Language as Constraint Field

Grammar, Repair, and Admissibility

in Semantic Manifolds

Flyxion
Independent Researcher

2026

Abstract

Grammar does not generate utterances. Grammar maintains the admissibility structure that makes utterances receivable. This essay develops a formal theory of language as a constraint field: a dynamical system in which utterances are operators on a discourse state space, grammaticality is a property of trajectories rather than objects, and linguistic competence is the ability to navigate admissibility manifolds rather than to apply production rules. The theory inverts the generative direction. Where Chomskyan grammar asks what rules produce well-formed sentences, the constraint-field approach asks what constraints make utterances receivable — a shift from production to admissibility as the organizing primitive. The framework is developed through three case studies. Arabic root-and-pattern morphology, reinterpreted as a repair ecology in which the triliteral root functions as a semantic attractor and the morphological patterns are repair operators maintaining coherence under contextual deformation. Spanish verbal agreement, reinterpreted as a multi-field coherence condition that native speakers satisfy by inhabiting the field rather than consulting rules. And conversation itself, reinterpreted as alternating CLIO projection and semantic repair, in which each turn maintains the discourse manifold by committing a trajectory that remains within the admissibility region. The mathematical framework connects directly to the Relativistic Scalar-Vector Plenum field theory (providing the field-dynamic substrate), CLIO projection dynamics (providing the discourse-reduction operator), and the Semantic Stability Theorem (providing the branch-coherence condition for grammatical constructions). Language death is repair collapse. Language change is manifold deformation. The gesture-before-symbol thesis establishes that meaning exists prior to its linguistic encoding, and language is the repair layer that preserves access to shared meaning against the attenuation introduced by distance, time, and individual variation.

A grammatical utterance is not one that was successfully generated. It is one that remains inside the discourse manifold.

The Wrong Question

The dominant tradition in formal linguistics — descending through Chomsky's generative grammar, the minimalist program, and their derivatives — takes as its central question: what rules generate well-formed sentences? The grammar is a production system. Grammaticality is a property of outputs. A sentence is grammatical if and only if it is generable by the grammar.

This question has been enormously productive. It forced the formalization of syntactic structure, revealed the existence of deep universals in language acquisition, and connected linguistics to computational theory. But it has a deep orientation problem: it begins with production rather than reception. It asks what mechanisms a speaker uses to generate sentences, not what conditions a discourse context imposes on receivable utterances.

The repair-equilibrium framework developed in companion papers suggests that this orientation is the source of several persistent difficulties in generative theory. If language is primarily a repair infrastructure — a system for maintaining shared semantic access against the attenuation introduced by distance, time, and individual variation — then the generative question is not the right first question. The right first question is: what constraints make an utterance receivable?

This shift from production to admissibility is not merely terminological. It changes what the primitive objects of the theory are, what counts as an explanation, and what the empirical targets of the theory should be. A generative grammar is a mechanism for producing sentences; its primitive objects are symbols, rules, and derivations. A constraint-field grammar is a dynamical system for maintaining discourse coherence; its primitive objects are fields, operators, and admissibility regions. The difference is the difference between asking how a sentence is built and asking why a sentence is receivable at this moment in this conversation by this community.

The Discourse Manifold

Definition 1 (Discourse state space)**.** Let X be a *discourse state space*: the space of all possible discourse configurations, where a configuration specifies the propositional content established so far, the discourse referents in play, the illocutionary commitments of each participant, the contextually salient alternatives, and the pragmatic norms active in the

exchange. A discourse state $x \in X$ is a complete specification of the conversational situation at a moment.

Definition 2 (Admissibility manifold)**.** The *discourse admissibility manifold* $\mathcal{A}(x) \subseteq X$ at discourse state x is the set of states reachable from x by utterances that maintain coherence: that are semantically interpretable in the context, that preserve the referential and illocutionary structure established by prior discourse, and that satisfy the pragmatic norms active in the exchange.

Definition 3 (Utterance as operator)**.** An utterance u acts as an operator

$$u : X \to X, \quad x \mapsto u(x),$$

transforming the discourse from state x to a new state $u(x)$. The utterance is *grammatically admissible* in state x if and only if $u(x) \in \mathcal{A}(x)$: the resulting discourse state is within the admissibility manifold of the current state.

Grammaticality is therefore a property of trajectories through discourse state space, not a property of sentences considered in isolation. The sentence "the dog bit the man" is admissible in some discourse states and inadmissible in others — not because its syntactic structure varies but because the admissibility manifold varies with context. "Colorless green ideas sleep furiously" is syntactically well-formed but semantically inadmissible in most discourse states because no standard interpretation maps it to a trajectory within the manifold of coherent discourse.

Theorem 1 (Dissolution of the Grammaticality/Acceptability Gap)**.** In the constraint-field framework, the generative distinction between grammaticality and acceptability is not a distinction between two properties of sentences but a distinction between two discourse contexts in which the same sentence is evaluated. More precisely: a sentence s that is "grammatical but unacceptable" in generative terms is, in constraint-field terms, a sentence that is admissible under the abstract grammatical constraint field $\mathcal{A}_G$ but inadmissible under the concrete discourse admissibility manifold $\mathcal{A}(x)$ at the current state x. A sentence that is "acceptable but ungrammatical" is admissible under $\mathcal{A}(x)$ but exits the boundaries of $\mathcal{A}_G$. The apparent gap arises because the generative framework holds the grammatical constraint field fixed and treats discourse context as noise; the constraint-field framework treats $\mathcal{A}(x)$ as the fundamental object and $\mathcal{A}_G$ as the degenerate case in which contextual constraints are minimized.

Worked examples. Example 1: Grammatical but unacceptable. "The horse raced past the barn fell" is syntactically well-formed (a reduced relative clause construction, fully derivable in any generative grammar). But in most discourse states x, $u(x) \notin \mathcal{A}(x)$: the interpretation that the utterance targets is not accessible without backtracking, be-

cause the initial parse trajectory "the horse raced" commits to a main- clause reading before the reduced relative can be recovered. The utterance is inadmissible not because it violates a rule but because the discourse state at the moment of parse failure has been driven outside $\mathcal{A}(x)$ by the garden-path effect. Grammatical competence (knowledge of $\mathcal{A}_G$) does not prevent the boundary crossing; the trajectory through discourse state space crosses the boundary before the full structure is assembled.

Example 2: Acceptable but ungrammatical. "Him and me went to the store" is widely stigmatized as ungrammatical (accusative pronouns in subject position violate Case assignment in standard generative accounts). But in informal registers, the utterance $u(x)$ lands inside $\mathcal{A}(x)$: it is fully interpretable, referentially clear, and pragmatically appropriate. What has happened is that the admissibility manifold $\mathcal{A}(x)$ for informal spoken English has expanded to include accusative subject pronouns, while the prescriptive constraint field $\mathcal{A}_G$ (constructed from formal written norms) has not. The "ungrammaticality" is a property of the mismatch between two admissibility manifolds, not a property of the sentence.

Example 3: Performance errors. Slips of the tongue, agreement attraction, and garden-path errors are all cases where the utterance operator u follows the locally strongest constraint gradient rather than the globally correct one — where $u(x)$ exits $\mathcal{A}(x)$ because the gradient of the constraint field was locally distorted by discourse salience. These are not distinct from grammaticality failures; they are the same phenomenon (trajectory exiting the admissibility manifold) occurring at different timescales and for different reasons. The generative competence/performance distinction is, in the constraint-field framework, the distinction between $\mathcal{A}_G$ (the idealized, context- free admissibility region) and $\mathcal{A}(x)$ (the actual, context-sensitive admissibility manifold at each discourse state). Both are admissibility manifolds; they differ in what constraints they encode. □

Syntax as Constraint Propagation

Syntactic structure, in the constraint-field framework, is the record of how admissibility constraints propagate through an utterance. A subject-verb agreement constraint is not a rule that is checked at a derivational step; it is a constraint that propagates from the subject constituent to the verb constituent through the admissibility field, requiring that the verb occupy a region of the field consistent with the person and number established by the subject.

Violation of agreement is not a failure to apply a rule. It is a trajectory that exits the admissibility manifold: the verb's position in the discourse state space is inconsistent with the constraints established by the subject. The discourse state $u(x)$ is outside $\mathcal{A}(x)$.

This reframing immediately explains several features of syntactic processing that are

puzzling on the generative account. Agreement attraction errors — in which speakers produce agreement with a local noun rather than the grammatical head — occur precisely when the local noun is more salient in the current discourse state, temporarily distorting the admissibility field. The error is not a failure of rule application but a trajectory that followed the locally strongest constraint gradient rather than the grammatically correct one. The gradient was distorted by the discourse context.

CLIO as Discourse Reduction

The CLIO operator — constraint-based limitation of interpretive output — provides the formal model of how utterances reduce the discourse state space.

At every conversational turn, the speaker faces a space X of possible meanings and must commit to an utterance that maps the current discourse state to a state within the admissibility manifold. This is a projection operation:

$$\pi_u : X \to \mathcal{A}(x), \quad \pi_u(X) = u(x).$$

The utterance selects, from the enormous space of possible meanings, a trajectory that lands in the admissible region. The listener performs the inverse:

$$\pi_u^{-1} : \mathcal{A}(x) \to X',$$

reconstructing from the utterance a trajectory in semantic space — not the unique trajectory the speaker intended, but one constrained to be consistent with the utterance and the discourse context.

Conversation is therefore alternating projection and reconstruction under shared constraint. The speaker projects from semantic space into the admissibility manifold. The listener reconstructs from the manifold back into semantic space. The gap between the projected trajectory and the reconstructed one is the fundamental source of communicative failure — not failure of information transmission in the Shannon sense, but failure of constraint alignment: the listener's reconstructed trajectory exits the admissibility manifold that the speaker was targeting.

Definition 4 (Conversational repair)**.** A *conversational repair* is an utterance r whose function is to return the discourse to the admissibility manifold after a failed transmission. Repair has the structure of the repair process from the repair ontology: $r : X \setminus \mathcal{A} \to \mathcal{A}$, mapping an inadmissible discourse state back into the admissibility region. Repair may target the speaker's prior utterance (self-initiated self-repair), the listener's interpretation (other-initiated other-repair), or the conversational trajectory (metalinguistic repair).

The repair-rate matching condition from the repair ontology applies directly: a conversation is stable if and only if the rate of repair r_R meets or exceeds the rate of semantic drift r_D:

$$r_R(x,t) \geq r_D(x,t) \quad \text{for all } x \in \mathcal{A},\ t \geq 0.$$

A conversation that exceeds its repair capacity — in which miscommunications accumulate faster than repair exchanges can correct them — will exit the admissibility manifold. The repair rate is a function of the participants' shared context, their metalinguistic resources, and the structural accessibility of repair in the discourse genre.

Theorem 2 (Admissibility Preservation)**.** Let $x_0 \in \mathcal{A}$ be an initial discourse state and let $u_1, u_2, \ldots, u_n$ be a sequence of utterance operators. The discourse trajectory

$$x_k = u_k(x_{k-1}), \quad k = 1, \ldots, n,$$

remains coherent if and only if $x_k \in \mathcal{A}(x_{k-1})$ for every k, equivalently $\rho(x_k) \geq 1$ for every state in the trajectory.

If there exists a smallest index m such that $\rho(x_m) < 1$, then $x_m \notin \mathcal{A}(x_{m-1})$: the discourse has exited the admissibility manifold. All subsequent discourse states $x_{m+1}, \ldots, x_n$ either depend on repair operators $r_k : X \setminus \mathcal{A} \to \mathcal{A}$ restoring the trajectory to $\mathcal{A}$, or they remain outside $\mathcal{A}$ and the conversation is incoherent from step m onward.

Proof. Discourse coherence at step k is defined as $x_k \in \mathcal{A}(x_{k-1})$. The trajectory remains coherent at every step iff this condition holds for all $k \in \{1, \ldots, n\}$, which by the definition of ρ is equivalent to $\rho(x_k) = S_{\text{support}}(x_k)/\theta(x_k) \geq 1$ for all k. If $\rho(x_m) < 1$ for some m, then $x_m \notin \mathcal{A}(x_{m-1})$ by definition. The utterance u_m has mapped the discourse to an inadmissible state. For $k > m$, the inherited starting state x_m is outside the admissibility manifold; subsequent operators either perform repair ($r : X \setminus \mathcal{A} \to \mathcal{A}$, restoring $\rho \geq 1$) or propagate the inadmissible state, producing further boundary violations. □

Remark 1. The Admissibility Preservation Theorem is the discourse-level analogue of the Semantic Stability Theorem's node-level condition $\rho(v) \geq 1$ and the branch stability index $\rho_{\min}(\mathcal{B}) \geq 1$. The difference is temporal: the Semantic Stability Theorem operates over a static tree structure; the Admissibility Preservation Theorem operates over a dynamic trajectory through discourse state space. The same repair ratio ρ governs both. The theorem unifies discourse coherence, conversational repair, CLIO projection, language change (as the evolution of $\mathcal{A}$ across time), and language death (as the condition $\rho_{\min} < 1$ propagating through the construction dependency network) under a single formal result.

Arabic: Root as Attractor, Pattern as Repair Operator

Arabic root-and-pattern morphology is the most formally transparent instantiation of the constraint-field approach to grammar, and its architecture repays careful analysis. The companion paper *Embodied Constraint and Minimal Assembly* develops a cost-distinctness model of Arabic script stabilization in which the primitive quantities are production cost C, perceptual distinctness D, and system stability $S = f(-C, D)$. That paper treats the template as a morphological operator $T : R \to W$ mapping a root to a word form. The present section performs the reinterpretation:

$$T : \mathcal{A}_t \to \mathcal{A}_{t+1},$$

where the template is a repair operator acting on a discourse admissibility manifold rather than a combinatorial generator acting on symbolic structures. The root is not a lexical seed; it is a semantic attractor. The word is not generated; it is stabilized. This shift from production to maintenance as the organizing primitive changes what requires explanation and what counts as explanation.

The Standard Account and Its Limitations

The standard morphological account treats the triliteral root as a semantic substrate and the pattern (or binyan) as a template that generates word forms by inserting the root consonants into fixed vowel-consonant frames. The root **k-t-b** (associated with writing) generates: *kataba* (he wrote), *kitāb* (book), *kātib* (writer, active participle), *maktūb* (written, passive participle), *maktab* (office, place of writing), *maktaba* (library), *kitāba* (writing, the act), and so on through Forms I–X of the binyan system.

On the generative account, the root is a lexical seed and the patterns are derivational rules. The morphological system is a productive generator of word forms. A word is well-formed if it is derivable from an attested root by an admissible pattern. *Embodied Constraint* formalizes this as $T = (C_1, C_2, C_3) \mapsto W$, where the template introduces fixed vocalic insertions, optional consonantal augmentation, and prefixation into root consonant slots. The generative capacity is constrained: $|\mathcal{W}| \leq |\mathcal{T}|$, with each template encoding a predictable semantic modulation.

This account is accurate as a description of the combinatorial productivity of the system. But it misses the dynamical question: why does the system have this architecture rather than another? Why does Arabic maintain a stable consonantal root that survives across phonological and morphological transformations? What work is the root doing that justifies the cost of preserving it? The cost-distinctness model answers this in static terms — the root is the motor-efficient substrate; dotting provides low-cost informational refinement — but does not explain why the system maintains the root's semantic

coherence across derivations over time.

The Root as Semantic Attractor

The constraint-field answer is that the root is a semantic attractor: a high- density node in the semantic field that exerts repair pressure on its derived forms, maintaining semantic coherence across contextual deformation.

Definition 5 (Root attractor and root support)**.** A triliteral root R defines a *semantic attractor field* $\Phi_R : \mathcal{L} \to \mathbb{R}_{\geq 0}$ over the lexical space $\mathcal{L}$, where $\Phi_R(w)$ measures the degree to which word w remains semantically anchored to R's core meaning field. The *root support* of a derived form w is:

$$S_R(w) = \sum_{p \in \text{Path}(R,w)} \alpha_p \, w_p,$$

where the sum runs over ancestral derivational paths from R to w, α_p is the attenuation along path p, and w_p is the admissibility weight of path p. A derived form w is *semantically alive* with respect to its root when $S_R(w) \geq \theta_w$, where θ_w is the minimum root support required for w to participate coherently in the root's semantic field.

The patterns (morphological templates) are, on this account, repair operators: structured deformations of the root that adapt it to specific grammatical contexts while preserving semantic continuity with the root field. Form I (*façala*: basic verb) is the identity repair — minimal deformation, maximum root visibility. Form II (*façççala*: intensive/causative) intensifies the root meaning with gemination. The gemination operator $G(C_i) = C_iC_i$ defined in *Embodied Constraint* is, in the constraint-field reinterpretation, not just a syllabic transformation but a semantic intensification repair: by doubling C_2, the pattern increases Φ_R at the derived form, strengthening its connection to the root field rather than merely modifying syllabic architecture. Form X (*istafçala*: estimative/requestive) introduces the most complex deformation while maintaining the root's semantic core through the invariant consonantal skeleton $\pi(R) = (C_1, C_2, C_3)$.

Proposition 1 (Pattern as constraint-preserving deformation)**.** Each morphological pattern P is a repair operator:

$$P : \mathcal{A}(R) \to \mathcal{A}(R, G),$$

mapping the root's semantic admissibility region to the intersection of the root's semantic field with the grammatical context G (verb, noun, active participle, place noun, etc.). A pattern is productive when the intersection $\mathcal{A}(R) \cap \mathcal{A}(G)$ is non-empty: there exist forms that simultaneously satisfy the root's semantic constraints and the grammatical context's requirements. The root-order invariance $\pi(R) = (C_1, C_2, C_3)$ proved in *Embodied Constraint* is, in the constraint-field framework, the formal condition for semantic

continuity: the root attractor's identity is preserved across morphological deformations precisely because the consonantal skeleton is invariant.

The system's resistance to lexical attrition is a direct consequence of the root-as-attractor structure. The root provides continuous repair pressure on its derived forms, maintaining semantic coherence against the attenuation that would otherwise disconnect a word from its semantic ancestry.

In languages without a visible root system, this attenuation runs to completion unchecked. The general mechanism is: a derived word w undergoes repeated use in semantic contexts that are progressively distant from its etymological source. Each such use increases the attenuation α along the derivational path from the source to w, reducing $S_R(w)$ incrementally. When the source is invisible in the word's form — because the morphological connection has been obscured by phonological change, borrowing, or orthographic opacity — no repair pressure maintains the semantic connection. The word's semantic admissibility region drifts from the source field entirely. The word persists as a stored structure carrying no active connection to its etymological origin. In Arabic, the consonantal skeleton in the orthography continually performs a repair function: every occurrence of a word containing **k-t-b** activates the root field, maintaining $S_R(w) \geq \theta_w$ for the entire family. The writing system is doing semantic repair work that English orthography, with its opacity to morphological structure, cannot do.

The Incremental Augmentation Principle as Repair Efficiency

Embodied Constraint establishes an Incremental Augmentation Principle formalized as:

$$\frac{I(A(X)) - I(X)}{C(A)} \geq \kappa,$$

where $C(A)$ is the production cost of augmentation A, I is informational distinctness, and $\kappa > 0$ is an efficiency threshold. The principle states that augmentations persist historically when they deliver high informational differentiation at low motor cost.

In the constraint-field reinterpretation, this principle is a special case of the repair efficiency condition: an augmentation persists because it successfully reduces S_{RSVP} at the affected nodes — it narrows the accessible future volume of interpretation — without requiring high repair cost. The dot on **b** (*ba*) versus undotted **t/th** (*ta*/*tha*) achieves $C(A) \approx \epsilon$ while achieving $D(g_1, g_2) \geq D_{\text{min}}$: a minimal repair operation that prevents the interpretation from exiting the admissibility manifold into an ambiguous region. The sukūn — the marker of vowel absence — is in this framework an explicit null-state repair: it prevents the default assumption of vowel continuation from generating inadmissible interpretations. The shaddah (gemination marker) is a temporal extension repair: a single graphic mark that modifies syllabic structure across two adjacent sylla-

bles by encoding consonantal doubling, maintaining admissibility at the morphological boundary without introducing new graphemic primitives.

The repair interpretation explains why the augmentation set $\mathcal{A}$ stabilizes at precisely the inventory it does. The stability criterion $\partial S/\partial g_i \approx 0$ in *Embodied Constraint* is, in constraint-field terms, the condition that no further repair operations are available that would increase the system's admissibility without exceeding the repair cost threshold. The script has reached a local repair equilibrium: an attractor state of the repair dynamics in which the system is simultaneously motor-efficient, perceptually discriminable, and semantically coherent.

Dotting as Boundary Preservation

The diacritical points (nuqat) that distinguish letter pairs sharing a rasm (consonantal skeleton) — the **d/dh** pair sharing one base, the **r/z** pair sharing another, the **s/sh** pair sharing a three- toothed form — are not merely phonemic distinguishers. In the constraint-field framework, they are boundary-preservation operators: graphic marks that prevent manifold folding in the interpretation space.

Embodied Constraint models this as $g_1 = b$, $g_2 = d(b)$ where d is the dot augmentation satisfying $C(d) \approx \epsilon$ and $D(g_1, g_2) \geq D_{\text{min}}$. The constraint-field reinterpretation is: the dot reduces S_{RSVP} at the graphemic node by collapsing the accessible future volume of interpretation from a branching ambiguous state into a single determinate trajectory. Without the dot, the interpretation space bifurcates at each ambiguous letter; the dot performs a CLIO projection, collapsing the branching into a single admissible path.

The **r**/**z** pair is especially informative. *Embodied Constraint* notes that the skeletal form is identical while the dot encodes a manner contrast involving sustained turbulent airflow (the voiced alveolar fricative versus the dental stop). In constraint-field terms, the dot marks the boundary between two distinct regions of the admissibility manifold: the stop-manner region and the fricative-manner region. Without the dot, an utterance containing this letter occupies a point equidistant from both boundary regions — a maximally unstable repair position, requiring the listener to perform external disambiguation from context. The dot is the writer's repair operation, performed in advance of transmission, that moves the interpretation away from the boundary and into the interior of the correct manifold region.

The non-connecting letters (*alif, dal, dhal, ra, zayn, waw*) function as semantic fault barriers: strategic discontinuities in the script that arrest ambiguity propagation. *Embodied Constraint* models these as constraint insertions that preserve minimum perceptual distance $D_{\text{join}} \geq D_{\text{min}}$ by imposing $J(g_i, g_{i+1}) = 0$ for specific g_i. The constraint-field reinterpretation: these letters are points in the word where the admissibility manifold

would fold if joining continued uninterrupted — where adjacent rasm forms, joined, would collapse into a single interpretive region that the listener cannot disambiguate. The non-connecting letter is not a motor convenience; it is a repair intervention built into the orthographic infrastructure, preventing the interpretation from crossing an admissibility boundary by forced segmentation.

The positional allograph system — the reduction from context-sensitive ligatures ($\mathcal{O}(n^2)$ possible forms) to positional forms ($\mathcal{O}(4n)$) — is in *Embodied Constraint* a dimensional reduction that decreases glyph-form entropy H while preserving cursive continuity. In the constraint-field framework, this is the system achieving a stable repair equilibrium: the high-entropy ligature system required too much disambiguation repair from readers and too much production precision from writers; the positional system achieves the same admissibility guarantees at lower repair cost. The compression from H_{context} to $H_{\text{positional}}$ is not a loss but an efficiency gain in the repair dynamics.

Lexicalization as Repair Collapse

The constraint-field framework gives a precise account of lexicalization — the process by which a derived form becomes semantically independent of its root — as a special case of repair collapse.

When a derived form w undergoes repeated use in contexts that are semantically distant from the root field Φ_R, the attenuation α along the path from R to w increases. The root support $S_R(w)$ declines. When $S_R(w)$ crosses the threshold θ_w, the form exits the root's semantic admissibility region:

$$S_R(w) < \theta_w \implies w \notin \mathcal{A}(R).$$

At this point w is semantically dead with respect to the root: it continues to exist as a lexical item but no longer participates in the root's repair network. It has become an isolated stored structure, carrying the consonantal form of the root without the semantic support that the root field provides.

This is the formal mechanism of lexicalization. The Arabic word *wazīr* (minister, vizier) derives from the root **w-z-r** (to bear a burden), but in ordinary use it no longer activates the burden- bearing semantic field. The root support has fallen below threshold. The word persists as a stored structure; the root attractor no longer repairs the semantic connection.

The Arabic system's resistance to lexicalization — relative to non-root- based languages — is a direct consequence of the orthographic repair infrastructure. The consonantal skeleton prevents the attenuation from completing: as long as the root is visible in the word's written form, the root attractor maintains some repair pressure, slowing the de-

scent of $S_R(w)$ below θ_w. This is the constraint-field interpretation of *Embodied Constraint*'s central thesis: the Arabic script's layered architecture (skeletal forms, dotting, diacritics, positional regularization) stabilizes as a constrained symbolic equilibrium because each layer is a repair operation that maintains the admissibility of the next layer's output. The script does not merely encode language; it continuously repairs the semantic coherence of the language it encodes.

Spanish: Multi-Field Coherence

Spanish verbal morphology illustrates a different mode of the same constraint- field structure: not the root-attractor system of Arabic but the multi-field coherence condition that arises when multiple independent constraint dimensions must be simultaneously satisfied.

Agreement as Field Intersection

The Spanish verbal agreement system requires the verb to be simultaneously consistent with person, number, tense, aspect, mood, and (in some constructions) gender. Traditional grammar describes this as rule application: the verb selects the form that satisfies all the relevant features of the subject. The constraint-field approach describes it as field intersection: the verb must occupy a region of the discourse state space that lies within the admissibility regions simultaneously imposed by each active constraint dimension.

The form *hablo* (I speak, present indicative) is not the result of applying a set of rules to the verb *hablar* and the features [1st person, singular, present, indicative]. It is the unique attractor in the intersection of the following constraint fields:

$$\begin{aligned}
\mathcal{A}(\text{1st person}) &: \text{forms marked for speaker reference} \\
\mathcal{A}(\text{singular}) &: \text{forms unmarked for plurality} \\
\mathcal{A}(\text{present}) &: \text{forms occupying the present tense basin} \\
\mathcal{A}(\text{indicative}) &: \text{forms within the realis mood region}
\end{aligned}$$

The correct form is the one that lies in $\bigcap_i \mathcal{A}(d_i)$, where d_i ranges over active constraint dimensions. An agreement error is a form that lies outside this intersection: it satisfies some constraint dimensions but not all, and the resulting discourse state exits the admissibility manifold.

Native Competence as Field Inhabitation

The constraint-field account explains an otherwise puzzling asymmetry between native and non-native speaker behavior. Native speakers of Spanish report that correct agreement *feels* right and incorrect agreement *feels* wrong — they do not consult a rule, they sense a violation. Non-native speakers, by contrast, often describe themselves as checking rules.

On the generative account this is a performance/competence distinction: native speakers have internalized the grammar and apply it automatically. The constraint- field account gives a different explanation: native speakers inhabit the discourse admissibility manifold. They perceive agreement errors as boundary crossings — as states outside the admissibility region — because they have calibrated their perceptual sensitivity to the field's geometry through years of exposure. They are not applying rules; they are detecting field violations.

Non-native speakers who report consulting rules are operating outside the field: they know the grammar's constraints as explicit propositions but have not yet developed the perceptual sensitivity to the field's geometry. The developmental trajectory of second-language acquisition is, on this account, the trajectory of field inhabitation: the learner progressively calibrates their sensitivity to the constraint field until explicit rule-consultation becomes unnecessary.

Proposition 2 (Competence as field sensitivity)**.** Linguistic competence in language L is the capacity to detect whether a discourse trajectory lies within $\mathcal{A}_L$, the admissibility manifold of L, without explicit computation of the constraints defining $\mathcal{A}_L$. This capacity is acquired through exposure to the distribution of admissible and inadmissible trajectories in L, which calibrates the learner's sensitivity to the field geometry.

Subjunctive as Admissibility Shift

The Spanish subjunctive mood is particularly illuminating because it cannot be described as a rule without appealing to semantic and pragmatic conditions that exceed the scope of formal syntax. The subjunctive is required in embedded clauses under verbs of desire, doubt, denial, and emotion — but not under verbs of assertion or perception. The syntactic distribution of the subjunctive is not formally characterizable without reference to the epistemic and affective commitments of the main clause.

The constraint-field account handles this naturally. The indicative occupies the realis region of the discourse admissibility manifold: the region in which the embedded proposition is asserted as obtaining in the actual or projected world. The subjunctive occupies the irrealis region: the region in which the embedded proposition is presented as desired, doubted, denied, or emotionally evaluated rather than asserted.

Verbs of desire, doubt, and denial require the subjunctive because they shift the embedded clause into the irrealis region of the admissibility manifold. The main clause verb is a field operator that transforms the constraint structure of the embedded clause: *quiero que vengas* (I want you to come) applies the desire operator, shifting the embedded clause from the realis to the irrealis region. Using the indicative (*quiero que vienes*) is a boundary crossing: it places the embedded proposition in the realis region, which is incompatible with the desire frame established by the main clause.

Conversation as Real-Time Boundary Maintenance

Turn-Taking as Constraint Relay

A conversation is a sequence of CLIO projections and repair operations in which participants alternate in applying utterances as operators on the shared discourse state space. Turn-taking is not merely a coordination mechanism for orderly speaker alternation; it is the structure through which constraint relay is maintained.

Each speaker turn is a commitment: the speaker applies an utterance operator u that commits the discourse to a new state $u(x)$. This state must lie within $\mathcal{A}(x)$ to be receivable. But it also establishes the constraint structure for the next speaker's turn: the next speaker inherits $u(x)$ as their starting state and must produce an utterance that maintains coherence from $u(x)$ rather than from x.

Turn-taking failures — interruptions, topic hijacks, non-sequiturs — are not merely violations of conversational norms. They are constraint structure violations: they apply utterance operators to discourse states other than the inherited state, producing trajectories that are inadmissible given the established discourse context. The conversational repair that follows such failures (the interrupted speaker resuming, the hijacked topic being reinstated, the non-sequitur being flagged) is the conversation's repair process at work: returning the discourse trajectory to the admissibility manifold.

Genre as Admissibility Architecture

Different discourse genres — casual conversation, academic lecture, legal testimony, religious ritual, intimate disclosure — maintain different admissibility manifolds. The same utterance may be admissible in one genre and inadmissible in another, not because the sentence is syntactically different but because the genre establishes different constraint structures.

Genre is therefore an admissibility architecture: a persistent constraint structure that shapes the discourse manifold for a class of communicative events. Participants entering a genre do not merely follow its conventions; they enter its admissibility manifold

and calibrate their utterance operators to the field's geometry. A participant who produces a genre-inappropriate utterance has not violated a convention; they have crossed a boundary. The conversational repair that follows (reframing, correction, discomfort) is the genre's repair infrastructure responding to the boundary crossing.

Language Change as Manifold Deformation

Constructions and Their Repair Ratios

Applying the Semantic Stability framework to grammatical constructions yields a precise account of language change as manifold deformation.

Definition 6 (Constructional repair ratio)**.** For a grammatical construction c, the *constructional repair ratio* is:

$$\rho(c) = \frac{S(c)}{\theta(c)},$$

where $S(c)$ is the *communicative support* of c — the degree to which the construction is actively reinforced through successful use in communicative exchanges — and $\theta(c)$ is the *maintenance threshold* of c — the minimum support required for c to remain within the admissibility manifold of the language.

A construction is grammatically alive when $\rho(c) \geq 1$. It is undergoing attrition when $\rho(c) < 1$: the communicative support has fallen below the maintenance threshold, and the construction is in the process of exiting the language's admissibility manifold.

Language change, on this account, is the evolution of the admissibility manifold through the differential attrition and emergence of constructions. A construction that consistently produces admissible discourse states accumulates repair support and deepens its basin in the manifold. A construction that fails to produce admissible states loses support and approaches the boundary. When $\rho(c)$ crosses 1, the construction undergoes the same failure-front propagation established for semantic trees: forms that depended on the failing construction for their own admissibility lose support, and the collapse propagates through the grammatical dependency structure.

Sound Change as Repair of Articulatory Effort

Sound change is the clearest case of manifold deformation as repair. The pressures driving sound change are well-understood: assimilation (adjacent sounds influence each other), lenition (intervocalic consonants weaken), merger (phonemically distinct sounds become indistinguishable in certain contexts), and chain shifts (systematic reorganizations of the vowel or consonant space). These are all repair operations applied to the articulatory- perceptual constraint manifold.

Assimilation is repair of articulatory discontinuity: the effort required to produce a sequence of sounds with very different articulations is reduced by partially assimilating adjacent sounds to each other. The result is a discourse trajectory that remains within the articulatory admissibility manifold with lower articulatory cost.

Lenition is repair of constraint excess: intervocalic consonants carry minimal distinguishing load in many phonological contexts, and the effort required to maintain their full stop or fricative articulation exceeds the constraint maintenance benefit. The lenited form costs less effort while remaining within the perceptual admissibility manifold of the listener.

Grammaticalization is repair of semantic underspecification: a construction that repeatedly appears in contexts where its meaning is semantically bleached — used for a grammatical function rather than its lexical meaning — accumulates the grammatical reading as an admissible trajectory and loses the lexical reading as the context consistently selects against it. The construction's admissibility region shifts from the lexical to the grammatical manifold.

Language Death as Repair Collapse

The Semantic Stability Theorem, applied to language as a whole, gives the following account of language death.

A language is a repair equilibrium: a constraint-field system in which the admissibility manifold $\mathcal{A}_L$ is maintained by the continuous repair activity of the speech community. Every successful communicative exchange is a repair operation that reinforces the constructions it employs. Every failed exchange that is subsequently repaired strengthens the repair infrastructure. The language persists as long as the repair rate meets or exceeds the degradation rate.

Language death is repair collapse: the failure of the repair rate to meet the degradation rate, producing a failure-front propagation from peripheral constructions toward the core vocabulary and grammar.

The branch stability index $\rho_{\min}(\mathcal{L})$ — the minimum repair ratio across all constructions in the language — is the survival condition. When $\rho_{\min}(\mathcal{L}) < 1$, at least one construction has crossed the admissibility boundary, and the failure-front dynamics of the Semantic Stability Theorem apply.

Corollary 1 (Language death as repair failure)**.** A language undergoes irreversible collapse when $\rho_{\min}(\mathcal{L}) < 1$ for a connected subgraph of the constructional dependency network. The collapse initiates at the peripheral constructions (those with the smallest repair surplus $S(c) - \theta(c)$), propagates inward through the dependency structure,

and terminates when either external support is introduced or the core vocabulary falls below its admissibility threshold.

The Critical Repair Divergence theorem applies here with particular force. As a language approaches the admissibility boundary — as $\rho_{\min}$ approaches 1 from above — the repair intensity required to maintain the remaining constructions diverges. Late-stage language revitalization is extremely costly not because the language is inherently difficult to revitalize but because the system is operating near its admissibility boundary, where $\iota \to \infty$. The cost of maintaining each surviving construction approaches infinity as the surrounding repair network collapses.

This also explains why revitalization requires reconstructing the repair infrastructure rather than merely reintroducing the content. A dictionary of a dying language contains the nodes $\ell(v)$ but not the constraint infrastructure $S_R(w) \geq \theta_w$ that makes each node semantically alive. A child who learns vocabulary from a dictionary but does not acquire the discourse manifold is not a speaker of the language; they are a user of stored structures disconnected from the living repair equilibrium.

RSVP as Field-Theoretic Synthesis

The discourse manifold, the utterance operator, the repair ratio, and the Admissibility Preservation Theorem are sufficient to account for the phenomena examined in the preceding sections. The Relativistic Scalar-Vector Plenum framework does not introduce new empirical claims about language; it provides a field-theoretic substrate in which the discourse-level quantities acquire continuous geometric descriptions. RSVP appears here as a consequence of the framework established through the case studies, not as a starting assumption.

Definition 7 (RSVP-discourse correspondence)**.** The RSVP field triple $(\Phi, v, S_{\text{RSVP}})$ is identified with the discourse field by the following explicit equalities, forced by the requirement that the RSVP stability condition coincide with the Admissibility Preservation Theorem:

$$\Phi(x) = S_{\text{support}}(x) \quad \text{(semantic density = support reaching state } x\text{)} \tag{1}$$

$$|v_{ij}| = \alpha(i,j) \quad \text{(vector magnitude = constraint transmission coefficient)} \tag{2}$$

$$S_{\text{RSVP}}(x) = -\log \rho(x) \quad \text{(entropy field = log-inverse repair ratio)} \tag{3}$$

The identification $S_{\text{RSVP}} = -\log \rho$ is the structurally significant one. It has three immediately verifiable properties: when $\rho(x) > 1$ (the discourse state is inside the admissibility manifold), $S_{\text{RSVP}} < 0$ — the entropy is negative, encoding the constraint surplus that keeps the state interior. When $\rho(x) = 1$ (exactly at the boundary), $S_{\text{RSVP}} = 0$. When

$\rho(x) < 1$ (outside the manifold), $S_{\text{RSVP}} > 0$ — positive entropy signals the expansion of interpretation volume beyond what the constraint structure can contain. The Admissibility Preservation Theorem translates directly: discourse coherence is maintained iff $S_{\text{RSVP}}(x_k) \leq 0$ for all steps k.

Proposition 3 (RSVP admissibility condition)**.** Under the identification above, an utterance u is admissible in discourse state x if and only if $S_{\text{RSVP}}(u(x)) \leq 0$.

Proof. $S_{\text{RSVP}}(u(x)) = -\log \rho(u(x)) \leq 0$ iff $\rho(u(x)) \geq 1$ iff $u(x) \in \mathcal{A}(x)$. □

The scalar field $\Phi = S_{\text{support}}$ provides the field geometry of the discourse. High Φ at a discourse state means richly constrained interpretation: a topic-initial position, a loaded term, a well-established presupposition. Low Φ means loosely constrained interpretation: an open-ended question, a genre-boundary crossing, an unfamiliar register. The vector field magnitude $|v_{ij}| = \alpha(i, j)$ is the constraint transmission coefficient between adjacent discourse states: how much of the constraint established at i reaches j. Low $|v|$ on an edge means high attenuation — culturally distant participants, mismatched registers, or the communicative equivalent of a non-connecting letter.

Orthographic Conventions as Generalized Repair Infrastructure

The Arabic analysis reveals a generalization that extends to all writing systems and to many non-orthographic symbolic conventions. The structural features that the Arabic system uses to maintain its admissibility manifold — dots, non-connecting letters, positional regularization, diacritic layers — all function as repair infrastructure built into the medium.

Once this is seen, the same structure becomes visible elsewhere. Word spacing in most alphabetic scripts is a non-connecting operator: a mandatory segmentation that prevents ambiguity from propagating across word boundaries, exactly as the non-connecting letters in Arabic prevent ambiguity from propagating within words. Without word spaces, a Latin-script text becomes interpretively ambiguous at every boundary: *therapist* versus *the rapist*, *nowhere* versus *now here*. The space is not a convention; it is a boundary-preservation operator that maintains $D_{\text{join}} \geq D_{\text{min}}$ at word boundaries.

Punctuation is turn-taking infrastructure for written discourse: the period terminates a clause bubble, the comma maintains syntactic connection across a pause, the colon announces a constraint-commitment (what follows will be constrained by what preceded), the quotation mark shifts the admissibility manifold to a nested register. Capitalization distinguishes the start of a new repair equilibrium (sentence-initial, proper name) from continuation within an existing one.

Paragraph breaks, section headings, and chapter divisions are higher-order fault bar-

riers: points at which the admissibility manifold resets to a lower constraint density, allowing readers to re-enter the document's semantic field without carrying all prior constraints. A text without paragraph breaks requires readers to maintain the entire prior constraint structure in working memory; the break is a repair operation that reduces the cognitive cost of re-entering the manifold.

In the RSVP field interpretation: every orthographic convention that reduces ambiguity is a point intervention that lowers S_{RSVP} at the affected node. Every convention that marks a boundary (word space, period, paragraph break) is a CLIO projection that collapses the branching interpretation space into a determinate trajectory. Writing systems are not representations of speech; they are layered repair infrastructures for discourse admissibility, and their conventions are the materialized residue of accumulated repair operations that have been found efficient enough to warrant standardization.

Remark 2 (Companion paper). The analysis of writing systems as boundary-preservation technologies is sufficiently rich to constitute a separate inquiry. A companion essay — *Repair Infrastructures: Writing Systems as Boundary-Preservation Technologies* — would extend the Arabic analysis to Latin script conventions, to ideographic and logographic systems, to the history of punctuation, and to digital typography. The present section establishes the theoretical connection; the companion would develop the empirical breadth.

Gesture Before Symbol: The Root Level

The gesture-before-symbol thesis — that meaning is prior to its linguistic encoding, and that symbols are repair layers applied to pre-symbolic meaning — specifies the position of the discourse manifold in the overall hierarchy of semantic structure.

Definition 8 (Semantic hierarchy)**.** The semantic hierarchy from meaning to utterance has the following structure:

$$\begin{aligned}
\text{Gesture} &: \text{pre-symbolic meaning grounded in shared attention and action} \\
&\downarrow \\
\text{Constraint field} &: \text{the admissibility structure that makes utterances receivable} \\
&\downarrow \\
\text{Root attractor} &: \text{semantic anchor maintaining coherence across derivations} \\
&\downarrow \\
\text{Morphological repair} &: \text{pattern operators adapting the root to grammatical context} \\
&\downarrow \\
\text{Utterance} &: \text{the committed trajectory within the admissibility manifold}
\end{aligned}$$

Each level of the hierarchy is a repair layer applied to the level above it. Grammar does not generate meaning from rules; grammar repairs the transmission of pre-symbolic meaning into the social medium of language, maintaining semantic coherence against the attenuation introduced by the arbitrariness of the symbol, the distance between interlocutors, and the variation in individual context.

This hierarchy inverts the generative direction: rather than meaning being produced by the grammar, meaning is the source, and the grammar is the infrastructure that maintains access to it. A language is not a machine for producing sentences. A language is a repair ecology that preserves shared access to the pre-symbolic meaning that speakers bring to the conversation.

The consequence for language acquisition is significant. A child acquiring language is not learning to generate sentences; they are calibrating their sensitivity to the discourse admissibility manifold. The critical period for language acquisition is not the period during which grammar rules are internalizable but the period during which the field geometry can be calibrated at the perceptual level. A child who acquires a language during the critical period inhabits its admissibility manifold; a child who acquires it afterward must compensate with explicit rule knowledge, which is a less efficient repair infrastructure.

Conclusion

The constraint-field theory of language proposed in this essay makes three central claims that distinguish it from the generative tradition.

The first is that the primitive object is the discourse admissibility manifold, not the sentence or the derivation. Grammar is the constraint structure that defines the manifold; syntax is the record of how those constraints propagate through utterances; grammaticality is a property of trajectories rather than objects.

The second is that linguistic competence is field sensitivity, not rule knowledge. Native speakers inhabit the admissibility manifold; they detect boundary crossings as violations rather than computing rule failures. The developmental trajectory of acquisition is the trajectory of field inhabitation. The asymmetry between native and non-native speaker behavior is the asymmetry between manifold inhabitation and explicit constraint computation.

The third is that language change is manifold deformation driven by the differential repair rates of grammatical constructions. Sound change, semantic change, and grammaticalization are all repair operations on the articulatory, semantic, and grammatical constraint manifolds respectively. Language death is repair collapse. Language revital-

ization requires reconstructing the repair infrastructure, not reintroducing the content.

The Arabic root-and-pattern system is the most formally transparent instantiation of this framework: the root as semantic attractor providing continuous repair pressure on its derived forms, the patterns as repair operators adapting the root to grammatical context, the orthography as a repair infrastructure that prevents lexicalization by keeping the root visible at every use. Spanish verbal agreement is an instance of multi-field coherence: the correct form is the unique attractor in the intersection of simultaneously active constraint fields, and native competence is the perceptual sensitivity to that intersection's geometry.

Grammar does not generate utterances. Grammar maintains the admissibility structure that makes utterances receivable. The utterance is not the output of a production system; it is the committed trajectory of a repair process that seeks to transmit meaning through a medium — the social-linguistic field — that continuously introduces attenuation. Every successful utterance is a repair. Every successful conversation is a shared repair equilibrium, maintained in real time by participants who have calibrated their sensitivity to the same constraint field.

The river is not the water. The grammar is not the sentence. The grammar is the work that makes the sentence possible.

Flyxion / Independent Researcher. 2026.