admissibility geometry and distinguishability theory, in which the same mathematical structure — distinctions, reachability, orientation, admissibility — recurs across language, memory, repair, and scientific explanation. Negation, on this view, is not an addition to logic. It is a transformation of the distinction space within which logic becomes possible.

1 The Operator-Count Failure

The simplest theory of negation processing is also the most natural one. If a sentence contains a negative expression, it is harder to process than the corresponding positive sentence. If it contains two negative expressions, it is harder still. Processing cost increases monotonically with the number of negative operators encountered during parsing. Call this the *operator-count hypothesis*.

The operator-count hypothesis has the virtue of simplicity and the vice of being wrong in precisely the cases that matter most theoretically. Three bodies of experimental evidence establish its failure, and the pattern of failure is informative: it is not random noise around a correct trend but a systematic deviation that points toward a different underlying variable.

1.1 The Double-Negation Asymmetry

If processing cost tracked operator count, then a sentence containing two negative expressions should be more costly than a sentence containing one. The prediction is monotone: each additional negation adds a fixed increment of difficulty. The experimental literature on double negation does not support this prediction.

Sentences of the form *it is not the case that none of the circles are red* involve two scope-taking negative expressions and require the reasoner to compute a positive conclusion by composing two reversals. Under operator-count, this double-reversal configuration should be the most expensive. Under an orientation account, the two reversals cancel: $N(N(\omega)) = \omega$, returning the field to its original orientation. The field does not count operators; it tracks net orientation. The prediction from an orientation account is therefore that double negation should not cost twice as much as single negation, and may in some configurations cost less, because the inferential field has returned to a familiar expansion regime.

This is what the experimental record shows. The cost of double negation is not additive. In some participant populations and task configurations, sentences with two negations are verified faster than sentences with one, when syntactic complexity is controlled. Operator-count has no account of this. Orientation reversal explains it directly.

1.2 Domain Sensitivity of Monotonicity Cost

A more precise failure of operator-count comes from experiments that manipulate monotonicity independently of the number of negative expressions. Consider the contrast between *more than half* and *fewer than half*. Both are quantifiers. Neither is a negative word in the lexical sense. Yet they induce different entailment profiles: *more than half* is upward-entailing in its restrictor, while *fewer than half* is downward-entailing. An operator-count model, applied to these expressions, predicts no difference in processing cost, since no negative operator has been added or removed.

The experimental evidence contradicts this prediction. Downward-entailing quantifiers of this kind impose a measurable processing cost relative to their upward-entailing counterparts, even when the sentences are matched for length, frequency, syntactic structure, and the absence of any overt negative morpheme. The source of the cost is the direction of inferential expansion, not the presence of a negative word.

More decisive still is the finding from the 2024 monotonicity-domain experiments. In those studies, the relevant experimental variable is not the number of downward-entailing operators but the structure of the monotonicity domain: whether a given expression occurs within a region whose inferential orientation is uniformly downward, or near a boundary where orientation shifts. The finding is that processing cost tracks domain structure. Two configurations with equal operator counts but different domain structures produce different costs. Two configurations with unequal operator counts but equivalent domain structures produce equivalent costs. This result cannot be accommodated by operator-count in any form. It requires a variable that tracks the geometry of inferential regions, not the inventory of lexical items.

1.3 The Persistence of Negation in Delayed Verification

A third failure of operator-count concerns the temporal profile of negation effects. The operator-count hypothesis is naturally interpreted as a parsing hypothesis: negative operators introduce additional computational steps during incremental sentence processing, and those steps consume time. On this interpretation, if sufficient processing time is provided before a verification judgment is required, the cost should diminish or disappear. The parser has had time to complete its work.

Delayed verification experiments test this prediction by introducing a temporal gap between sentence presentation and the onset of the verification image. The operator-count prediction is that negation effects should attenuate with delay. The experimental finding is the opposite: negation effects are *sustained* across delays. A negated sentence does not become easier to verify simply because more time has passed. The cost persists into the verification phase even when parsing has long concluded.

This result is incompatible with a purely incremental parsing account of negation cost. It suggests that negation does not merely insert additional steps into the parse; it changes the structure of the working-memory representation that is passed forward to verification. In the framework developed here, this is the *persistence cost*: the residual distortion of the inferential field that remains in the representation after processing. Negation changes the orientation of the stored representation, and orientation is not automatically restored by the passage of time.

1.4 What the Failures Share

The three failures of operator-count share a common structure. In each case, the quantity that predicts processing cost is not the number of negative expressions but a property of the inferential environment those expressions create: the net orientation of the field after all reversals have been applied, the geometry of domain boundaries, and the persistence of orientation change in the stored representation. These are properties of a field, not properties of a count. The theory required to explain them must be a field theory.

2 Inferential Fields

The failures documented in the preceding section converge on a single diagnostic: the relevant variable is not a count of operators but a property of the inferential environment those operators construct. To make this precise, we need a mathematical object that can represent inferential environments, support a notion of orientation, and admit a well-defined transformation corresponding to negation. We introduce this object here.

2.1 Distinctions and Reachability

Let X be a set of *distinctions*. Intuitively, a distinction is a difference that the inferential system can represent and act upon: the difference between red and not-red, between all and some, between the circle being inside the square and outside it. We do not require X to carry propositional structure at the outset. Propositions, on the view developed here, are derivative objects constructed from distinctions; the geometry of the distinction space is prior.

Definition 2.1 (Inferential reachability)**.** *A binary relation $R \subseteq X \times X$ is an* inferential reachability relation *on X if it satisfies the following conditions:*

(i) Reflexivity*: $R(x, x)$ for all $x \in X$.*

(ii) Transitivity*: if $R(x, y)$ and $R(y, z)$ then $R(x, z)$.*

(iii) Non-triviality*: there exist $x, y \in X$ such that $R(x, y)$ and $\neg R(y, x)$.*

Condition (iii) ensures that R is not the trivial equivalence relation. Inferential motion has direction: some distinctions are reachable from others without the converse holding. This directedness is what makes orientation a meaningful concept in the space.

For a subset $A \subseteq X$, define the *inferential expansion* of A under R as

$$R(A) = \{y \in X : \exists\, x \in A,\ R(x, y)\}.$$

This is the set of distinctions reachable from any member of A in one inferential step. The behavior of $R(A)$ as A varies is what monotonicity theory studies, and it is where orientation enters.

2.2 The Orientation Field

Classical monotonicity theory classifies expressions as upward- or downward-entailing and treats this classification as a binary, lexically determined property. The operator-count failures suggest that this classification is insufficient: what matters is not the sign of an individual expression but the orientation of the inferential environment that expression inhabits. We represent this environment with a continuous field.

Definition 2.2 (Orientation field)**.** *An* orientation field *on* (X, R) *is a continuous function* $\omega : X \to [-1, +1]$ *satisfying:*

(i) Upward regime*: if* $\omega(x) > 0$ *and* $A \subseteq B \subseteq X$*, then* $R(A) \subseteq R(B)$*, so inferential expansion preserves set inclusion.*

(ii) Downward regime*: if* $\omega(x) < 0$ *and* $A \subseteq B \subseteq X$*, then* $R(B) \subseteq R(A)$*, so inferential expansion reverses set inclusion.*

(iii) Boundary*: if* $\omega(x) = 0$*, then* x *lies on the boundary between upward and downward regimes, and the local expansion behavior of* R *is indeterminate.*

The choice of $[-1, +1]$ rather than $\{-1, +1\}$ is deliberate. It allows ω to vary continuously across the distinction space, so that regions near orientation boundaries are represented as having intermediate orientation values rather than being forced into a binary classification. This continuity is what permits a well-defined notion of orientation curvature in the next subsection.

2.3 Admissibility

Not every distinction in X is available for inferential use at every point in processing. Licensing conditions, syntactic constraints, and contextual restrictions determine which regions of the distinction space are open to a given inferential operation. We represent this with an admissibility structure.

Definition 2.3 (Admissibility structure)**.** *An* admissibility structure *on* (X, R, ω) *is a function* $A : X \to \{0, 1\}$ *such that* $A(x) = 1$ *if the distinction* x *is currently admissible for inferential use and* $A(x) = 0$ *otherwise. An* admissible region *is a connected component of* $A^{-1}(1)$*.*

Admissibility is not truth. A distinction may be admissible without being verified, and verified without being admissible in the relevant inferential context. The separation between admissibility and truth is essential to the account of negative polarity items developed in Section 4: an NPI is sensitive to admissibility conditions, not directly to truth conditions.

Definition 2.4 (Inferential field)**.** *An* inferential field *is a quadruple* $\mathcal{F} = (X, R, \omega, A)$ *where* X *is a set of distinctions,* R *is an inferential reachability relation on* X*,* $\omega : X \to [-1, +1]$ *is an orientation field, and* $A : X \to \{0, 1\}$ *is an admissibility structure.*

2.4 Negation as Orientation Transformation

With the inferential field in hand, we can give negation a precise geometric meaning.

Definition 2.5 (Negation operator)**.** *The* negation operator $N : \mathcal{F} \to \mathcal{F}$ *is the map that acts on the orientation field by pointwise reversal:*

$$N(\omega)(x) = -\omega(x) \quad \textit{for all } x \in X,$$

while leaving X*,* R*, and* A *unchanged.*

Several consequences follow immediately.

Proposition 2.6 (Double negation)**.** $N \circ N = \mathrm{id}_{\mathcal{F}}$*.*

Proof. $(N \circ N)(\omega)(x) = N(-\omega(x)) = -(-\omega(x)) = \omega(x)$ for all $x \in X$. □

Proposition 2.7 (Orientation of upward and downward environments)**.** *A linguistic environment is upward-entailing if and only if the corresponding region of* $\mathcal{F}$ *satisfies* $\omega > 0$*. It is downward-entailing if and only if the corresponding region satisfies* $\omega < 0$*.*

Proof. By conditions (i) and (ii) of Definition 2.2, the expansion behavior of R under set inclusion is determined by the sign of ω. Upward entailment is defined by the preservation of set inclusion under inferential expansion; downward entailment by its reversal. These correspond exactly to the upward and downward regimes of ω. □

Proposition 2.6 is already theoretically significant. It establishes that the field-theoretic account of negation predicts the cancelation of double negation as a structural theorem, not as an empirical stipulation. The field does not count negative operators; it tracks the net orientation that results from their composition. Two reversals return the field to its original state. The non-additivity of double-negation cost documented in the experimental literature is therefore a consequence of the geometry, not a puzzle requiring separate explanation.

2.5 Orientation Curvature and Processing Cost

The domain-sensitivity result from the 2024 monotonicity experiments requires a further quantity: not merely the sign of ω at a point but the rate at which ω changes across inferential regions. We define this as orientation curvature.

Definition 2.8 (Orientation curvature)**.** *The* orientation curvature *at $x \in X$ is*

$$\kappa_\omega(x) = |\nabla\omega(x)|\,,$$

where ∇ denotes the gradient with respect to the topology on X induced by R.

Intuitively, $\kappa_\omega(x)$ measures how rapidly the inferential orientation changes in the neighborhood of x. In a uniformly upward-entailing region, ω is constant and positive, so $\kappa_\omega = 0$. In a uniformly downward-entailing region, ω is constant and negative, so $\kappa_\omega = 0$ again. High curvature occurs at and near orientation boundaries, where ω transitions between positive and negative values.

Definition 2.9 (Processing cost)**.** *The* processing cost *at $x \in X$ is*

$$C(x) = \alpha \cdot \mathbf{1}[\omega(x) < 0] + \beta \cdot \kappa_\omega(x),$$

where $\alpha, \beta > 0$ are empirically determined weights, $\mathbf{1}[\omega(x) < 0]$ is the indicator function for the downward regime, and $\kappa_\omega(x)$ is orientation curvature.

The first term captures the cost of operating in a reversed inferential regime: downward-entailing environments impose a baseline cost because the ordinary expansion direction of inference is no longer available. The second term captures the additional cost of processing near domain boundaries, where the local orientation is changing and neither a pure upward nor a pure downward inference strategy is applicable.

This cost function separates three empirically distinct quantities that previous accounts have conflated:

1. *Polarity cost*: the cost associated with scalar negativity, which in the present account is absorbed into the admissibility structure A rather than ω directly, since polarity is a property of the scalar position of an expression rather than the orientation of its inferential environment.

2. *Monotonicity cost*: captured by the first term, the baseline cost of operating in a downward-entailing region regardless of how that region was constructed.

3. *Boundary cost*: captured by the second term, the additional cost of processing near orientation domain boundaries, which is the quantity most directly measured by the 2024 monotonicity-domain experiments.

The persistence cost identified in delayed-verification studies does not appear in $C(x)$ directly, because persistence is a property of the representation passed forward to the verification stage rather than a cost incurred during field navigation. We address persistence separately in Section 5.

2.6 Negation Distortion

The final quantity introduced in this section connects the field-theoretic account to the broader program in distinguishability geometry. In that program, a transformation is characterized not only by its action on objects but by the distortion it introduces in the reachability structure of those objects. We define the analogous quantity for negation.

Definition 2.10 (Negation distortion)**.** *The* negation distortion *at $x \in X$ is*

$$D_N(x) = d_R(x, N(x)),$$

where d_R is the reachability pseudometric on X defined by

$$d_R(x, y) = \inf\{n \in \mathbb{N} : R^n(x) \ni y \text{ and } R^n(y) \ni x\},$$

with $d_R(x, y) = \infty$ if no such n exists.

Negation distortion measures how far a distinction x is displaced in the reachability structure of X when the orientation of the field is reversed. A distinction that remains easily reachable from its negated counterpart has low distortion; a distinction that becomes unreachable or reaches only after many inferential steps has high distortion.

The global distortion of the negation operator is then

$$\|D_N\| = \int_X D_N(x)\, d\mu(x),$$

where μ is a suitable measure on X reflecting the prior probability or salience of distinctions. This global quantity is what the overall processing cost of negation in a given inferential context approximates. The three cost components now have a unified representation: polarity cost, monotonicity cost, and boundary cost all contribute to $\|D_N\|$ through their respective effects on d_R. Persistence cost is the residual of D_N that has not been resolved by the time the representation is passed to the verification stage.

3 The Orientation-Reversal Principle

The formal machinery of Section 2 permits a precise statement of the central theoretical claim of this paper. We state it as a principle rather than a theorem because its full justification draws on both the mathematical structure of inferential fields and the empirical record reviewed in Section 1. The principle unifies the four experimental phenomena — double-negation non-additivity, domain-sensitive monotonicity cost, sustained verification effects, and negative polarity item licensing — as consequences of a single underlying geometric quantity.

Theorem 3.1 (Orientation-Reversal Principle). *Let $\mathcal{F} = (X, R, \omega, A)$ be an inferential field and let $T : \mathcal{F} \to \mathcal{F}$ be a semantic operator. The cognitive processing cost of T is determined not by the count of negative expressions introduced by T but by two separable quantities:*

(i) the magnitude of net orientation reversal induced by T, measured by $\int_X |\omega(x) - T(\omega)(x)|\, d\mu(x)$, and

(ii) the curvature of admissibility domain boundaries in the region of $\mathcal{F}$ traversed during processing, measured by $\int_X \kappa_\omega(x)\, d\mu(x)$ over the traversed region.

Proof sketch. We proceed by showing that operator count is neither necessary nor sufficient to determine processing cost, and that quantities (i) and (ii) are jointly sufficient within the field-theoretic framework.

Operator count is not sufficient. By the double-negation result (Proposition 2.6), two orientation reversals compose to the identity transformation. A sentence containing two downward-entailing operators therefore induces a net orientation change of zero, since $N \circ N = \text{id}$. Under operator count, the predicted cost is twice the single-negation cost. Under the orientation-reversal measure (i), the predicted cost is zero for the orientation component, plus boundary costs incurred during the traversal. The experimental record supports the orientation prediction over the operator-count prediction: double-negation cost is non-additive and can fall below single-negation cost under controlled conditions.

Operator count is not necessary. By Proposition 2.7, downward-entailing environments correspond to regions where $\omega < 0$, regardless of whether a negative morpheme is present. Quantifiers such as *fewer than half* induce downward orientation without containing any negative word. The processing cost associated with such quantifiers — documented in the monotonicity literature — is therefore a cost of orientation, not a cost of operator occurrence. Operator count, which registers zero for these expressions, fails to predict the observed cost.

Quantities (i) and (ii) are jointly sufficient. The cost function $C(x) = \alpha \cdot \mathbf{1}[\omega(x) < 0] + \beta \cdot \kappa_\omega(x)$ captures the baseline cost of reversed orientation and the additional cost of domain boundaries respectively. Integrating over the traversed region yields a total cost that depends on the geometry of the inferential path, not on the inventory of operators along it. This accounts for: the non-additivity of double negation (net reversal is zero); the domain-sensitivity of monotonicity cost (cost tracks boundary curvature, not operator occurrence); and the differential cost of downward-entailing quantifiers without negative morphemes (orientation cost with zero operator count). The persistence effect requires an additional claim about the representation passed to verification, which is addressed in Section 5, but is consistent with the present framework since $D_N(x) > 0$ whenever orientation has been reversed and the distortion has not been resolved. □

3.1 Four Phenomena as Consequences

We now show explicitly how the four experimental phenomena documented in Section 1 follow from the Orientation-Reversal Principle.

3.1.1 Double-Negation Non-Additivity

Under the operator-count hypothesis, the cost of a doubly negated sentence S^{--} satisfies $C(S^{--}) = 2 \cdot C(S^{-})$, where S^{-} is the singly negated sentence. Under the Orientation-Reversal Principle, the relevant quantity is the net orientation after composition of operators. Since $N \circ N = \text{id}$, the net reversal induced by S^{--} is zero, and the orientation component of cost vanishes. The [Oresidual cost is purely the boundary cost incurred during the processing of the intermediate state, which lies at the orientation boundary $\omega = 0$. This boundary cost is real — the system does pass through a region of high curvature — but it is bounded above by the single boundary traversal, not multiplied by two.

The empirical prediction is therefore:

$$C(S^{--}) < 2 \cdot C(S^{-}),$$

and in populations or task configurations where boundary traversal cost is low relative to orientation cost, the stronger prediction

$$C(S^{--}) < C(S^{-})$$

may hold. Both predictions are supported by the experimental record and both are consequences of the geometry.

3.1.2 Domain-Sensitive Monotonicity Cost

The 2024 monotonicity-domain experiments show that processing cost tracks domain structure rather than operator occurrence. Under the Orientation-Reversal Principle, this is expected. The cost function $C(x)$ depends on $\omega(x)$ and $\kappa_\omega(x)$, both of which are properties of the inferential region in which processing occurs, not of the count of operators that constructed that region. Two configurations with equal operator counts but different domain structures will have different values of κ_ω and therefore different costs. Two configurations with unequal operator counts but equivalent domain structures will have equivalent costs.

More precisely, consider two sentences S_1 and S_2 such that:

- S_1 contains one downward-entailing operator and processing occurs uniformly within a single downward-entailing domain.
- S_2 contains two downward-entailing operators but processing traverses a boundary between an upward and a downward domain.

Operator count predicts $C(S_2) > C(S_1)$. The Orientation-Reversal Principle predicts $C(S_2) > C(S_1)$ only if the boundary cost of S_2 exceeds the difference in baseline orientation cost, which depends on the specific geometry of the domains. The empirical result — that S_2 can be more costly than predicted by its operator count and S_1 less costly — follows naturally from the curvature term $\beta \cdot \kappa_\omega(x)$.

3.1.3 Sustained Verification Effects

The delayed-verification result presents a challenge to any account that treats negation cost as a purely incremental parsing phenomenon. Under the Orientation-Reversal Principle, the sustained effect has a natural explanation via negation distortion. When a negated sentence is processed, the orientation of the resulting representation is reversed relative to the corresponding positive sentence. This reversed representation is stored in working memory and passed to the verification stage. The distortion $D_N(x) = d_R(x, N(x))$ quantifies how far the representation has been displaced in reachability terms. This distortion does not automatically resolve with the passage of time; it persists in the representation until the verification process has explicitly re-oriented the field with respect to the scenario being compared.

The prediction is that the sustained cost should correlate with the magnitude of D_N for the relevant distinctions, and that it should diminish only when the verification process has completed the reachability comparison — not merely when parsing time has elapsed.

3.1.4 Negative Polarity Item Licensing

The analysis of negative polarity item licensing requires the admissibility structure A of the inferential field and is developed in full in Section 4. Here we note only the consequence of the Orientation-Reversal Principle that is relevant to licensing: an NPI is not sensitive to the presence of a negative operator but to the orientation and admissibility of the local inferential region. Two environments that are equally negative in operator count but differ in orientation or admissibility will differ in NPI licensing. This prediction separates the present account from both standard downward-entailment accounts, which license NPIs by operator type, and from polarity accounts, which license them by scalar position.

3.2 The Principle as Unification

The Orientation-Reversal Principle is not merely a restatement of the four phenomena in geometric language. It provides a single quantity — the distortion of inferential orientation and its curvature across admissibility domains — from which all four phenomena are derivable. This is a stronger claim than a descriptive unification. It implies that the four phenomena are not independently explained by four separate mechanisms but are four measurements of the same underlying geometric variable.

Experiments designed to manipulate one of the four phenomena while holding the geometric variable constant should attenuate or eliminate the effect. Experiments designed to manipulate the geometric variable directly — by constructing sentences that traverse domain boundaries without using negative morphemes, for instance — should produce costs not predicted by any account that relies on negative operator occurrence.

These are testable predictions. We return to them in Section 6.

4 Negative Polarity Items as Admissibility Probes

Negative polarity items present the licensing problem in its sharpest form. An expression such as *ever*, *any*, *yet*, or *lift a finger* is grammatical in some environments and ungrammatical in others, and the distribution of grammatical environments does not reduce to the mere presence of a negative word. The sentences in (1) illustrate the basic pattern.

1. (a) Nobody has ever been to that island. [grammatical]
   (b) Somebody has ever been to that island. [ungrammatical]
   (c) Few students have ever read that book. [grammatical]
   (d) Many students have ever read that book. [ungrammatical]
   (e) I doubt that she has ever been there. [grammatical]
   (f) I believe that she has ever been there. [ungrammatical]

The standard generalization, due to Ladusaw, is that NPIs are licensed in downward-entailing environments. *Nobody*, *few*, and the complement of *doubt* are downward-entailing; *somebody*, *many*, and the complement of *believe* are not. This generalization is largely correct as a descriptive matter, but it leaves two theoretical questions unanswered. First, why should downward entailment be the licensing condition? What is the connection between inference reversal and the grammaticality of *ever*? Second, how does the licensing condition interact with syntactic scope and domain structure, given the 2024 finding that processing cost tracks domains rather than operators?

The inferential field framework answers both questions by recasting NPI licensing as an admissibility probe.

4.1 Admissibility and Orientation as Independent Conditions

The key theoretical move is to separate two conditions that previous accounts have run together. In the standard downward-entailment account, a single condition — the DE-ness of the licensing environment — determines grammaticality. In the field-theoretic account, two conditions must be jointly satisfied.

Definition 4.1 (NPI probe)**.** *An* NPI probe *is a function $P : X \to \{0,1\}$ defined by*

$$P(x) = \begin{cases} 1 & \text{if } A(x) = 1 \text{ and } \omega(x) < 0, \\ 0 & \text{otherwise.} \end{cases}$$

An NPI expression is licensed at position x in an inferential field $\mathcal{F} = (X, R, \omega, A)$ if and only if $P(x) = 1$.

The two conditions are:

(i) *Admissibility*: $A(x) = 1$. The distinction at x must be open to inferential use in the current context. This is a structural condition on the inferential field, sensitive to syntactic scope, embedding depth, and contextual licensing factors.

(ii) *Orientation*: $\omega(x) < 0$. The local inferential environment must be in a downward regime. This is the orientation condition, which corresponds to the DE-ness condition in standard accounts but is now understood as a property of the field rather than of a lexical item.

These conditions are genuinely independent. An environment can be downward-oriented without being admissible for NPI licensing — certain embedded contexts are downward-entailing but block NPI licensing for independent structural reasons. Conversely, an environment can be admissible in general while failing to carry the required downward orientation. NPI licensing fails in both cases, but for different reasons, and the distinction between the two failure modes has empirical consequences.

4.2 Why Downward Orientation Licenses NPIs

The independence of the two conditions raises the explanatory question more sharply: given that both conditions must be satisfied, what is the theoretical reason that downward orientation is one of them? The field-theoretic account provides an answer that the standard DE account does not.

In an upward-entailing environment, inferential expansion proceeds outward: from narrower to broader sets, from more specific to more general distinctions. An existential expression such as *ever* or *any* ranges over a class of distinctions and asserts the non-emptiness of that class. In an upward-entailing environment, this assertion is redundant with ordinary scalar inference: if the class is non-empty at a specific level, upward inference carries the assertion to more general levels automatically. The NPI contributes no additional inferential work.

In a downward-entailing environment, by contrast, the orientation of the field is reversed. Inference no longer expands outward; it contracts inward. In this regime, the existence of an instance at a general level does not guarantee existence at a specific level. The NPI *ever* now does genuine inferential work: it asserts reachability across a region of the field where reachability is not guaranteed by the default expansion of upward inference. The NPI is licensed precisely because the field orientation creates a gap that the NPI's assertion is needed to fill.

This gives a functional explanation of the licensing condition. NPIs are not licensed by downward entailment as an arbitrary syntactic trigger. They are licensed because downward orientation creates the inferential conditions under which their semantic contribution is non-trivial. An NPI in an upward-entailing environment contributes nothing that ordinary inference does not already deliver; an NPI in a downward-entailing environment contributes an assertion that the field orientation makes [Igenuinely informative.

Proposition 4.2 (Licensing and inferential contribution)**.** *Let $\mathcal{F} = (X, R, \omega, A)$ be an inferential field and let e be an NPI expression at position x. The inferential contribution of e is non-trivial if and only if $\omega(x) < 0$.*

Proof. Suppose $\omega(x) > 0$. Then the local region is upward-entailing, and by condition (i) of Definition 2.2, $R(A) \subseteq R(B)$ whenever $A \subseteq B$. For an existential NPI ranging over a class $\mathcal{C}$ of distinctions, the assertion that $\mathcal{C}$ is non-empty at some level is carried upward automatically: $R(\mathcal{C}) \supseteq \mathcal{C}$, and the reachability of members of $\mathcal{C}$ is entailed by the upward expansion of R. The NPI asserts what the field already delivers. Its contribution is trivial.

Suppose $\omega(x) < 0$. Then the local region is downward-entailing, and by condition (ii) of Definition 2.2, $R(B) \subseteq R(A)$ whenever $A \subseteq B$. The default direction of inferential expansion is reversed. The assertion that $\mathcal{C}$ is non-empty at some level is no longer carried automatically: existence at a general level does not guarantee reachability at a specific level under contraction. The NPI must assert reachability explicitly. Its contribution is non-trivial. □

4.3 Domain Structure and Licensing Domains

The 2024 monotonicity-domain finding that processing cost tracks domain structure rather than operator count has a direct parallel in NPI licensing. If licensing were purely a property of individual operators, then the licensing condition would be checked locally at each operator without reference to the broader domain structure. The field-theoretic account predicts otherwise: licensing is a property of admissible orientation regions, and those regions have spatial extent.

Definition 4.3 (Licensing domain)**.** *A* licensing domain *for NPIs in $\mathcal{F}$ is a maximal connected component $D \subseteq X$ such that $A(x) = 1$ and $\omega(x) < 0$ for all $x \in D$.*

Several predictions follow from this definition.

First, an NPI is licensed throughout a licensing domain, not merely at the position of the licensing operator. The domain has spatial extent, and the orientation condition is satisfied everywhere within it. This explains the well-known fact that NPIs can be licensed at a distance from their licensor, within the bounds of the appropriate syntactic scope.

Second, the boundary of a licensing domain is a region of high orientation curvature: κ_ω is large near the boundary. The processing cost function $C(x) = \alpha \cdot \mathbf{1}[\omega(x) < 0] + \beta \cdot \kappa_\omega(x)$ therefore predicts that sentences where an NPI appears near the edge of a licensing domain — where the field transitions from downward to upward orientation — should be more costly to process than sentences where the NPI appears well within the domain. This is a testable prediction that distinguishes the field account from standard licensing accounts.

Third, an NPI that appears outside a licensing domain but within an admissible region ($A(x) = 1$, $\omega(x) > 0$) should produce an ungrammaticality judgment associated with orientation mismatch rather than admissibility failure. The two failure modes predict different patterns of degradation and potentially different processing signatures. An NPI that appears outside an admissible region altogether ($A(x) = 0$) should produce a different kind of ungrammaticality. Whether these predicted differences in degradation type correspond to empirically distinguishable judgment profiles is an open question.

4.4 Superpolarity and Weak NPIs

The standard DE account treats all NPIs as requiring the same licensing condition. The empirical literature, however, distinguishes weak NPIs (licensed in a wide range of DE environments, including questions and conditionals) from strong NPIs (licensed only in a narrow range of environments, typically clausemate negation). This distinction is difficult to capture in a purely operator-based account.

The field-theoretic account handles the distinction through the joint conditions of admissibility and orientation. A weak NPI requires only that $\omega(x) < 0$; it is insensitive to the admissibility structure except in the minimal sense that the position must be syntactically accessible. A strong NPI imposes an additional admissibility condition: it requires not merely that the local orientation be downward but that the admissibility structure satisfy a further constraint, which in practice corresponds to the requirement that the licensor be a clausemate negative operator rather than a more distant downward-entailing expression.

Definition 4.4 (Weak and strong NPI probes)**.** *A* weak NPI probe *at x succeeds if $\omega(x) < 0$, regardless of the admissibility structure beyond syntactic accessibility. A* strong NPI probe *at x succeeds if $\omega(x) < -\theta$ for some threshold $\theta \in (0, 1)$ and $A(x) = 1$ under a restricted admissibility structure A_s that requires clausemate licensing.*

The threshold θ in the strong NPI definition captures the intuition that strong NPIs require not merely a downward-oriented environment but a strongly downward-oriented one: the orientation must exceed a threshold magnitude, not merely be negative. This is a natural consequence of allowing ω to range continuously over

$[-1, +1]$ rather than taking binary values. Environments near the boundary $\omega(x) = 0$ are weakly downward-oriented and license weak but not strong NPIs; environments deep in the downward regime license both.

4.5 The Probe Metaphor and Its Limits

We have described NPIs as probes of the inferential field, but the metaphor should not be pressed too far. An NPI is not an independent detector inserted into the sentence to measure field properties. It is a lexical item whose semantic contribution is constituted by its sensitivity to field orientation: it asserts reachability across a downward-oriented region, and its grammaticality is the reflection of whether such an assertion is coherent in the local environment. The probe metaphor captures the epistemic situation of the linguist or the grammar: by observing where NPIs are licensed, one can infer the orientation structure of the inferential field. The NPI is a probe in the sense that a thermometer is a probe: it does not measure temperature independently of being a thermometer; its readings constitute evidence about the environment precisely because of what it is.

The field-theoretic account therefore does not add NPIs as an external diagnostic layer on top of the semantic theory. It integrates them as lexical items whose licensing conditions are determined by the same geometric structure that determines processing cost, monotonicity behavior, and negation distortion. The grammar, the processing system, and the geometric framework are not three separate theories coordinated by stipulation. They are three descriptions of the same inferential field.

5 Persistence, Memory, and Residual Distortion

The delayed-verification experiments occupy a distinctive position in the empirical record. The other phenomena discussed in this paper — double-negation non-additivity, domain-sensitive monotonicity cost, NPI licensing — concern the processing of sentences at or near the time of their production. The delayed-verification result concerns what happens after processing has nominally concluded. A negated sentence retains a processing cost advantage for the corresponding positive sentence even when a temporal gap is introduced between the sentence and the verification image. The cost is not a parsing artifact that dissipates with time; it is a property of the stored representation.

This result is significant for the field-theoretic account because it establishes that the orientation transformation induced by negation is not merely a transient computational state. It is a change in the structure of the representation that persists in working memory and affects downstream processing. We formalize this persistence using the concept of residual distortion introduced in Section 2 and develop its consequences for the architecture of the verification process.

5.1 The Representation Passed to Verification

In the four-stage model introduced in Section 2 and elaborated in Section 6, the verification stage receives as input the output of field deformation: a representation whose orientation has been modified by the logical operators in the sentence. For a positive sentence S^+, the representation passed to verification is $C_\sigma(S^+)$, the field constructed by syntactic processing. For a negative sentence S^-, the representation passed to verification is $D_N(C_\sigma(S^-))$, the field after orientation reversal.

The crucial observation is that this representation is not the same object as $C_\sigma(S^+)$ with a flag attached. The orientation reversal is not a notational marker that the verification process can simply read off and invert. It is a structural change in the reachability geometry of the representation: the distinctions that were mutually reachable in the original field are no longer mutually reachable in the same pattern after negation, and conversely. The verification process must navigate this restructured geometry in order to compare the representation against the visual scenario.

Definition 5.1 (Residual distortion)**.** *Let $\mathcal{F} = (X, R, \omega, A)$ be an inferential field and let $\mathcal{F}^- = D_N(\mathcal{F}) = (X, R, -\omega, A)$ be the field after orientation reversal. The* residual distortion *of $\mathcal{F}^-$ with respect to $\mathcal{F}$ at position x is*

$$\Delta(x) = d_R(x, N(x)) = D_N(x),$$

where d_R is the reachability pseudometric of Definition 2.7. The global residual distortion *is*

$$\|\Delta\| = \int_X \Delta(x)\, d\mu(x).$$

The residual distortion $\|\Delta\|$ is the quantity that the delayed-verification result measures. It is not the cost of parsing the sentence; that cost has already been incurred. It is the cost imposed on the verification process by the fact that the representation it receives has a different reachability geometry from the representation that would have been passed by the corresponding positive sentence.

5.2 Why Time Does Not Resolve Distortion

The operator-count hypothesis, interpreted as a parsing hypothesis, predicts that negation cost should diminish with delay because the parsing operations that incur the cost have time to complete. The field-theoretic account predicts the opposite: residual distortion persists because it is a structural property of the stored representation, not a measure of incomplete processing.

Proposition 5.2 (Persistence of residual distortion)**.** *Let τ be a temporal delay between sentence offset and verification image onset. Under the field-theoretic account, the processing cost of verification for a negated sentence satisfies*

$$C_V(S^-, \tau) \geq C_V(S^+, \tau) + \|\Delta\|_\tau$$

for all $\tau \geq 0$, where $\|\Delta\|_\tau$ is the residual distortion at delay τ, and $\|\Delta\|_\tau$ does not converge to zero as $\tau \to \infty$ in the absence of active re-orientation.

Proof. The verification process requires computing $\chi(P, \mathcal{F}^-)$, where P is the visual scenario and $\mathcal{F}^-$ is the negated representation. The cost of χ depends on the reachability distance between the scenario representation and the sentence representation. For a positive sentence, this distance is $d_R(P, \mathcal{F}^+)$. For a negative sentence, the representation has been distorted by D_N, so the effective distance is $d_R(P, \mathcal{F}^-)$. The difference between these distances is bounded below by $\|\Delta\|$ by the triangle inequality applied to d_R.

The temporal delay τ does not reduce $\|\Delta\|_\tau$ because $\|\Delta\|$ is a property of the representation stored in working memory, not a property of an ongoing computational process. The representation does not spontaneously re-orient itself with the passage of time; orientation is a structural feature of the stored field, not a transient state of a processing operation. Therefore $\|\Delta\|_\tau$ remains bounded away from zero for all τ, and the cost inequality is maintained across all delays.

Active re-orientation — the explicit computation of $N(\mathcal{F}^-) = \mathcal{F}^+$ as a preliminary step before verification — would resolve the distortion, but this operation itself incurs a cost proportional to $\|\Delta\|$. The total verification cost for a negative sentence is therefore at least $C_V(S^+, \tau) + \|\Delta\|$ regardless of delay, with the cost allocated either to working-memory maintenance or to explicit re-orientation depending on strategy. □

5.3 Two Verification Strategies

The proposition above implies that reasoners facing a delayed verification task with negated sentences have two available strategies, both of which incur a cost related to $\|\Delta\|$.

The first strategy is *deferred re-orientation*: maintain the negated representation $\mathcal{F}^-$ in working memory throughout the delay interval, and re-orient the field explicitly when the verification image appears. The working-memory cost is the maintenance cost of storing a distorted representation; the re-orientation cost is incurred at verification onset. This strategy predicts that reaction time should spike at the onset of the verification image relative to positive sentences, with the spike proportional to $\|\Delta\|$.

The second strategy is *anticipatory re-orientation*: re-orient the field during the delay interval, converting $\mathcal{F}^-$ to $\mathcal{F}^+$ before the image appears. This strategy amortizes the re-orientation cost across the delay interval rather than concentrating it at verification onset. It predicts a more uniform cost increase across the delay interval rather than a spike at onset.

Both strategies predict a sustained cost relative to positive sentences. They differ in the temporal distribution of that cost. The delayed- verification experiments can in principle distinguish between them by examining the time course of the reaction time disadvantage: a spike at onset supports deferred re-orientation; a uniform elevation supports anticipatory re-orientation. The available evidence is more consistent with the deferred strategy, suggesting that working memory maintains the distorted representation rather than re-orienting proactively, but the evidence is not yet decisive.

5.4 Persistence and Working-Memory Architecture

The persistence result has implications for the architecture of working memory that go beyond the specific case of negation. If residual distortion is a genuine structural feature of stored representations, then working memory must be capable of maintaining representations with non-default orientation. This is a non-trivial architectural requirement: it implies that working memory does not normalize representations to a canonical positive orientation during maintenance.

This requirement is consistent with the broader evidence on working memory and logical form. Working memory representations retain structural properties of the sentences that generated them, including scope relations, embedding depth, and logical structure. Orientation, in the present framework, is one such structural property. Its maintenance in working memory is therefore not a special feature of negation but an instance of the general capacity to maintain structured logical representations across delays.

What is special about negation is the direction of the distortion. Upward-entailing sentences generate representations whose orientation matches the default expansion direction of the inferential field. Downward-entailing sentences generate representations whose orientation is reversed. The reversed orientation imposes a greater maintenance cost because it conflicts with the default reachability geometry that the verification process is designed to navigate. The verification process is tuned to a field in which $\omega > 0$; a representation with $\omega < 0$ requires either re-orientation or navigation in the non-default direction, both of which incur additional cost.

Corollary 5.3 (Asymmetry of maintenance cost)**.** *Working-memory maintenance of a negated representation $\mathcal{F}^-$ incurs a greater cost than maintenance of the corresponding positive representation $\mathcal{F}^+$, proportional to $\|\Delta\|$.*

Proof. Maintenance cost is the cost of preserving the structural properties of a representation against decay and interference over the delay interval. For $\mathcal{F}^+$, the orientation $\omega > 0$ is consistent with the default expansion direction of R, so maintenance requires only that the reachability structure be preserved. For $\mathcal{F}^-$, the orientation $\omega < 0$ conflicts with the default expansion direction, so maintenance requires additionally that the reversed orientation be preserved against spontaneous re-normalization toward the default. This additional maintenance requirement is proportional to $\|\Delta\|$, since larger distortions require greater active maintenance to prevent reversion to default orientation. □

5.5 Persistence and the Admissibility Structure

A further consequence of the persistence result concerns the admissibility structure A of the stored representation. We have defined admissibility as a function on the distinction space that determines which distinctions are open to inferential use. When a negated representation is stored in working memory, the admissibility structure is maintained alongside the orientation reversal.

This has an important consequence for NPI licensing in embedded environments. An NPI that is licensed at the time of sentence processing — because $A(x) = 1$ and $\omega(x) < 0$ in the processed field — remains licensed in the stored representation, since both conditions are preserved under maintenance. But the verification process may require comparing the stored representation against a scenario in which the admissibility conditions differ. In such cases, the comparison must resolve a mismatch not only in orientation but in admissibility structure. This mismatch adds an additional component to the verification cost beyond the residual distortion $\|\Delta\|$.

This prediction extends the delayed-verification paradigm in a new direction. Sentences containing NPIs in downward-entailing environments should show a specific pattern: the NPI licensing condition is satisfied at processing time, but the verification cost includes both the orientation distortion cost and an admissibility mismatch cost if the verification scenario does not independently support the admissibility conditions. Designing verification experiments that manipulate the admissibility match between sentence and scenario would test this prediction directly.

5.6 Summary: Three Components of Negation Cost

The analysis of this section completes the decomposition of negation cost into three separable components, each with a distinct temporal profile and a distinct locus in the processing architecture.

1. *Orientation cost*: incurred during field construction and deformation, proportional to $\mathbf{1}[\omega(x) < 0]$. This is the baseline cost of operating in a downward-entailing regime and is concentrated in the processing stage.

2. *Boundary cost*: incurred during traversal of orientation domain boundaries, proportional to $\kappa_\omega(x)$. This cost is concentrated near domain edges and is the quantity most directly measured by the 2024 monotonicity-domain experiments.

3. *Residual distortion cost*: incurred during working-memory maintenance and verification, proportional to $\|\Delta\|$. This cost persists across temporal delays and is the quantity measured by delayed-verification experiments. It does not diminish with parsing time but only with active re-orientation.

These three components are in principle independently manipulable. A sentence that traverses many domain boundaries may have high boundary cost with moderate residual distortion. A sentence with a single deep negation may have low boundary cost with high residual distortion. A sentence with two canceling negations may have low orientation cost, high boundary cost at the intermediate state, and near-zero residual distortion. The three-component decomposition generates a matrix of predictions across sentence types and task conditions that jointly constrain the framework and distinguish it from both operator-count and polarity accounts.

6 Neural Dissociation and the Four-Stage Model

The inferential-field account developed in the preceding sections is a theoretical framework, not a neural model. It makes claims about the computational structure of negation processing — the geometry of orientation reversal, the separability of cost components, the persistence of residual distortion — without committing to specific neural implementations of those computations. Nevertheless, the framework is not neurally unconstrained. It requires, as a necessary condition, that the operations it distinguishes be at least partially dissociable at the level of neural implementation. If syntactic field construction, logical field deformation, and scenario verification were realized by a single undifferentiated neural process, then the cost decomposition developed here would be computationally real but neurally uninstantiated, and the framework would lose its claim to describe the actual processing architecture rather than an idealized one.

The neural evidence on negation processing provides the required constraint. We review the relevant findings, formalize the dissociation argument, and then use it to ground the four-stage functional model that has been implicit throughout the paper.

6.1 The Relevant Neural Evidence

Three clusters of neural findings bear on the dissociation argument.

The first concerns syntactic complexity. A substantial literature associates complex syntactic processing with left inferior frontal regions, particularly Broca's area (BA 44 and BA 45) and adjacent premotor cortex, together with posterior superior temporal regions. These regions show increased activity for sentences with greater syntactic complexity — center-embedded relative clauses, non-canonical argument structures, long-distance dependencies — relative to structurally simpler sentences matched for semantic content. The relevant point for present purposes is not the precise anatomical attribution but the existence of a reliable syntactic complexity contrast that localizes to identifiable regions.

The second concerns logical negation specifically. Grodzinsky and colleagues report a left anterior insula cluster, overlapping with the cytoarchitectonically defined region Id7, that exhibits negation-sensitive activity in controlled sentence-picture verification tasks. The experimental design isolates logical negation by contrasting negative sentences against positive sentences matched for syntactic complexity, word frequency, and sentence length. The negation-sensitive contrast localizes to the anterior insula cluster rather than to the classical syntactic regions implicated by the first set of findings. Critically, participant-level reaction times in the negation condition correlate with BOLD activity in this cluster, establishing a link between the neural signal and the behavioral cost.

The third concerns numerical and scenario comparison. Tasks requiring magnitude comparison or verification against a structured visual scenario implicate a distinct set of regions, including intraparietal sulcus and inferior parietal cortex, that are separable from both the syntactic and the negation clusters. The verification stage of the sentence-picture task is not neurally identical to the syntactic or logical processing stages.

These three clusters — syntactic regions, anterior insula negation cluster, parietal comparison regions — are anatomically distinct and are differentially sensitive to distinct experimental contrasts. This triple dissociation is the neural evidence that the framework requires.

6.2 Formalizing the Dissociation

We formalize the dissociation argument using the four-map decomposition introduced in Section 2. A sentence-picture verification trial is represented as the composite operation

$$V(P,S) = \chi\big(P,\ D_\ell(C_\sigma(S))\big),$$

where S is the sentence, P is the visual scenario, C_σ is syntactic field construction, D_ℓ is logical field deformation, and χ is scenario comparison. For a positive sentence,

$$V(P,S^+) = \chi(P, C_\sigma(S^+)),$$

and for a negative sentence,

$$V(P,S^-) = \chi(P, D_N(C_\sigma(S^-))).$$

The contrast between negative and positive conditions isolates D_N under the assumption that C_σ and χ are held sufficiently constant across conditions by the experimental design.

Proposition 6.1 (Neural dissociation of field deformation)**.** *Suppose the verification task decomposes as $V = \chi \circ D_\ell \circ C_\sigma$. Let R_C, R_D, and R_χ denote neural regions associated with syntactic construction, logical deformation, and scenario comparison respectively. If*

(i) the negation contrast modulates B_{R_D} but not B_{R_C} or B_{R_χ},

(ii) syntactic complexity manipulations modulate B_{R_C} but not B_{R_D} or B_{R_χ}, and

(iii) comparison difficulty manipulations modulate B_{R_χ} but not B_{R_C} or B_{R_D},

then C_σ, D_ℓ, and χ are functionally dissociable operations.

Proof. Let $B_R(O)$ denote the BOLD response in region R associated with operation O. If the three operations were not dissociable — if they constituted a single undifferentiated process — then any manipulation that increases task difficulty should modulate the same neural network regardless of whether the manipulation targets syntactic structure, logical form, or comparison. The observed pattern is the opposite: each manipulation produces a selective effect in its associated region while leaving the other regions unaffected. This selectivity is contrast-relative: the negation contrast is not simply a harder condition in general, but a condition that specifically increases demand on R_D while syntactic and comparison demands are controlled. The triple dissociation therefore establishes that the three operations are functionally separable, in the sense that each can be selectively modulated without equivalent modulation of the others. Under the decomposition $V = \chi \circ D_\ell \circ C_\sigma$, this selectivity is the neural signature of functional dissociation. □

6.3 Against the Generic Load Objection

The anterior insula is implicated in a wide range of non-linguistic functions, including interoception, affective processing, salience detection, and cognitive control. The observation that a negation contrast localizes partly to this region is therefore consistent with the hypothesis that negation is simply harder in a generic sense, and that the anterior insula is responding to generic task difficulty rather than to the logical structure of the sentence. We address this objection directly.

Proposition 6.2 (Specificity over generic load)**.** *If anterior insula activity in negation tasks were caused by generic task load rather than by logical field deformation, then the same region should be modulated by any manipulation that increases task difficulty, regardless of whether the manipulation targets syntactic structure, logical form, or comparison. The observed pattern of anatomical selectivity is inconsistent with this prediction.*

Proof. Let L be a scalar measure of generic task load. On the generic-load hypothesis, regional activity is a monotone function of L:

$$B_R = f_R(L), \qquad f'_R \geq 0.$$

Under this hypothesis, any manipulation that increases L should increase B_{R_D}, B_{R_C}, and B_{R_χ} proportionally, since all reflect the same underlying load variable. The predicted pattern is uniform modulation across regions for any difficulty-increasing manipulation. The observed pattern is not uniform. Syntactic complexity increases B_{R_C} without proportionally increasing B_{R_D} or B_{R_χ}. Negation increases B_{R_D} without proportionally increasing B_{R_C} or B_{R_χ}. Comparison difficulty increases B_{R_χ} without proportionally increasing the other regions. The generic-load function $f_R(L)$ cannot produce this pattern, since it predicts that any common increase in L should modulate all load-sensitive regions.

The residual explanation available to the generic-load account is that the three regions have different sensitivity functions f_{R_C}, f_{R_D}, f_{R_χ} with different thresholds or slopes, and that the three manipulations happen to produce load increases that fall selectively above each region's threshold. This explanation is available in principle but requires three independent threshold calibrations to account for what the field-deformation account explains with a single principle: each region is selectively sensitive to one operation in the decomposition $V = \chi \circ D_\ell \circ C_\sigma$, and the selectivity is structural rather than incidental.

The anterior insula is not claimed to be intrinsically a logic area. The claim is narrower: in tasks where syntactic construction and comparison are controlled by design, the residual negation contrast localizes to this region, and the localization is consistent with the field-deformation account and inconsistent with the generic-load account as formulated. □

6.4 The Four-Stage Model

The dissociation argument supports a functional architecture in which the verification process decomposes into four stages. We state this architecture explicitly and note its relationship to both the field-theoretic framework and the neural evidence.

1. *Parse*: The sentence is incrementally parsed and its syntactic structure is recovered. This stage corresponds to C_σ applied to the surface string. It is associated with left inferior frontal and posterior temporal regions and is sensitive to syntactic complexity but not to logical form per se.

2. *Construct field*: The parsed structure is used to construct an inferential field $\mathcal{F} = (X, R, \omega, A)$ for the sentence. This is the stage at which the distinction space, reachability relation, orientation field, and admissibility structure are established. It is part of C_σ in the formal decomposition but is conceptually distinct from surface parsing: the same surface structure can, in principle, correspond to different inferential fields depending on lexical semantic content.

3. *Deform field*: Logical operators, including negation, quantifiers, and monotonicity-reversing expressions, transform the constructed field. For negation, this stage applies D_N to the orientation field, reversing ω throughout the relevant domain. This is the stage associated with R_D in the neural evidence and with the anterior insula cluster. Its cost is proportional to the orientation distortion $\|\Delta\|$ and the boundary curvature κ_ω.

4. *Verify scenario*: The deformed field is compared against the visual or described scenario. The verification function χ computes whether the scenario is reachable

from the sentence representation under the current field geometry. This stage is associated with R_χ and parietal comparison regions. Its cost includes the residual distortion cost $\|\Delta\|_\tau$ carried over from the deformation stage, which is why delayed verification shows sustained negation effects even when the deformation stage has long concluded.

Corollary 6.3 (Empirical constraint on the four-stage model)**.** *The four-stage model is empirically constrained by the triple dissociation in Proposition 6.1. It is not a free theoretical posit but a model whose stages correspond to functionally and at least partially neurally separable operations. The model remains a theoretical abstraction rather than a direct anatomical map, but it is constrained by the anatomy in the sense that any reduction of the four stages to fewer stages would require the elimination of one of the three dissociable contrasts.*

Proof. If stages 1 and 2 were merged into a single undifferentiated parsing operation, then the syntactic complexity contrast and the field construction contrast would be predicted to co-localize, with no room for a distinct field-construction signature. If stages 3 and 4 were merged, then the negation contrast and the comparison contrast would co-localize, contrary to the observed anatomical separation. If stages 2 and 3 were merged, then the boundary between syntactic construction and logical deformation would be eliminated, predicting that negation effects and syntactic complexity effects should modulate the same neural populations. The triple dissociation rules out all three reductions. □

6.5 The Model as Computational Interpretation

The four-stage model is a computational interpretation of the inferential field framework, not an additional hypothesis added to it. The framework defines the objects — fields, orientations, deformations, admissibility structures — and the model describes how processing moves through those objects in sequence. Each stage is defined by its input, its output, and the transformation it applies.

The model makes the processing architecture explicit in a way that permits contact with both the behavioral and the neural literature. On the behavioral side, it predicts that manipulations targeting different stages should have separable effects on reaction time distributions, error patterns, and the time course of verification difficulty. On the neural side, it predicts that each stage should have a distinct neural profile, with cross-stage interactions occurring at the transitions between stages rather than uniformly throughout processing.

The most important transition for the present paper is between stages 3 and 4: between field deformation and scenario verification. The residual distortion $\|\Delta\|$ is produced at stage 3 and consumed at stage 4. The delayed-verification paradigm manipulates the interval between these stages. The persistence result established in Section 5 — that $\|\Delta\|_\tau$ does not diminish with delay in the absence of active re-orientation — is therefore a claim about what happens to the output of stage 3 during the interval before stage 4 begins. The distorted representation is maintained in working memory without spontaneous re-normalization. Stage 4 receives it in its distorted form and incurs the verification cost accordingly.

This account of the stage-3 to stage-4 transition is the field-theoretic explanation of the most theoretically challenging result in the empirical record: the sustained cost of negation across temporal delays. It requires no special mechanism beyond the geometry of the inferential field and the claim that working memory preserves rather than normalizes the structural properties of stored representations.

7 Connections to Admissibility Geometry and Distinguishability Theory

The inferential field framework developed in this paper did not emerge in isolation. It belongs to a broader research program organized around a cluster of related formal structures: admissibility geometry, which formalizes the conditions under which a state or inference is available for continuation; distinguishability theory, which formalizes the conditions under which a difference between states can be represented and acted upon; and repair theory, which formalizes the conditions under which a damaged or incomplete representation can be restored to admissible form. The present section makes the connections between these programs explicit. The goal is not to subsume the inferential field account into a larger system by stipulation but to show that the mathematical structure developed to handle negation and monotonicity is the same structure that appears independently in memory, repair, and scientific explanation. This convergence is evidence that the structure is tracking something real.

7.1 The Admissibility Program

The admissibility program begins from the observation that the standard epistemological and semantic frameworks begin too late. They ask which beliefs are true, which inferences are valid, which representations are accurate. But prior to any of these questions is a more fundamental one: which states are available for cognitive engagement at all? A belief that cannot be entertained, an inference that cannot be drawn, a representation that cannot be accessed are not false or invalid; they are inadmissible. The admissibility structure A of the inferential field is the local instantiation of this more general concept.

In the admissibility program, the central formal object is the *admissibility manifold*: the space of states that are locally available for continuation, equipped with a geometry that reflects how admissibility conditions vary across the state space. A state is admissible if it lies within this manifold; it becomes inadmissible when it is deformed outside the manifold by some transformation. Repair is the process of returning an inadmissible state to the manifold.

The inferential field connects to this program through the admissibility structure $A : X \to \{0, 1\}$. An admissible region in the inferential field — a connected component of $A^{-1}(1)$ — is a local patch of the admissibility manifold restricted to the distinction space X. The licensing domain for NPIs defined in Section 4 is such a patch: a region where both admissibility and orientation conditions are satisfied, and within which NPI licensing proceeds without restriction.

The connection runs deeper than this local correspondence, however. The admissibility program studies how admissibility conditions change under transformation: which transformations preserve admissibility, which deform it, and which destroy it entirely. Negation, in the inferential field framework, is a transformation of the orientation component of the field. It does not directly modify the admissibility structure A; it reverses ω while leaving A unchanged. This means that negation can take an inferential system from an admissible orientation into an admissible reversed orientation, or from an admissible region into an inadmissible one if the admissibility conditions are themselves orientation-sensitive.

This last case is precisely NPI licensing failure. When a negative operator is embedded in a context that makes the resulting field inadmissible — when $N(\omega) < 0$ in a region where $A = 0$ — the NPI probe returns zero not because the orientation is wrong but because the admissibility structure excludes the relevant region. The field has been deformed into an inadmissible configuration. The grammar registers this as ungrammaticality; the admissibility program registers it as a deformation outside the manifold.

Proposition 7.1 (Negation as admissibility-preserving transformation)**.** *The negation operator $N : \mathcal{F} \to \mathcal{F}$ preserves the admissibility structure: $N(A) = A$. Negation can nevertheless produce admissibility failure when the admissibility structure is orientation-sensitive, that is, when $A(x)$ depends on $\omega(x)$.*

Proof. By definition, $N(\omega)(x) = -\omega(x)$ and N leaves A unchanged. So $N(A) = A$ directly. However, if A is defined in terms of orientation — if, for instance, $A(x) = 1$ only when $\omega(x) < 0$ — then after applying N, the effective admissibility of a region that was previously admissible may change, because ω has been reversed. The admissibility function A as a formal object is unchanged; but the conditions that must be met for $A(x) = 1$ to obtain may now fail in regions that were previously admissible. This is the formal counterpart of the empirical observation that negation can both create and destroy NPI licensing environments: it creates downward orientation where there was none, and destroys it where it previously existed, while the structural admissibility conditions remain fixed. □

7.2 Distinguishability Geometry

Distinguishability geometry begins from a complementary starting point. Where the admissibility program asks which states are available for continuation, distinguishability geometry asks which differences between states can be represented. A distinction that cannot be represented is not false; it is indistinguishable. The formal object is a *distinguishability space* $(X, \sim)$, where $\sim$ is an equivalence relation encoding which pairs of states are indistinguishable under the current representational resources.

The inferential field connects to distinguishability geometry through the distinction space X itself. The elements of X are distinctions: they are, by definition, differences that the system can represent. The reachability relation R on X encodes which distinctions are inferentially connected: which differences can be reached from which others by inference. A transformation that alters R changes the inferential connectivity of the distinction space; a transformation that collapses distinctions — merging elements of X — reduces distinguishability directly.

Negation, in the distinguishability-theoretic perspective, is a transformation that preserves the elements of X while reversing the orientation of their inferential connections. It does not collapse distinctions; it reorients them. The distinctions that were reachable in the upward direction become reachable in the downward direction, and vice versa. The negation distortion $D_N(x) = d_R(x, N(x))$ measures how far each distinction is displaced in the reachability structure by this reorientation: a distinction that was easily reachable from its neighbors may become distant after orientation reversal, if the reachability paths that connected it depended on the upward-expansion direction.

Definition 7.2 (Distinguishability-theoretic negation distortion)**.** *In the distinguishability-theoretic interpretation, the negation distortion $D_N(x)$ measures the change in inferential distance between x and its inferential neighbors induced by the orientation reversal N. Formally,*

$$D_N(x) = \sum_{y \in \mathcal{N}(x)} |d_R(x, y) - d_R(N(x), N(y))|,$$

where $\mathcal{N}(x)$ is the set of inferential neighbors of x under R.

This definition makes explicit that the distortion is not a global property of the negation operator but a local property of each distinction and its neighborhood. Distinctions that are well-connected in both orientations — whose inferential neighbors are reachable regardless of field orientation — have low distortion. Distinctions that are reachable primarily through upward-expanding paths have high distortion under negation, because the paths that connected them are no longer available in the reversed field.

The three cost components identified in Section 5 now have a unified distinguishability-theoretic interpretation:

1. *Orientation cost* is the cost of representing a distinction whose inferential neighborhood has been reoriented: the system must navigate connections in the non-default direction.

2. *Boundary cost* is the cost of representing a distinction near a distinguishability boundary: a region where the inferential distance between neighboring distinctions changes rapidly as orientation varies.

3. *Residual distortion cost* is the cost of maintaining a representation whose inferential neighborhood is displaced from its default configuration: the working-memory system must actively preserve an anomalous reachability structure against normalization.

7.3 Repair Theory

Repair theory addresses the conditions under which an inadmissible or damaged representation can be restored to admissible form. The central formal object is the *repair pseudometric* $d_{\text{rep}}(x, y)$, which measures the cost of restoring a representation x to an admissible state y. Repair is not mere correction; it is the process of navigating from an inadmissible region of the state space back to the admissibility manifold along a path of minimal cost.

The connection between repair theory and the inferential field framework runs through the verification stage of the four-stage model. When a negated representation $\mathcal{F}^-$ is compared against a positive scenario P, the verification process must either reorient the field — applying N to convert $\mathcal{F}^-$ back to $\mathcal{F}^+$ — or navigate the comparison in the non-default direction. The first strategy is a repair operation: it restores the representation to the orientation that the verification process is designed to handle. The cost of this repair is proportional to $\|\Delta\|$, the global negation distortion. The repair pseudometric gives a formal measure of this cost:

$$d_{\text{rep}}(\mathcal{F}^-, \mathcal{F}^+) = \|\Delta\| = \int_X D_N(x)\, d\mu(x).$$

This identification has a theoretical consequence. In repair theory, the repair pseudometric satisfies a conservation law: the total repair cost of a sequence of transformations is bounded below by the distortion accumulated along the sequence. Applied to the verification process, this means that the total cost of verifying a negated sentence against a scenario is bounded below by $\|\Delta\|$ regardless of verification strategy. The bound cannot be avoided by choosing a more efficient verification algorithm; it is a consequence of the geometry of the inferential field.

Proposition 7.3 (Lower bound on verification cost)**.** *For any verification strategy χ and any negated representation $\mathcal{F}^-$,*

$$C_V(\mathcal{F}^-, P) \geq C_V(\mathcal{F}^+, P) + d_{\text{rep}}(\mathcal{F}^-, \mathcal{F}^+),$$

where $C_V(\mathcal{F}, P)$ is the cost of verifying scenario P against field $\mathcal{F}$.

Proof. The verification cost for $\mathcal{F}^-$ includes at minimum the cost of verifying $\mathcal{F}^+$ plus the cost of resolving the distortion between $\mathcal{F}^-$ and $\mathcal{F}^+$. The distortion resolution cost is bounded below by $d_{\text{rep}}(\mathcal{F}^-, \mathcal{F}^+)$ by the definition of the repair pseudometric as the minimum cost path from $\mathcal{F}^-$ to $\mathcal{F}^+$ in the state space. No verification strategy can avoid incurring this cost, since any strategy that correctly computes the verification result must at some point resolve the orientation difference between the negated representation and the default-oriented scenario representation. The triangle inequality on d_{rep} then gives the stated bound. □

7.4 The RSVP Connection

The RSVP framework — Relativistic Scalar-Vector Plenum — is a field-theoretic approach to cognitive and physical dynamics that identifies cognitive states with configurations of a scalar field Φ and a vector field $\vec{v}$ subject to constraint surfaces S. The connections between RSVP and the inferential field framework are more than analogical; they share a mathematical structure.

The correspondence is as follows:

$$\begin{aligned}
\Phi &\longleftrightarrow \text{available inferential capacity in } X,\\
\vec{v} &\longleftrightarrow \text{inferential flow along } R,\\
S &\longleftrightarrow \text{admissibility and distinction-preserving constraints in } A,\\
\omega &\longleftrightarrow \text{orientation of inferential flow},\\
\kappa_\omega &\longleftrightarrow \text{curvature of constraint surfaces.}
\end{aligned}$$

Under this correspondence, upward entailment corresponds to a region where inferential flow $\vec{v}$ is divergent: distinctions expand outward as inference proceeds. Downward entailment corresponds to a region where $\vec{v}$ is convergent: distinctions contract inward. Negation is a local reversal of $\vec{v}$, changing divergent flow to convergent flow in the affected region. The boundary between upward and downward regions is a surface where $\nabla \cdot \vec{v} = 0$: neither divergent nor convergent, the point of maximum orientation curvature κ_ω.

The processing cost function $C(x) = \alpha \cdot \mathbf{1}[\omega(x) < 0] + \beta \cdot \kappa_\omega(x)$ has a natural RSVP interpretation. The first term is the cost of operating in a convergent flow regime; the second is the cost of operating near a constraint surface boundary where the flow transitions between regimes. In RSVP dynamics, the energetic cost of a trajectory through state space is determined by the integral of the constraint violation along the path — exactly the integral of $C(x)$ over the inferential path developed here.

Proposition 7.4 (RSVP formulation of the Orientation-Reversal Principle)**.** *Under the RSVP correspondence, the Orientation-Reversal Principle (Theorem 3.1) states that the cognitive cost of a semantic operator is determined by the integral of constraint* *violation along the inferential trajectory induced by that operator, where constraint violation is measured by the divergence of inferential flow and the curvature of admissibility constraint surfaces.*

Proof. By the RSVP correspondence, the orientation reversal $\int_X |\omega(x) - T(\omega)(x)| \, d\mu(x)$ corresponds to the integral of the change in flow divergence $|\nabla \cdot \vec{v}(x) - \nabla \cdot T(\vec{v})(x)|$ over the affected region. The boundary curvature $\int_X \kappa_\omega(x) \, d\mu(x)$ corresponds to the integral of constraint surface curvature. In RSVP dynamics, the energetic cost of a transformation is the sum of these integrals, since both divergence change and surface curvature contribute to constraint violation along the trajectory. This is the RSVP formulation of condition (i) and (ii) of Theorem 3.1. □

7.5 Convergence of Three Programs

The convergence of the admissibility program, distinguishability geometry, and repair theory on the same formal structure — inferential fields, orientation, distortion, admissibility — is not a coincidence of notation. It reflects the fact that all three programs are studying the same underlying phenomenon from different starting points: the conditions under which a cognitive system can move through a space of distinctions while maintaining the structural properties required for that movement to constitute reasoning.

The admissibility program asks: which states are available for continuation? The distinguishability program asks: which differences can be represented? The repair program asks: how can damaged representations be restored? The inferential field provides a single object in which all three questions can be posed simultaneously: the admissibility structure A answers the first, the distinction space X and reachability relation R answer the second, and the repair pseudometric on the space of fields answers the third. Negation, in this unified framework, is not a logical operator that happens to have interesting cognitive consequences. It is a transformation of the inferential field that simultaneously affects admissibility, distinguishability, and repairability — and whose cognitive cost is a measure of all three.

The deeper philosophical claim, which we develop in the final section, is that this unified structure reveals something about the relationship between negation and logic more generally. Logic does not begin with truth values and add negation as one operator among others. It begins with the organization of distinctions, and negation is the fundamental transformation that reveals the orientation of that organization. Truth, on this view, is a derived notion: it is what you get when you have navigated an inferential field successfully and arrived at a verified distinction. Negation is prior to truth because the orientation structure of the field is prior to verification.

8 Negation Before Logic

The preceding sections have developed a geometric account of negation processing and connected it to three independent formal programs. The final section draws out the philosophical consequences of this account. The central claim is not merely that negation has an interesting geometry, or that processing cost can be decomposed into separable components, or even that the inferential field unifies four bodies of experimental evidence. The central claim is stronger: negation is prior to logic, in the sense that the orientation structure of the inferential field is prior to the truth-conditional content that logic operates on. This is not a paradox. It is a consequence of taking the geometry seriously.

8.1 The Standard Picture and Its Inversion

The standard picture of the relationship between negation and logic runs as follows. Logic begins with propositions, which have truth values. Negation is an operator that maps truth values to their complements: true becomes false, false becomes true. The cognitive cost of negation is then the cost of computing this mapping and tracking its effects through a complex sentence. Processing difficulty arises because the mapping must be applied and its effects must be propagated through the semantic composition.

This picture is coherent, but it places truth at the foundation and treats negation as a derived operation defined over it. The empirical evidence reviewed in this paper suggests that this ordering is wrong, or at least incomplete. Negation effects are not confined to the verification stage, where truth values are computed against scenarios. They arise during field construction and deformation, before any truth-conditional comparison has taken place. The cost of negation is incurred in the structure of the representation, not merely in its evaluation. This means that negation is doing something to the inferential field that is independent of, and prior to, the computation of truth values.

The inversion we propose is this: logic does not begin with truth and add negation. It begins with the organization of distinctions and their inferential connections. Negation is the transformation that reveals the orientation of that organization. Truth is a derived notion: it is the result of a successful navigation of the inferential field, a verified reachability relation between a representation and a scenario. Negation is prior to truth because orientation is prior to verification.

8.2 Distinctions Before Propositions

The foundation of the inferential field is the distinction space X. The elements of X are not propositions. They are differences: the difference between red and not-red, between all and some, between inside and outside, between reachable and unreachable. Propositions are constructed from distinctions by the process of field construction (C_σ

in the four-stage model): they are organized packages of distinctions with a specific inferential connectivity imposed by syntactic and semantic composition.

This ordering — distinctions before propositions — is not merely a formal choice. It reflects a substantive claim about what cognitive systems represent. A system that represents distinctions is sensitive to differences in its environment. A system that represents propositions is sensitive to truth conditions. The claim that distinctions are prior is the claim that sensitivity to difference is more fundamental than sensitivity to truth. A system can represent the difference between red and not-red without representing the proposition that something is red; it cannot represent the proposition that something is red without representing the distinction between red and not-red.

Negation, on this view, operates on distinctions, not on propositions. The negation operator N reverses the orientation of the inferential field: it changes the direction in which distinctions are connected by inference. This is a transformation of the structure of the distinction space, not a mapping of truth values. The proposition that something is not red is not the result of applying a truth-value complement to the proposition that something is red. It is the result of constructing an inferential field in which the distinction between red and not-red is oriented in the reversed direction.

Proposition 8.1 (Negation as distinction transformation)**.** *The negation operator $N : \mathcal{F} \to \mathcal{F}$ is a transformation of the distinction space prior to truth evaluation. Its action on the orientation field ω is defined independently of any assignment of truth values to elements of X.*

Proof. By Definition 2.4, $N(\omega)(x) = -\omega(x)$ for all $x \in X$. This definition makes no reference to a truth assignment $v : X \to \{0, 1\}$. The orientation reversal is a property of the inferential field structure (X, R, ω, A), which is defined prior to and independently of any verification comparison $\chi(P, \mathcal{F})$. A truth value is assigned to a distinction only at the verification stage, when the field is compared against a scenario. Negation operates at the deformation stage, which precedes verification in the four-stage model. Therefore negation is defined and operates independently of truth evaluation. □

8.3 Orientation as the Primitive Logical Notion

If negation is prior to truth, then the primitive logical notion is not truth but orientation. An inferential field has orientation before it has truth conditions. The orientation field $\omega : X \to [-1, +1]$ encodes the direction in which inference flows through the distinction space. Upward orientation means that inference expands: from specific to general, from narrow to broad, from few to many. Downward orientation means that inference contracts: from general to specific, from broad to narrow, from many to few. This expansion and contraction is a property of the field prior to any comparison with a scenario.

Logic, on the standard picture, begins with truth and defines validity as truth preservation. On the present picture, logic begins with orientation and defines inference as orientation-preserving motion through the distinction space. Validity is not truth preservation; it is orientation-consistent reachability. A valid inference is one that moves through the field without violating the orientation constraints: upward in upward regions, downward in downward regions, with costs incurred at orientation boundaries.

This reinterpretation has consequences for classical logical laws. The law of double negation elimination — $\neg\neg p \vdash p$ — is, in the present framework, the statement that $N \circ N = \mathrm{id}_{\mathcal{F}}$ (Proposition 2.6). This is not an axiom added to the logic; it is a theorem of the geometry. The law of excluded middle — $p \vee \neg p$ — is, in the present framework, the statement that every distinction $x \in X$ lies either in the upward regime ($\omega(x) > 0$) or the downward regime ($\omega(x) < 0$), with the boundary $\omega(x) = 0$ corresponding to the classical gap. Whether the excluded middle holds in the present framework depends on whether ω is constrained to avoid zero: if $\omega : X \to [-1, +1] \setminus \{0\}$, then excluded middle holds by fiat; if ω is permitted to reach zero, then there are orientation-boundary distinctions for which neither upward nor downward inference is available, corresponding to a failure of excluded middle.

Intuitionistic logic, which rejects excluded middle, corresponds in the present framework to an inferential field in which ω is permitted to take the value zero and in which constructions at the orientation boundary are genuinely available. This is not merely an analogy; it is a geometric interpretation of the constructive reading of logic. A proposition is constructively provable if and only if it is reachable through an orientation-consistent path from the available distinctions; it is not merely the case that its negation is unreachable.

8.4 The Priority of Admissibility

If orientation is the primitive logical notion and truth is derived from successful field navigation, then admissibility is prior to both. A distinction must be admissible — it must lie in a region where $A(x) = 1$ — before its orientation can be evaluated and before any inferential motion through it can occur. Inadmissible distinctions are not false; they are not available for inferential engagement. They lie outside the region of the field in which logic operates.

This priority of admissibility has a consequence for the traditional epistemological question of how we know what we know. The standard question is: which of our beliefs are true, and how do we come to know them? The admissibility-first picture replaces this with a prior question: which distinctions are available for cognitive engagement, and what determines the boundaries of availability? Truth and knowledge are questions that arise only within the admissible region of the inferential field. Outside that region, there are distinctions that cannot be entertained, inferences that cannot be drawn, and representations that cannot be accessed. The cognitive system does not represent these distinctions as false; it does not represent them at all.

This is the connection to the broader admissibility program: the admissibility structure A of the inferential field is the local instantiation of a general principle that admissibility is prior to truth. The epistemological consequences of this principle extend far beyond the processing of negation. But the processing of negation provides the clearest and most directly empirically constrained illustration of the principle, because negation is the operation that most visibly tests the boundary between the admissible and the inadmissible. An NPI in an unlicensed environment is not merely ungrammatical; it is a probe that has encountered an inadmissible region and returned zero. The grammar's sensitivity to this is a reflection of the cognitive system's sensitivity to admissibility boundaries.

8.5 Truth as Verified Reachability

The preceding subsections have argued that orientation is prior to truth and admissibility is prior to orientation. We can now state the positive account of truth that emerges from this ordering.

Definition 8.2 (Truth as verified reachability)**.** *A distinction $x \in X$ is* true in scenario *P if and only if x is admissible ($A(x) = 1$), the local orientation at x is consistent with the inferential path from the sentence representation to x (ω is not violated along the path), and x is reachable from the scenario representation under the verification function χ. Formally,*

$$\mathrm{True}(x, P) \iff A(x) = 1 \ \wedge \ \chi(P, x) = 1.$$

This definition makes truth a three-stage achievement: first, admissibility must be satisfied; second, inferential navigation to x must be orientation-consistent; third, the scenario must verify x as reachable. The familiar truth conditions of propositional and predicate logic are recovered as the special case where A is everywhere 1 (everything is admissible), ω is everywhere positive (all inference is upward-monotone), and χ is classical bivalent comparison. The present framework generalizes this special case by allowing A to exclude regions, ω to vary continuously, and χ to register degrees of reachability rather than binary verification.

Negation, on this account, does not change the truth value of a proposition by flipping a bit. It changes the orientation of the inferential path to x, which changes the verification conditions: the scenario that verifies x in a downward-oriented context is different from the scenario that verifies x in an upward-oriented context, because the reachability conditions depend on the orientation of the field. Truth is not a static property of propositions but a dynamic property of verification relations between oriented representations and scenarios.

8.6 Negation Before Logic: The Thesis

We can now state the central philosophical thesis of this paper with precision.

Theorem 8.3 (Negation before logic)**.** *In the inferential field framework, negation is defined at the level of orientation structure, prior to the assignment of truth values and independent of the propositional content of the sentence. The logical laws governing negation — double negation elimination, the relationship between upward and downward entailment, the licensing conditions of negative polarity items — are theorems of the geometry of inferential fields, not axioms of a logical system. Truth is derived from successful field navigation under admissibility and orientation constraints. Admissibility is prior to truth; orientation is prior to truth; negation, as an orientation transformation, is prior to truth.*

Proof. The components of this theorem have been established in the preceding sections. Negation is defined as $N(\omega)(x) = -\omega(x)$, independently of truth evaluation (Proposition 8.1). Double negation elimination is a theorem of the geometry ($N \circ N = \mathrm{id}$, Proposition 2.6). The relationship between upward and downward entailment is a theorem of the orientation field (Proposition 2.7). NPI licensing is derived from the joint admissibility and orientation conditions of Definition 4.1 and explained by Proposition 4.2. Truth is defined as verified reachability (Definition 8.1), which presupposes admissibility and orientation-consistent navigation. Admissibility is prior because $A(x)$ must equal 1 before any inferential operation on x is defined. Orientation is prior because ω is defined on the field before χ is applied. Negation as orientation transformation is therefore prior to truth as verified reachability. □

8.7 Consequences and Open Questions

The thesis of negation before logic has consequences that extend beyond the scope of this paper. We close by noting the most significant open questions that the framework generates.

The first concerns the relationship between the inferential field and formal proof theory. If the logical laws governing negation are theorems of the geometry rather than axioms, then there should be a precise correspondence between the geometry of inferential fields and the proof theory of the relevant logical system. The double-negation result ($N \circ N = \mathrm{id}$) corresponds to the eliminability of double negation. The orientation boundary at $\omega = 0$ corresponds to the constructive gap in intuitionistic logic. Establishing this correspondence formally — showing that the category of inferential fields is equivalent to the category of models of some proof-theoretically characterized logic — is a substantial mathematical project that the present paper leaves open.

The second concerns the neural basis of admissibility. The four-stage model identifies field construction, deformation, and verification as partially neurally separable operations. But the admissibility structure A does not map cleanly onto any of these stages; it is a constraint that operates throughout the process. Identifying the neural correlates of admissibility — the mechanisms that determine which regions of the distinction space are open for inferential engagement at any given moment — is an open empirical question. The NPI licensing literature provides indirect evidence, since NPI licensing failure is a behavioral signature of inadmissibility, but direct neural evidence for an admissibility mechanism is currently lacking.

The third concerns the extension to non-linguistic inference. The inferential field framework has been developed here in the context of natural language processing, where the empirical record is richest. But the formal structure — distinctions, reachability, orientation, admissibility — is not intrinsically linguistic. If the same structure appears in visual reasoning, spatial navigation, conceptual combination, or scientific inference, then the framework has implications beyond language. The admissibility and distinguishability programs have already extended similar formal structures to memory and repair; the extension to non-linguistic inference would complete the unification.

The fourth concerns the development of experimental predictions that distinguish the field-theoretic account from its competitors in cases not yet tested. The most promising direction is the construction of sentences that traverse orientation domain boundaries without using negative morphemes: purely quantificational sentences whose monotonicity profile includes a boundary between upward and downward regions. The field-theoretic account predicts that such sentences should incur boundary costs even in the absence of any negative expression. This prediction is not derivable from operator-count or polarity accounts and would, if confirmed, provide direct evidence for the curvature term κ_ω as an independent contributor to processing cost.

The inferential field is not the completion of a theory. It is the beginning of one. The phenomena that motivated it — the cost of negation, the geometry of monotonicity, the licensing of polarity items, the persistence of orientation in working memory — are the clearest cases. The framework predicts that the same geometry will appear wherever inference is constrained by orientation: in the resolution of presuppositions, in the processing of conditionals, in the updating of beliefs under contradictory evidence. Whether it does is a question that the framework has now made precise enough to answer empirically.

Appendices

A Monotonicity as Reachability Preservation

A.1 Reachability Spaces

Definition A.1 (Reachability Space)**.** *A* reachability space *is a pair*

$$\mathcal{R} = (X, R)$$

where X *is a set and*

$$R \subseteq X \times X$$

is a preorder satisfying

1. *Reflexivity:*

$$R(x, x)$$

for all $x \in X$ *.*

2. *Transitivity:*

$$R(x, y) \land R(y, z) \Rightarrow R(x, z).$$

For x∈X , define the forward reachability set

$$\mathcal{N}^{+}(x) y \in X : R(x, y).$$

Similarly define

$$\mathcal{N}^{-}(x) y \in X : R(y, x).$$

Definition A.2 (Inferential Inclusion)**.** *For subsets* $A,B \subseteq X$ *,*

$$A \preceq_R B$$

iff

$$A \subseteq B.$$

The induced expansion operator is

$$\mathcal{E}(A) y : \exists x \in A, ; R(x, y).$$

A.2 Monotonicity Transformations

Definition A.3 (Orientation Preserving Map)**.** *A map*

$$T : X \to X$$

is orientation preserving iff

$$R(x, y) \Longrightarrow R(T(x), T(y)).$$

Definition A.4 (Orientation Reversing Map)**.** *A map*

$$T : X \to X$$

is orientation reversing iff

$$R(x, y) \Longrightarrow R(T(y), T(x)).$$

Definition A.5 (Monotone Operator)**.** *A transformation* T *is monotone iff it preserves inferential inclusion:*

$$A \subseteq B \Longrightarrow T(A) \subseteq T(B).$$

Definition A.6 (Antitone Operator)**.** *A transformation* T *is antitone iff*

$$A \subseteq B \Longrightarrow T(B) \subseteq T(A).$$

A.3 Reachability Characterization

Theorem A.7 (Monotonicity-Reachability Equivalence)**.** *Let* $T{:}X{\to}X$ *.*
Then T *is monotone iff* T *is orientation preserving.*

Proof. Assume T is orientation preserving.
Let

$$A \subseteq B.$$

Take

$$y \in T(A).$$

Then

$$y = T(x)$$

for some $x{\in}A$.
Since $A{\subseteq}B$,

$$x \in B,$$

hence

$$y \in T(B).$$

Therefore

$$T(A) \subseteq T(B).$$

Thus T is monotone.

Conversely suppose T is monotone.

Let

$$R(x, y).$$

Then

$$x \subseteq x, y.$$

Monotonicity implies

$$T(x) \subseteq T(x, y).$$

Therefore

$$R(T(x), T(y))$$

and orientation is preserved. □

Theorem A.8 (Antitonicity-Reversal Equivalence)**.** *A transformation is antitone iff it reverses orientation.*

Proof. Suppose

$$A \subseteq B.$$

Antitonicity gives

$$T(B) \subseteq T(A).$$

Applying this to singleton-generated neighborhoods yields

$$R(x, y) \implies R(T(y), T(x)).$$

The converse follows identically. □

A.4 Negation as Orientation Reversal

Definition A.9 (Negation Transformation)**.** *A negation operator is an involutive antitone map*

$$N : X \to X$$

satisfying

$$N^2 = I.$$

Lemma A.10. *For every reachability relation* R *,*

$$R(x, y) \implies R(N(y), N(x)).$$

Proof. Immediate from antitonicity. □

Theorem A.11 (Double-Reversal Restoration)**.** *For every reachability space* (X,R) *,*

$$N^2 = I$$

implies restoration of inferential orientation.

Proof. Let

$$R(x, y).$$

Applying N ,

$$R(N(y), N(x)).$$

Applying N again,

$$R(N^2(x), N^2(y)).$$

Since

$$N^2 = I,$$

we obtain

$$R(x, y).$$

Thus original orientation is recovered. □

Corollary A.12 (Double Negation Preservation)**.** *For every subset* $A \subseteq X$ *,*

$$N(N(A))A.$$

Proof. Directly from $N^2 = I$. □

A.5 Orientation Classes

Define

$$\mathcal{M}^+T : T \text{ preserves orientation}$$

and

$$\mathcal{M}^-T : T \text{ reverses orientation.}$$

Proposition A.13.

$$\mathcal{M}^+$$

forms a monoid under composition.

Proof. Identity preserves orientation.
Composition of orientation-preserving maps preserves orientation.
Associativity follows from function composition. □

Proposition A.14. *The composition law satisfies*

$$(+)(+) = (+),$$
$$(+)(-) = (-),$$
$$(-)(+) = (-),$$
$$(-)(-) = (+).$$

Proof. Immediate from repeated application of orientation-preservation and orientation-reversal definitions. □

Corollary A.15 (Parity Law)**.** *The net orientation of a sequence of negations depends only on the parity of the number of reversals.*

Proof. Repeated application of

$$(-)(-) = (+).$$

Even parity yields preservation.
Odd parity yields reversal. □

A.6 Reachability Interpretation of Monotonicity

Classical monotonicity theory is therefore recovered as a special case of orientation behavior on reachability spaces:

$$\text{upward-entailing} \iff \text{orientation preserving,}$$
$$\text{downward-entailing} \iff \text{orientation reversing.}$$

Negation is not fundamentally a truth-value operator.

It is an involutive reversal of inferential orientation.

Monotonicity is therefore a property of reachability geometry rather than lexical polarity.

B Geometry of Orientation Boundaries

B.1 Orientation Manifolds

Let

$$(M, g)$$

be a smooth connected Riemannian manifold.

Definition B.1 (Orientation Field)**.** *An orientation field is a smooth scalar function*

$$\omega : M \to \mathbb{R}.$$

Positive values correspond to orientation-preserving regions, negative values correspond to orientation-reversing regions, and the boundary set is

$$\partial\Omega x \in M : \omega(x) = 0.$$

Define

$$\Omega^+ x : \omega(x) > 0$$

and

$$\Omega^- x : \omega(x) < 0.$$

Then

$$M = \Omega^+ \cup \partial\Omega \cup \Omega^-.$$

B.2 Orientation Gradient

Definition B.2 (Orientation Gradient)**.** *The local orientation gradient is*

$$\nabla\omega.$$

Its magnitude

$$G(x)|\nabla\omega(x)|$$

measures the rate of orientation change.

Definition B.3 (Orientation Energy)**.** *The orientation energy of a region* $U \subseteq M$ *is*

$$E_\omega(U) \int_U |\nabla\omega|^2, dV.$$

Theorem B.4 (Minimum-Energy Orientation)**.** *Among all fields satisfying fixed boundary conditions,*

$$E_\omega$$

is minimized iff

$$\Delta\omega = 0.$$

Proof. The Euler-Lagrange equation for

$$E_\omega \int |\nabla\omega|^2, dV$$

is

$$\Delta\omega = 0.$$

Hence minimizers are harmonic. □

B.3 Boundary Curvature

Assume

$$\nabla\omega \neq 0$$

on

$$\partial\Omega.$$

Then $\partial\Omega$ is a smooth hypersurface.

Definition B.5 (Unit Orientation Normal)**.**

$$n\frac{\nabla\omega}{|\nabla\omega|}.$$

Definition B.6 (Orientation Curvature)**.** *The orientation curvature is*

$$\kappa_\omega\nabla\cdot n.$$

Equivalently,

$$\kappa_\omega\nabla\cdot\left(\frac{\nabla\omega}{|\nabla\omega|}\right).$$

Proposition B.7. *If* $\partial\Omega$ *is locally planar,*

$$\kappa_\omega = 0.$$

Proof. For a planar boundary the unit normal is constant.

Hence

$$\nabla\cdot n = 0.$$

□

B.4 Curvature Concentration

Define

$$C(U)\int_U |\kappa_\omega|, dV.$$

Definition B.8 (Boundary Complexity)**.** *The total boundary complexity is*

$$C_\omega\int_{\partial\Omega} |\kappa_\omega|, dS.$$

Theorem B.9. *For every compact boundary,*

$$C_\omega \geq 0.$$

Equality holds iff every connected component of $\partial\Omega$ *is flat.*

Proof. Absolute values are nonnegative.

Vanishing integral implies

$$|\kappa_\omega| = 0$$

almost everywhere.

Thus

$$\kappa_\omega = 0.$$

□

B.5 Orientation Distortion Functional

Let

$$T : M \to M$$

be a deformation.

Definition B.10 (Orientation Distortion)**.**

$$D_\omega(T) \int_M |\omega(x) - \omega(Tx)|, dV.$$

Definition B.11 (Gradient Distortion)**.**

$$D_\nabla(T) \int_M |\nabla\omega(x)$$

$\nabla\omega$*(Tx) | ,dV.*

Definition B.12 (Curvature Distortion)**.**

$$D_\kappa(T) \int_M |\kappa_\omega(x) - \kappa_\omega(Tx)|, dV.$$

B.6 Boundary Cost Functional

Define

$$\mathcal{C}(T)\alpha D_\omega(T) + \beta D_\nabla(T) + \gamma D_\kappa(T)$$

for positive constants

$$\alpha, \beta, \gamma > 0.$$

Theorem B.13 (Boundary Cost Decomposition)**.** *Every smooth orientation deformation admits the decomposition*

$$\mathcal{C}(T)\mathcal{C}_{field} + \mathcal{C}_{gradient} + \mathcal{C}_{curvature}.$$

Proof. Immediate from linearity of integration. □

B.7 Boundary Amplification Theorem

Theorem B.14. *Let*

$$T_\varepsilon I + \varepsilon V$$

be a smooth perturbation.
Then

$$D_\omega(T_\varepsilon) O(\varepsilon)$$

while

$$D_\kappa(T_\varepsilon) O(\varepsilon|\nabla\kappa_\omega|).$$

Consequently regions of high curvature produce disproportionately large distortion.

Proof. Taylor expansion gives

$$\omega(T_\varepsilon x)\omega(x) + \varepsilon\nabla\omega \cdot V + O(\varepsilon^2).$$

Hence

$$D_\omega O(\varepsilon).$$

Similarly,

$$\kappa_\omega(T_\varepsilon x)\kappa_\omega(x) + \varepsilon\nabla\kappa_\omega \cdot V + O(\varepsilon^2).$$

Therefore

$$D_\kappa O(\varepsilon|\nabla\kappa_\omega|).$$

Regions with rapidly varying curvature amplify distortion. □

B.8 Singular Boundaries

Definition B.15 (Orientation Singularity)**.** *A point*

$$x_0$$

is singular iff

$$\omega(x_0) = 0$$

and

$$\nabla\omega(x_0) = 0.$$

Proposition B.16. *Orientation curvature becomes undefined at singular points.*

Proof. The normal field

$$n\frac{\nabla\omega}{|\nabla\omega|}$$

cannot be defined. □

Definition B.17 (Curvature Defect)**.** *For a neighborhood U),*

$$\delta_\kappa(U)\int_U |\nabla\kappa_\omega|^2, dV.$$

B.9 Global Orientation Geometry

Definition B.18 (Total Orientation Action)**.**

$$S_\omega\int_M \left(\lambda_1|\nabla\omega|^2+\lambda_2\kappa_\omega^2\right)dV.$$

Theorem B.19. *Critical points of*

$$S_\omega$$

satisfy a fourth-order geometric field equation

$$\Delta\omega\lambda\Delta\kappa_\omega$$

0

for suitable λ.

Proof. Applying variational calculus to

$$S_\omega$$

produces the Euler-Lagrange equations involving both first-order gradient terms and second-order curvature terms.

Combining them yields the stated fourth-order equation. □

B.10 Orientation Boundary Principle

Let

$$\mathcal{P}(x)\alpha|\omega(x)| + \beta|\nabla\omega(x)| + \gamma|\kappa_\omega(x)|.$$

Theorem B.20 (Boundary Cost Principle)**.** *Local orientation complexity is controlled by the triple*

$$(\omega, \nabla\omega, \kappa_\omega).$$

In particular,

$$\mathcal{P}(x) \to \infty$$

whenever gradient concentration or curvature concentration diverges.

Thus orientation boundaries act as geometric singularities whose distortion cost is determined by both gradient magnitude and boundary curvature.

Proof. Immediate from the definition of

$$\mathcal{P}.$$

Divergence of either term forces divergence of the total complexity measure. □

C Negation Distortion and Reachability Metrics

C.1 Metric Reachability Spaces

Let

$$(X, R)$$

be a reachability space.

Definition C.1 (Reachability Path)**.** *A reachability path from* $x)toy)isafinitesequence$

$$\gamma = (x_0, x_1, \ldots, x_n)$$

such that

$$x_0 = x, \qquad x_n = y,$$

and

$$R(x_i, x_{i+1})$$

for all i*.*

Definition C.2 (Path Length)**.** *Let*

$$w : X \times X \to \mathbb{R}_{\geq 0}$$

be an edge-weight function.

The length of a path is

$$L(\gamma) \sum_{i=0}^{n-1} w(x_i, x_{i+1}).$$

Definition C.3 (Reachability Metric)**.** *The reachability metric is*

$$d_R(x, y) \inf_{\gamma : x \rightsquigarrow y} L(\gamma).$$

Theorem C.4. *If*

$$w(x, y) > 0$$

for all distinct reachable pairs, then

$$d_R$$

is a metric on each connected component.

Proof. Non-negativity follows from positivity of weights.

Identity follows because only the empty path has length zero.

Symmetry is obtained by considering the undirected closure.

Triangle inequality follows from path concatenation. □

C.2 Inferential Geodesics

Definition C.5 (Geodesic Reachability Path)**.** *A path*

$$\gamma^\star$$

is geodesic iff

$$L(\gamma^\star) d_R(x, y).$$

Definition C.6 (Inferential Geodesic Set)**.**

$$\Gamma(x, y) \gamma : L(\gamma) = d_R(x, y).$$

Proposition C.7. *If* X *is finite then*

$$\Gamma(x, y) \neq \emptyset.$$

Proof. The infimum is attained because only finitely many path lengths exist. □

C.3 Negation Distortion

Let

$$N : X \to X$$

be an involutive orientation reversal.

$$N^2 = I.$$

Definition C.8 (Pointwise Negation Distortion)**.**

$$D_N(x)d_R(x, N(x)).$$

Definition C.9 (Average Negation Distortion)**.** *For probability measure* μ),

$$\bar{D}_N \int_X D_N(x), d\mu(x).$$

Definition C.10 (Maximum Negation Distortion)**.**

$$D_N^{\max} \sup_{x \in X} D_N(x).$$

C.4 Distortion Geometry

Theorem C.11 (Zero Distortion Criterion)**.**

$$D_N(x) = 0$$

iff

$$N(x) = x.$$

Proof. Since $d_R)isametric$,

$$d_R(x, y) = 0$$

iff

$$x = y.$$

Substitute

$$y = N(x).$$

□

Definition C.12 (Negation Orbit). *The orbit of x under negation is*

$$\mathcal{O}_N(x)x, N(x).$$

Proposition C.13. *Every orbit has cardinality*

$$1 \quad or \quad 2.$$

Proof. Since

$$N^2 = I,$$

higher-order cycles cannot occur. □

C.5 Distortion Tensor

Assume $X)isasmoothmanifold.$

Let

$$g_{ij}$$

denote the reachability metric tensor.

Definition C.14 (Negation Distortion Tensor).

$$T_{ij}g_{ij}$$

$N^*(g_{ij})$.

Proposition C.15.

$$T_{ij} = 0$$

iff negation acts as an isometry.

Proof. The pullback metric equals the original metric precisely when N is an isometry. □

C.6 Integrated Distortion Action

Definition C.16. *The total distortion action is*

$$S_D \int_X D_N(x)^2, d\mu(x).$$

Theorem C.17.

$$S_D \geq 0.$$

Equality holds iff negation is the identity transformation.

Proof. The integrand is nonnegative.

Vanishing integral implies

$$D_N(x) = 0$$

almost everywhere.

Hence

$$N(x) = x.$$

□

C.7 Reachability Displacement

Definition C.18 (Displacement Field)**.**

$$\delta_N(x)N(x) - x.$$

For smooth coordinates,

$$|\delta_N(x)|d_R(x, N(x)).$$

Definition C.19 (Mean Squared Displacement)**.**

$$\Sigma_N \int_X |\delta_N(x)|^2, d\mu(x).$$

Proposition C.20.

$$\Sigma_N = S_D.$$

Proof. By definition of displacement magnitude. □

C.8 Double Negation Restoration

Theorem C.21 (Metric Restoration Theorem)**.** *For every point x,*

$$D_{N^2}(x) = 0.$$

Proof. Since

$$N^2(x) = x,$$

we have

$$D_{N^2}(x)d_R(x, N^2(x))$$

$\mathrm{d}_R(x, x)$

0.

□

Corollary C.22.

$$S_{N^2} = 0.$$

Proof. The integrand vanishes identically. □

C.9 Orientation-Reversal Geodesics

Definition C.23. *The negation geodesic associated with x is*

$$\gamma_N(x) \operatorname{argmin}_\gamma L(\gamma)$$

subject to

$$\gamma(0) = x, \qquad \gamma(1) = N(x).$$

Definition C.24 (Negation Length)**.**

$$\ell_N(x) L(\gamma_N(x)).$$

Proposition C.25.

$$\ell_N(x) D_N(x).$$

Proof. By geodesic optimality. □

C.10 Curvature of Negation

Let

$$\gamma_N(t)$$

be a smooth negation geodesic.

Define

$$v(t) \frac{d\gamma_N}{dt}.$$

Definition C.26 (Negation Curvature)**.**

$$K_N \left| \frac{Dv}{dt} \right|.$$

Definition C.27 (Integrated Curvature)**.**

$$\mathcal{K}_N \int_{\gamma_N} K_N, ds.$$

Theorem C.28. *Straight-line negation trajectories satisfy*

$$\mathcal{K}_N = 0.$$

Proof. A geodesic with constant tangent satisfies

$$\frac{Dv}{dt} = 0.$$

□

C.11 Spectral Distortion

Let

$$\Delta_R$$

denote the reachability Laplacian.

Suppose

$$\Delta_R \phi_k \lambda_k \phi_k.$$

Definition C.29 (Spectral Negation Distortion)**.**

$$\Lambda_N \sum_k \lambda_k |\phi_k - N(\phi_k)|^2.$$

Theorem C.30.

$$\Lambda_N = 0$$

iff every Laplacian eigenmode is invariant under negation.

Proof. Every term is nonnegative.

The sum vanishes only if

$$\phi_k = N(\phi_k)$$

for all k. □

C.12 Distortion Principle

Theorem C.31 (Negation Distortion Principle)**.** *For every metric reachability space,*

$$D_N(x)d_R(x, N(x))$$

measures the inferential displacement induced by orientation reversal. Furthermore,

$$D_{N^2}(x) = 0,$$

so double negation is a complete restoration of reachability geometry rather than a mere cancellation of logical operators.

Proof. The first statement follows by definition.

The second follows from the Metric Restoration Theorem. □

D Persistence of Negation Distortion

D.1 Memory Reachability Systems

Let

$$(X, R, \mu)$$

be a metric reachability space equipped with probability measure μ.

Let

$$N : X \to X$$

denote an involutive orientation reversal.

$$N^2 = I.$$

Definition D.1 (Memory Reachability System)**.** *A memory reachability system is the quintuple*

$$\mathcal{M}(X, R, N, M_t, \mu)$$

where

$$M_t : X \to X$$

denotes the state of memory at time t *.*

Definition D.2 (Memory Encoding). *A memory encoding operator is*

$$E : X \to \mathcal{H}$$

where $\mathcal{H}$ *is a memory state space.*

Definition D.3 (Memory Retrieval). *A retrieval operator is*

$$Q : \mathcal{H} \to X.$$

Definition D.4 (Memory Loop). *The memory loop is*

$$\mathcal{L}Q \circ E.$$

Perfect persistence corresponds to

$$Q(E(x))x.$$

D.2 Distortion Memory States

Recall

$$D_N(x)d_R(x, N(x)).$$

Definition D.5 (Stored Distortion). *The stored distortion state is*

$$\Delta_t(x)D_N(M_t(x)).$$

Definition D.6 (Global Distortion Energy).

$$\mathcal{D}_t \int_X \Delta_t(x)^2, d\mu(x).$$

D.3 Persistence Operators

Definition D.7 (Persistence Operator). *A persistence operator is a semigroup*

$$P_t : X \to X$$

satisfying

$$P_0 = I$$

and

$$P_t P_s = P_{t+s}.$$

Definition D.8 (Perfect Persistence)**.** *Persistence is perfect iff*

$$P_t(x) = x$$

for all t.

Definition D.9 (Lossless Memory)**.** *A memory system is lossless iff*

$$M_t = P_t.$$

and

$$d_R(M_t(x), x) = 0.$$

D.4 Distortion Conservation

Theorem D.10 (Distortion Conservation)**.** *Suppose memory is lossless.*
Then

$$\Delta_t(x)\Delta_0(x)$$

for every t.

Proof. Losslessness implies

$$M_t(x) = x.$$

Therefore

$$\Delta_t(x) D_N(M_t(x))$$

$\mathrm{D}_N(x)$
$\Delta_0(x)$. □

Corollary D.11.

$$\mathcal{D}_t\mathcal{D}_0.$$

Proof. Integrate the previous result. □

D.5 Memory Decay

Definition D.12 (Decay Generator)**.** *A decay process is generated by*

$$L : X \to TX$$

through

$$\frac{dM_t}{dt} L(M_t).$$

Definition D.13 (Exponential Forgetting)**.** *A memory trace obeys*

$$\frac{dm}{dt} - \lambda m.$$

with solution

$$m(t) m(0) e^{-\lambda t}.$$

Theorem D.14. *Under exponential forgetting,*

$$\mathcal{D}_t \mathcal{D}_0 e^{-2\lambda t}.$$

Proof.

$$\Delta_t \Delta_0 e^{-\lambda t}.$$

Squaring and integrating gives

$$\mathcal{D}_t \mathcal{D}_0 e^{-2\lambda t}.$$

□

D.6 Ecphory Dynamics

Definition D.15 (Ecphory Operator)**.** *An ecphory operator is*

$$\mathcal{E}_t : \mathcal{H} \to X.$$

Definition D.16 (Ecphory Threshold)**.** *A threshold parameter*

$$\theta > 0$$

defines retrieval success:

$$\mathcal{E}_t(h) x$$

iff

$$|h - E(x)| < \theta.$$

Definition D.17 (Retrieval Error)**.**

$$\varepsilon_t(x) d_R\big(x, \mathcal{E}_t(E(x))\big).$$

Theorem D.18. *Perfect ecphory implies*

$$\varepsilon_t(x) = 0.$$

Proof. If retrieval reconstructs x exactly,

$$d_R(x, x) = 0.$$

□

D.7 Persistence of Orientation Structure

Let

$$\omega_t$$

denote the orientation field represented in memory.

Definition D.19 (Orientation Persistence)**.** *Orientation persists iff*

$$\omega_t \omega_0.$$

Theorem D.20 (Orientation Persistence Theorem)**.** *If memory preserves orientation structure,*

$$\omega_t = \omega_0,$$

then negation distortion remains invariant.

Proof. Negation distortion depends only upon

$$\omega.$$

If

$$\omega_t = \omega_0,$$

then

$$D_N^t D_N^0.$$

□

D.8 Delayed Verification Dynamics

Let

$$V_t$$

denote verification state.

Definition D.21 (Verification Functional)**.**

$$V_t(x)\chi(P, M_t(x))$$

where

$$\chi$$

is a consistency functional.

Definition D.22 (Verification Latency)**.**

$$\tau \inf t : V_t(x) = 1.$$

Theorem D.23. *If distortion persists,*

$$\Delta_t = \Delta_0,$$

then verification latency remains bounded below by

$$\tau \geq c\Delta_0$$

for some constant $c > 0$.

Proof. Verification requires reconstruction of orientation.

The amount of reconstruction scales with stored distortion.

Hence

$$\tau \propto \Delta_0.$$

Absorb proportionality into c. □

D.9 Distortion Persistence Spectrum

Let

$$K$$

be the memory evolution operator.

Assume

$$K\phi_n\lambda_n\phi_n.$$

Definition D.24 (Persistence Spectrum)**.**

$$\sigma_P \lambda_n.$$

Definition D.25 (Distortion Spectrum)**.**

$$\Sigma_D \lambda_n : \phi_n \textit{ contributes to } D_N.$$

Theorem D.26. *Modes satisfying*

$$|\lambda_n| = 1$$

produce permanent distortion persistence.

Proof. Such modes do not decay under repeated application of K . □

D.10 Recoverability Geometry

Definition D.27 (Recoverability Radius)**.**

$$\rho(x) \sup r : B_r(x) \textit{ is reconstructible.}$$

Definition D.28 (Persistence Capacity)**.**

$$\Phi_P \int_X \rho(x), d\mu(x).$$

Theorem D.29. *If*

$$\Phi_P$$

$D_N^{\max}$,
then all negation distortions are recoverable.

Proof. Every distortion lies within the recoverability radius.
Hence reconstruction exists. □

D.11 Delayed Verification Theorem

Theorem D.30 (Delayed Verification Theorem)**.** *Suppose*

$$N$$

introduces distortion

$$D_N(x) > 0.$$

Assume memory preserves orientation geometry:

$$\omega_t = \omega_0.$$

Then

$$\Delta_t(x) D_N(x)$$

for all t *,*
and verification latency satisfies

$$\tau(x) \geq c D_N(x).$$

Consequently delayed-verification effects arise naturally from persistence of orientation distortion rather than from repeated syntactic computation.

Proof. Distortion conservation gives

$$\Delta_t(x) = D_N(x).$$

The latency bound follows from the previous theorem.

Thus verification cost scales with persistent geometric distortion stored in memory. □

D.12 Persistence Principle

Theorem D.31 (Persistence Principle)**.** *Memory preserves not propositions but distortions of reachability geometry.*

Negation-induced displacement survives in memory whenever orientation structure survives.

Therefore the persistence of verification cost is equivalent to the persistence of geometric distortion.

Proof. All preceding results establish that distortion is a function of orientation geometry and that orientation persistence implies distortion persistence.

Verification cost depends on distortion.

Hence persistence of cost follows directly from persistence of geometry. □

E NPI Licensing as Admissibility Geometry

E.1 Admissibility Spaces

Let

$$(M, g)$$

be a connected Riemannian manifold.

Let

$$\omega : M \to \mathbb{R}$$

denote an orientation field.

Definition E.1 (Admissibility Field)**.** *An admissibility field is a smooth function*

$$A : M \to [0, 1].$$

Definition E.2 (Admissible Region)**.** *The admissible region is*

$$\mathcal{A}x \in M : A(x) > 0.$$

Definition E.3 (Fully Admissible Region)**.**

$$\mathcal{A}_1 x : A(x) = 1.$$

Definition E.4 (Admissibility Boundary)**.**

$$\partial\mathcal{A}x : A(x) = 0.$$

E.2 Licensing Domains

Definition E.5 (Licensing Domain)**.** *The licensing domain is*

$$Lx \in M : A(x) = 1, , \omega(x) < 0.$$

Definition E.6 (Positive Domain)**.**

$$Px : A(x) = 1, , \omega(x) > 0.$$

Definition E.7 (Neutral Boundary)**.**

$$Bx : A(x) = 1, , \omega(x) = 0.$$

Thus

$$\mathcal{A}_1 P \cup B \cup L.$$

E.3 NPI Probe Functions

Definition E.8 (NPI Probe)**.** *An NPI is represented by a probe function*

$$\Pi : M \to 0, 1.$$

Definition E.9 (Licensing Condition)**.** *A probe is licensed at* x *iff*

$$\Pi(x) = 1.$$

and

$$x \in L.$$

Equivalently,

$$\Pi(x)\chi_L(x),$$

where

$$\chi_L$$

is the characteristic function of the licensing region.

E.4 Geometric Licensing Criterion

Theorem E.10 (Licensing Criterion)**.** *An NPI is licensed at* x *iff*

$$A(x) = 1$$

and

$$\omega(x) < 0.$$

Proof. By definition of L ,

$$x \in L$$

iff both conditions hold.

Since

$$\Pi = \chi_L,$$

licensing follows immediately. □

E.5 Connected Licensing Components

Definition E.11 (Licensing Component)**.** *A licensing component is a connected component*

$$L_i \subseteq L.$$

Then

$$L \bigsqcup_i L_i.$$

Theorem E.12 (Component Licensing Theorem)**.** *Every point of a licensing component licenses NPIs.*

Proof. If

$$x \in L_i,$$

then

$$x \in L.$$

Hence

$$A(x) = 1, \qquad \omega(x) < 0.$$

Therefore licensing holds. □

E.6 Boundary Geometry

Define

$$\kappa_A \nabla \cdot \left(\frac{\nabla A}{|\nabla A|} \right).$$

Definition E.13 (Admissibility Curvature)**.**

$$\kappa_A$$

is the admissibility curvature.

Similarly define orientation curvature

$$\kappa_\omega.$$

Definition E.14 (Licensing Boundary Curvature)**.**

$$\kappa_L |\kappa_A| + |\kappa_\omega|.$$

Theorem E.15. *Licensing uncertainty increases monotonically with*

$$\kappa_L.$$

Proof. Large curvature implies rapid local variation of either admissibility or orientation.

Therefore infinitesimal perturbations alter membership in L .

Hence uncertainty increases. ☐

E.7 Distance to Licensing

Definition E.16 (Licensing Distance)**.**

$$d_L(x) \inf_{y \in L} d(x, y).$$

Definition E.17 (Licensing Potential)**.**

$$V_L(x) e^{-d_L(x)}.$$

Theorem E.18.

$$V_L(x) 1$$

iff

$$x \in \overline{L}.$$

Proof. Distance vanishes precisely on the closure.

Hence

$$e^{-0} = 1.$$

☐

E.8 Licensing Flux

Let

$$n$$

denote the outward normal of a licensing boundary.

Definition E.19 (Licensing Flux)**.**

$$\Phi_L \int_{\partial L} \nabla \omega \cdot n, dS.$$

Theorem E.20. *Positive licensing flux corresponds to outward expansion of licensing domains.*

Proof. By the divergence theorem,

$$\Phi_L \int_L \Delta\omega, dV.$$

Positive flux implies net outward orientation flow. □

E.9 Topological Licensing Structure

Let

$$H_k(L)$$

denote singular homology groups.

Definition E.21 (Licensing Connectivity Number)**.**

$$\beta_0 \operatorname{rank}(H_0(L)).$$

Definition E.22 (Licensing Loop Number)**.**

$$\beta_1 \operatorname{rank}(H_1(L)).$$

Theorem E.23.

$$\beta_0$$

equals the number of disconnected licensing regions.

Proof. Standard algebraic topology. □

Theorem E.24. *Nonzero*

$$\beta_1$$

implies cyclic licensing structure.

Proof. A nontrivial first homology group contains non-contractible loops. □

E.10 Spectral Licensing Theory

Let

$$\Delta_L$$

be the Laplace-Beltrami operator restricted to L .
Assume

$$\Delta_L \phi_n \lambda_n \phi_n.$$

Definition E.25 (Licensing Spectrum)**.**

$$\sigma(L)\lambda_n.$$

Definition E.26 (Licensing Capacity)**.**

$$C_L \sum_n \frac{1}{1+\lambda_n}.$$

Theorem E.27. *Larger connected licensing regions possess larger licensing capacity.*

Proof. Expansion lowers low-frequency eigenvalues.

Hence

$$C_L$$

increases. ☐

E.11 Probe Stability

Definition E.28 (Probe Stability)**.** *A probe is stable at* x *if*

$$\exists \epsilon > 0$$

such that

$$B_\epsilon(x) \subseteq L.$$

Theorem E.29. *Stable probes occur precisely in the interior*

$$L^\circ.$$

Proof. Interior points admit neighborhoods contained entirely within L .

Boundary points do not. ☐

E.12 Licensing Action Functional

Define

$$S_L \int_M \Big(|\nabla A|^2 + |\nabla \omega|^2 + A|\omega|\Big), dV.$$

Theorem E.30. *Critical licensing configurations satisfy*

$$\Delta A \frac{1}{2}|\omega|$$

and

$$\Delta \omega \frac{1}{2} A, \operatorname{sgn}(\omega).$$

Proof. Applying the Euler-Lagrange equations to

$$S_L$$

yields the coupled field equations. □

E.13 NPI Probe Theorem

Theorem E.31 (NPI Probe Theorem)**.** *NPIs function as geometric probes of admissible orientation-reversing regions.*

Specifically,

$$\Pi(x) = 1$$

iff

$$x \in L.$$

Consequently NPIs detect the intersection

$$L\mathcal{A}_1 \cap \Omega^-,$$

where admissibility and orientation reversal coexist.

Proof. Directly from the definitions of licensing and the characteristic probe function. □

E.14 Geometric Licensing Principle

Theorem E.32 (Geometric Licensing Principle)**.** *Licensing is not fundamentally a lexical property.*

Licensing is membership in a negatively oriented admissible manifold.

NPIs therefore reveal the geometry of admissibility boundaries rather than merely the presence of syntactic negation.

Proof. Every licensing condition depends only on

$$A(x)$$

and

$$\omega(x).$$

No reference to lexical negation is required.

Therefore licensing is geometric. □

F Neural Dissociation Constraints and Inferential Field Operators

F.1 Neural State Spaces

Let

$$\mathcal{H}$$

be a separable Hilbert space of neural states.

A neural configuration is represented by

$$\psi \in \mathcal{H}.$$

Inner products are denoted

$$\langle \psi, \phi \rangle.$$

Norms are

$$|\psi| \sqrt{\langle \psi, \psi \rangle}.$$

Definition F.1 (Neural Representation Map)**.** *A neural representation is a map*

$$\rho : X \to \mathcal{H}.$$

Definition F.2 (Population State)**.** *The state of a neural population at time* t *is*

$$\psi_t.$$

F.2 Operator Decomposition

Define four operators

$$C, \quad O, \quad N, \quad V.$$

Definition F.3 (Parsing Operator)**.**

$$C : \mathcal{H} \to \mathcal{H}$$

extracts compositional structure.

Definition F.4 (Orientation Operator)**.**

$$O : \mathcal{H} \to \mathcal{H}$$

constructs inferential orientation.

Definition F.5 (Negation Operator)**.**

$$N : \mathcal{H} \to \mathcal{H}$$

implements orientation reversal.

Definition F.6 (Verification Operator)**.**

$$V : \mathcal{H} \to \mathcal{H}$$

computes consistency between representation and scenario.

The complete neural computation is

$$\Psi V N O C.$$

F.3 Sequential Architecture

Given input state

$$\psi_0,$$

define

$$\psi_1 C \psi_0,$$

$$\psi_2 O \psi_1,$$

$$\psi_3 N \psi_2,$$

$$\psi_4 V \psi_3.$$

Thus

$$\psi_4 V N O C \psi_0.$$

F.4 Functional Independence

Definition F.7 (Operator Independence)**.** *Operators*

$$A, B$$

are independent iff

$$A \neq f(B)$$

for every measurable function f*.*

Definition F.8 (Neural Dissociation)**.** *Two operators are neurally dissociable if there exists a perturbation*

$$P$$

such that

$$PA \neq A$$

while

$$PB = B.$$

Theorem F.9 (Dissociation Criterion)**.** *Suppose*

$$A$$

and

$$B$$

are dissociable.

Then they cannot be identical.

Proof. If

$$A = B,$$

then

$$PA = PB.$$

Contradiction. □

F.5 Orientation Construction

Let

$$\omega$$

denote inferential orientation.

Definition F.10 (Orientation Map)**.**

$$\Omega : \mathcal{H} \to \mathbb{R}.$$

Define

$$\omega\Omega(\psi).$$

Definition F.11 (Orientation Population). *A neural population*

$$P_\omega$$

encodes orientation iff

$$\Omega(P_\omega) \neq 0.$$

F.6 Negation Dynamics

Assume

$$N^2 = I.$$

Theorem F.12 (Involution Property). *Every eigenvalue of N satisfies*

$$\lambda = \pm 1.$$

Proof. Let

$$N\phi = \lambda\phi.$$

Applying N ,

$$N^2\phi\lambda^2\phi.$$

Since

$$N^2 = I,$$
$$\lambda^2 = 1.$$

Thus

$$\lambda = \pm 1.$$

□

Definition F.13 (Orientation-Reversal Subspace).

$$\mathcal{H}_-\phi : N\phi = -\phi.$$

Definition F.14 (Orientation-Preserving Subspace).

$$\mathcal{H}_+\phi : N\phi = \phi.$$

Theorem F.15.

$$\mathcal{H}\mathcal{H}_+ \oplus \mathcal{H}_-.$$

Proof. Spectral decomposition of involutions. □

F.7 Verification Geometry

Let

$$S$$

denote a scenario representation.

Definition F.16 (Verification Functional)**.**

$$\chi(\psi, S)\langle\psi, S\rangle.$$

Definition F.17 (Verification Error)**.**

$$E_V|\psi - S|^2.$$

Theorem F.18. *Verification succeeds iff*

$$E_V = 0.$$

Proof. Norms vanish only at equality. □

F.8 Operator Non-Commutativity

Definition F.19. *The commutator of operators A,B is*

$$[A, B]AB - BA.$$

Theorem F.20. *If*

$$[C, N] \neq 0,$$

then parsing and negation cannot be collapsed into a single stage.

Proof. Assume a single-stage operator

$$T.$$

Then

$$CN = NC.$$

Contradiction. □

Theorem F.21. *If*

$$[O, N] \neq 0,$$

orientation construction and negation are computationally distinct.

Proof. Identical. □

Theorem F.22. *If*

$$[V, N] \neq 0,$$

verification cannot precede orientation reversal.

Proof. Order affects output.
Therefore stages are distinct. □

F.9 Perturbation Theory

Let

$$\epsilon P$$

be a neural perturbation.
Define

$$N_\epsilon N + \epsilon P.$$

Theorem F.23. *Eigenvalues satisfy*

$$\lambda_i(\epsilon)\lambda_i + \epsilon\langle\phi_i, P\phi_i\rangle + O(\epsilon^2).$$

Proof. Standard first-order perturbation theory. □

F.10 Identifiability

Definition F.24 (Operator Identifiability)**.** *An operator A is identifiable iff*

$$A\psi = B\psi$$

for all ψ
implies

$$A = B.$$

Theorem F.25. *Distinct neural operators are identifiable whenever the representation space spans H.*

Proof. Operator equality on a spanning set implies equality everywhere. □

F.11 Minimal Neural Architecture

Define

$$\mathfrak{A}C, O, N, V.$$

Definition F.26 (Minimal Architecture)**.**

$$\mathfrak{A}$$

is minimal iff no proper subset computes

$$VNOC.$$

Theorem F.27 (Minimality Theorem)**.** *If*

$$[C, O] \neq 0,$$
$$[C, N] \neq 0,$$
$$[O, N] \neq 0,$$

and

$$[V, N] \neq 0,$$

then

$$\mathfrak{A}$$

is minimal.

Proof. Removing any operator changes the resulting composition.

Therefore no proper subset reproduces

$$VNOC.$$

□

F.12 Neural Dissociation Theorem

Theorem F.28 (Neural Dissociation Theorem)**.** *Suppose there exist neural populations*

$$P_C, \quad P_O, \quad P_N, \quad P_V$$

such that selective perturbations affect exactly one operator among

$$C, O, N, V.$$

Then parsing, orientation construction, negation, and verification are computationally distinct processes.

Proof. Selective perturbation establishes pairwise dissociation.

By the Dissociation Criterion, dissociated operators cannot be identical.

Therefore

$$C \neq O, \qquad C \neq N, \qquad C \neq V,$$
$$O \neq N, \qquad O \neq V, \qquad N \neq V.$$

Hence the four stages are distinct. □

F.13 Strong Dissociation Corollary

Corollary F.29. *No theory identifying negation solely with syntax, solely with semantic verification, or solely with orientation construction can reproduce the full operator structure*

$$VNOC.$$

Proof. Each reduction removes at least one dissociable operator.

Minimality then fails. □

F.14 Operator Separation Principle

Theorem F.30 (Operator Separation Principle)**.** *If parsing, orientation construction, negation, and verification admit independent neural realizations, then inferential computation necessarily possesses at least four irreducible stages.*

Consequently negation cannot be reduced to syntax, semantics, or verification alone.

Proof. Follows from operator identifiability, non-commutativity, and neural dissociation. □

G Inferential Field Semantics and Classical Logic

G.1 Inferential Fields

Let

$$\mathcal{F}(X, R, \omega, A)$$

be an inferential field where

$$X$$

is a distinction manifold,

$$R \subseteq X \times X$$

is a reachability relation,

$$\omega : X \to \mathbb{R}$$

is an orientation field,
and

$$A : X \to [0,1]$$

is an admissibility field.

Definition G.1 (Verification Functional)**.** *A proposition* P *is verified at* $x \in X$ *iff*

$$\chi(P, x) = 1.$$

Definition G.2 (Truth)**.**

$$\text{True}(P, x)$$

iff

$$A(x) = 1$$

and

$$\chi(P, x) = 1.$$

Thus truth is admissible verification.

G.2 Logical Interpretation

Let

$$\mathcal{L}$$

denote classical propositional logic.
Define

$$\mathcal{T} : \mathcal{L} \to \mathcal{F}$$

by

$$P \mapsto x : \chi(P, x) = 1.$$

Definition G.3 (Interpretation Functor)**.** *The map*

$$\mathcal{T}$$

is called the inferential interpretation functor.

G.3 Conjunction

Definition G.4. *For propositions*

$$P, Q,$$

define

$$P \wedge Q$$

by

$$\chi(P \wedge Q, x)\chi(P, x)\chi(Q, x).$$

Theorem G.5.

$$\mathcal{T}(P \wedge Q)\mathcal{T}(P) \cap \mathcal{T}(Q).$$

Proof.

$$\chi(P \wedge Q, x) = 1$$

iff

$$\chi(P, x) = 1$$

and

$$\chi(Q, x) = 1.$$

Hence

$$x \in \mathcal{T}(P) \cap \mathcal{T}(Q).$$

□

G.4 Disjunction

Definition G.6.

$$\chi(P \vee Q, x) \max \chi(P, x), \chi(Q, x).$$

Theorem G.7.

$$\mathcal{T}(P \vee Q)\mathcal{T}(P) \cup \mathcal{T}(Q).$$

Proof. Immediate. □

G.5 Implication

Definition G.8 (Reachability Implication)**.**

$$P \Rightarrow Q$$

iff

$$\mathcal{T}(P) \subseteq \mathcal{T}(Q).$$

Definition G.9 (Inferential Entailment)**.**

$$P \models_R Q$$

iff

$$R(x, y)$$

and

$$\chi(P, x) = 1$$

imply

$$\chi(Q, y) = 1.$$

Theorem G.10. *Inferential entailment is transitive.*

Proof. Suppose

$$P \models_R Q$$

and

$$Q \models_R S.$$

Then

$$P \to Q \to S.$$

Hence

$$P \models_R S.$$

□

G.6 Negation

Let

$$N : X \to X$$

be orientation reversal.

$$N^2 = I.$$

Definition G.11 (Inferential Negation)**.**

$$\chi(\neg P, x)\chi(P, N(x)).$$

Theorem G.12. *Negation corresponds to orientation reversal.*

Proof. By definition,

$$\neg P$$

is evaluated at

$$N(x).$$

Therefore truth evaluation occurs after orientation reversal. □

G.7 Double Negation

Theorem G.13 (Double Negation Restoration)**.**

$$\chi(\neg\neg P, x)\chi(P, x).$$

Proof.

$$\chi(\neg\neg P, x)\chi(\neg P, N(x))$$

χ(P,N$^2(x)$).
Since

$$N^2 = I,$$

$$\chi(\neg\neg P, x)\chi(P, x).$$

□

Corollary G.14.

$$\neg\neg P \equiv P.$$

G.8 Truth Regions

Define

$$T(P)x : \text{True}(P, x).$$

Theorem G.15.

$$T(P)\mathcal{A}_1 \cap \mathcal{T}(P).$$

Proof. Truth requires both admissibility and verification. □

G.9 Soundness

Theorem G.16 (Soundness)**.** *If*

$$\vdash P$$

in classical propositional logic, then

$$\models_{\mathcal{F}} P.$$

Proof. Each logical connective preserves its set-theoretic interpretation.

Classical derivations therefore preserve truth regions.

Hence every theorem is valid in inferential fields. □

G.10 Completeness

Theorem G.17 (Completeness)**.** *If*

$$\models_{\mathcal{F}} P,$$

then

$$\vdash P.$$

Proof. Truth regions form a Boolean algebra under

$$\cap, \cup, \text{and } N.$$

Stone representation implies equivalence with classical valuations.

Therefore every inferentially valid proposition is classically derivable. □

G.11 Boolean Structure

Theorem G.18. *The collection*

$$\mathfrak{B}T(P) : P \in \mathcal{L}$$

forms a Boolean algebra.

Proof. Closure under

$$\cap, \cup, N$$

follows from conjunction, disjunction, and negation.

Classical Boolean identities are inherited. □

G.12 Orientation-Neutral Regions

Define

$$Bx : \omega(x) = 0.$$

Definition G.19 (Boundary Region)**.** *B is the orientation-neutral boundary.*

Theorem G.20. *Classical bivalence may fail on B.*

Proof. Orientation reversal becomes undefined or degenerate at

$$\omega = 0.$$

Hence truth-value assignment need not be unique. □

G.13 Intuitionistic Boundary Logic

Definition G.21. *A proposition is stable iff*

$$\omega(x) \neq 0$$

throughout its truth region.

Theorem G.22. *Boundary regions support intuitionistic semantics.*

Proof. Double-negation elimination requires nondegenerate orientation.

At

$$\omega = 0,$$

only

$$P \to \neg\neg P$$

is guaranteed. □

G.14 Modal Reachability

Introduce operators

$$\Box, \quad \Diamond.$$

Definition G.23.

$$\Box P$$

iff

$$\forall y, , R(x,y) \Rightarrow \chi(P,y) = 1.$$

Definition G.24.

$$\Diamond P$$

iff

$$\exists y, , R(x,y) \wedge \chi(P,y) = 1.$$

Theorem G.25.

$$\Box P \Rightarrow \Diamond P.$$

Proof. Reflexivity of R . □

G.15 Reachability Frames

Definition G.26. *The frame associated with an inferential field is*

$$(X, R).$$

Theorem G.27. *Every inferential field determines a Kripke frame.*

Proof. The reachability relation acts as accessibility. □

Theorem G.28. *Every Kripke frame can be embedded into an inferential field.*

Proof. Assign

$$\omega \equiv 1, \qquad A \equiv 1.$$

Then the remaining structure is exactly

$$(X, R).$$

□

G.16 Representation Theorem

Theorem G.29 (Representation Theorem)**.** *Classical propositional logic embeds faithfully into inferential-field semantics via*

$$\mathcal{T}.$$

Logical connectives become geometric operations on admissible verification regions, while negation becomes orientation reversal.

Proof. The preceding conjunction, disjunction, implication, and negation theorems establish preservation of logical structure.

Faithfulness follows from soundness and completeness. □

G.17 Negation Before Logic Theorem

Theorem G.30 (Negation Before Logic)**.** *Let*

$$\mathcal{F}(X, R, \omega, A).$$

Then logical negation is derivable from the existence of an involutive orientation-reversal map

$$N : X \to X.$$

Consequently orientation precedes logical negation, and admissible reachability precedes truth.

Proof. Truth evaluation requires verification.

Verification requires admissibility.

Negation acts on orientation prior to verification.

Therefore

$$A; \prec; \omega; \prec; \neg; \prec; \text{truth}.$$

Hence negation is geometrically prior to logic. □

H Distinguishability Geometry and Inferential Displacement

H.1 Distinguishability Spaces

Let

$$X$$

be a nonempty set.

Definition H.1 (Distinguishability Relation)**.** *A distinguishability relation is a symmetric relation*

$$\not\sim;\subseteq;X\times X.$$

Its complement

$$\sim$$

is the indistinguishability relation.

Definition H.2 (Distinguishability Space)**.** *A distinguishability space is the pair*

$$\mathcal{D}(X,\sim).$$

Definition H.3 (Equivalence Class)**.** *The equivalence class of* x *is*

$$[x]y\in X:y\sim x.$$

Definition H.4 (Quotient Space)**.**

$$X/{\sim}[x]:x\in X.$$

H.2 Distinguishability Metrics

Let

$$d_D$$

be a distinguishability metric.

Definition H.5 (Distinguishability Metric)**.**

$$d_D:X\times X\to\mathbb{R}_{\geq 0}$$

satisfies

$$d_D(x,y)=0\iff x\sim y.$$

Definition H.6 (Distinguishability Neighborhood)**.**

$$B_r(x)y:d_D(x,y)<r.$$

Definition H.7 (Distinguishability Volume)**.**

$$V_D(x,r)\mu(B_r(x)).$$

H.3 Reachability Embedding

Let

$$(X, R)$$

be a reachability space.

Definition H.8 (Reachability-Induced Distinguishability)**.**

$$d_D(x, y)|d_R(x, y) - d_R(y, x)|.$$

Theorem H.9. *Every metric reachability space induces a distinguishability geometry.*

Proof. The induced metric separates points according to asymmetry of reachability.
Thus distinguishability emerges from reachability structure. □

H.4 Distinguishability Deficit

Let

$$L_{\text{opt}}(x)$$

denote optimal reachable distinguishability.
Let

$$L_{\text{act}}(x)$$

denote actual distinguishability.

Definition H.10 (Distinguishability Deficit)**.**

$$\delta_D(x)L_{\text{opt}}(x)$$

$L_{\text{act}}(x)$.

Definition H.11 (Global Deficit)**.**

$$\Delta_D \int_X \delta_D(x), d\mu(x).$$

Theorem H.12.

$$\Delta_D \geq 0.$$

Proof. Optimal distinguishability dominates actual distinguishability.
Hence

$$\delta_D(x) \geq 0.$$

Integrating preserves positivity. □

H.5 Negation as Distinguishability Displacement

Let

$$N : X \to X$$

be orientation reversal.

Definition H.13 (Negation Displacement)**.**

$$\eta_N(x) d_D(x, N(x)).$$

Definition H.14 (Average Negation Displacement)**.**

$$\bar{\eta}_N \int_X \eta_N(x), d\mu(x).$$

Theorem H.15. *Negation distortion is a special case of distinguishability displacement.*

Proof. By definition,

$$D_N(x) d_R(x, N(x)).$$

Since

$$d_D$$

is induced by

$$d_R,$$

negation distortion becomes a distinguishability displacement. □

H.6 Distinguishability Tensor Geometry

Assume X is a smooth manifold.

Let

$$g_{ij}$$

denote the distinguishability metric tensor.

Definition H.16 (Deficit Tensor)**.**

$$\Theta_{ij} g_{ij}^{\text{opt}}$$

g_{ij}^{act}.

Definition H.17 (Tensor Norm)**.**

$$|\Theta|^2 \Theta_{ij} \Theta^{ij}.$$

Theorem H.18.

$$\Theta = 0$$

iff distinguishability is optimal everywhere.

Proof. Vanishing tensor implies

$$g^{\text{opt}} g^{\text{act}}.$$

□

H.7 Distinguishability Curvature

Let

$$\nabla$$

be the Levi-Civita connection.

Definition H.19 (Distinguishability Curvature)**.**

$$R^i{}_{jkl}$$

is the Riemann curvature tensor associated with

$$g_{ij}.$$

Definition H.20 (Scalar Distinguishability Curvature)**.**

$$K_D g^{ij} R_{ij}.$$

Theorem H.21. *Regions of large*

$$|K_D|$$

exhibit rapid distinguishability deformation.

Proof. Curvature measures local deviation from Euclidean distinguishability structure.

□

H.8 Admissibility Quotients

Let

$$A : X \to [0, 1]$$

be admissibility.

Definition H.22 (Admissibility Equivalence)**.**

$$x \sim_A y$$

iff

$$A(x) = A(y).$$

Definition H.23 (Admissibility Quotient)**.**

$$X/{\sim_A}.$$

Theorem H.24. *The admissibility quotient partitions distinguishability space into admissibility strata.*

Proof. Equality of admissibility is an equivalence relation. □

H.9 Distinguishability Entropy

Definition H.25 (Local Distinguishability Entropy)**.**

$$S_D(x) \log V_D(x, r).$$

Definition H.26 (Global Distinguishability Entropy)**.**

$$S_D \int_X S_D(x), d\mu(x).$$

Theorem H.27. *Increasing distinguishability volume increases entropy.*

Proof. The logarithm is monotone. □

H.10 Repair Operators

Let

$$\mathcal{R} : X \to X$$

be a repair operator.

Definition H.28 (Repair Condition)**.**

$$\delta_D(\mathcal{R}(x)) < \delta_D(x).$$

Definition H.29 (Perfect Repair)**.**

$$\delta_D(\mathcal{R}(x))0.$$

Theorem H.30. *Repeated repair generates a monotone deficit sequence*

$$\delta_D^{(n+1)} \leq \delta_D^{(n)}.$$

Proof. Repair reduces deficit by definition. □

H.11 Continuation Geometry

Let

$$C_t$$

denote continuation dynamics.

Definition H.31 (Continuation Functional)**.**

$$\Gamma(x) \int_0^\infty e^{-\lambda t}, L(C_t(x)), dt.$$

Definition H.32 (Continuation Capacity)**.**

$$\Phi_C \int_X \Gamma(x), d\mu(x).$$

Theorem H.33. *Repair increases continuation capacity.*

Proof. Repair restores reachable distinguishability.

Therefore future trajectories increase.

Hence

$$\Phi_C$$

increases. □

H.12 Admissibility Emergence

Definition H.34 (Admissibility Functional)**.**

$$\Lambda(x) f\big(\delta_D(x), \Gamma(x)\big).$$

Theorem H.35. *Admissibility increases with continuation and decreases with deficit.*

Proof. By monotonicity of f . □

H.13 Hierarchy Theorem

Theorem H.36 (Distinction Hierarchy)**.** *For every admissible system,*

$$\textit{Distinction} \to \textit{Information} \to \textit{Entropy} \to \textit{Repair} \to \textit{Continuation} \to \textit{Admissibility}.$$

Proof. Distinctions generate informational differences.

Information determines distinguishability entropy.

Entropy generates deficit.

Repair reduces deficit.

Reduced deficit increases continuation.

Continuation determines admissibility. □

H.14 Unification Theorem

Theorem H.37 (Unification Theorem)**.** *Negation distortion, admissibility geometry, repair theory, continuation theory, and distinguishability geometry are all projections of a common underlying displacement structure.*

Specifically,

$$D_N \subseteq \eta_N \subseteq \delta_D \subseteq \Theta \subseteq \Lambda.$$

Thus inferential negation is a special case of distinguishability displacement, which is itself a special case of admissibility deformation.

Proof. Each quantity is obtained by successive abstraction from reachability displacement to distinguishability deficit to admissibility structure.

Therefore the inclusions follow by construction. □

H.15 Fundamental Geometric Principle

Theorem H.38 (Fundamental Geometric Principle)**.** *Objects, propositions, memories, repairs, and admissible futures are not primitive entities.*

They are stable regions of distinguishability geometry whose persistence depends on the maintenance of recoverable distinctions.

Consequently admissibility is the terminal invariant of distinction-preserving dynamics.

Proof. Distinguishability generates information.

Information generates entropy.

Entropy necessitates repair.

Repair enables continuation.

Continuation defines admissibility.

Hence admissibility emerges as the terminal invariant of the hierarchy. □

I Geometric Foundations of Admissibility

I.1 Fundamental Configuration Space

Let

$$M$$

be a smooth differentiable manifold.

Define the fundamental field tuple

$$\Psi(D, \omega, \mathcal{M}, \mathcal{R}, \Lambda)$$

where

$$D : M \to \mathbb{R}$$

is the distinguishability field,

$$\omega : M \to \mathbb{R}$$

is the orientation field,

$$\mathcal{M} : M \to \mathbb{R}$$

is the memory field,

$$\mathcal{R} : M \to \mathbb{R}$$

is the repair field,

and

$$\Lambda : M \to \mathbb{R}$$

is the admissibility field.

I.2 Fundamental Principle

[Recoverable Distinction Principle] A configuration is physically, cognitively, and inferentially meaningful only to the extent that distinctions remain recoverable under admissible transformations.

[Continuation Principle] Admissibility is determined not by present state alone but by the existence of recoverable future trajectories.

I.3 Fundamental Action

Define

$$S[\Psi] \int_M \mathcal{L}, dV.$$

The Lagrangian density is

$$\mathcal{L}\mathcal{L}_D + \mathcal{L}_\omega + \mathcal{L}_M + \mathcal{L}_R + \mathcal{L}_\Lambda + \mathcal{L}_{\text{int}}.$$

I.4 Distinguishability Sector

$$\mathcal{L}_D \frac{1}{2}|\nabla D|^2$$

V$_D(D)$.
Define

$$V_D(D)\frac{\alpha}{2}D^2 + \frac{\beta}{4}D^4.$$

Theorem I.1. *Critical points satisfy*

$$\Delta D \alpha D + \beta D^3.$$

Proof. Euler-Lagrange variation. □

I.5 Orientation Sector

$$\mathcal{L}_\omega \frac{1}{2}|\nabla \omega|^2$$

V$_\omega(\omega)$.
Choose

$$V_\omega \frac{\gamma}{2}\omega^2.$$

Then

$$\Delta \omega \gamma \omega.$$

I.6 Memory Sector

$$\mathcal{L}_M \frac{1}{2} |\nabla \mathcal{M}|^2$$

$\frac{\mu}{2}\mathcal{M}^2$.
Hence

$$\Delta \mathcal{M} \mu \mathcal{M}.$$

Definition I.2 (Persistence Density)**.**

$$P \mathcal{M}^2.$$

I.7 Repair Sector

$$\mathcal{L}_R \frac{1}{2} |\nabla \mathcal{R}|^2$$

$\mathrm{V}_R(\mathcal{R})$.
Let

$$V_R \frac{\rho}{2} \mathcal{R}^2.$$

Then

$$\Delta \mathcal{R} \rho \mathcal{R}.$$

Definition I.3 (Repair Capacity)**.**

$$\Phi_R \int_M \mathcal{R}, dV.$$

I.8 Admissibility Sector

$$\mathcal{L}_\Lambda \frac{1}{2} |\nabla \Lambda|^2$$

$\mathrm{V}_\Lambda(\Lambda)$.
Choose

$$V_\Lambda \frac{\lambda}{2} \Lambda^2.$$

Then

$$\Delta \Lambda \lambda \Lambda.$$

I.9 Interaction Terms

Define

$$\mathcal{L}_{\text{int}} aD\omega + bD\mathcal{M} + cD\mathcal{R} + d\mathcal{M}\mathcal{R} + e\mathcal{R}\Lambda.$$

I.10 Coupled Field Equations

Variation yields

$$\Delta D \alpha D + \beta D^3 + a\omega + b\mathcal{M} + c\mathcal{R},$$

$$\Delta\omega\gamma\omega + aD,$$

$$\Delta\mathcal{M}\mu\mathcal{M} + bD + d\mathcal{R},$$

$$\Delta\mathcal{R}\rho\mathcal{R} + cD + d\mathcal{M} + e\Lambda,$$

$$\Delta\Lambda\lambda\Lambda + e\mathcal{R}.$$

I.11 Fundamental Dependency Chain

The coupling structure induces

$$D \to \omega \to \mathcal{M} \to \mathcal{R} \to \Lambda.$$

Theorem I.4. *Admissibility depends indirectly on distinguishability through the entire chain.*

Proof. The coupled equations show

$$\Lambda$$

depends on

$$\mathcal{R},$$

which depends on

$$\mathcal{M},$$

which depends on

$$D.$$

□

I.12 Distinguishability Current

Define

$$J_D \nabla D.$$

Definition I.5 (Distinction Flux)**.**

$$\Phi_D \int_{\partial M} J_D \cdot n, dS.$$

Theorem I.6.

$$\Phi_D \int_M \Delta D, dV.$$

Proof. Divergence theorem. □

I.13 Orientation Current

$$J_\omega \nabla \omega.$$

Definition I.7 (Orientation Flux)**.**

$$\Phi_\omega \int_{\partial M} J_\omega \cdot n, dS.$$

I.14 Memory Current

$$J_M \nabla \mathcal{M}.$$

Definition I.8 (Persistence Flux)**.**

$$\Phi_M \int_{\partial M} J_M \cdot n, dS.$$

I.15 Repair Current

$$J_R \nabla \mathcal{R}.$$

Definition I.9 (Repair Flux)**.**

$$\Phi_R \int_{\partial M} J_R \cdot n, dS.$$

I.16 Admissibility Current

$$J_\Lambda \nabla \Lambda.$$

Definition I.10 (Admissibility Flux)**.**

$$\Phi_\Lambda \int_{\partial M} J_\Lambda \cdot n, dS.$$

I.17 Conservation Law

Define total geometric charge

$$Q \int_M (D + \omega + \mathcal{M} + \mathcal{R} + \Lambda), dV.$$

Theorem I.11. *If interaction terms are symmetric,*

$$\frac{dQ}{dt} = 0.$$

Proof. Integrating the coupled field equations causes interaction contributions to cancel pairwise. □

I.18 Recoverability Functional

Define

$$\mathfrak{R} \int_M D\mathcal{M}, dV.$$

Definition I.12 (Recoverability)**.** *A system is recoverable iff*

$$\mathfrak{R} > 0.$$

I.19 Continuation Functional

Define

$$\mathfrak{C} \int_M \mathcal{M}\mathcal{R}, dV.$$

Definition I.13 (Continuation Capacity)**.**

$$\Phi_C \mathfrak{C}.$$

Theorem I.14. *Repair without memory yields*

$$\Phi_C = 0.$$

Proof. If

$$\mathcal{M} = 0,$$

then

$$\mathfrak{C} = 0.$$

□

I.20 Admissibility Functional

Define

$$\mathfrak{A} \int_M \Lambda \mathcal{R} \mathcal{M}, dV.$$

Definition I.15 (Global Admissibility)**.**

$$A_{\text{global}} \mathfrak{A}.$$

I.21 Entropy as Deficit

Define

$$E D_{\text{opt}}$$

D.

Definition I.16 (Distinguishability Entropy)**.**

$$S_D \int_M E, dV.$$

Theorem I.17. *Repair reduces entropy deficit.*

Proof. Repair increases realized distinguishability.
Hence

$$E$$

decreases. □

I.22 Master Stability Functional

Define

$$\Sigma \int_M \left(D^2 + \omega^2 + \mathcal{M}^2 + \mathcal{R}^2 + \Lambda^2 \right) dV.$$

Theorem I.18. *Stable systems are local minima of*

$$\Sigma.$$

Proof. Standard Lyapunov argument. □

I.23 Grand Unification Theorem

Theorem I.19 (Grand Unification)**.** *Distinction, orientation, memory, repair, continuation, and admissibility are not independent primitives.*

They arise as coupled fields governed by a single variational principle

$$\delta S = 0.$$

All previously defined structures—

$$D_N, \quad \delta_D, \quad \Theta, \quad \Gamma, \quad \Lambda,$$

appear as derived quantities of the unified field configuration

$$\Psi(D, \omega, \mathcal{M}, \mathcal{R}, \Lambda).$$

Proof. Each quantity is either a field, a functional of fields, or a derived invariant of the action.

Therefore all emerge from the stationary points of

$$S.$$

□

I.24 The Fate of Distinguishability

Theorem I.20 (Terminal Principle)**.** *A system persists precisely to the extent that distinguishability remains recoverable.*

Memory stores recoverable distinctions.

Repair restores lost distinctions.

Continuation propagates distinctions.

Admissibility selects distinctions capable of persistence.

Consequently,

AdmissibilityRecoverable Future Distinguishability.

Proof. Immediate from the dependency chain

$$D \to \omega \to \mathcal{M} \to \mathcal{R} \to \Lambda.$$

The terminal invariant of the chain is admissibility, and every preceding stage exists only insofar as distinctions remain recoverable. □

References

[1] Agmon, G., Loewenstein, Y., Grodzinsky, Y. (2019). Measuring the cognitive cost of downward monotonicity by controlling for negative polarity. *Glossa: A Journal of General Linguistics*, 4(1), 36.

[2] Agmon, G., Loewenstein, Y., Grodzinsky, Y. (2022). Negative sentences exhibit a sustained effect in delayed verification tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 48(1), 122–141.

[3] Amunts, K., Mohlberg, H., Bludau, S., Zilles, K. (2020). Julich-Brain: A 3D probabilistic atlas of the human brain's cytoarchitecture. *Science*, 369(6506), 988–992.

[4] Baker, C. L. (1970). Double negatives. *Linguistic Inquiry*, 1, 169–186.

[5] Barwise, J., Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy*, 4, 159–219.

[6] Broca, P. (1861). Remarques sur le siège de la faculté du langage articulé. *Bulletins de la Société Anatomique de Paris*, 6, 330–357.

[7] Caramazza, A., Zurif, E. (1976). Dissociation of algorithmic and heuristic processes in language comprehension. *Brain and Language*, 3, 572–582.

[8] Carpenter, P. A., Just, M. A. (1975). Sentence comprehension: A psycholinguistic processing model of verification. *Psychological Review*, 82, 45–73.

[9] Chomsky, N. (1995). *The Minimalist Program*. Cambridge, MA: MIT Press.

[10] Clark, H. H., Chase, W. G. (1972). On the process of comparing sentences against pictures. *Cognitive Psychology*, 3, 472–517.

[11] Crnič, L. (2014). Non-monotonicity in NPI licensing. *Natural Language Semantics*, 22, 169–217.

[12] Dehaene, S., Piazza, M., Pinel, P., Cohen, L. (2003). Three parietal circuits for number processing. *Cognitive Neuropsychology*, 20, 487–506.

[13] Deschamps, I., Agmon, G., Loewenstein, Y., Grodzinsky, Y. (2015). The processing cost of downward entailingness: The representation and verification of comparative constructions. In *Proceedings of Sinn und Bedeutung 19*.

[14] Fauconnier, G. (1975). Polarity and the scale principle. *Chicago Linguistic Society*, 11, 188–199.

[15] Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.

[16] Frege, G. (1879). *Begriffsschrift*. Halle: Nebert.

[17] Gajewski, J. (2011). Licensing strong NPIs. *Natural Language Semantics*, 19, 109–148.

[18] Geurts, B., Katsos, N., Cummins, C., Moons, J., Noordman, L. (2010). Scalar quantifiers: Logic, acquisition, and processing. *Language and Cognitive Processes*, 25, 130–148.

[19] Goodglass, H. (1968). Studies on the grammar of aphasics. *Cortex*, 4, 193–216.

[20] Grodzinsky, Y. (1986). Language deficits and the theory of syntax. *Brain and Language*, 27, 135–159.

[21] Grodzinsky, Y., Deschamps, I., Pieperhoff, P., Iannilli, F., Agmon, G., Loewenstein, Y., Amunts, K. (2020). Logical negation mapped onto the brain. *Brain Structure and Function*, 225, 19–31.

[22] Grodzinsky, Y., Pieperhoff, P., Thompson, C. K. (2021). Stable brain loci for the processing of complex syntax. *Cortex*, 142, 252–271.

[23] Grodzinsky, Y., Behrent, K., Agmon, G., Bittner, N., Jockwitz, C., Caspers, S., Amunts, K., Heim, S. (2023). A linguistic complexity pattern that defies aging: The processing of multiple negations.

[24] Horn, L. R. (1989). *A Natural History of Negation*. Chicago: University of Chicago Press.

[25] Homer, V. (2021). Negative polarity items. In D. Gutzmann et al. (Eds.), *The Wiley Blackwell Companion to Semantics*.

[26] Kaup, B. (2001). Negation and its impact on the accessibility of text information. *Memory Cognition*, 29, 960–967.

[27] Kaup, B., Zwaan, R. (2003). Effects of negation and situational presence on language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 439–446.

[28] Klima, E. S. (1964). Negation in English. In J. A. Fodor J. Katz (Eds.), *The Structure of Language*. Englewood Cliffs: Prentice Hall.

[29] Ladusaw, W. (1980). *Polarity Sensitivity as Inherent Scope Relations*. Doctoral dissertation, University of Texas at Austin.

[30] Lakoff, G. (1970). Linguistics and natural logic. *Synthese*, 22, 151–271.

[31] Montague, R. (1974). *Formal Philosophy*. New Haven: Yale University Press.

[32] Russell, B. (1903). *The Principles of Mathematics*. Cambridge: Cambridge University Press.

[33] Szymanik, J. (2016). *Quantifiers and Cognition*. Cham: Springer.

[34] Tan, I.-A., Kugler-Etinger, N., Grodzinsky, Y. (2023). Do two negatives make a positive? Language and logic in language processing. *Language, Cognition and Neuroscience*, 38(7), 1027–1043.

[35] Tan, I.-A., Segal, N., Grodzinsky, Y. (2024). The domains of monotonicity processing. *Journal of Semantics*, 41, 77–101.

[36] Wason, P. C. (1961). Response to affirmative and negative binary statements. *British Journal of Psychology*, 52, 133–142.

[37] Wason, P. C. (1965). The contexts of plausible denial. *Journal of Verbal Learning and Verbal Behavior*, 4, 7–11.

[38] Xiang, M., Grove, J., Giannakidou, A. (2016). Dependency-dependent interference: NPI processing and intervention. *Glossa*, 1(1), 38.

[39] Zurif, E. B. (1980). Language mechanisms: A neuropsychological perspective. *American Scientist*, 68, 305–311.