

Zentzumen-fluxu jarraituak izan arren, animalien portaerak ekintza diskretuak erakusten ditu. Etenkadura horrek emergentzia ahul bat eskatzen du: konstante biofisikoen moldatutako espektro-kanalak, erabaki motor kategorikoetara eramaten dutenak. Hala ere, bilaterioen ahalmen konposizionalak, inferentziales eta kontrafaktuales emergentzia sendoago bat eskatzen dute, irudikapen-espazio ez-kommutatiboetan gauzatzen dena. Esparru honek aljebra lineala eta lotura oszilatarioa uztartzen ditu, sentsazioetik pentsamendura doan eboluzio-ibilbidea azaltzeko.

Giltza-Hitzak: Pentsamenduaren hizkuntza. Sintaxia. Ez-kommutatibitatea. Ziurgabetasunaren Printzipioa. (Kognizioaren) Eboluzioa. Pertzepzio-ekintza begizta.

Contra un flujo sensorial continuo, la conducta animal muestra acciones discretas. Tal discontinuidad requiere una *emergencia débil*: canales sensoriales configurados por constantes biofísicas que implican decisiones motrices categóricas. Pero los bilaterios presentan capacidades composicionales, inferenciales y contrafactuals que requieren *emergencia fuerte*, posible en espacios representacionales no-conmutativos. Este marco conceptual conecta álgebra lineal y sincronización oscilatoria para una trayectoria evolutiva de la sensación al pensamiento.

Palabras Clave: Lenguaje-del-pensamiento. Sintaxis. No-conmutatividad. Principio de Incertidumbre. (Evolución de la) Cognición. Bucle Percepción-acción.

Face à un flux sensoriel continu, le comportement animal se manifeste par des actions discrètes. Cette discontinuité implique une *émergence faible* : des canaux sensoriels configurés par des constantes biophysiques qui induisent des décisions motrices catégorielles. Cependant, les bilatériens possèdent des capacités compositionnelles, inférentielles et contrefactuelles qui nécessitent une *émergence forte*, possible dans des espaces de représentation non commutatifs. Ce cadre conceptuel relie l'algèbre linéaire et la synchronisation oscillatoire pour décrire une trajectoire évolutive allant de la sensation à la pensée.

Mots-clés: Langage-de-la-pensée. Syntaxe. Non-commutativité. Principe d'incertitude. (Évolution de la) Cognition. Boucle Perception-action

Uriagereka, Juan; Orús, Román: Weak & Strong Emergence in Language & Mind

Mobilis in Mobili: Weak & Strong Emergence in Language & Mind

Uriagereka, Juan (1); Orús, Román (2)

(1) Linguistics and School of Languages, Literatures & Cultures. University of Maryland. 1401 Marie Mount Hall (or 4225 Juan Ramón Jiménez Hall). College Park, MD, 20742.

juan@umd.edu

(2) Donostia International Physics Center. Manuel de Lardizabal, 4. 20018 – Donostia-San Sebastián.

roman.orus@dipc.org

<http://doi.org/10.61879/riev7022zkia202503>

Reccep.: 2025-09-04

Acept.: 2025-12-24

BIBLID [eISSN 2952-4180 (2025), 70: 2]ABSTRACT

1. Setting up the Issue

The appearance of macroscopic biological functions (or life itself), from complex molecular interactions, and of structured animal behaviors (or mind itself), from unknown conditions in brains, are often considered candidates for emergent characteristics: when new properties cannot be predicted from a reductionist analysis of the system's fundamental parts. Demonstrating strong emergence in biophysics would require showing that emergent properties cannot be explained by a reductionist account (Wilson 2021). A weaker form of emergence is concerned, instead, with difficult-to-predict complex behaviors, although in principle they may be derivable from their lower-level components (see Phillips et al. 2009 on cell biophysics). While discussions about emergence became systematic in the 19th century, they took a deeper turn with the advent of particle physics and its long-range implications.

Dating back to Aristotle, even if emergence was a background concept, the prevailing mechanistic worldview held the universe to be fully understood through its parts. But the rise of thermodynamics made (weak) emergence a bona-fide scientific topic. The irreversible arrow of time in relation to the [second law](#), with a concept of [entropy](#) that cannot be found in the principles governing individual atoms, was the next step. Entropy, however, can still be addressed through [statistical mechanics](#). In contrast, in an [entangled](#) particle pair the properties of the whole are fundamentally non-separable from those of its parts: relevant correlations cannot be explained by analyzing the particles individually or any pre-existing local information. In a sense, the *relationship* between the particles is then more fundamental than the particles themselves, a description of the system's [quantum state](#) requiring a single [wavefunction](#) for the whole [ensemble](#). This is a more radical form of non-reducibility and a candidate for strong emergence.

Customarily, philosophers link strong emergence to what looked like uniquely human capacities: consciousness, creativity, and language. This of course is puzzling: in that the human species only evolved within the last half a million years. Current research, however, has been consistently showing how many animal species possess remarkable cognitive abilities, once thought to be exclusively human (Gallistel & King 2009; see Godfrey Smith 2016 for an argument that the relevant animal emergence is possibly [Precambrian](#)). In the present context, the idea is explored that animal minds may indeed require a form of strong emergence, particularly if they are governed by a logic that is complex enough for them to carry [inferences](#).

Consider a daily human scenario: suppose you normally leave your car in the spot nearest the garage exit. After your daily business, you mindlessly return to pick up the vehicle where you left it, this being routine. Connectionist models reproduce this quite effectively, for neurons that fire together wire together. But today there was a traffic jam and you arrived late, so you had to go to the first garage spot that was empty – in a different floor. When returning to the vehicle, you experienced a tension: you unconsciously went to where you normally park, until you backtracked. Computational theorists of mind rightfully ask systemic connectionists how that inference is even posed, if all we have is neurons firing together after having wired. While one's "first instinct" is to (erroneously) go to the familiar place, an animal mind also seems capable of bypassing that, through illative reasoning, in the process carrying crucial information forward in time: enough to self-correct and attain a different outcome, even if not wired into neural pathways...

That elementary scenario emphasizes two different cognitive states. On one hand, animals use forms of [Hebbian learning](#): for routine tasks that constitute their repetitive habits. On the other, many animals, and certainly humans, also seem capable of breaking said habits. It is because of that creativity that Chomsky's

critique of behaviorism – as a model of the language faculty – led to Fodor’s (1975) extension of the [Computational Theory of Mind](#) (CTM): as a theory of preexisting knowledge. That tension recalls one between empiricism and rationalism, a debate that gained new momentum with the evolutionary synthesis: Species (not just individuals) may in some form accumulate “innate knowledge”. The devil, though, is in the details: to claim that some forms are somehow innate can be implemented multifariously. Are they so because of some condition encoded in “Bauplans” of thought processes, at the level of the underlying molecular biology? Is this an even more abstract physico-mathematical condition on the very possibility of thought processes? Is anything else neurophysiologically possible and relevant to the emergence?

Famously, Fodor took [connectionism](#) to be a reiteration of primitive associationism. At that time, he also predicted how connectionism “will pass”. That assertion has proven wrong, so we find it useful to consider whether there is any point in attempting to reconcile these two central approaches to emergence. For that, we must understand the quarrel that the garage example illustrates: how distributed representations and learning through connections could capture the structure for human (animal) cognition and underlying cognitive processes of a *systematic*, *productive*, and *transparent* sort. It is those conditions that support a claim that thought is, ultimately, a [\(formal\) language](#), in the sense of [generative grammar](#).

In what follows, we delve into the issue of mind emergence, attempting to reconcile the dominant rationalist and empiricist positions on the field. We will do so by invoking the sort of algebra that, while being relevant to the study of language and cognitive psychology – and certainly behind the recent explosion of [Large Language Models](#) – is also foundational in understanding the quantum architecture. We do this in stages, from acoustics up to more abstract realities needed to account for the syntax/semantics interface. Our move is to start in [Fourier analysis](#) (FA), emphasizing the arguable relevance of the [Fourier Transform](#) (FT) in a variety of so-called [representations](#), which we show relevant to other animals too. But once we argue that central point, we can approach formal and substantive realities in the language faculty not just from an algebraic perspective, but furthermore one that could be analyzed in terms of a [Hilbert \(vector\) space](#).

That is a more radical move, which takes us from classical spectral conditions, where the narrow [Uncertainty Principle](#) is fundamental, to more architectural conditions: where a generalized version of this principle obtains with radical consequences. It is then that we review the possibility of transitioning from classical to quantum waves, which has consequences beyond the debate between connectionism and the CTM. We conclude with some reflections on how we expect these notions to be relevant not just to the language faculty, but for any animal mind.

2. An approach based on Algebra

To implement the CTM’s symbolic manipulation, a system of mental representation must possess [generative power](#): a finite set of primitives combined by recursive operations yielding hierarchical structures. The [Language of Thought](#) (LoT) hypothesis instantiates such a system. Whether modeled as a [phrase-structure grammar](#), a [lambda-calculus](#), a [combinatory algebra](#), or a [categorical logic](#), all these formalisms share the minimal expressive architecture required for [compositional](#) thought. Connectionist models too are algebraic, as they rely on vector spaces and matrix transformations, with learning treated as a [stochastic Markovian dynamical process](#) over weight states. Hebbian rules are based on linear algebra (synaptic weights following [associative memory](#) algebra), while more advanced versions with [deep learning](#) use nonlinear [activation](#)

[functions](#) that rely on [gradient-based](#) algebra, with [recurrent neural networks](#) introducing dynamical systems algebra (Smolensky 1990). As early as Marcus's 2001 *The Algebraic Mind*, researchers were arguing for hybrid architectures for neural networks to instantiate rule-like behavior, the algebraic tension being between discrete generative structure (LoT/CTM) and continuous associative structure (connectionism). While this intuition – as amply discussed in Smolensky & Legendre 2005 – is certainly relevant now, we can explore more complex algebraic ideas too.

In a connectionist network, each synaptic weight acts as a kind of dynamically adjusted coefficient, governing how much one unit's output influences another. Such a habitual learning relies on an algebra for continuous dynamical systems, a point of tension for CTM models that require arbitrarily presumed discrete symbols. Therein the rub, though: the symbols that need to be presumed, whose nature is as necessary as far from obvious, from our present understanding of neurophysiology. Do these decompose into more elementary features? If so, how do those relate to neurophysiology, so that they can be represented in minds or bound to brain processes? And even if discrete symbols are aggregates of more elementary features (like nasality, syllabicity, and the like), one must consider whether there is any “geometry” to those – with some more fundamental than others, as classically sketched in McCarthy 1988. Of course, as Descartes and Fermat taught us, geometric problems can be expressed algebraically, in a deeper sense.

There are core features in phonology (like *syllabicity*) that seem essential to that level of representation, in ways that others (like, say, *nasality*) appear less so. While no language lacks syllabic distinctions, some are deprived of nasal [phonemes](#) (Makah in the US Northwest coast, Central Rotokas in Papua New Guinea, etc.). It is possible to treat fundamental features like $[\pm\text{syllabic}]$ (determining syllable architecture) as formal, still considering substantive, subservient on the formal ones, such attributes as implement nasality, aspiration, or other such nuances. Basically, the formal features define a broad, underlying, space of possibilities, within which additional specifications implement further distinctions. Formal features should be universal and independent of cultural instantiation, while substantive features, though available to learners throughout, may vary not just across cultures: they also appear to be specific to modalities, like voicing in phonology, but not in signed languages, where substantive phonological features may involve, instead, hand shapes or facial expressions. Although said features offer a significant chance of being identified in terms of brain conditions – like motor instructions to activate voice onset, in the vocal folds – the focus of the present exercise is more on the formal features.

Abstract formal features are arguably replicated across levels, thus extending from phonology to syntax. Modeling them on phonological features like $[\pm\text{syllabic}]$ or $[\pm\text{consonantal}]$, as in Chomsky & Halle 1968, Chomsky 1974 proposed “categorical” features $[\pm V]$ or $[\pm N]$ within the syntactic level of representation, as abstract and universal. The strongest argument for a formal nature is to remove *any* substantive residue of the feature analysis at each level. For instance:

- (1) a. $[\pm X]$ instantiated as $[\pm\text{CONSONANTAL}]$ or $[\pm N]$ b. $[\pm Y]$ instantiated as $[\pm\text{SYLLABIC}]$ or $[\pm V]$

If (1) holds, a direct question is whether X/Y, whatever they may be, also relate to one another or how a correlation can be established. For an instantiation via (CONSONANTALITY, SYLLABICITY), it is plausible to take values in one such attribute to *inversely correlate* with the other (the more consonantal an element, the least syllabic, and vice-versa). Building on such correspondences, we may be able to develop an algebraic apparatus to understand a geometry for formal features.

Compare, for instance, a [back vowel](#) like /u/ and a [voiceless stop](#) like /k/, to ponder to what extent a brain may process their information differently. Unlike vowels residing in well-defined acoustic representational spaces, a voiceless stop depends less on an auditory than somatosensory-motor loop. Resonance from varying tongue position allows us to distribute phonemes from front to back stops, transitioning from [alveolar](#) to [velar fricatives](#), on to the [glides](#) (front /y/, back /w/), then moving into [vocalic](#) territory: from front /i/, through mid /a/, to back /u/. Although further [articulators](#) may arise, within the [vocal tract](#) space, pairs as follows obtain:


- (2) a. Glide/vowel pairs: (y, i) (w, u)
- b. Lax/tense vowel pairs: (ɪ, i) (ɔ, o)
- c. Stop/fricative consonant pairs: (t, ʃ) (k, x)

In all these, the paired [manner of articulation](#) is virtually identical, though not the realization of each [segment](#): those to the left are shorter, laxer, more punctual, while those to the right are longer, tenser, more continuant – that distinction holds in music too: a string can be plucked ([pizzicato](#)) or glid over ([arco](#)). These sorts of ideas can be made precise through FA and the FT.

3. Fourier Analysis and the Fourier Transform

A (t, ω) correlation, for time t and [angular frequency](#) ω , exists for “wave dependent” entities: the shorter their duration, the larger spread of angular frequencies it takes to approximate, for instance, a relevant sound, and conversely: a precise frequency cluster is ascertained in longer durations. That is because the Fourier Transform (FT) describes a function localized in time in terms of the correlated wave frequencies (in the case of sound, from air pressure) that add up to it. This presumes that a repeated intensity in a medium over time, $g(t)$, is approximated by sums of simpler, sinusoidal, trigonometric functions (Fourier Analysis, FA). The FT implements an integration mapping from the original space of $g(t)$, to separate, in the transformed function, each component wave that adds up to $g(t)$, where properties of the original function are easier to typify than in the original function space. Figure 1 illustrates how a signal reduces to simpler sinusoidal waves of different frequencies and how its expression in the time domain differs from the one in the frequency domain, spiking at frequencies corresponding to sinusoidal decomposition.

Figure 1: View of a signal in the time and frequency domains, from [Kinder Chen's Denoising Data with Fast Fourier Transform](#)



Per FA, we generalize the wave-life in the [unit circle](#) to any wave; correlated variables t, ω describe the recurring period, expressed by the trigonometric relations as in Figure 2 (for a sine wave, equal relations holding for a cosine wave in the complementary [phase](#)). Sine/cosine relations within the unit cycle are algebraically expressed as in Figure 3.

Figure 2: Sine curve and the unit cycle, capture from a [tutorial by Arkadi Etkin](#).





Figure 3: From Wikipedia, [Euler's formula](#)



Function $g(t) e^{-2\pi\omega t}$ maps $g(t)$'s cyclic nature into a unit circle at frequency ω . The FT is an [integral](#) of that expression derived over t : In turn, FT of $g(t)$ yields a function $\hat{g}(\omega)$ over frequency, whose output is a [complex number](#) corresponding to the strength of the frequency in the signal. We can, thus, also express (3a) inversely as in (3b), to integrate over ω instead of t , thus converting a signal in the frequency domain ω to one in the time domain t :

$$(3) \text{ a. } \int_{t_1}^{t_2} g(t) e^{-2\pi\omega t} dt \quad \text{b. } \int_{\omega_1}^{\omega_2} \hat{g}(\omega) e^{-2\pi\omega t} d\omega$$

In order to get accuracy for t in (3a), we then need as large a sum of frequencies over times as possible (a [periodic function](#) is only approximated by a sequence of sums in FA, so the more of these we have, the more accurate the approximation is); conversely, to get accuracy for ω in (3b), we need as large a sum of times over frequencies as we can get.

4. Basic Uncertainty Conditions

The mathematical relationship between the duration of a [wave packet](#) and its frequency bandwidth is summarized in the uncertainty statement in (4):

$$(4) (\Delta t)^2 (\Delta \omega)^2 \geq 1/4, \text{ so } \Delta t \Delta \omega \geq 1/2$$

Frequency expresses a wave's periodicity per time unit – the rate at which the wave occurs. Time magnitudes are positive, just as frequency magnitudes. (We may describe wave direction as clockwise or the opposite through sign polarity, as in the conventional exponents in (3), but the magnitude of the frequency is still positive.) Uncertainty in relation $(\Delta t, \Delta \omega)$ arises because the expressions in (3) ([wave equations](#)) are FTs of one another (time, frequency being [conjugate variables](#)). A way to quantify these random variables is through their [variance](#) (squared [standard deviation](#) from the [mean](#) $\Delta t, \Delta \omega$). Although well-known algebraic ways get (4) from multiplying (3a) and (3b), all that matters in the present context are *the converse relations* in (5), expressing an inverse correlation for the standard deviation of t with regards to that of ω and vice-versa:

$$(5) \Delta \omega \geq 1/2 \Delta t \text{ and } \Delta t \geq 1/2 \Delta \omega.$$

This leads to Heisenberg's [Uncertainty Principle](#) (UP) at the [Planck scale](#).

While not discovering the mathematical property in (4)/(5) (which had been recognized earlier), Heisenberg realized some of its physical significance for subatomic particles (see section 9). He interpreted “spread”, in a quantum mechanical wavefunction's spatial and [momentum](#) dimensions, as the inherent imprecision in simultaneously measuring those physical properties: a precise measurement of one requires maximum uncertainty in the other. The connection presumes the [de Broglie relation](#) taking momentum to be proportional to spatial frequency of the [matter wave](#) at the Planck scale, thus turning the mathematical constraint above into a statement about the limits of reality and measurement at the quantum level. The inequalities in (5) entail the [complementary](#) behavior of elements in dual states, as particles or waves: this [wavefunction collapse](#) stems from Max Born's interpretation of the wave as carrying the probability of a particle's existence at a given point, as defined only through [measurement ensembles](#), each with a different possible outcome. The act of measurement, empirically observed to result in a definite quantity, is taken to remove this probabilistic state and, with it, the wave state itself. But let's not yet force uncertainty in such terms: statement (5) is relevant to perfectly classical waves, like the acoustic ones.

To make that concrete for classical waves, it is useful to conceptualize sound processing prior to our mammalian [cochlea](#), in more basic conditions than those implied for phonemes. Many fish sense sound through sensitivity to water particle motion, detected by hair cells in their inner ears; for instance, the [Lusitanian toadfish](#) is capable of distinguishing a continuous frequency (a “tone”) from a broadband transient (a “click”) – an ability that is present, also, in insects with [tympanal organs](#) that contain mechanoreceptors sensitive to a range of frequencies and amplitude modulations. As in the operation of Edison's [phonograph](#), the fundamental mechanism of sound detection in such (cochlea-less) organisms relies on sheer mechanics. Sound waves push against a biological membrane, a physical displacement that directly opens [ion channels](#) on the receptor cells, the influx of which changes the [electrical potential](#) of the cell, generating a neural signal.

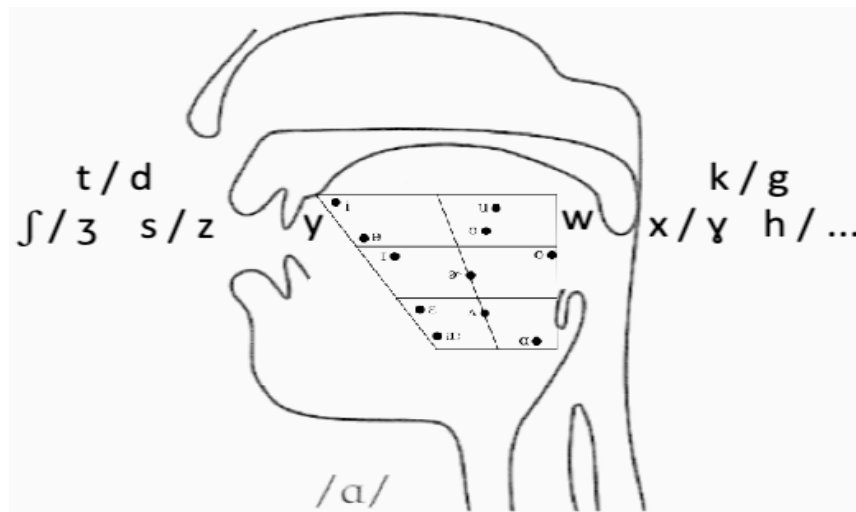
Before any brain processing, sound detection in most organisms relies on a mechanical step, through mechanoreceptors that get “translated” into neuronal firing patterns. These organisms rely on the time-frequency duality in the sound wave of the signal by detecting (i) continuous oscillations (thus focusing on

frequency conditions) or (ii) a sharp sudden onset and rapid offset (thus focusing on temporal conditions). For a click, they primarily encode the onset and temporal pattern of this transient event, with an emphasis on its timing; a tone, instead, is encoded by sustained firing that mirrors the frequency of the oscillation (phase-locking), the emphasis then being on a consistent temporal rhythm. Simple organisms have downstream neurons that detect the presence, amplitude, and timing of the physical vibration, lacking a “frequency-band detector”; for that, a cochlea is needed, as we discuss next.

5. Phonological Correlations

Edison’s analog recording seized a narrow segment of the frequency spectrum. The [condenser microphone](#) and [vacuum amplifiers](#) captured a wider range, converting sound into an electrical signal. These tools allowed engineers to manipulate spectral content, which a cochlea does, too, by separating frequencies along the [basilar membrane](#). The spectrograph, developed for analyzing voice transmissions, decomposed sound waves into frequency components over time – the phonetic structure of speech. Of course, presence of information within a signal is meaningless to an organism lacking the mechanism to extract it; in chordate evolution, the biological receiver of a signal evolved to perform a sophisticated analysis thereof. Cochleated organisms use spectral decomposition to access frequency-based information in parallel, providing a wealth of data about sound quality (fundamental frequency for voice pitch, [formants](#) as resonant frequencies of a vocal tract, etc.). An organism without a cochlea detects fundamental frequency as a vibration pattern, but it cannot simultaneously detect the F1/F2 ratio that defines vowel quality (Figure 4).

Figure 4: Phonemic scaffolding within the vocal tract



A FT analyzes an entire signal over all time to determine its frequency components with infinite precision, but no precision in time – while an animal cochlea must analyze sound *in real-time*. In fact, since t and ω are correlated variables, the standard FT poses just as much of an infinitude issue in the frequency domain as it does in the time domain. The [short-time Fourier transform](#) (STFT, operating over a fixed-size time window) or a [wavelet transform](#) (WT, operating over a variable-size time window) are constrained applications designed to work within physically relevant, band-limited, and time-dependent contexts, like mammal audition. In effect, this imposes a trade-off between t and ω resolution in the cochlea: (i) the basilar membrane is a

traveling wave system analyzing ω within a short, localized t window; (ii) near the [base](#), the cochlea uses t resolution but broader ω tuning at high frequencies; in turn (iii) near the [apex](#), it has sharper ω tuning but coarser t resolution, at low frequencies. These, thus, amount to non-linear filtering characteristics for frequency-dependent resolution.

Although other animals lack speech-specific brain modules, as shown in a literature summarized in ten Cate 2025, some can apply auditory processing capabilities as sketched (to distinguish vowel formants) and temporal resolution (to distinguish rapid consonant onsets). This suggests that the ability to perform the necessary time-frequency analysis is shared across chordates with a specialized inner ear, based on a cochlea or [basilar papilla](#). In actual speech systems, however, stable states are very systematic beyond vocalic situations ascertaining the ω the aspect, and narrower t windows presumed for the punctual consonants: they also include “compromise” states with information about both correlated variables – if more imprecisely, in such a setup. This arguably presumes that relevant variables (engaged in the STFT or WT analogues) stand in an uncertain correlation in the sense of (4)/(5). Pertinent in such terms is how many simultaneous analyses are carried in the [primary auditory cortex](#), at the level of individual neurons (David & Shamma 2013), which affects detailed sound categorization for speech. That seems relevant to pairs as in (2), with various intermediate positions also identified:

- (6) a. Glide/vowels combine into diphthongs: (y, i): [yj], or semivowels (y, i): /ji/
b. Stop/fricatives combine into affricates: (t, s): tʃ, (k, x): kx

While generalizations as in (2)/(6) may correspond to unrelated [allophones](#), it is worth pondering how they may stem from correlated variables (t , ω), with acoustics as discussed (Stevens 1998). This provides a tool for analyzing geometry patterns underlying the feature space in Figure 4, with relevant dimensions allowing for further articulatory or acoustic constraints.

William Idsardi notes through personal communication how, although feasible via a (t , ω) correlation, “inverse affricates” (from a fricative to a stop state) are unattested; [diphthongs](#), in contrast, can be rising (glide first) or falling (glide last). So even if speech arose from an evolutionarily older system (enough for other animals to process distinctive features), humans have evolved specialized nuances beyond acoustic/articulatory correlates, into patterns generalizing in ways that are still being studied. Among human specificities we find the [Perceptual Magnet Effect](#), the highly automatic [McGurk Effect](#) of audiovisual integration, or mechanisms to integrate fragmented acoustic cues spread over time within the signal, as arising in [coarticulation](#). The human cochlea, in sum, breaks the speech signal into distinct frequency channels and the [brainstem](#) and [auditory cortex](#) manage to process parallel frequency streams to detect formant ratios, regardless of speaker pitch (or pitch variations across languages) – with spectro-temporal receptive fields analyzing cochlear outputs as rich as one may imagine (Elhilali et al. 2005).

But while all sorts of spectro-temporal integration are possible, limitations may emerge from the uncertainty relation in (5). For example, these systems may be able to utilize the equivalent of narrowband vs. broadband spectrographs, each more sensitive to the correlated variables; temporal scales for each sort of analysis would then be significantly different. Some of these possibilities are part of what is already established in acoustic, phonetic, and phonological studies. For instance, stop consonants present a closure and release, with the release being practically instantaneous and only visible on a wideband spectrogram; in contrast, narrow band spectrograms are good for slowly changing signals, a precise frequency representation being useful for pitch – formant transitions may be too fast and frequency-indeterminate to track accurately with narrowband spectrograms. The point is: if some of these limitations are purely physical, they ought not to

require duplication *for each of the pairs* in (2) or (6). This is to be kept in mind when we move to conceptual patterns, as we will witness a similar duality: between proto-concepts present in other animals and sophisticated linguistic nuances that may be human-specific.

6. Ballistic and Tetanized Gestures in Vocal Learning

The rare animal trait of [vocal learning](#) involves imitating complex sounds (Lattenkamp & Vernes 2018), with songbirds as a model organism for understanding its mechanisms and evolution (Goller & Shizuka 2018). This is a highly orchestrated process including perception mechanisms, action processes, and feedback loops between the two systems, all worth reflecting upon.

In the primary auditory area (called Field L in birds), a spatial array optimizes some regions for narrow spectral tuning (frequency sensitivity) and others for temporal tuning (rapid changes over time), thus initially separating spectral and temporal information. Higher auditory areas (grouped together in Figure 5 as an auditory module AUD) and [HVC](#) (high vocal center) contain neurons that are highly selective for the bird's own song, responding to specific combinations of spectral and temporal patterns. The auditory system thus performs an initial separation of information that helps manage the time-frequency uncertainty inherent in an FT. Within Field L, some neurons are “tuned” to slow temporal frequencies (longer sounds like vowels, requiring sustained response); others, to fast temporal frequencies (rapid changes like consonant onsets, requiring transient responses). This initial neural separation means the brain has distinct neural “handles” for different types of acoustic information, which act as complex feature detectors, forming the basis of the auditory template used for song learning.

The motor system then mirrors that specialization in its output architecture. The link between perception and action occurs in the specialized [song control system](#). HVC neurons project to the robust nucleus of the arcopallium (RA) and exhibit sparse bursts of spikes at specific times during the song. This internal motor command forms a “time-locked” pattern that persists even if the bird is deafened: a mechanism for the timing of vocal gestures, as RA neurons project to brainstem motor neurons that control the vocal organs ([syrinx](#) and [respiratory muscles](#)). High-level (frequency-rich) features are likely encoded by sustained activity patterns in RA, which lead to [tetanic](#) tonic muscle contractions; rapid (time-sensitive) features are encoded, instead, by the precise timing and [ballistic](#) bursts generated by HVC and executed by RA.

But the precise mapping from perception to action happens through an auditory feedback loop, which reinforces motor gestures that match the desired spectro-temporal features of the song template – the auditory template stored in NCM/HVC containing both spectral and temporal information. The act of vocal learning involves the [anterior forebrain pathway](#) (AFP), which uses auditory feedback to refine the HVC-RA motor commands. When the bird produces a sound, the AFP compares its spectro-temporal features to the memorized template. This feedback system ensures that if, in a juvenile hypothesis, there is a spectral error, motor commands are adjusted, and if there is a temporal error, timing commands in HVC are modified via the AFP loop. This allows a vocal output to match a relevant input through sensorimotor practice, after feedback loops.

Figure 5




Figure 5 [Adapted from Amador et al. 2017] Schematic of songbird’s song system and sensory pathways, comprised of the song motor pathway (SMP, black arrows) and anterior forebrain pathway (AFP, lavender arrows). In turn: [1] Activity originates at **HVC** and projects downstream to the **robust nucleus of the arcopallium** RA, which in turn projects to the **dorsomedial intercollicular nucleus**, DM, in the midbrain and to brainstem nuclei: the **tracheosyringeal portion of the hypoglossal nucleus**, nXIIIts (whose motor neurons innervate the syringeal muscles), the **nucleus Retroambiguus**, RA, and the **nucleus Parambigualis**, PAm, which control expiration/inspiration, respectively. [2] The SMP involves recurrent motor pathway connecting both DM and PAm indirectly to HVC via the **nucleus Uvaeformis**, Uva. [3] The AFP presents an indirect pathway from HVC to RA, crucial for song learning and adult song maintenance. [4] HVC receives auditory information from two additional pathways (brown arrows). In one, auditory information is transmitted through Uva to HVC directly and indirectly via Nif; the other pathway sends the auditory input through the **nucleus Ovoidalis**, Ov, which projects to highly-interconnected nuclei dedicated to auditory processing (**Field L** and the **caudal medial nidopallium**, NCM, represented as “AUD”, which also includes the **caudal mesopallium**, CM). (The figure also abbreviates the **nucleus interfascialis of the nidopallium** as Nif, the **lateral magnocellular nucleus of the anterior nidopallium** as LMAN, the **dorsal lateral nucleus of the medial thalamus** as DLM, and the **caudal mesopallium** as CM.)

But a core puzzle is still how a nervous system converges on such a task, unless sensory and motor systems speak in commensurable “units”. Vocal learning can only work if the animal manages to successfully map spectral (slowly varying) auditory features onto tonic/tetanized motor gestures, and temporal (rapidly varying) auditory features onto phasic/ballistic motor gestures, so that the sensorimotor loop has natural gradients to climb in the feedback process. To understand how convergence becomes possible, beneath neural circuitry one must ponder the biophysical substrate. Interestingly, the vocal-learning trait only seems to

appear in lineages with highly specialized cochleae (or cochlea-like organs) – see Costalunga et al. 2024 for perspective.

The cochlea physically implements a time-frequency decomposition embodying Fourier uncertainty, so that, even before learning, a juvenile receives the song through partitioned (slower or faster) auditory channels. High-frequency sensitivity amounts to short integration windows and poor spectral precision, while low-frequency sensitivity boils down to long integration windows and rich spectral resolution. This gears evolution towards favoring slow channels for [spectral envelopes](#) and fast channels for temporal edges instead. Concomitantly, the motor system comes with broad regimes determined by muscle physiology and [motoneuron](#) time constants: it presents integrative tonic pathways and transient phasic pathways, corresponding to the mechanical possibilities of the syrinx and respiratory system. Evolutionary and developmental mechanisms ([molecular gradients](#), [subtype-specific adhesion molecules](#), etc.) bias sensorimotor connectivity, so that slow sensory channels preferentially connect to slow motor pools and fast sensory channels, to fast motor pools. The system thus begins in a lawful map from the physics of sound to the neurophysiology of behavior. Spectral features recruit sustained/tetanized gestures and temporal features, ballistic gestures, as a physical property of mechanical wave propagation.

Evolution seems to have solved the convergence problem by aligning biophysics on both ends. At a sensory level, the cochlea generates slow/fast spectro-temporal channels via the physics of FT (in some limited variety like STFT or WT); at a motor level, musculature and its premotor pools support slow/tonic vs. fast/phasic regimes, as muscles and motoneurons also obey physical time constants. Developmental genetics then biases slow-to-slow and fast-to-fast linkages, so that the result is for the embodied physics in the cochlea to create an information geometry that motor circuits can evolve to mirror. So long as the lineages where this matching took place were good enough to evolve an initial vocal imitation, specialized nuclei (HVC, RA, etc.) could be invoked thereafter: to evolve on top of this scaffold, thus pushing fidelity as relevant. In sum, the cochlea is the prerequisite enabling the initial step in this convergence possible, with refinements making vocal learning as astonishing as witnessed, for instance, in the [lyrebird](#).

7. Generalizing beyond Vocal Learning and the specificities of FT

While different in neural circuitry and anatomical structures, vertebrates and invertebrates alike rely on a distinction between fast, feed-forward and slower, feedback-dependent movements. The cellular and molecular mechanisms of muscle contraction and synaptic transmission are also conserved (Büschges & Ache 2025). Perceptually, it is not just [echolocation](#) clicks vs. social tones in, say, bats, that matter; equally relevant is the evolutionary response from their prey: [katydids](#) have hearing organs with tonotopically arranged sensory neurons. The female's auditory neurons respond to the narrow-band frequency of her own species' song (in the 20-40 kHz range); but that same auditory system also detects the broader and higher-frequency clicks of predatory bats (a couple of hundred kHz in some bats; see Symes et al. 2020). That said, compare moths that differentiate between the low-intensity (distant) pulsed clicks used by bats for initial detection and the rapid [feeding buzz](#) (close-range) of clicks as a bat closes in for the kill (ter Hofstede & Ratcliffe 2016). Unlike t and f , time and amplitude are not subject to a FT, so it is impossible to transform between independent variables that do not represent the same information (Figure 5).

Figure 6: Frequency determines cycles per time unit, reflecting t dependence for the wave, while amplitude sets its vertical height, showing t independent.

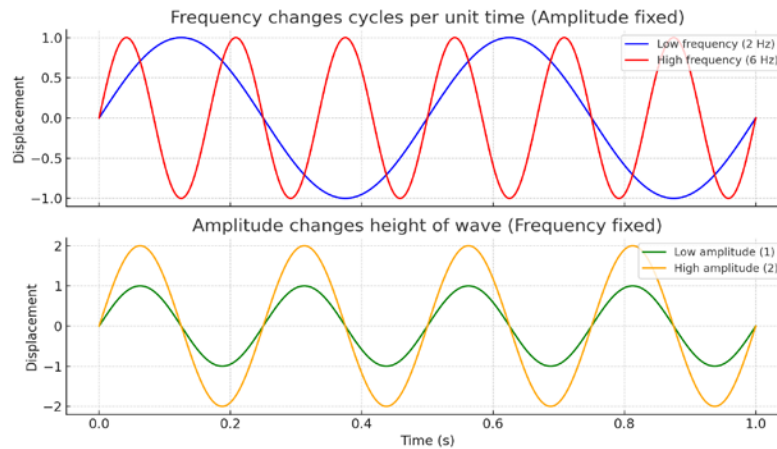


Diagram created with assistance from ChatGPT.

Those comments are meant in two different methodological regards. Positively, now that experimental technology permits, it should be pertinent to study the general dynamics this piece is concerned with in animals for which anatomy is less elaborate and we have less ethical limitations with. At the same time, negatively, we are treading in terms of a mathematical transformation at the core of the UP, which only applies to correlated variables (like time and frequency). Once this clarification is presumed, putatively correlated FT variable pairs may, in fact, shed light on what to look for beyond specific levels of representation (like speech, phonology, other “externalization” mechanisms for language as in signed languages, semantic and pragmatic representations, or whatever else may be relevant to animal cognition).

Those questions presuppose a related one: How does an analog, continuous, spectrally structured physical processes give rise to discrete, categorical, decision-level mental events? Perception is spectral, but acting forces discreteness: a nematode turning left vs. right, approaching vs. withdrawing, crawling vs. swimming, must commit to a categorical output, which requires a decision point so that the [hydrostatic musculature](#) executes one program at a time. Mechanisms for this are known in the *C. elegans* [chemotaxis](#) circuit: gradual increases in a (chemical, thermal, metabolic) attractant entails more runs; gradual decreases, more turns. While the (modality specific) sensory system is continuous, the (modality general) motor output is discrete: the run/turn switch implemented by a mutually inhibitory circuit that performs [winner-take-all dynamics](#).

Among 16 pairs of [amphid](#) sensory neurons in *C. elegans*, several are [chemosensory](#), including the [AWC](#) pair. Left and right AWC neurons begin as symmetric, but during development undergo stochastic lateralization: one becomes AWC(ON), responding to odor decrements; the other, AWC(OFF), responding to odor increments. The two are initially electrically coupled through a [gap-junctional](#) network and spontaneous fluctuations give one neuron a slight activity advantage, triggering a calcium-dependent inhibitory cascade that suppresses the partner. The more active neuron becomes AWC(OFF), while the inhibited one becomes AWC(ON) (Chuang et al. 2007). This ON/OFF opponency increases the system’s coding bandwidth, allowing the animal to detect odor concentrations and compare them bilaterally.

Nematodes learn mapping modifiers in more abstract ways too. For instance, [NaCl](#) is attractive if previously paired with food, but aversive if matching starvation (Jang et al. 2019). Moreover, in patterned [optogenetic](#) stimulations delivered to the AWC neurons, in an environment with no chemical gradient – mimicking how a neuron would fire if the animal followed a chemical gradient – Kocabas et al. 2012 expressed a [channelrhodopsin](#) variant in AWC, shining light on it only when the nematode moves in a direction, then altered light intensity to simulate more or less attractant. Although these gradients exist *only inside the sensory neuron*, not as a chemical gradient in the environment, they still produce correct run-and-turn decisions, just as if the attractant were physically present. These experiments reveal how the worm’s navigational logic is not tied to the absolute chemical concentration: the animal performs an abstract form of chemotaxis in a featureless arena and altering stimulation patterns reverses behavior.

Taking a step back, a spectral-to-categorical evolutionary process, from perceptual to action conditions, requires a pipeline as follows:

- (7) i. A continuous input space with structured statistics and a biophysical filter bank that produces multiscale, band-limited, channels.
- ii. A distributed encoding of channels that can be combined to synthesize or discriminate fine-grained signals, with methods to tune weights based on co-occurrence and predictive success.
- iii. Nonlinear decision dynamics that compress continuous evidence into categorical motor choices, and a loss function to evaluate actions by consequence, for relevant categories to stabilize.

The perceptual magnet effect (Feldman et al. 2009) is a manifestation of this (experience biases the [manifold](#) toward prototypes), with the algorithms generalizing beyond particular sensory inputs or corresponding actions. Practically, this requires a high-dimensional sensory stream $x(t)$, time-scale constraints to implement a multiscale decomposition, and a behavior space with discrete outputs (motor programs) whose success produces an evaluative signal.

Pipelines as in (7) must have happened over the eons for species, an empirical question being how much managed to get compiled and, thus, to be (epi)genetically transmissible. Unlike whatever may be going on within individuals (by adjusting synapses, membrane properties, gene expression programs, etc.), evolutionary change across lineages emerges by selecting variants that bias what an individual is likely to learn. Evolution finds a way to transmit ion-channel repertoires, [transcription factor networks](#), recurrent motifs like [cortical columns](#) or [mushroom bodies](#), [neuromodulatory](#) architectures, and so on: more inductive biases than full-fledged solutions, as seen for vocal learning in section 6. All of this implies hardwired architectures and dynamics.

That said, while the weak emergence implementing analog-to-categorical transition in the perception-action dichotomy may harvest proto-symbolic states from discrete decision-making, it is far from yielding even the proto-logic that a bona-fide computation presumes. Although the form of computation involving categorical phonological objects like the vowels in Figure 4, with elaborate correlations as in (2)/(6) and the like, may be restricted to human speech, it would be surprising if other elaborate animal mental behaviors fail to pose similar challenges. In short, although we have seen a way of compiling weak-emergent spectral-to-categorical structures, nothing in that yields, in particular, compositionality. Therefore, something stronger must emerge within development or evolution – or possibly a different physical substrate becomes necessary.

8. Formal vs. Substantive Realities in different Levels of Representation

The use of substantive distinctive features is common in linguistics, either to ground those on interface conditions or because a general account has not been established yet. Formal distinctive features are usually invoked to encode the structural constraints that enable syntactic processes, particularly if there is computational theory behind them. To some extent, whether a feature is deemed formal or substantive may depend on whether an explanation has been found for how it contributes to the system's internal coherence – or it must be blamed on its interactions. In that regard, though hard to establish, feature-geometries are interesting: first, they presume an abstract reality independent of external grounding; second, once a geometry is empirically settled on, it is possible to find an algebra for which that geometry obtains.

Feature theory in phonology started in Trubetzkoy (1939). Distinctive features were seen as articulatory descriptors in Jakobson, Fant & Halle 1952, thus assumed to correspond to phonetic substance. After Chomsky 1965 – chapter 1 of which raises the formal/substantive distinction – features were taken to be psychologically real. By Clements & Hume 1995, much of the structure in phonology was being analyzed as formal, hierarchy then seen as algebraic constraints on feature interactions. Hale & Reiss 2008 postulate that all genuine phonological features are formal, at which point bona-fide feature geometry stems from a feature algebra. The theoretical shift is unsurprising: early geometry was tied to land measurement, although Euclid attempted a formal systematization by way of his [postulates](#). Gauss showed how the [parallel postulate](#) is hard to justify, which through Descartes's work on coordinate systems led to [Hilbert's axiomatization](#). Then geometry became formal and, eventually, linear algebra generalized it into vector spaces. By non-Euclidean geometry discarding the last “substantive” postulate, the field moved in an abstract direction where it became a collection of formal constraints on a model.

What is “substantive” may depend on the interface; what is “formal”, on the generative core. In phonology, where distinctive features were first postulated, the migration from substance to form has been radical. For other levels of representation, discussion continues. Consider, in this regard, correlated formal variables as [canonical conjugates](#) of one another. This is straightforward in standard spectral conditions, but these are far from obvious in human language (beyond acoustics). It must be kept in mind, however, that, as such, Heisenberg's UP holds of correlated variables – leading to [complementarity](#) – as an algebraic reality *with no substantive commitment*: if we find any such pair, the same formal result obtains regardless of the substance of the interaction: whether acoustic, visual, or even more abstract ones, as we consider next.

If a formal regularity, what holds for consonantal or syllabic dimensions should also hold for syntactic or other dimensions, as expressed in generalization (1) in section 2. A classification as in (8i), via the systematization in Chomsky & Halle 1968, allows us to sort out syllabic consonants (8a) vs. more familiar non-syllabic ones (8b) or, among the vocalic phonemes, those that appear in syllabic nuclei (8c) from those like glides that do not (8d). But, then, in the formal terms already discussed, it is somewhat reasonable to express the still substantive (8ii) in the algebraic fashion in (8iii), which Martin et al. 2019 deploy for their “fundamental assumption”. Algebraic features like those in (8iii) allow us to plug relevant distinctions – again, if of a spectral sort, given what we have been discussing – into the Fourier formulas in (3a)/(3b).

- (8) i. a. [+CONSONANTAL, +SYLLABIC] b. [+CONSONANTAL, -SYLLABIC] c. [-CONSONANTAL, +SYLLABIC] d. [-CONSONANTAL, SYLLABIC]
ii. [±CONSONANTAL, ±SYLLABIC]
iii. [±1, ±i]

Euler's formula (9a), to express complex [exponentiation](#) through [rotation](#) in the [complex plane](#), shows how multiplication by $e^{i\theta}$ corresponds a counterclockwise turn by θ [radians](#) around the unit circle (as seen already in Figure 3), which for $\theta = \pi$ yields [Euler's identity](#) (9b):

$$(9) \text{ a. } e^{i\theta} = \cos \theta + i \sin \theta \qquad \text{b. } e^{i\pi} = \cos \pi + i \sin \pi = -1 + i(0) = -1$$

The [exponential function](#) e^x has the fundamental property of being its own [derivative](#), a self-replicating behavior that extends to complex numbers by translating into circular motion: a deep connection between [exponential growth](#), rotation, and [oscillation](#) (as fundamental to [signal processing](#)). Deriving Euler's formula $e^{-2\pi i \omega t}$ for time t in the $g(t)$ integral (between times t_1 and t_2), or for frequency ω in the inverse $\hat{g}(\omega)$ integral (between frequencies ω_1 and ω_2) has the (uncertainty) consequences discussed for pairs as in (2)/(6) or any other such correlations; again, so long as they are of a spectral sort, as is the case for signals – therein the rub, though. In what sense is it legitimate to speak of spectral conditions beyond acoustics?

Phonological theory is an example of a scientific domain that transitioned from substantive to formal primitives. This shift is also an evolution in explanatory adequacy and a way to illustrate the difference between a substance-driven weak emergence and stronger formal emergence that, in being structure-driven, is constraint-determined. The weak emergence already discussed only gives us a kind of quantization, through action constraints from stable attractors in sensorimotor maps and concomitant categorical perception. It hardly achieves any significant logic – not even negation. Useful though a taxonomy of oppositions is, as a structure of perceptual contrasts grounded in continuous physical substrate, it simply yields no operations over symbols. The question is whether the generalized version of the UP can help in this.

9. A Formal (stronger) Emergence in Historical Perspective

[Heisenberg](#) did not come up with the UP when using only observables to represent them as arrays indexed by transitions (instead of continuous variables over hidden trajectories), stipulating how multiplying such arrays is non-commutative, to prevent the simultaneous specification of observables to arbitrary precision. It was [Born](#) and [Jordan](#) who realized the arrays could be seen as matrices with non-commutative products, establishing the [canonical commutation](#) in (10a) for particle position (x) and momentum (p) at the [Planck scale](#) (here, \hbar). A definite value for an observable means the system's fundamental state (vector $|\psi\rangle$ in a Hilbert space, an [eigenstate](#)) is left unchanged when acted upon by a particular operator \hat{A} for the observable in point.

$$(10) \text{ a. } px - xp = i\hbar$$

b. more generally, $[\hat{A}, \hat{B}] = i\hbar I$ (for I the identity matrix)


c. If $[\hat{A}, \hat{B}]|\psi\rangle = 0$, then they do not admit a complete basis of common eigenstates.

Born and Jordan may not have fully grasped this generalization either, until the algebra for non-commuting operators became clear with [von Neumann](#)'s Hilbert space formalism, but their (10a) is one example of a general [non-commutative](#) algebra of observables.

In classical physics, quantities like x and p are numbers or functions, where multiplication order is irrelevant ($px = xp$); but at the Planck scale observables correspond to linear operators ([self-adjoint matrices](#)) on a

Hilbert space. The fact that their commutator is a non-zero value $i\hbar$ implies that the order in which the operators are applied (or measurements are performed) fundamentally changes the result. So: the dynamical laws as observed at this scale are best expressed as algebraic relations among non-commuting objects, with Born's statistical interpretation of the wavefunction providing probabilistic meaning to non-commutativity (the quantum wave Ψ is a complexed-value probability amplitude): $|\Psi(x, t)|^2$ represents the [probability density](#) of finding a given particle at a specific, observable, position and time (Figure 6). Then non-commuting observables cannot have simultaneously sharp values. In general, the set of all possible eigenstates for an operator entails that any arbitrary state of the system can be expressed as a [superposition](#) (sum) of these eigenstates (forms a [complete basis set](#) in the Hilbert space).

Figure 7: (From the College Sidekick Study Guide's introduction to chemistry, the [De Broglie wavelength propagation](#). Propagation of de Broglie waves in 1 dimension (the real part of the complex amplitude is blue and the imaginary part is green; top: plane wave, bottom: wave packet.). The probability (shown as the color opacity) of finding the particle at a given point x is spread out like a waveform, with no definite position of the particle. As the amplitude increases above zero the curvature decreases, so the amplitude decreases again, and vice versa – the result is an alternating amplitude: a wave.



If $|\psi\rangle$ could exist as a joint eigenstate as in (11a) below, then (11b) should obtain; but applying the commutator to $|\psi\rangle$ yields (11c), which makes sense only under conditions as in (11d). Now, the zero-state vector is not a sensible physical condition; nor is it to equate \hbar to zero, as we can positively measure that constant – in fact, were Planck's constant zero, the result would be a classical universe, with a commutative algebra for all observables.

- (11) a. $A|\psi\rangle = a|\psi\rangle, B|\psi\rangle = b|\psi\rangle$
 b. $[A, B]|\psi\rangle = (AB - BA)|\psi\rangle = (ab - ba)|\psi\rangle = 0$
 c. $i\hbar|\psi\rangle \neq 0$
 d. Sensible only with the zero-state vector or with $\hbar = 0$

So: in a universe with a measurable Planck constant, in conditions as discussed, for any pair whose commutator is a non-zero multiple of the identity, no mutual eigenstates exist.

Pauli's introduction of [spin matrices](#) was the earliest use of finite-dimensional non-commutative operators in quantum theory, with the algebra in (12a). The indexes in that expression range over coordinate labels $\{x, y, z\}$, as relevant to the matrices $\sigma_x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, $\sigma_y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}$, and $\sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$. In turn, ϵ_{ijk} is the [Levi-Civita symbol](#), which in these conditions yields only values as follows: when any particular index repeats, $\epsilon_{ijk} = 0$; if all three indices are unique and in alphabetical order when read cyclically, $\epsilon_{ijk} = 1$; and if all three indices are unique but in reverse alphabetical order, $\epsilon_{ijk} = -1$. The final value of the commutator is $2i$ times the value of ϵ_{ijk} multiplied by whatever happens to be the Pauli matrix *not* used in the commutator $[\sigma_i, \sigma_j]$. Of all $3^3 (= 27)$ combinations arising from the possible permutations of the three indices (i, j, k) in (12a), 21 have at least one repeated index and yield $\epsilon_{ijk} = 0$; three are cyclic permutations of $\sigma_x, \sigma_y,$ and $\sigma_z,$ yielding $\epsilon_{ijk} = 1$; and the remaining three are anti-cyclic permutations yielding $\epsilon_{ijk} = -1$.

$$(12) \text{ a. } [\sigma_i, \sigma_j] = 2i \epsilon_{ijk} \sigma_k$$

$$\text{ c. Spin operator } S_i := \frac{\hbar}{2} \sigma_i$$

$$\text{ b. } \Delta S_x \Delta S_y \geq \frac{\hbar}{2} | \langle S_z \rangle |$$

Defining [spin operators](#) as in (12b) implies uncertainty relations for spin components (12c). That has little to do with Heisenberg's initial use of uncertainty: it arises from operator structure in a finite-dimensional Hilbert space. Pauli's insight was that, rather than about measurement disturbance in FT conditions, uncertainty is a structural consequence of [non-abelian symmetry](#).

That was the seed of a generalized uncertainty principle (GUP), later stated by [Schrödinger](#):

$$(13) \Delta A \Delta B \geq \frac{1}{2} | \langle [A, B] \rangle |$$

With a GUP whose structure is linear-algebraic, Hilbert space notions ([projectors](#), [orthogonality](#), [tensor-product compositionality](#), [gate-like transformations](#)) become pertinent, with much relevance to symbol manipulations: from negation to conditionals, through recursion leading to inference, all of which can be defined under such conditions. This is, then, a way for a theory (regardless of its subject matter) to transition into presenting binary distinctive features with privative/formal specifications in the form of feature matrices within geometrical hierarchies and, if generative rules are specified, a full-blown formal apparatus. In Schrödinger's GUP statement about operators in a Hilbert space, if these lack a constant with [dimensions of action](#), nothing in the inequality requires Planck's \hbar to appear. This is consistent, for instance, with the Pauli matrices expressed in the relations as in (12a), with no \hbar , although the spin operators as in (12b) may have it – so \hbar enters if the operators are defined as physical generators with units of action for our universe. Another way of saying this is that the GUP is “quantum” not in the sense of microscopic physics, but in *belonging to a non-commutative operator algebra* that quantum theory uses.


Several researchers are exploring non-commutative structures in language and cognition. Already mentioned in section 2 was Smolensky's work on tensor-product variable binding, the first to show how symbolic structures can be embedded in high-dimensional vector spaces. Phillips & Wilson 2010 exposed how connectionist models need tensor products to express compositionality. Within so-called [quantum cognition](#), Busemeyer & Bruza 2012 address situations where answering one question influences a response to the

next, developing quantum-inspired models of inference using Hilbert-space geometries. Within [decision theory](#), Pothos & Busemeyer 2009 have quantum probability models account for order effects and inferential reasoning. Coecke et al. 2010 use categorical tools from quantum theory to model meaning compositionality. In the context of a [conceptual-spaces](#) framework, Gärdenfors and Zenker 2023 offer geometric refinements directly relevant to Hilbert-like representational schemes. From a more mathematical viewpoint, Marcolli & Port 2015 show how certain classes of graph grammars can associate to a Lie algebra, whose structure is reminiscent of the [Lie algebra extensions](#) of [quantum field theory](#). These efforts all involve the mathematical architecture associated with quantum theory in studying representation, inference, and compositional structure in cognition.

10. Pauli Matrices in Syntactic Structure?

Just as it is not capricious to interpret (8i, ii) in algebraic fashion (8iii), one can treat Chomsky's 1974 foundational pair $[\pm N, \pm V]$, from section 2, in that algebraic fashion. Consider, for illustration, a graphic instantiation of syntactic operations with formal objects as in (14):

(14) a. N: $[1, -i]$, b. V: $[-1, i]$, c. A: $[1, i]$, d. P: $[-1, -i]$



Michael Jarret proposed (15) to us via personal communication, presuming a scalar, multiplicative, basis for [labeling](#) the syntactic merger of a categorial nucleus (a so-called head) and its syntactic dependent. In this form, the [multiplication is entry-wise](#). The syntactic merge relation is taken to be [anti-symmetrical](#), and thus possible in [reflexive](#) conditions (a head with itself). Moreover, it is postulated that nouns self-merge (Guimaraes 2000), so the product of Chomsky's (14a) by itself is projection $[1, -1]$. Also, it is assumed that syntactic labels correspond to the multiplication of relevant entries, producing "twin" categories like $[1, -1]$ and $[-1, 1]$ with identical labels. The same is true for all other products, yielding projections with labels $-1, 1, i$, and $-i$, empirically determined to correspond to NP, AP, VP, and PP, respectively, because of generalizations as follows. While (16a) is assumed as a *cognitive postulate*, the rest of the "selection" statements follow from the formal system, presuming Chomsky's (1974) substantive assumptions.

- (16) a. Only nouns self-merge (by hypothesis).
 b. Only nouns and adjectives take PPs as complements (here $-i$ categories).
 c. Only verbs and pre/postpositions take NPs as complement (-1 categories).
 d. Only NPs and PPs enter tail recursion (*pictures of enemies of war of...*).

We can then up the formal ante, with Chomsky's features as matrix diagonals. In this instantiation, projection labels for each category in (17) (with subindices as in (18)) are matrix [determinants](#); presumed products, [matrix multiplications](#); and the product results in (18) – starting with the self-multiplication of (17b) – yield an [Abelian group](#) for the items in Jarret's recursive directed graph in (15), in matrix guise, half of whose elements (**boldfaced**) are in the [Pauli group](#). Systemic consequences ensue of presuming that algebraic base. Off the bat, one need not stipulate the "selectional" restrictions in (16): they are a consequence of the

algebra, once (16a) is presumed as an anchoring statement. That is the foundation axiom (mapping to semantic representations, effectively as a [meaning postulate](#)). Human languages present richer dependencies, for instance exhibiting verbal [complementation](#) beyond (16c), [bi-clausal dependencies](#) of various sorts, and also so-called [functional categories](#) of an abstract sort; but Orús et al. 2017 show how to extend the group in (18) to its mirror image by multiplying its elements by Pauli's matrix $\sigma_x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, which results in another group for matrix multiplication including 32 elements from all possible 2x2 [diagonal and off-diagonal matrices](#) with entries corresponding to the four projection labels $\pm 1, \pm i$. This group, which Orús et al. 2017 call the Chomsky/Pauli group G_{CP} , properly contains the Pauli group (as relevant to [quantum information theory](#), see Rieffel & Polak 2011).

$$(17) \text{ a. V: } \begin{bmatrix} -1 & 0 \\ 0 & +i \end{bmatrix} \quad \text{b. N: } \begin{bmatrix} 1 & 0 \\ 0 & -i \end{bmatrix} \quad \text{c. P: } \begin{bmatrix} -1 & 0 \\ 0 & -i \end{bmatrix} \quad \text{d. A: } \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}$$

$$(18) \text{ a. } \begin{bmatrix} -1 & 0 \\ 0 & -i \end{bmatrix} / \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}, \text{ b. } \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} / \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}, \text{ c. } \begin{bmatrix} -1 & 0 \\ 0 & i \end{bmatrix} / \begin{bmatrix} 1 & 0 \\ 0 & -i \end{bmatrix}, \text{ d. } \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix} / \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

Orús et al. also show how the Pauli group can be expressed via the Chomsky matrices in (17). With the elements of such a group, it is simple to define a Hilbert space. This has rich consequences for [long-range correlations](#) in grammar that we will not be reviewing in this context, which were central in motivating generative grammar. While the approach is ultimately very related both to Smolensky & Legendre 2005 and the [DisCoCat model](#) in Coecke et al. 2010, the present system stems from foundational generative assumptions, as discussed in section 8, and selection restrictions as presumed in (16), which neither of those models is concerned with. Note, also, that if one is going to articulate a Hilbert space around syntactic features, there simply is nothing more basic than involving Chomsky's $[\pm N, \pm V]$ categorial features – and doing so in terms of the most elementary algebraic form of $[\pm 1, \pm i]$, as discussed in (8iii). In that sense, the present group-theoretic approach can be seen as coming from a different, more syntactically oriented, tradition, although it is of course encouraging that models stemming from radically different presuppositions may end up in comparable formal conclusions.

11. The CTM v. Connectionism Bout Redux

The real numbers form a [field](#) \mathbb{R} with properties for addition and multiplication: [commutativity](#), [associativity](#), [distributivity](#), [closure](#), plus the presence of [identity](#) and [inverse](#) elements. These obtain in the extension to the complex field \mathbb{C} from \mathbb{R} , which maintains foundational conditions while adding new structure through conjugation: the [complex conjugate](#) introduces a symmetry that allows us to control over magnitude and phase as needed for [wave representation](#); also, in \mathbb{C} , instead of relying on [linear scaling](#), rotations arise and magnitudes that combine multiplicatively. Interestingly, a characteristic is lost from \mathbb{R} to \mathbb{C} : there is, at that level, no natural ordering as in the reals. The question is whether such a representational expansion makes the move a compelling candidate for extending the foundation of connectionist models – in a direction the CTM demands – without sacrificing the stability that scalar algebra provides. There is a sense in which the [imaginary unit \$i\$](#) (which arises from algebraic manipulation of objects in \mathbb{R} , but is not in \mathbb{R} , instead constituting a gateway to \mathbb{C}) is effectively “a new symbol” that emerges from resources in \mathbb{R} , without having to postulate it ex nihilo. Presuming the reality of Chomsky's $[\pm N, \pm V]$ is, in that regard, quite different from presuming,

instead, the existence of $[\pm 1, \pm i]$: while $\pm i$ emerges from algebraically manipulating ± 1 , there is no sense in which the substantive $\pm V$ stems from $\pm N$.


Invoking i thus amounts to a qualitative shift that introduces a distinct symbolic dimension into an otherwise scalar system. This move arises when the logic of the algebra itself is pushed, instead of inserting extraneous symbolic elements of a substantive sort. That said, i does act like a “symbol”. First, qualitatively: Elements in \mathbb{R} differ only by magnitude, as they exist on a continuum; but i is orthogonal to that entire array, its introduction signaling a new dimension of representation. Second, i is an operator in its own right; when applied to a real, it rotates it into a new representational space. Third, a duality persists: while i is tightly linked to the real system via its defining relation ($i^2 = -1$), it itself never reduces to a real number. Finally, the element implies inferential structure in its very existence: since i and real-valued coefficients combine, one thus gains the capacity to encode what amounts to *layers* of structure, as central to the CTM.

Computationalists may protest that i is a poor sort of symbol and the hierarchical structure it introduces, too meagre to run the calculations to recognize, say, that the car is not parked where expected! But [Hamilton](#) showed how there is plenty of algebraic structure beyond the complex plane at higher-dimensional transformations, where multiple simultaneous rotations occur. His breakthrough was to introduce [quaternionic](#) units j, k alongside i , each rotationally defined:

$$(19) i^2 = j^2 = k^2 = ijk = -1$$

These j and k are neither real nor imaginary, but qualitatively different: just like i introduced a new “symbolic” element when seeking $\sqrt{-1}$, so too quaternions introduce two more such elements when we push the algebra to accommodate their dynamics (Figure 7). So: the algebraic path iteratively compels us to introduce qualitatively new symbols to resolve its own formal demands.

Figure 8: From the Wikipedia entry on quaternions, the first image is the [Cayley Q8 graph](#) showing multiplication cycles by i, j and k . In the second, red arrows represent multiplication by i , and the green arrows represent multiplication by j .



Just as we paused to examine what is preserved from \mathbb{R} to \mathbb{C} ; so too must we ask what happens when we go from complex numbers to Hamilton’s quaternions. Given (19), associativity and distributivity are preserved for addition and multiplication in the [hyperplane](#) where the quaternions live, and they still have clear additive and multiplicative identity elements, as well as inverses for the nonzero quaternions. However, commutativity breaks down, a major structural change in the algebraic structure that reflects a profound shift in how the system encodes information. While in \mathbb{R} and \mathbb{C} multiplication is fundamentally scaling and rotation in a fixed plane – inherently commutative processes – for quaternions multiplication tracks sequences of rotations in

three-dimensional space, where the order of rotation matters (Figure 7). In addition, note that, by proceeding into a hyperplane, the inferential structure implied by the very existence of i in \mathbb{C} , from manipulating elements in \mathbb{R} , indicates that, with effectively generating j and k from manipulations in \mathbb{C} , we immediately obtain two new layers of inferential structure.

Quaternions are isomorphic with matrices as presumed in the Jarret graph in (15) (Zee 2016): i, j and k , behave like the Pauli matrices. (Keep in mind that, by postulate (16a) and definition (16b), Pauli's σ_z is the starting point of an algorithmic instantiation of the Jarret graph.) The Pauli matrices satisfy the same (non-commutative) multiplication rules as the quaternion units:

$$(20) \sigma_x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \sigma_y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

$$(21) \sigma_x \sigma_y = i \sigma_z, \sigma_y \sigma_z = i \sigma_x, \sigma_z \sigma_x = i \sigma_y$$

If we identify i, j, k with $i \sigma_x, i \sigma_y, i \sigma_z$, we recover the quaternion algebra in (19) (Figure 8). Just like quaternion conjugation, conjugation by a Pauli matrix corresponds to a transformation, and any quaternionic transformation can be written as a 2×2 complex matrix (an element of the [special unitary group](#) $SU(2)$). So: quaternions behave like rotation operators in three dimensions, just like Pauli matrices generate rotations, and conjugation by i in the quaternion group is equivalent to rotating around the x-axis in the Pauli matrix representation; the “instability” of j and k under conjugation mirrors how Pauli matrices transform under similarity operations (see (12), section 9).

Figure 9. Quaternion units (i, j, k) correspond to Pauli matrices ($\sigma_x, \sigma_y, \sigma_z$) up to factors of i , reflecting an isomorphic algebraic structure and shared role in representing rotations within a Hilbert space.




Diagram generated with assistance from ChatGPT.

The 32 element G_{CP} (which properly contains the 16 element Pauli group and also the Abelian group conforming the Jarret graph) exhibits the hierarchical algebraic constraints we track within the quaternion group: some transformations preserve structure (those with real determinants), others induce sign flips or rotations in higher-dimensional spaces. The present linguistic project seeks to postulate “parts of speech” (extended from substantive elements like nouns, verbs, and so on, to functional elements like auxiliaries, determiners, and the like) from the elements of the G_{CP} . Phrasal structure is projected in relation to “twin”

categories emerging from interactions as in the Jarret graphs and extensions to other matrices within the G_{CP} . Whatever the ultimate algebraic nature of the G_{CP} s, and the Hilbert spaces that can be constructed from its interactions, it preserves the implicational structure in the quaternion group. This is important in terms of giving the CTM proponents another inferential layer, which – once again – is not based on substantive external systems (for semantics or even pragmatics), but is generated in formal terms. Not only does it, then, come “for free” within the algebra, but it is universally inviolable.

Since the number of algebraic properties, as we understand them, is known to be finite, and some fall apart as we march formally into higher formulations, at some point we no longer even have an algebraic system by proceeding in the fashion outlined. Full algebraic stability is present at \mathbb{R} ; at \mathbb{C} we lose ordering, even if we retained all other field properties; by the time the quaternions enter the picture, we lose commutativity; we obtain [octonions](#) from the quaternions but losing general associativity – and if the system takes us into [sedenions](#) in like fashion, zero divisors appear, yielding uninterpretable results. In sum: as new dimensions arise, the algebraic scaffolding weakens. A construction that progressively discards algebraic properties eventually reduces the formal system to something that no longer behaves like a familiar algebra.

The point of the present exercise has been to argue, first, that there is use in the mental system for elements like i , which presuppose a qualitative change in formal tools. Second, we made representational use of the extra formal support, so that $\pm i$ becomes foundational. Moreover, it was argued that the algebraic conditions that take us to an element like the imaginary i (seeking roots for all elements in the presumed algebraic space) yield elements like j or k into the higher dimensional quaternion hyperspace. If, again, we allow for the representational use of the $\pm j$, $\pm k$ elements – as the Jarret graph in (11) presumes – we will be engaging an elaborate implicational structure for logical inference. The quaternion group is isomorphic to the Pauli group, which is known to carry the quantum gates. $SU(2)$, a Lie group, is the geometric object of interest (the group of relevance to our operations), the vector space tangent to which at the identity element is the [Lie algebra \$su\(2\)\$](#) , which describes the group’s local dynamics. Consider next why this matters.

12. What should be Expected in an Animal LoT?

When pondering a putative LoT for a bilaterian, one must of course go beyond signals in any external sense, by the nature of “thought”. Then again, although the study of signals is the impetus behind much of what we have been analyzing, formally there is no need to tie oneself to such strictures, in two different ways – one more radical than the other.

Mathematically, “spectral” can mean eigenmodes of any linear operator associated with an input space. One could then utilize eigenfunctions of the sensory manifold, whose eigenmodes are the natural “basis functions” for an agent’s data. Probability distributions live naturally on manifolds, with characteristic functions as FTs. Beyond a cochlea producing time-frequency channels, for instance, animals may present multiscale sensory apparatus to induce a spectral decomposition, via the eigenmodes of the sensory covariance. When those eigenmodes align with motor primitives (muscle biomechanics, central pattern generators, etc.), the same kind of pre-alignments we considered for vocal learning could obtain in different conditions. The mathematics generalizes: replace “time-frequency” by “data-manifold-spectrum” and the same technical tools are relevant. That would be as weak an emergence as we saw, with whatever properties vocal learning may exhibit under that analysis. Moreover, if actual nematodes, moths, octopuses, chordates,

or whatever, tend to encounter similar situations as they hatch, swim around, seek food, escape predators, attract mates, and all that, it is even possible that genuine spectral conditions in terms of the relevant frequencies arise for such circumstances, correlating with proto-concepts that may stabilize under such circumstances. At present, this is not really known.

In the paleolithic conditions for language evolution, babies in human communities would have faced structured daily and seasonal routines closer to those of higher apes than what we expect from present societies. Well-known developmental milestones (Bavin & Naigles 2015) make the exploration of a time-frequency framing worth testing, even at present. From the perspective of corpora – kept only since we have written records, within historical times – language occurrences may appear haphazard or cluster into paradigmatic uses that seem far from systematic, productive, or transparent, let alone periodic. In such terms, any FA of bursty data we may obtain will produce either broad spectra or misleading sidebands (recall Figure 5 for amplitude vs. frequency). But this is just an artifact of our data gathering. To be sure, it will be impossible to approximate paleolithic conditions for our ancestors, but at the level of granularity we are now considering, even conclusions reached for other higher apes (or any animal we can study) would be significant in terms of the emergence of proto-concepts. These may end up being every bit as specific or general, but at least analyzable, as the nuances of vocal learning in different species.

However, although their descriptive power arose from Fourier dualities and the fact that measurements ride on uncertainty trade-off, uncertainty consequences in Hilbert spaces are of a more architectural sort. It is this aspect that takes us into a putative strong emergence, to go from proto-concepts, whatever those may be, to at least a *proto-logical* apparatus to sanction inference – the real bite of LoT. Stemming from that spirit within the generative enterprise, our approach thus takes it to be a key how an animal brain may carry the sort of information that permits logical inference. Therefore, it asks whether it is efficient (for nature) to get categorical compositionality as a testable (naturally evolvable and developmentally acquirable) way for concepts to interact. It is for that reason that our system bets on invoking a wave-dual interpretation involving time and frequency, first, as in the weak-emergence examples discussed for acoustics. Although, at some point, a stronger form of emergence is needed, both to go into full-fledge phonology in the case of human language and, more generally, for any animal that presents a serious LoT.

Theoretical decisions in this regard are more than just philosophical. If distributional vectors are treated as state vectors in a Hilbert space, for instance, the amplitude dimension corresponds naturally to association strength (its geometric/probabilistic aspect). But recovering full spectral power, one could identify complementary observables whose bases are related by FTs. The interpretation that the present approach suggests to the human language faculty is, in that sense, ontologically radical: by seeking Fourier duals consistent with the presumed lexical bases of syntactic knowledge on one hand and, one hopes, the full-fledged generative syntax on the other. In short, those will have to pair quantities analogous to time and frequency and, also, quantities analogous to position and momentum in particle physics. We have already speculated on psycholinguistic/developmental correlates of time/frequency in LoT (abstractly comparable to energy in particle physics), for acquirable (proto-)concept representations. In turn, correlates of position/wavelength (abstractly comparable to momentum?) ought to model syntax.

Although we will not develop this idea here, the motivation for the present approach, as sketched in Martin & Uriagereka 2014, is that there is, indeed, meaning to correlating syntactic position (in a [phrase marker](#)) with displacement in the sense of generative grammar. In the vector spaces this program sanctions, an observable slot in a syntactic configuration is represented in a positional basis, while the displacement correlate is represented, instead, in a basis associated to a functional extension of the Abelian group as in

(15), in section 10, to matrices in the G_{CP} with entries “swapped around” both horizontally and vertically (to mirrored anti-diagonal matrices). In that extension, so-called “chains” are state superpositions arising from tensor products built from the substantive vocabulary in (15), which upon interpretation at the interfaces of Phonetic Form (PF) and Logical Form (LF) collapse into observable configurations, an analogue of [Born’s rule](#). This is conceptually related to the “trees” in Coecke et al.’s 2010 DisCoCat, but the syntax is more consistent with the generative enterprise. The approach allows for modeling fundamental properties of human syntax, like the ability to generate formal occurrences from lexical types, tokenized into observable structures. Current understanding of quantificational logic is impossible without such occurrences, which have otherwise not been successfully modeled. Then again, an interesting question is whether these arise in other animals too, which is also unknown.

13. Towards Some Conclusions

Neuroimaging techniques reveal animal brains thriving in waves; the issue is how their foundational properties relate to information carried at diverse levels of representation. In the case of language, where binding conditions are easier to ascertain, the study of brain rhythms has a long experimental and theoretical pedigree. Interest in oscillations as computational is more recent, as synthesized in Giraud & Poeppel 2012, even more so with extensions to hierarchical structure (Ding et al. 2016, Matchin & Hickok 2020, Murphy 2024). But the sound assumption is that relevant dynamics arise as a function of the parsed unit, syntax presuming underlying phonemes, etc. As one ponders intuitions about semantic notions like the rigidity of names, on one end, and the relative fluidity of predicative descriptions on the other, it is sound to expect a treatment of lexicosemantic notions in terms of representation that could also be described through waves, but experimental results in this direction may take a while. In the terms sketched here, one could pursue them more intrusively with other animals, although it is still hard to pin down what exactly to test – for instance, whether they may understand at least part-whole relations.

Assuming that a CTM need not mean postulating arbitrary symbols ex nihilo, our model “rides the wave” of algebraic structure emerging from connectionist processes, if enriched in ways that we think we need for modeling brains. This took us to a weak emergence, whereby discrete categories are constructed from continuous sensory statistics by biologically plausible mechanisms triggered by action conditions, imposed on the animal by its anatomy and physiology. That, however, takes us only so far: it may yield proto-symbols, but hardly any interactions thereof, as needed for proto-logic. For that, we explored a stronger emergence, stemming from generalizations arising within the language faculty as presently understood. Our move boils down to: a) presuming that the levels of linguistic representation have a more articulated algebraic structure than is reputed to them, and b) that postulated “jumps” in algebraic dimensionality constitute a way to limit mappings between levels, as certain homomorphisms – which preserve algebraic structure – are more reasonable than alternatives. The mathematics thus generalizes: from spectral nuances arising for classical waves, as relevant at least to vocal learning, to more elaborate conditions that map the transition in particle physics from Heisenberg’s UP to a GUP, as represented in the Pauli analysis of spin, or further generalizations into quantum field theory, where the algebraic structure forces non-commutativity (Lie algebras with their associated conditions).

Information-processing operations that fail to commute are more common in biological systems than it may seem at first, long before animals evolved cognition. Order-dependent [transcription regulatory networks](#),

Calcium-dependent [synaptic plasticity](#), or [olfactory processing](#) suggest that evolution had the raw material for non-commutative processes that we have extended to animal cognition. We needed this because weak emergence only gives us discontinuities, relevant to proto-symbols as best; but if the LoT hypothesis is even remotely correct, in a physicalist approach its component parts must interact via operations (comparison, negation, combination, etc.) that do not commute (e.g., negating before categorizing \neq categorizing before negating). Given a rich set of non-commuting representational operators, a Hilbert space is a natural mathematical setting for their informational interaction.

Whatever a LoT ultimately is, systematicity, productivity, and even transparency can be observed in creatures as distant from us as cephalopods. If so, it is no stretch to argue that whatever took place in evolution – in that scenario, in pre-Cambrian conditions – managed to summon rather elaborate emergences to manipulate information, classically or otherwise. This is not to say that other animals can perform actual algebraic exercises or engage in generalized quantification: only that they have the mental ingredients to do so, in evolutionary terms. Evidently, if the representational machinery required for a LoT already is computational, the issue then becomes less whether other animals have access to recursive thoughts than whether they can share them with conspecifics. The remarkable human PF (presupposing vocal learning) must have been a piece of the evolutionary puzzle in our lineage, as humans externalize their LoT in a way other animals seem less willing or able to. This has all sorts of consequences for social intelligence, rich culture, and so on, in turn, posing the question of how different LoT is in humans vis-à-vis the faculty of language as we experience it through the LF level of representation.

These are all important questions, but empirical ones at that. Hopeful pre-verbal babies or tragic post-verbal seniors, aside from those suffering from aphasias and even more dramatic conditions, clearly make one separate “language” from “thought”, at least in impressionistic terms. But such categories are pre-theoretical, the nuance being in the details of what representations each of these states presuppose. Our contribution is less to think of that in algebraic terms, which others have suggested already, than to imply that going there need not entail giving up the CTM/LoT edifice. Quite the opposite: if we deepen the presuppositions in such theories in terms of the wealth of knowledge we have from generative grammar, the result is a more elegant version of the interaction between lexicons and grammars, in ways that may be found in other species too, *mutatis mutandis*. The human species seems to have found further mobility within that moving ensemble, exploiting the very dynamicity that makes it possible – hence, the Vernean title of the piece.

MOBILIS IN MOBILI
N

References

- Amador, A., S. Boari, & G. Mindlin. 2017. "From perception to action in songbird production: Dynamics of a whole loop". In *Curr. Op. Sys. Biol.* 3, 30-35
- Bavin, E. & L. Naigles. 2015. *The Cambridge Handbook of Child Language*. Cambridge: Cambridge University Press.
- Büschges, A. & J. Ache. 2025. "Motor control on the move: from insights in insects to general mechanisms." In *Physiological Review* 105-3, pp. 975-1031.
- Busemeyer, J. & P. Bruza. 2012. *Quantum Models of Cognition and Decision*. Cambridge: Cambridge University Press.
- Chomsky, N. 1974. "The Amherst Lectures." University of Paris VII.
- Chomsky, N. & M. Halle. 1968. *The Sound Pattern of English*. New York: Harper & Row.
- Chuang, C.-F., VanHoven, M., Fetter, R., Verselis, V., & C. Bargmann. 2007. "An innexin-dependent cell network establishes left-right neuronal asymmetry in *C. elegans*." *Cell*, 129(4): 787-799.
- Clements, G., & E. Hume. 1995. "The Internal Organisation of Speech Sounds." In *The Handbook of Phonological Theory* (J. Goldsmith ed.), Blackwell.
- Coecke, B., M. Sadrzadeh & S. Clark. 2010. "Mathematical foundations for a compositional distributional model of meaning." In *Linguistic Analysis*. 36-1, 345-384.
- Costalunga, G., D. Vallentin, & J. Benichov. 2024. "A neuroethological view of the multifaceted sensory influences on birdsong." In *Curr. Op. Neurob.* 86. <https://doi.org/10.1016/j.conb.2024.102867>
- David, S. & S. Shamma. 2013. "Integration over Multiple Timescales in Primary Auditory Cortex." in *The Journal of Neuroscience*, 33-49:19154 -19166.
- Ding, N., L. Melloni, H. Zhang, X. Tian, & D. Poeppel. 2016. "Cortical tracking of hierarchical linguistic structures." In *Nature Neuroscience* 19-1, 158-164.
- Elhilali, M., D. Klein, J. Fritz, J. Simon, & S. Shamma. 2005. The enigma of cortical responses: slow yet precise. In *Auditory Signal Processing: Physiology, Psychoacoustics, and Models*. New York: Springer.
- Feldman, N., T. Griffiths, & J. Morgan. 2009. "The Influence of Categories on Perception: Explaining the Perceptual Magnet Effect as Optimal Statistical Inference." In *Psych. Review* 116-4, 752-782.
- Fodor, J. 1975. *The Language of Thought*. Cambridge: Harvard University Press.
- Gallistel, R. & A. King. 2009. *Memory and the Computational Brain: Why Cognitive Science Will Transform Neuroscience*. Hoboken: Blackwell Wiley.
- Gärdenfors, P., & F. Zenker. 2023. *Representing Information: A Conceptual Spaces Approach*. Berlin: Springer.
- Giraud, A., & D. Poeppel. 2012. "Cortical oscillations and speech processing." In *Nat. Neur.* 15-4, 511-517.
- Godfrey-Smith, P. 2016. *Other Minds: The Octopus, the Sea, and the Deep Origins of Consciousness*. New York: Farrar, Straus, & Giroux. Malden, MA: Wiley-Blackwell.
- Goller, M. & D. Shizuka 2018. "Evolutionary origins of vocal mimicry in songbirds." In *Evol. Let.*, 2-4.1, 417-426.
- Guimarães, M. 2000. "In Defense of Vacuous Projections in Bare Phrase Structure," In *UMWPL* 9.
- Hale, M., & C. Reiss. 2008. *The Phonological Enterprise*. Oxford: Oxford University Press.
- Hebb, D. 1949. *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley and Sons.
- ter Hofstede, H. & J. Ratcliffe. 2016. "Evolutionary escalation: the bat-moth arms race Free." In *Jour. Exper. Biol.* 219-11: 1589-1602.
- Idsardi, W. 2022. Underspecification in time. *Canadian Journal of Linguistics*. 67, 1-13.
- van der Hulst, H. Forthcoming. *A Mind for Language: An Introduction to the Innateness Debate*. Cambridge: Cambridge University Press.
- Jang, M. S., Toyoshima, Y., Tomioka, M., Kunitomo, H., & I. Yuichi. 2019. "Multiple sensory neurons mediate starvation-dependent aversive navigation in *Caenorhabditis elegans*." In *PNAS* 116(37):18673-18683

- Jakobson, R., G. Fant, & M. Halle. 1952. *Preliminaries to Speech Analysis: The Distinctive Features and Their Correlates*. Cambridge: MIT Press.
- Kocabas, A., C.-H. Shen, Z. Guo, & S. Ramanathan. 2012. "Controlling interneuron activity in *Caenorhabditis elegans* to evoke chemotactic behavior." *Nature*, 490, 273-277.
- Lattenkamp, E. & S. Vernes. 2018. "Vocal learning: a language-relevant trait in need of a broad cross-species approach." In *Current Opinion in Behavior. Sci.*, 21, 209-215
- Le Du, S., 2022. "A gentle introduction to Quantum Natural Language Processing." In [arXiv:2202.11766v1](https://arxiv.org/abs/2202.11766v1)
- Marcolli, M. & A. Port. 2015. "Graph Grammars, Insertion Lie Algebras, and Quantum Field Theory". In *Mathematics in Computer Science* 9-4, 391-408.
- Marcus, G. 2001. *The Algebraic Mind Integrating Connectionism and Cognitive Science*. Cambridge, MA MIT Press.
- Martin, R. & J. Uriagereka. 2014. Chains in Minimalism. In: Kosta, P., Franks, S. Radeva-Bork, T. & Schürcks, L. (eds.), *Minimalism and Beyond: Radicalizing the interfaces*. Amsterdam: Benjamins.
- Martin, R., R. Orús, & J. Uriagereka. 2019. "Towards Matrix Syntax". *Catalan Jour. Ling.*, Special Issue, 27-44.
- Matchin, W. & G. Hickock. 2020. "The Cortical Organization of Syntax". In *Cer. Cortex* 30-3:1481-1498.
- McCarthy, J. 1988. "Feature Geometry and Dependency: A Review." *Phonetica* 45, 2-4, pp. 84-108.
- Murphy, E. 2024. "ROSE: A neurocomputational architecture for syntax." In *Jour. Neur.* 70, doi: <https://doi.org/10.1016/j.jneuroling.2023.101180>
- Orús, R., R. Martin, & J. Uriagereka. 2017. "Mathematical Foundations of Matrix Syntax." [Arxiv.org/abs/1710.00372](https://arxiv.org/abs/1710.00372).
- Paulk, A. C., Y. Zhou, P. Stratton, L. Liu, & B. Van Swinderen. 2013. "Multichannel brain recordings in behaving *Drosophila* reveal oscillatory activity and local field potential coherence." *Jour. Neur.*, 33-13, 5934-43.
- Phillips, R., J. Kondev, J. Theriot, & H. García. 2009. *Physical Biology of the Cell*. New York: Garland.
- Phillips, S. & W. Wilson, W. 2010. Categorical compositionality: A category theory explanation for symbolic structure in cognition. *PLOS Comp. Biol.*, 6-8, doi: [10.1371/journal.pcbi.1000858](https://doi.org/10.1371/journal.pcbi.1000858)
- Pothos, E., & J. Busemeyer. 2009. "A quantum probability explanation for violations of "rational" decision theory." *Psych. Rev.*, 116-4, 165-185.
- Rieffel, E. & W. Wolfgang Polak. *Quantum Computing: A Gentle Introduction*. Cambridge: The MIT Press.
- Sainsbury, R. 2001. "Two ways to smoke a cigarette." *Ratio* 14: 386-406.
- Stevens, K. 1998. *Acoustic Phonetics*. Cambridge: MIT Press.
- Smolensky, P. 1990. "Tensor Product Variable Binding and the Representation of Symbolic Structures in Connectionist Systems". In *Art. Intell.* 46 (1-2): 159-216.
- Smolensky, P., & G. Legendre, G. (2005). *The Harmonic Mind*. Cambridge, MA: MIT Press.
- Symes, L., S. Martinson, C. Kerman, & H. ter Hofstede. 2020. "Sheep in wolves' clothing: prey rely on proactive defences when predator and non-predator cues are similar." In *Proc. Roy. Soc. B*, 287.
- ten Cate, Carel. 2025. "Animal Analogs of Speech Perception." In *Oxford Research Encyclopedia of Linguistics*, DOI 10.1093/acrefore/9780199384655.001.0001/acrefore-9780199384655-e-1080.
- Trubetzkoy, N. S. 1939. "Grundzüge der Phonologie." *Travaux du Cercle Linguistique de Prague* 7.
- Uriagereka, J. 2023. "Correlated attributes: Toward a labeling algorithm of complementary categorial features." In *Front. Lang. Sci.*, *Neurobiology of Language: Syntax, the brain, and linguistic theory*.
- Wilson, J. 2021. *Metaphysical Emergence*. Oxford: Oxford University Press.
- Zee, A. 2016. *Group Theory in a Nutshell for Physicists*. Princeton: Princeton University Press.