

The Persistence Hierarchy

Adjective Order as Reachability Gradient

Flyxion

Independent Researcher

June 2026

Abstract

Every native English speaker knows that *a small red wooden box* sounds right while *a wooden red small box* sounds wrong. Almost none can explain why. The standard linguistic account lists a fixed sequence—opinion, size, age, shape, colour, origin, material, purpose, noun—but offers no principled explanation of why this sequence and not another. This paper proposes that adjective order encodes a *persistence hierarchy*: the ordering reflects how deeply each property is coupled to the identity and continued existence of the object it modifies. Properties that survive more perturbations of the object sit closer to the noun; properties that are observer-dependent, contextual, or easily altered sit further away. This account derives the standard ordering from a single explanatory principle rather than memorising a chart. It connects naturally to the Spanish *ser/estar* distinction, which grammaticalises the same hierarchy through verb choice rather than word order, and to a broader architectural principle visible in writing systems: that stable symbolic substrates support low-cost refinements, with high-value distinctions encoded closest to the base and low-value annotations pushed outward. The Arabic script's consonantal skeleton and diacritic layer instantiate this logic in orthographic form; the English adjective sequence instantiates it in positional order. A cross-linguistic analysis of Arabic nominal modification further reveals that the hierarchy connects to the classical Arabic philosophical distinction between *jawhar* (substance, essence) and *carad* (accident, contingent property)—a distinction developed by Al-Farabi and Ibn Sina in which the English adjective-order rule participates implicitly. The persistence hierarchy is further interpretable as a reachability gradient in the sense of admissibility theory: properties closer to the noun have higher reachability divergence, meaning that losing the distinction they mark forecloses more futures for agents who depend on accurate object representation. Native speakers can navigate the persistence hierarchy without knowing the chart exists because they are tracking a genuine feature of the world—the differential stability of object properties—rather than applying an arbitrary syntactic convention.

1. The Chart Nobody Knows

Ask a native English speaker whether *a small red wooden box* sounds natural. They will say yes immediately. Ask whether *a wooden red small box* sounds natural. They will say no, and they will feel the wrongness viscerally, the way a wrong note in a familiar melody feels wrong.

Now ask them why.

Almost none can answer. They know the rule in the sense that they reliably follow it, but they do not know that they know it. The rule has never been explicitly taught to most native speakers. It is acquired implicitly, the way children acquire the phonology of their language: through exposure, without instruction, and without conscious access to the underlying principle.

The rule in question is the English adjective order constraint. When multiple adjectives modify a noun, they must appear in a fixed sequence:

Position	Category	Examples
1	Determiner	a, the, this, those
2	Quantity	four, some, many
3	Opinion	adorable, awful, delicious
4	Size	big, tiny, massive
5	Age	old, young, ancient
6	Shape	round, square, oval
7	Colour	red, blue, faded
8	Origin	Italian, eastern, Greek
9	Material	wooden, silk, metal
10	Purpose	dining, sleeping, racing
	Noun	box, table, car

This ordering is not a stylistic preference. It is a grammatical constraint. Violations are not merely inelegant; they are ungrammatical in the same sense that subject-verb disagreement is ungrammatical. A native speaker who produces *a wooden red small box* has made an error, even if they cannot identify it as such when asked.

The standard linguistic account describes the sequence but does not explain it. Why should opinion precede size? Why should material precede purpose but follow origin? Why should colour sit between shape and origin rather than elsewhere? The chart is a description of a pattern. It is not a theory of why the pattern exists.

This paper proposes such a theory.

2. The Persistence Hierarchy

2.1. The basic idea

The ordering of adjectives in English reflects the degree to which each property is coupled to the identity and continued existence of the object it modifies. Properties that are constitutive of the object—that survive many perturbations without the object ceasing to be itself—sit close to the noun. Properties that are observer-relative, contextual, or easily altered without changing the object’s identity sit farther away.

Call this the *persistence hierarchy*. Formally:

Definition 2.1 (Persistence of a Property). Let O be an object and P a property attributed to O . The *persistence* of P with respect to O is the measure of the set of perturbations of O under which P remains true:

$$\text{Pers}(P, O) = \mu(\{\phi \in \Phi : P(\phi(O)) = \mathbf{true}\}),$$

where Φ is the space of possible perturbations of O and μ is a measure on Φ reflecting the prior probability of each perturbation.

Definition 2.2 (Persistence Depth). The *persistence depth* of adjective category \mathcal{C} is the average persistence of properties in \mathcal{C} across objects and domains:

$$\text{Depth}(\mathcal{C}) = \mathbb{E}_{O, P \in \mathcal{C}}[\text{Pers}(P, O)].$$

The Persistence Hierarchy Hypothesis is then:

Proposition 2.3 (Persistence Hierarchy). *Adjectives in English are ordered such that adjective category \mathcal{C}_1 precedes \mathcal{C}_2 in the prenominal sequence if and only if $\text{Depth}(\mathcal{C}_1) < \text{Depth}(\mathcal{C}_2)$.*

In plain language: the less persistent the property, the farther from the noun. The more persistent the property, the closer to the noun.

2.2. Deriving the standard ordering

Let us work through the standard sequence and verify that persistence depth increases monotonically from opinion to noun.

Opinion (position 3). Opinion adjectives encode the speaker’s evaluative response to the object. *Beautiful, adorable, awful*. These are maximally observer-relative: the same object is beautiful to one observer and ugly to another. The same observer

may revise their opinion tomorrow. Opinion properties are highly sensitive to perturbations of the *observer* rather than the object. Their persistence with respect to the object itself is minimal. Hence they sit furthest from the noun.

Size (position 4). Size is more objective than opinion but highly context-relative. *Big* relative to what? A big mouse is smaller than a small elephant. Size also changes: objects shrink, expand, are cut down or built up. The property is less observer-relative than opinion but more mutable than material constitution. Persistence is low-to-moderate.

Age (position 5). Age is more stable than size but not fully constitutive. An old table was once a young table. The same object passes through age predicates over time. Age is a temporal property that the object accumulates rather than a structural property that defines it. Persistence is moderate: the property will change, but more slowly than opinion or contextual size.

Shape (position 6). Shape is more stable than age in many cases. The round table remains round across many perturbations—moving it, painting it, ageing it. Shape changes require structural intervention. Persistence increases.

Colour (position 7). Colour sits between shape and origin. It is more stable than opinion (there is an objective fact about the colour of the table) but less stable than material (you can repaint the table without changing what it is made of). Repainting is a less drastic intervention than reshaping, and reshaping is less drastic than replacing the material entirely. The persistence ordering follows: shape < colour < material is wrong; the correct ordering is shape slightly more stable, then colour, then origin and material. The chart confirms this: shape (6) precedes colour (7) which precedes origin (8) and material (9).

Origin (position 8). Origin is historically fixed. An Italian table was made in Italy and cannot retrospectively have been made in France. Origin is a relational property fixed at creation. It cannot be changed by any perturbation of the object that leaves the object in existence. Persistence is high.

Material (position 9). Material is constitutive in the strongest sense. A wooden table is made of wood throughout its existence as that table. You can paint it, scratch it, move it across the world—it remains wooden. To change the material is to destroy the table and create a different object. Material properties survive almost all perturbations short of destruction. Persistence is near-maximal.

Purpose (position 10). Purpose sits adjacent to the noun because it specifies what the object is *for*, which is intimately connected to what kind of thing it is. A dining table is not merely a table that happens to be used for dining: the purpose is built into the categorisation. Purpose properties are relational but stable: a dining table does not cease to be a dining table when it is not in use. Persistence is high, reflecting the tight coupling between function and identity.

The noun (position 11). The noun is the maximally persistent residue: it is the compressed identity of the object itself. All other properties are modulations of this core.

2.3. Why the ordering is felt rather than known

Native speakers acquire the persistence hierarchy implicitly because it tracks a genuine feature of the world rather than an arbitrary convention. The hierarchy of property-persistence is available to any cognitive system that represents objects over time and across perturbations. Children who have learned that wooden things remain wooden when painted, but beautiful things can stop being beautiful without any physical change, have implicitly acquired the persistence ordering of those two property types. The grammatical rule is downstream of this cognitive representation, not upstream of it.

This explains an otherwise puzzling asymmetry: children acquire adjective order before they can explain it, ESL learners can explain it before they acquire it fluently. The native speaker is navigating the persistence hierarchy; the ESL learner is consulting a map of it. Both arrive at the same sequence, but by different routes, and the native speaker's route is faster and more reliable in production because it is grounded in object cognition rather than metalinguistic rule-following.

3. The Spanish Parallel: Ser and Estar

Spanish provides independent evidence for the persistence hierarchy by grammaticalising it through a different means: the choice between two copulas.

English uses a single verb *to be* for all predications. Spanish uses two: *ser* and *estar*. The traditional description contrasts permanent properties (*ser*) with temporary ones (*estar*), but this description is incomplete and generates famous counterexamples. A better description is that *ser* encodes properties treated as constitutive or identity-defining, while *estar* encodes properties treated as contingent, resultant, or context-specific.

Ser (constitutive)	Estar (contingent)
La mesa <i>es</i> de madera.	La mesa <i>está</i> rota.
The table is (made of) wood.	The table is broken.
Él <i>es</i> médico.	Él <i>está</i> enfermo.
He is a doctor.	He is sick.
La silla <i>es</i> redonda.	La silla <i>está</i> sucia.
The chair is round.	The chair is dirty.
El cielo <i>es</i> azul.	El cielo <i>está</i> nublado.
The sky is blue (by nature).	The sky is cloudy (right now).

The *ser/estar* distinction is a grammaticalisation of persistence depth. *Ser* selects for high-persistence properties: material constitution, professional identity, geometric shape, natural colour. *Estar* selects for low-persistence properties: broken states, illness, dirt, current atmospheric conditions.

The famous puzzle cases resolve naturally under this analysis. *Él está muerto* (he is dead, using *estar*) seems to violate the permanent/temporary rule since death is permanent. But death is a resultant state, the outcome of a process, and *estar* is the verb of resultant and contextual states. The persistence-hierarchy account handles this without stipulation.

Similarly, *la fruta está verde* (the fruit is green, not yet ripe) uses *estar* even though colour is a relatively stable property. But here the speaker is treating the greenness as a contingent stage in the fruit's development—it will ripen—rather than as a constitutive colour property. The choice of *estar* signals that the colour is being construed as low- persistence in this context. The same property can be construed as high- or low-persistence depending on speaker intent, and *ser/estar* grammaticalises that construal. English adjective order lacks this flexibility: material always precedes colour in English because the ordering is fixed in the grammar rather than available for construal-sensitive adjustment.

This cross-linguistic comparison suggests that the persistence hierarchy is a universal feature of human object cognition that different languages grammaticalise in different ways. English encodes it as positional order. Spanish encodes it as copula selection. Other languages may encode it through other means: classifier systems, evidentiality hierarchies, or nominal modification patterns.

4. Adjective Order as Reachability Gradient

The persistence hierarchy connects naturally to the reachability framework developed in a companion paper [1]. That paper introduces *reachability divergence* $R_D^\theta(X, Y)$ as the measure of the set of futures accessible from meaning X but not meaning Y in domain θ . The central claim is that linguistic distinctions survive when collapsing them would foreclose futures a community needs to keep open.

Adjective order can be interpreted as a spatial encoding of reachability divergence. When an adjective A with high persistence is placed close to the noun, it occupies a position that signals: the distinction A /not- A controls which object you are dealing with, and therefore controls which futures are available to you. When an adjective B with low persistence is placed far from the noun, it signals: the distinction B /not- B is peripheral to the object's identity and controls relatively few futures.

More precisely:

Definition 4.1 (Adjectival Reachability Divergence). Let O be an object, P a property, and θ a domain of action. The *adjectival reachability divergence* of P with respect to O in θ is

$$R_D^\theta(P, O) = \mu^\theta \left(\mathcal{A}_{O=P}^\theta \triangle \mathcal{A}_{O \neq P}^\theta \right),$$

where $\mathcal{A}_{O=P}^\theta$ is the admissible future region for an agent who correctly knows that O has property P , and $\mathcal{A}_{O \neq P}^\theta$ is the admissible future region for an agent who correctly knows that O lacks P .

The Persistence Hierarchy Hypothesis then becomes a special case of a more general reachability claim:

Proposition 4.2 (Reachability Gradient). *Adjective order in English encodes adjectival reachability divergence: adjective category C_1 precedes C_2 in the prenominal sequence if and only if $R_D^\theta(C_1, O) < R_D^\theta(C_2, O)$ on average across objects O and domains θ .*

The connection to persistence depth is immediate: high-persistence properties tend to have high reachability divergence because they are more constitutive of the object's identity, and correct knowledge of constitutive properties opens more futures than correct knowledge of peripheral properties.

Consider a concrete illustration. An agent navigating a woodworking domain needs to know whether the material is wood or metal before selecting tools, fasteners, finishing products, and joining methods. The distinction wooden/non-wooden has high R_D^θ in this domain: knowing it correctly opens the correct set of futures; not knowing it (or holding a merged representation that does not mark it) forecloses those futures. By contrast, knowing whether the object is *beautiful* controls almost no futures in the woodworking domain. The carpenter proceeds identically

regardless. The reachability divergence of *beautiful* / *not-beautiful* is near zero in this domain.

The adjective-order rule places *wooden* near the noun and *beautiful* far from it. The positional encoding tracks the reachability significance of the distinction.

4.1. Why violations feel wrong

The reachability gradient account makes a specific prediction about grammaticality judgements. Adjective-order violations should feel worse when the misplaced adjective has higher reachability divergence. Specifically:

Proposition 4.3 (Violation Severity). *The perceived ungrammaticality of an adjective-order violation is a monotonically increasing function of the adjectival reachability divergence of the displaced adjective.*

This predicts that moving a material adjective inward (away from the noun to a position that implies lower persistence) should produce stronger ungrammaticality judgements than moving an opinion adjective outward (to a position implying higher persistence than it merits). Informally:

Violation	Type	Predicted severity
a red wooden small box	colour before material misplaced	moderate
a wooden beautiful box	material before opinion	severe
a beautiful wooden box	opinion before material	mild (near-canonical)
a small wooden old box	size before material/age reordered	moderate

The violation *a wooden beautiful box* should be perceived as severely ungrammatical because it places the high-reachability material adjective in a low-reachability position, misrepresenting the persistence hierarchy of the object’s properties. The violation *a beautiful wooden box* is near the canonical order (opinion, then material) and should be perceived as only mildly deviant or acceptable, because opinion naturally precedes material in the standard sequence.

This prediction is testable against existing psycholinguistic data on adjective order. Several studies have collected grammaticality ratings for adjective-order violations [11, 7, 9, 4]; the reachability gradient account predicts a correlation between rated ungrammaticality and the persistence depth of the displaced adjective.

4.2. Implicit navigation without explicit knowledge

The finding that native speakers feel adjective-order violations without being able to explain them is predicted by the reachability account but would be puzzling on a purely syntactic account. If adjective order were an arbitrary grammatical

convention, there would be no reason to feel it rather than merely know it. Arbitrary conventions can be learned and applied without affective response.

But if adjective order tracks the persistence hierarchy, and the persistence hierarchy reflects genuine facts about how object properties interact with future action, then violation of adjective order is a misrepresentation of object structure. Agents who have learned to navigate the world using objects as tools will have developed sensitivity to the persistence hierarchy through ordinary cognition. The grammatical rule is felt because it encodes something real that object cognition already tracks.

This is analogous to why mathematical proofs feel compelling to trained mathematicians even before they have fully verified every step: the proof tracks genuine structure, and the trained mind can detect that tracking even partially. A grammaticality judgement on adjective order is a compressed evaluation of whether the property ordering matches the persistence hierarchy of the object being described.

5. Nested Structure and the Spherepop Interpretation

There is a natural interpretation of the persistence hierarchy in terms of nested process structure. Consider the noun not as a static object but as the compressed residue of a generative process: the table is the outcome of selecting material, imposing shape, applying finish, assigning function, and locating in a context.

Each adjective layer represents a shell in the process that generated the object. Properties at inner shells are constitutive of the process outcome. Properties at outer shells are features added after the constitutive process is complete.

Shell	Property type
Innermost	Material: constitutive of what the object is Origin: historically fixed at creation Purpose: built into the object's functional identity
Middle	Colour: applied after shaping Shape: imposed on material Age: accumulated over time
Outermost	Size: relational and context-dependent Opinion: observer-relative and mutable

In this picture, the noun is the innermost bubble in a nested structure. Each adjective is a layer around that bubble. The layers are ordered by how closely they participate in the generative process that produces the object. Inner layers can be thought of as programming the object's identity; outer layers annotate it from the outside.

The standard adjective-order sequence then reads as a traversal of this nested structure from outside in: the reader or listener encounters the most contingent, observer-relative properties first, then progressively closer to the constitutive core, arriving at the noun as the innermost identity.

This traversal order has a functional interpretation. In language processing, early adjectives set up the context of evaluation and constrain the space of possible referents coarsely. Later adjectives narrow the referent space more finely. Opinion tells you roughly what evaluative class the object falls into. Size narrows it further. By the time you reach material, you are very close to uniquely identifying the object type. The noun then closes the identification. The ordering is cognitively efficient: it moves from high-entropy, context-setting information to low-entropy, identity-fixing information.

This efficiency is not separable from the persistence hierarchy. It is the same hierarchy viewed from two directions. Moving from observer to object is simultaneously moving from contingent to constitutive, from high-entropy to low-entropy, from mutable to persistent.

6. Script Architecture and the Incremental Augmentation Principle

The persistence hierarchy in adjective order is not an isolated phenomenon. It belongs to a broader family of structures in which a stable substrate supports layered, low-cost refinements that increase informational resolution without requiring wholesale redesign. Writing systems exhibit this architecture with unusual clarity, and examining it strengthens the theoretical basis for the persistence hierarchy account.

6.1. Stable substrate and minimal refinement

Consider how the Arabic consonantal script handles the problem of phonemic distinction. Early manuscripts were written in a consonantal skeleton—the *rasm*—in which many letters that are today graphically distinct shared identical base shapes. The forms now written as **b**, **t**, and **th** were graphically identical. Disambiguation relied on context.

Over time, a solution emerged that did not abandon the efficient skeletal layer: systematic dotting. A single dot placed above or below the shared base form was sufficient to differentiate multiple phonemes. The informational return of this augmentation is disproportionately large relative to its cost. A dot is a minimal motor gesture, yet it yields a categorical distinction.

The structural principle this illustrates—call it *incremental augmentation*—is that symbolic systems under constraint tend to introduce new distinctions through minimal additions to existing structure rather than through wholesale replacement.

The base layer is preserved because it is already efficient; the augmentation is accepted because its cost is low relative to its informational benefit.

Definition 6.1 (Incremental Augmentation). Let X be a base symbolic structure and A a minimal augmentation operator. The augmentation is *efficient* if

$$C(A) \ll C(X) \quad \text{while} \quad I(A(X)) - I(X) \gg 0,$$

where C is production cost and I is informational distinctness. An efficient augmentation achieves high discriminative gain at low additional effort.

The dotted Arabic letters are a paradigm case. The base form b (bare skeletal shape) and the augmented form $d(b)$ (dot added above) satisfy $C(d) \approx \varepsilon$ while $D(b, d(b)) \geq D_{\min}$: minimal cost, maximal discriminative payoff.

6.2. The adjective hierarchy as augmentation sequence

The persistence hierarchy is a cognitive instance of the same principle. The noun is the base form—the maximally compressed representation of the object. Each adjective layer is an augmentation that adds information at increasing cost as you move outward from the noun toward the observer.

Material augmentation (“wooden”) has high informational return because it is constitutive: knowing it correctly opens the full set of futures associated with the object type. Its informational leverage is high. Opinion augmentation (“beautiful”) has low informational return because it is observer-relative: knowing it correctly adds little to your navigational capacity relative to knowing the material and origin. Its leverage is low.

The ordering opinion \rightarrow material is therefore an ordering from low-leverage augmentation to high-leverage augmentation as you approach the noun. The noun is the base; the adjective sequence is a series of augmentations arranged so that the highest-leverage refinements sit closest to what they refine.

This is the same efficiency logic as the dot. Put the high-value distinctions closest to the structure they modify. Put the low-value annotations further away, where they can be dropped if context permits without losing navigational capacity.

6.3. Positional regularization as compression

Writing systems that allow fully context-sensitive joining face a combinatorial explosion: if the shape of each letter depends on both its left and right neighbours, the number of possible forms grows as $\mathcal{O}(n^2)$ for pairwise dependencies. Arabic solved this by reducing context-sensitive variation to four predictable positional allographs—isolated, initial, medial, and final—which scale as $\mathcal{O}(4n)$, linear rather

than quadratic. The combinatorial space is compressed by replacing local variation with a global positional rule.

Adjective order performs exactly this compression for nominal modification. Rather than permitting every possible ordering of adjectives and resolving ambiguity pragmatically, English fixes the positions so that the parsing load is eliminated. A native speaker does not need to consider whether *wooden red small* might mean something different from *small red wooden*: the positional encoding renders one of them ungrammatical and removes the interpretive computation entirely. The grammar pays a small cost in flexibility to achieve a large reduction in processing load—the same trade as the script makes when it moves from context-sensitive ligatures to positional allographs.

6.4. Non-connecting letters as controlled interruption

Arabic scripts face a second optimization problem: maximizing cursive continuity reduces production cost (fewer pen lifts) but risks perceptual collapse (letters become visually indistinguishable when merged into uninterrupted flow). The solution is selective interruption: a small set of non-connecting letters terminates the cursive sequence at exactly the points where continued joining would erode discriminability. These letters function as structured boundary markers, paying a local cost in motor continuity to preserve global legibility.

Adjective order has an analogue. The constraint is not uniform: it applies more rigidly to some category pairs than others, and native speakers show graded judgements rather than categorical rejection across all violations. The rigidity is highest precisely where the positional encoding matters most for disambiguation—where without the fixed order, the persistence hierarchy of the object would be misrepresented. Where it matters less (two high-similarity-persistence categories adjacent to each other), flexibility appears. The grammar interrupts only where interruption is necessary to preserve the hierarchy, not uniformly throughout.

6.5. The general pattern

Across these parallels—augmentation efficiency, positional compression, controlled interruption—the same structural logic recurs. Systems under embodied constraint stabilize not by maximizing any single variable (pure motor efficiency, pure discriminability, pure flexibility) but by achieving local optima where small adjustments no longer improve the balance between competing pressures. The persistence hierarchy in adjective order is one such equilibrium: it is the ordering that maximizes navigational information delivered per unit of positional slot, given the constraint that the slots must be linearly ordered and that the noun anchors the right edge.

This suggests that the English adjective-order rule is not an accident of grammatical history but a convergent solution to a constraint satisfaction problem that

writing systems, phonological systems, and nominal modification systems all face in their own medium.

7. Arabic and the Persistence Hierarchy

Arabic provides an independent and typologically distant illustration of the persistence hierarchy. Egyptian Arabic, like Modern Standard Arabic, does not grammaticalise adjective order as rigidly as English, but the hierarchy is nonetheless visible in how speakers and analysts articulate the relationship between adjectives and the nouns they modify.

Consider the following analysis, offered by a native Arabic speaker reflecting on the English adjective-order phenomenon:

el-fekra di acma' šwayya men eš-šarh et-taqlīdi li-tartīb es-sifāt fil-ingilīzi.

THE-IDEA THIS DEEPER A-LITTLE FROM THE-EXPLANATION TRADITIONAL OF-ORDER THE-ADJECTIVES IN-ENGLISH

“This idea is a little deeper than the traditional explanation of adjective order in English.”

et-tafsīr illi ana šāyfo inn tartīb es-sifāt biyyackis daragit irtibāt es-sifa bi-huwiyyit eš-šay' nafsu.

THE-EXPLANATION THAT I SEE-IT THAT ORDER THE-ADJECTIVES REFLECTS DEGREE ATTACHMENT THE-ADJECTIVE TO-IDENTITY THE-THING ITSELF

“The explanation I see is that adjective order reflects the degree to which an adjective is attached to the identity of the thing itself.”

kull mā kānit es-sifa a'rab li-tabīcit eš-šay' wi-gawharu, kull mā 'irribit min el-ism.

EVERY TIME WAS THE-ADJECTIVE CLOSER TO-NATURE THE-THING AND-ESSENCE-ITS, EVERY TIME IT-CAME-CLOSER TO THE-NOUN

“The closer the adjective is to the nature and essence of the thing, the closer it comes to the noun.”

huwwa miš biytabbi' qācida nahwiyya wācya, lakinnu biyitcāmīl mac taba'āt mukhtalifa min il-irtibāt biš-šay' nafsu.

HE NOT APPLIES RULE GRAMMATICAL CONSCIOUS, BUT-HE DEALS WITH LAYERS DIFFERENT OF ATTACHMENT TO-THE-THING ITSELF

“He is not consciously applying a grammatical rule; rather, he is dealing with different layers of attachment to the thing itself.”

Several things are worth noting in this formulation.

The gloss makes constituency visible. The interlinear Arabic translation makes explicit a feature of the persistence hierarchy that English obscures: the morphological transparency of the attachment relation. Arabic *irtibāt* (attachment, connection) combined with *huwīyyit eš-šay'* (identity of the thing) gives the phrase *daragit irtibāt es-sifa bihuwīyyit eš-šay'*—literally, the degree of the adjective's attachment to the identity of the thing. The construct state (*idāfa*) in Arabic grammaticalises the possession-and-attachment relation directly: *huwīyyit eš-šay'* is “the thing's identity” as a syntactic unit. The persistence hierarchy is therefore expressible in Arabic with greater morphological directness than in English.

The observer/object distinction is grammatically marked. The Arabic formulation distinguishes between *ra'y il-murā'ib* (the observer's opinion, literally “view of the observer”) and *huwīyyit eš-šay'* (the identity of the thing). This is not merely a conceptual distinction; the Arabic morphosyntax makes it a syntactic one. The possessive construction in each case foregrounds whose perspective is operative. Opinion adjectives belong to the observer; constitutive adjectives belong to the thing. The persistence hierarchy is already a distinction between observer-relative and object-relative properties, and Arabic grammaticalises this distinction at the morphological level.

Arabic adjectival agreement reflects persistence depth. In Arabic, adjectives agree with their nouns in gender, number, definiteness, and case. The agreement morphology creates a tight binding between noun and modifier. Critically, adjectives in Arabic follow the noun rather than preceding it: *el-tawla el-khašabiyya* (the table the-wooden, “the wooden table”). This post-nominal position means that Arabic encodes the hierarchy differently from English—not through prenominal ordering but through the sequence of post-nominal modifiers when multiple adjectives appear.

Studies of Arabic nominal modification suggest that when multiple adjectives follow a noun in Arabic, the ordering preferences mirror the persistence hierarchy: constitutive modifiers (material, origin) tend to appear closer to the noun while contingent modifiers (opinion, temporary state) appear farther away. The hierarchy is the same; the directionality is reversed because the noun precedes rather than follows its modifiers.

The verbal gloss reveals the process structure. The Arabic reflection uses the verb *biyyackis* (reflects, mirrors) to describe what adjective order does to the persistence hierarchy. The choice of this verb is significant: a mirror is a faithful representation of structure, not a stipulation. The implication is that adjective order is not imposing an order on properties but revealing an order that already exists in the relationship

between properties and object identity. This is precisely the claim of the persistence hierarchy account: the grammatical rule is downstream of the cognitive structure, not constitutive of it.

The key Arabic term: *jawhar*. The word *jawhar* (essence, substance, jewel) appears in the phrase *tabīcit eš-šay' wi-gawharu* (the nature of the thing and its essence). In classical Arabic philosophy, *jawhar* is the technical term for substance in the Aristotelian sense—the category of being that persists through change and constitutes the identity of a thing. Its appearance here is not accidental. The intuition that material and constitutive properties sit closest to the noun because they are closest to the *jawhar* of the object is a direct inheritance from the Arabic philosophical tradition's engagement with Aristotelian substance metaphysics.

This is a deeper connection than a cross-linguistic parallel. It suggests that the persistence hierarchy in adjective order is the grammatical encoding of a metaphysical distinction—between substance and accident, *jawhar* and *carad*—that has been articulated in the Arabic philosophical tradition since at least Al-Farabi and Ibn Sina. The English adjective-order rule, which most English speakers have never analysed, encodes a distinction that medieval Arabic philosophers developed elaborate theoretical machinery to describe.

The ESL learner consulting a chart is memorising a rule whose philosophical depth they may not suspect. The native speaker following it implicitly is navigating a distinction between substance and accident that Aristotle, translated into Arabic and refined across centuries of Islamic philosophy, made one of the central problems of metaphysics.

8. Cross-Linguistic Predictions

If the persistence hierarchy is a universal feature of human object cognition, it should leave traces in languages other than English, even when those languages do not grammaticalise adjective order as rigidly.

Several predictions follow:

Languages with flexible adjective order should show persistence effects in neutral contexts. In languages where adjective order is grammatically unconstrained, speakers should nevertheless prefer orders that match the persistence hierarchy in unmarked, neutral utterances. Deviations from the hierarchy should be interpretable as marked: an adjective moved outward from its canonical position may signal that the speaker is treating the property as more contingent than usual.

Copula systems should reflect the hierarchy. Languages with multiple copulas should align their copula distinctions with the persistence hierarchy. Spanish

ser/estar is one case. Igbo *bù/ dì* makes a similar distinction. Arabic *kana/laysa* have related distribution patterns. The prediction is that the high-persistence side of any copula distinction should select for the property types that sit closer to the noun in English adjective order.

Classifier systems should track material and shape. Languages with nominal classifiers typically classify by material and shape rather than by opinion or size. Mandarin measure words, Japanese classifiers, and Burmese classifiers all prioritise the constitutive properties that sit innermost in the English adjective-order sequence. This is expected if classifiers are grammaticalisations of the most identity-relevant property dimensions.

Adjective-order rigidity should correlate inversely with other persistence-encoding mechanisms. Languages that grammaticalise the persistence hierarchy through copula selection, classifier systems, or nominal morphology should show less rigid adjective order, because the hierarchy is already encoded elsewhere. Languages that lack these mechanisms should show more rigid adjective ordering as the primary grammatical encoding of the hierarchy.

These are testable predictions that connect the persistence hierarchy account to the typological literature on adjective ordering, copula systems, and nominal classification.

9. Implications for Language Acquisition

The persistence hierarchy account makes specific predictions about the acquisition of adjective order that differ from a purely syntactic account.

On a syntactic account, adjective order is a language-specific parameter that children must learn by exposure to their target language. The order is arbitrary relative to the world and must be memorised. Errors should be random with respect to property type: children should be as likely to misorder material before opinion as opinion before material.

On the persistence hierarchy account, the ordering reflects a non-linguistic cognitive structure that children develop through ordinary object interaction. Children who have learned that wooden objects remain wooden when painted, but beautiful objects stop being beautiful when damaged, have implicitly acquired the persistence ordering of those property types. Grammatical acquisition then consists in aligning this pre-existing cognitive hierarchy with the language-specific positional encoding.

This predicts that children's errors in adjective order should be *non-random*: errors involving high-persistence adjectives displaced to low-persistence positions should be rarer than the reverse, because the cognitive hierarchy already privileges

material and origin adjectives as identity-constitutive. Children may more readily err by placing opinion adjectives too close to the noun (treating a subjective evaluation as more constitutive than it is) than by placing material adjectives too far from the noun (treating a constitutive property as more contingent than it is).

The ESL learner presents the mirror case. They have the cognitive persistence hierarchy from their native-language experience and they have the chart from instruction. Their errors arise not from lacking the hierarchy but from failing to automate the mapping between the hierarchy and positional order. They must consciously consult the rule because the rule has not been integrated into the object-cognition process that the native speaker navigates implicitly. Fluent production requires not knowing the chart but being the chart.

10. Relationship to the Reachability Theory of Lexical Preservation

This paper is a companion to [1], which argues that linguistic distinctions survive when collapsing them would foreclose futures a community needs to keep open. The present paper extends this framework from the survival of lexical distinctions to the internal ordering of nominal modifiers.

The connection runs in both directions.

Adjective order as implicit stability. The persistence hierarchy is a clear case of a linguistic structure that is highly stable without being consciously maintained. No speaker enforces adjective order by consulting the chart. The rule survives because it tracks the persistence hierarchy, and the persistence hierarchy tracks genuine facts about object constitution that agents need for navigation. This supports the companion paper's claim that linguistic structures can be preserved by navigational pressure rather than communicative convention or conscious rule-following.

The persistence hierarchy as reachability ordering. Proposition 4.2 establishes that adjective order is a reachability gradient: position encodes reachability divergence. High-persistence properties have high reachability divergence and sit close to the noun. Low-persistence properties have low reachability divergence and sit far from the noun. The positional encoding is therefore a spatial representation of the same quantity—reachability divergence—that the companion paper shows governs the survival of lexical distinctions.

A unified picture. Both phenomena—which lexical distinctions survive compression, and in what order adjectives appear before a noun—are governed by the same underlying quantity: how much future access depends on the distinction being preserved. Lexical distinctions with high reachability divergence resist colexification. Adjectival properties with high reachability divergence sit close to the noun.

Both are expressions of a single principle: language preserves and foregrounds the distinctions that matter most for navigating the futures agents face.

This suggests that the reachability framework is not specific to lexical organisation but extends to syntactic organisation as well. The persistence hierarchy in adjective order is a syntactic implementation of the same principle that governs lexical stability: put the high-stakes distinctions where they cannot be lost.

11. Conclusion

The English adjective-order constraint is one of the most widely known examples of implicit grammatical knowledge: a rule that native speakers reliably follow without knowing it exists. The standard linguistic account lists the sequence but does not explain it.

This paper proposes that the sequence encodes a persistence hierarchy: the ordering from opinion to material reflects the ordering from least persistent to most persistent among the properties an adjective can express. Properties that survive more perturbations of the object, that are more constitutive of its identity, sit closer to the noun. Properties that are observer-relative, contextual, or easily altered sit further away.

The persistence hierarchy derives the standard adjective order from a single explanatory principle rather than a memorised chart. It explains why native speakers feel the order rather than merely knowing it: the order tracks a genuine feature of object constitution that cognition already represents. It explains the Spanish *ser/estar* distinction as a grammaticalisation of the same hierarchy through copula choice rather than word order. The same efficiency logic—stable base, high-value distinctions closest to the base, low-value annotations pushed outward—governs the Arabic consonantal skeleton and its diacritic layer, the compression of context-sensitive ligatures into positional allographs, and the selective use of non-connecting letters to interrupt cursive flow only where discriminability requires it. Adjective order and writing system design are convergent solutions to the same constraint satisfaction problem. The analysis of Arabic nominal modification further reveals that the hierarchy connects to the classical philosophical distinction between *jawhar* and *carad*—between substance and accident—a distinction that Al-Farabi and Ibn Sina developed into a rigorous theoretical framework while reflecting on the same Aristotelian intuitions that underlie the English adjective-order rule. The chart that English teachers draw on whiteboards is a pedagogical rediscovery of a metaphysical principle that medieval Islamic philosophy had already formalised. The account generates testable predictions about violation severity, cross-linguistic copula systems, classifier hierarchies, and acquisition errors.

The persistence hierarchy is further interpretable as a reachability gradient: ad-

jective position encodes how much future access depends on knowing the property correctly. This connects adjective order to a companion theory of lexical preservation, suggesting that both phenomena—which distinctions survive compression and how modifiers are ordered—are governed by the same underlying principle.

Language puts the things that matter most for navigation where they are hardest to lose.

References

- [1] Flyxion (2026). Why distinctions survive: A reachability theory of lexical preservation. *Independent Research Monograph*.
- [2] Ibn Sīnā (Avicenna) (c. 1027). *Kitāb al-Shifā'* (The Book of Healing). Partial English translation: A. M. Goichon (1938), *La distinction de l'essence et de l'existence d'après Ibn Sīnā*. Desclée de Brouwer.
- [3] Al-Fārābī (c. 930). *Kitāb al-Hurūf* (The Book of Letters). Edition and commentary: M. Mahdi (1969). Dar el-Mashreq.
- [4] Cinque, G. (1994). On the evidence for partial N-movement in the Romance DP. In G. Cinque et al. (Eds.), *Paths Towards Universal Grammar*, 85–110. Georgetown University Press.
- [5] Cinque, G. (2010). *The Syntax of Adjectives: A Comparative Study*. MIT Press.
- [6] Dixon, R. M. W. (1982). *Where Have All the Adjectives Gone?* Mouton.
- [7] Martin, J. E. (1969). Semantic determinants of preferred adjective order. *Journal of Verbal Learning and Verbal Behavior*, 8, 697–704.
- [8] Murphy, G. L. (2002). *The Big Book of Concepts*. MIT Press.
- [9] Scott, G.-J. (2002). Stacked adjectival modification and the structure of nominal phrases. In G. Cinque (Ed.), *Functional Structure in DP and IP*, 91–120. Oxford University Press.
- [10] Svenonius, P. (2008). The position of adjectives and other phrasal modifiers in the decomposition of DP. In L. McNally & C. Kennedy (Eds.), *Adjectives and Adverbs*, 16–42. Oxford University Press.
- [11] Whorf, B. L. (1956). *Language, Thought, and Reality*. MIT Press.