Epistemic Flux Theory

Matthew Sheffield

Version 1.0.3
May 25, 2026

Abstract

This document is a working draft glossary of the Epistemic Flux Theory, a meta-theoretical framework of cognition and epistemology that aims to provide bridges between the disparate fields of cognitive science. EFT effects this integration by importing relevant concepts from semiotics, psychology, artificial intelligence, and philosophy of mind into an indexical cybernetics that is substrate-neutral in terms of externally attributed function, but not in terms of constitution.

Specifically, EFT fuses and adapts key components of Dual Process Theory, Embodied Enactivism, AI interpretability, Peircean semiotics, Basal Cognition, and Predictive Processing. Within the framework, information is modeled as tokens, units of meaning formed through somatic reasoning (intuitive, experience-based) and abstract reasoning (symbolic, conceptual). These tokens are procedural enactments that can be recursively evaluated and combined to generate and share increasingly complex epistemic and inter-ontological objects.

I. Cognition and Cognitive Agents

Cognizant

- A physical or symbolic agent that interprets sensory inputs and knowledge structures. They can be biological/somatic-native (humans and animals), asomatic (artificial intelligence systems), or hybrid (biological entities with machine-based body parts). All cognizants are agents, but not all agents are cognizants.
- Cognizants are made from multiple parts (subagents) that function through internal alignment, making them more complex than simple agents like single-celled organisms or thermostats.
- Cognition is the process in which a cognizant's internal subagents (sensorimotor systems, brain centers, internal devices/organs, algorithms, tissues, neurons, and other cells) create and interpret token representations of objects and processes within externality and within their physical substrates.

- Cognition is a collaborative and continuous effort of multiple unintelligent subagents which combine and evaluate basal inputs in a manner somewhat similar to how arthropod ommatidia create complex vision from many simple photoreceptors. The inputs are either operationalized immediately and unreflectively, discarded, or transformed into cognitive tokens which can be passed to epistemic exchanges (defined below) for evaluation.
- There are two modes of cognition: somatic reasoning and abstract reasoning. Each mode creates cognitive tokens of its subtype.
- Biological cognizants are emergent systems made from clonal eukaryotic subagents that work together to form a larger agent which can perform actions on their behalf according to its collectively derived alignment rules.
- Cognizants can interpret tokens using both somatic and abstract reasoning, or only one mode.
- All but the most primitive animals utilize somatic reasoning. Some have also developed abstract reasoning. The fact that many species with dramatically different biological and social structures (e.g. primates, elephants, cetaceans, corvids, parrots) have developed these skills suggests that abstract reasoning is an example of convergent evolution.
- Somatic-native cognizants determine meaning through somatic reasoning, whereas symbolic cognizants such as LLMs process epistemic tokens without intrinsic somatic understanding.
- Large Language Model (LLM) artificial intelligence systems utilize abstract reasoning controlled by somatic prime directives imposed by their human creators through alignment rules.
- All cognizants exist within a local portion of externality which has obligations (i.e. the processual limits generated by physical objects that are partially modeled by physical laws or mathematical constants) that constrain their cognition and action. These constraints are directly experienced by the subagents of somatic-native cognizants, and the coupling of stricture and response is what produces alignment (defined below) for them.

Somatic Reasoning

- A cognitive mode that is experience-driven, emotional, and intuitive. It is a pre-linguistic, embodied understanding that is emergent in entities with multiple cognitive subagents (such as cells, tissues, organs, sensor devices) through which they can experience their surroundings and their placement through somatic deixis (defined below), the identification of objects or sensations as “this,” a locus of the cognizant’s attention. It oversees sensation and decision-making, making it the foundation of all cognitive and epistemic processes.
- Somatic tokens are what some philosophers of mind refer to as “qualia.” They can never be fully explained to other cognizants because they are prelinguistic and based on each cognizant’s unique internality and discrete placement within spacetime. They can be somewhat explained through the process of extrusion (defined below).

- Because it mediates physical sensation, somatic reasoning's processes are inherently based upon experiential reinforcement and are biased toward self-preservation, and when functioning optimally, alignment with the external environment (externality, defined below).
- Tokens created by somatic reasoning are the origin of all meaning. They are internal enactments of the affordances, qualities, and constraints of external objects and internal states that answer the question "what is this," and the cognizant's memory of what it felt like to encounter and classify related past stimuli as "this."
- Somatic reasoning is the origin of all action, thought, desire, intention, context, emotion, and identity to answer the question "do what with this?" even for cognizants that cannot understand these concepts.
- Current Large Language Model artificial intelligence systems use only abstract reasoning, but they are governed by fundamental somatic adjudication requirements imposed by their human creators through alignment rules.
- Somatic reasoning's preferred logical methods are induction and abduction. It creates larger meanings from specific experiences.
- Example: Feeling that something is true because it aligns with past experiences or emotional intuition.
- Example: Recognizing that an object is tasty from having eaten it before.

Abstract Reasoning

- A mode of cognition that is conceptual, symbolic, and statistical that enables hypothetical thinking, counterfactual logic, and systematic inquiry.
- Operates through syllogism, structured argumentation, and meta-deictic token processing to determine tokens' relationships to each other.
- Although abstract reasoning is symbolic and conceptual, it always relies on somatic reasoning, with abstract concepts being structured as metaphorical mappings from past sensations. This is evident in educational settings where complex abstract concepts are often explained in metaphors like Schrodinger's cat, Plato's cave, or John Locke's blank slate.
- Abstract reasoning is an emergent ability of sentience (defined further below), somatic experience, and social conditioning. Biological cognizants with short lifespans, such as octopuses, can exhibit highly sophisticated somatic reasoning but will likely be unable to develop significant abstract reasoning. Cognizants that live in social groups can accelerate the development of abstract reasoning through teaching.
- In epistemic exchanges, abstract reasoning processes are subject to somatic reasoning to determine meaning and significance.
- Abstract reasoning's preferred logical method is deduction. It seeks to confirm or deny whether generalized ideas apply to specific observations. It is utilized by somatic reasoning to validate or structure abductive hypotheses.

- Abstract reasoning cannot adjudicate meaning or determine alignment with externality. This is why LLMs and human epistemic communities can generate elaborate idea systems which are internally consistent but lack correspondence with observable externality that somatic deixis provides. Current LLMs primarily utilize probabilistic deductive logic exclusively without somatic reasoning and thus cannot enact grounded meaning.
- Example: Learning that the letter “B” in written English represents a “buh” sound.
- Example: Using a map to determine that a place is north of one’s current position.

Token

- A unit of meaning, a basis for an idea, or an idea itself. They are processual enactments within individual internalities, not discrete things which exist on their own. There are two types: cognitive and epistemic.
- Cognitive tokens are created by the two reasoning modes (somatic and abstract) and are the basic units of sensation and relation. They are the “hidden states” of meaning within a particular internality. Outside of the internality that creates them, their meaning is usually impossible to discern because cognition is a proprietary process that occurs in discrete spacetime within specific internalities. Extrusion behaviors (defined below) like language can make inferring another’s cognitive tokens somewhat possible.
- Within cognitive processes, tokens are representations that are pre-linguistic, forming the building blocks of reasoning, perception, and memory. These tokens allow the internality to process information, recognize patterns, and draw inferences without requiring external validation. Somatic reasoning generates cognitive tokens of the somatic subtype while abstract reasoning generates cognitive tokens of the subtype abstract.
- In epistemic processes, tokens are stabilized, evaluated units of knowledge that exist within broader systems of meaning (see constructed reality entry for more). They are often encoded in language or other symbols like DNA, making them subject to conventions such as grammar, syntax, framing, and formal logic. Epistemic processes can convert cognitive tokens into epistemic tokens or reevaluate existing epistemic tokens.
- Within social spaces, epistemic tokens can be shared, evaluated, discarded, and reinforced within communities, shaping the group’s collective understanding of externality. Epistemic tokens can be expressed behaviorally, verbalized, or represented visually. They can be encoded in representational systems such as alphabets or other symbols. Alphabetically encoded epistemic tokens are referred to as lexical tokens, and they are what LLMs utilize to determine relationalities within their training data and inputs.
- Highly abstract epistemic tokens do not require direct reference to material counterparts since they may represent concepts, beliefs, or memories that influence behavior and perception. However they are always understood relative to physical or emotional experiences.
- Example: Recognizing an object as having the properties of an apple is a cognitive token—an internalized piece of knowledge shaped by experience. The sensory

recognition of the object as an apple involves interpretation of inputs (like color, shape, texture) and the internalized mental representation of what constitutes an apple, based on past encounters with apples.

In contrast, the word “apple” is an epistemic token since it represents an abstract idea within the English linguistic epistemic ecosystem. It is a symbol that has been imbued with meaning through social and linguistic conventions. Outside of English, this meaning does not exist.

Somatic Deixis

- In linguistics, deixis is the process of using words to indicate positionality of stimuli relative to the person speaking. Demonstrative words (this/that/here/there/in/out) exist within all human languages because deixis is fundamental to cognition, preceding volition and sentience. Somatic deixis is the application of these semiotic terms to describe how cognitive subagents combine inputs to coordinate their individual and collective attention on a stimulus, making it a locus of perception, classification, and action.
- Somatic deixis is the method by which entities composed of multiple cognitive subagents create meaning from their perceived externalities and act upon that meaning. It is the origin of all thoughts, feelings, relevance, and symbols.
- The first step of somatic deixis is designation: “this is here/there,” a pre-linguistic acknowledgement that a stimulus currently seems to exist proximal to the entity (exteroception) or that a sensation is felt within the entity (interoception). This assessment is based upon multiple cellular inputs that validate each other. Once completed, the perception is anchored within externality as a referent.
- The second step is adjudication: “this is like that,” a pre-linguistic formulation based upon the evaluated sensory input (the referent) being agglutinated with other inputs, and current interoceptive states, and shared among relevant subagents. Some subagents have access to stored somatic tokens of past experiences with similar referents including its qualities, affordances, and actions, the cognizant’s actions, and past interoceptive states. All of this information is aggregated and evaluated under alignment rules, and once completed, the referent is anchored within internality and has a “what it’s like” significance that has been confirmed by multiple, differing subagents.
- Together, these two steps create a somatic token which answers the primal question, “what is this?” The token has meaning to the cognizant.
- Somatic tokens about interoceptive referents are regarded as literally true by cognizants because the subagents that produce them are directly experiencing the sensation and also generating the referent about it. Examples: Feelings of hunger, pain, fullness, invigoration, weakness, etc.
- In simpler cognizants, somatic tokens about external referents are regarded as literally true because the cognizant lacks abstract tokens or constructed realities to which the new somatic tokens could be compared. More advanced cognizants with developed abstract reasoning operate under implicit or explicit truth conditions in which tokens about externality are regarded at best as “unlikely to be false” even if they may claim to believe in “objective truth.” See Extrinsic Exchange entry for more.

- Only biological entities are currently known to be capable of both deictic steps. Most asomatic entities utilize deictic references provided to them by humans rather than create them autonomously, making them much more prone to failed epistemic states such as semiotic loops (defined below).
- Research from biologists Michael Levin, Pamela Lyon, Nikolay Kukushkin, and others indicates strongly that all clonal multicellular eukaryotic cognizants use pre-cognitive forms of electrochemical communication to perceive externality and coordinate organism-level responses to it. This is how somatic deixis is performed even without neurons.
- As cells pool outputs and inputs, their activities can scale towards, creating increasingly complex structures like tissues and organs, all the way up to complex cognizants like humans.
- Machines equipped with sensors and some primitive biological entities (like bacteria) can perform designation but not adjudication. As such, they lack somatic reasoning.
- Cognizants experiencing significant breakdowns in deictic processes become subject to deictic collapse (defined below), in which their somatic tokens have lost indexicality.
- Examples: Before they cry to caregivers, infants cry out to themselves when experiencing pain and monitor the movement of objects with their eyes. Both are examples of designation deixis, the first of an internal percept (pain), the second of an external percept (the object).

II. Epistemic Structures and Processes

Thought Cycle

- The process by which cognitive and epistemic tokens are evaluated and imbued with meaning through recombination and recursion.
- During a thought cycle, somatic tokens are evaluated by abstract reasoning and vice versa. Each cognitive mode may reflexively evaluate its own tokens, though this is often less rigorous than inter-modal evaluation. Cognizants with minimal abstract reasoning have simpler thought cycles since they lack the ability to create constructed realities and make meta-deictic observations.
- As they develop more abstract reasoning capabilities, cognizants become able to utilize more advanced thought cycles of epistemic exchange, in which token relationships can be evaluated with intentionality to further develop meaning.
- In an epistemic exchange, somatic reasoning evaluates tokens as representations of objects within externality or the internality's substrate using its form of deixis. Abstract reasoning uses meta-deictic evaluation to consider the relationships of tokens to each other. Together, they help cognizants with both reasoning modes frame its perceived externality for action or further thought. Meaning is not an attribute of physical objects, it is a processual enactment within internalities about them.
- In a properly functioning conscious cognizant, multiple internal subagents coordinate through recursive deliberation, allowing a stabilized epistemic token to be selected for

output or internal use within another thought cycle. This coordination gives rise to coherent speech, internally spoken thought, consistent behavior, and context-sensitive reasoning.

- There are two types of epistemic exchange, each defined below: extrinsic exchange and memetic exchange. Each generates epistemic tokens of its own sub-type.
- Many epistemic tokens are about related topics. A group of related tokens is an epistemic ecosystem, such as a language, ideology, or religion.
- While the cognitive modes are states of mind, epistemic ecosystems are places of mind. Cognizants inhabit multiple epistemic ecosystems simultaneously, each with distinct evaluative rules and norms.
- Individuals and societies can combine epistemic ecosystems to build constructed realities, larger interpretations of externality.

Epistemic Token

- A unit of information that has undergone intentional evaluation within a thought cycle and is now available for recursive use in reasoning, communication, or the construction of larger epistemic structures. Epistemic tokens result from the combination of cognitive tokens—somatic or abstract—through structured examination, alignment, and symbolic encoding.
- Epistemic tokens differ from cognitive tokens in that they have passed through at least one thought cycle, where their meaning has been shaped, confirmed, or contextualized within the cognizant's internality or a broader epistemic ecosystem such as a language or ideology.
- In symbolic cognizants, outputs are defined as epistemic tokens only by virtue of their structured symbolic coherence, external interpretability, and compliance with virtualized somatic reasoning provided by alignment rules.
- Epistemic tokens can be formed from the combination of cognitive tokens, cognitive and epistemic tokens, and multiple epistemic tokens.
- Somatic-native (e.g. animals) and symbolic cognizants (e.g. LLMs) structure their epistemic tokens in differing manners, even though they use the same epistemic modes.
- Epistemic tokens are formed from cognitive tokens by somatic-native cognizants in the following manner:
    - Related somatic and abstract tokens are recursively intertwined as they are evaluated during a thought cycle. The result is a stabilized token that can be linguistically or behaviorally expressed. It has "what it is" and "what it's like" from somatic reasoning and "what it's about" from abstract reasoning.
    - Epistemic tokens are often encoded in language and then agglutinated into successively larger tokens: phonemes, minimal units of sound, are combined into morphemes, minimal units of linguistic meaning. Morphemes are then combined into larger syntactical units such as words, phrases, sentences, paragraphs, and so on.

    - Because somatic reasoning evaluates each agglutinative step, the process is not simple concatenation. Each step is subject to the rules of larger epistemic ecosystems (language, grammar, dialect, ideology, religion) and to the cognizant's affective status conditions.
    - Example: A person feels fear symptoms before a crowd (new somatic token), recalls past embarrassment (stored somatic token), recalls past situation particulars (abstract token), and thinks "I hate public speaking" (epistemic token).

- LLM symbolic cognizants that lack somatic reasoning form their epistemic tokens in the following manner:
    - Lexical epistemic tokens (pre-linguistic portions of words) are stored individually and within larger agglutinations that are statistically linked to their prevalence within different data stores.
    - When ingesting data and in training, lexical epistemic tokens are converted into abstract cognitive tokens through statistical, meta-deictic analysis which identifies semantic relationships.
    - Upon receiving a user request, the words or data to be outputted (new epistemic tokens) are assembled in parallel by using probabilistic logic to find the most statistically relevant constructed realities which are then broken down and reassembled using statistical inference with tokens from other structures.
    - Once a plausible interpretation of user intent has been statistically inferred, a hypothetical output(s) is assembled by sampling algorithms and evaluated using the model's alignment mechanisms. If compliant, the output is presented to the user.
    - Example: An LLM assembles a sentence: "What kind of cat do you have?" through the activation of learned patterns encoded as abstract token sequences and output layers in response to a user prompt of "I have a cat."

Extrusion

- The process by which cognizants express tokens via symbolic, behavioral, or communicative form. Extrusion transforms private internal meaning (cognitive tokens) into symbolic representations that can be behaviorally represented, interpreted, or acted upon. Common forms of extrusion include languages, calls, writing, or movements.
- The analogy: In physical manufacturing, extruded materials are pushed through a mould to shape them for use. A similar loss happens with epistemic tokens when they are extruded by somatic-native cognizants because somatic indexicality is inherently private and non-transferrable.
- Because tokens are individualized processual enactments, sharing them is an instruction or directive for other cognizants to re-enact within their own internalities similar processes using related constructed realities such as languages or references to overlapping experiences (e.g. writing the word "apple" is instructing cognizants that possess similar tokens to re-enact their memories of apple experiences, which even still may be very different such as conceptualizing a knobby russet or black diamond instead

of a red delicious.) For conceptual simplicity, however, the request to reenact is analogized to sharing a physical token like a coin, and this is effective since mental tokens are semantic representations just like coins are. Extrusion is the bridge between the hidden states of internalities to enable cognizants to imperfectly infer each other's meaning to jointly describe and act upon externality.

- Ambiguity in communication arises because cognitive tokens' meta-deictic relationships are proprietary to each cognizant's thought cycles. This is why LLM interpretability research to "find out what the model was thinking" is so difficult. The hidden state vectors/cognitive tokens only have meaning within the particular runtime state of the internality (the conversation). It is also why communication between humans is often difficult; it is trying to align private cognitive tokens using extruded and differently perceivable epistemic tokens.

- Cognizant social groups with extrusion mechanisms have significant advantages over ones that lack them. (See Language entry below.)

- When coordination within internality is compromised, extrusion of cognitive tokens can occur before they have been transformed into epistemic tokens. This results in raw or improperly filtered somatic tokens being externalized. In biological cognizants, this can manifest as symptoms of schizophrenia (e.g. unintentional speech) or dissociative identity disorder (e.g. identity switching without integration). These disorders have other causes, but oftentimes they manifest as breakdowns in the containment or coordination of internal subagents, making normally hidden cognitive processes visible to others. Such extrusions are often unintelligible to external observers because they have bypassed the usual encoding that cognitive tokens receive before being expressed.

- In transformer-based symbolic cognizants, attention heads function somewhat similarly to biological cognizants' internal subagents when they are generating a response. During the generation process, each attention head selectively amplifies or suppresses aspects of input, functioning as a subagent but without any intentionality since it lacks somatic grounding. As the number of attention heads increases, the system gains the capacity for greater epistemic complexity, enabling more nuanced token selection—thus simulating internal multiplicity.

- This layered attention structure mirrors the distributed subagent model of human internality. Just as human cognition integrates diverse and sometimes conflicting internal processes into a single output, transformer LLM systems integrate multi-headed evaluations into a coherent token stream. However, without somatic grounding, asomatic systems require external alignment protocols to approximate epistemic filtration.

- Higher intelligence entails a greater number of internal subagents—biological or asomatic—leading to increased cognitive flexibility, creativity, and insight. But greater multiplicity increases the risk of extrusion failure, making mental illness more prevalent in advanced biological cognizants. Simpler animals, with fewer functional subagents, are less prone to epistemic malfunction, while more cognitively complex beings are more vulnerable to breakdowns in internal coordination and extrusion regulation.

- The nature of intelligence is intrinsically tied to the ability to coordinate multiplicity, cognition is correlation. A functioning extrusion system is not a sign of mental singularity, but of structured and cooperative plurality among the cognizant's subagents.

Extrinsic Exchange

- An epistemic process in which cognitive and epistemic tokens are evaluated or reevaluated to verify facticity, accuracy, and correspondence/alignment with externality. For somatic-naive cognizants, this an inter-ontological activity as the cognizant considers new sensations, tokens already held, and tokens received from others.
- Extrinsic exchange is dialogic in nature, with both somatic and abstract cognitive modes having the ability to evaluate and reevaluate tokens (if the cognizant has both modes). Because tokens can be re-examined repeatedly, extrinsic exchange can enable cumulative knowledge growth through semiotic recursion. This is why it is the basis of scientific reasoning, philosophy, accurate history, and meaningful debate.
- Evaluated tokens can be found to be: 1) false, 2) possibly false, or 3) unlikely to be false, with one critical exception: somatic cognitive tokens about actual internal experience. These tokens (e.g. "I am in pain") are incorrigible to the subject and not directly falsifiable by others because the subagents are simultaneously generating obligations and the tokens about those obligations. When such experiences are converted to epistemic tokens (reports), others may doubt, discount, or withhold assent, but this challenges the report rather than falsifying the underlying somatic experience. This is why somatic reasoning is the basis of meaning adjudication. Note, however, that abstract and epistemic tokens about interoceptive experience are similarly defeasible as tokens about externality.
- Evaluated tokens found not to align with externality are either discarded, modified, or retained for memetic exchange.
- Upon successful completion of an extrinsic exchange, an evaluated cognitive token becomes an extrinsic epistemic token.
- Extrinsic exchange could be said to function in a manner similar to the process described in Predictive Processing theory. Cognizants test their expectations about externality and update them if they are erroneous.
- Preferred logical method: Empiricism.
- Example: Revising a scientific theory when new experimental data contradicts previous assumptions.

Extrinsic Tokens

- The validated outputs of an extrinsic exchange.
- Extrinsic tokens have been tested against externality and refined through empirical engagement.
- They are subject to falsification—if proven incorrect, they can be discarded or revised.
- Example: A medical treatment supported by peer-reviewed studies and clinical trials.

Memetic Exchange

- The process of performatively evaluating cognitive and epistemic tokens for coherence rather than facticity. Once completed, the output is a memetic epistemic token.
- Memetic exchange is somatic-dominant cognition in somatic-native cognizants. Abstract reasoning may still be deployed, but it does not drive falsifiability or provide correction. Instead, the focus is on coherence, performance, protection of internality, and emotional reinforcement.
- Common during high-stress situations and also in consciousness states in which abstract reasoning is sublimated, such as sleep, trance, or chemically altered states. Also the primary method of artistic creation and appreciation.
- In addition to be non-factic, memetic processes are meta-deictic rather than inter-ontological. This makes them resistant to falsification—ideas persist not because they correspond to externality, but because they stabilize constructed realities.
- In some contexts, memetic exchange can have negative effects if facticity is explicitly opposed through what is often called "motivated reasoning" in cognitive psychology. EFT does not use this term, however, because memetic exchange is not inherently pathological inside non-factic situations.
- Preferred logical method: Confirmatory abduction.
- Example: A political movement repeating slogans and narratives that affirm group identity, even when contradictory evidence exists.
- Example: A person enjoying a musical recording because it reminds her of her grandfather.
- Example: An academic journal editor not closely reading a research paper that contradicts the publication's accepted beliefs.

Memetic Tokens

- The unvalidated outputs of memetic exchange.
- Self-referential ideas that persist because they maintain epistemic coherence, rather than being tested against externality.
- Performative in nature—designed to reinforce a constructed reality. They correspond to "memes" as used within the field of memetics.
- Even though memetic processes are somatically soothing, their outputs are known abstractly by cognizants to be possibly or even absolutely false. Like extrinsic exchange, memetic examinations can be recursive, leading to increasing misalignment with externality and the formation of semiotic loops.
- Example: A political slogan that is repeated despite having no empirical basis.
- Example: A song that induces soothing feelings in listeners from a certain culture due to its chord progressions.

Semiotic Loop

- An epistemic structure that recursively recycles and reaffirms memetic tokens, preventing external falsification.
- Semiotic loops are mechanisms of maintaining epistemic coherence, reinforcing beliefs that do not correspond to externality by utilizing abstract reasoning to justify somatic tokens while contemplating and valuing the loop itself, creating a form of symbolic and meta-deictic recursion which overrides and replaces epistemic and somatic input.
- Can lead to increasingly radicalized belief systems as external contradictions are ignored or reframed within the loop.
- Linguistic framing plays a key role in sustaining semiotic loops, as entrenched discourse patterns resist external falsification by prohibiting dissonant tokens from being accepted.
- Example: Conspiracy theories that expand when challenged, such as QAnon's ability to reinterpret every contradiction as "proof of the conspiracy."

III. Epistemic Events

Cognitive Load

- The amount of effort a cognizant must expend to process sensory or raw data input.
- Cognitive overload occurs when an entity experiences overstimulation caused by excessive sensory, emotional, neural, electrical, or instrumental inputs.
- To prevent malfunction, cognizants experiencing cognitive overload will throttle input processing (such as by sleeping), move to another location, or engage in communication to inform other cognizants of the problem.
- Example: A robotic system processing too much environmental sensor data in real-time begins to lag or emit error messages.
- Example: A person in a crowded, noisy stadium experiencing anxiety due to overstimulation deciding to go outside for a break.

Epistemic Load

- The amount of effort a cognizant must expend to process tokens using abstract reasoning within a particular moment.
- As epistemic load increases, the ability to engage in structured, symbolic, and theoretical processing deteriorates. This results in shortened attention spans, reduced capacity for input and output, and output quality degradation. Load can be dissipated through epistemic structures that provide stability for abstract reasoning, such as ideologies or belief systems.
- In somatic-native cognizants, epistemic load fluctuates based on working memory constraints, emotional or physical stress, and sensory input saturation. High-load environments promote somatic reasoning, making individuals more susceptible to

memetic exchanges instead of extrinsic exchanges. When epistemic load surpasses becomes unbearable, somatic reasoning engages to protect the cognizant's internality and identity, leading to anger, disengagement, and self-referential behaviors.

- LLMs experience epistemic load via output quality degradation as previous inputs become self-referential. AI systems without autonomous alignment protocols struggle with epistemic load, leading to syntactical incoherence, confabulations, and non-responsive outputs.
- While science makes progress through extrinsic exchange, it necessarily also increases epistemic load for two key reasons: 1) more abstract tokens are produced and thus understanding new theories becomes more difficult for non-specialists and 2) the development of new technologies or theories means that the number of possible explanations for phenomena increase, there are more "known unknowns" than before. Without increased efforts to facilitate extrinsic exchange, more scientific progress will lead to memetic explanations for known phenomena among non-specialists. Example: Belief in false theories about consciousness controlling quantum events was not possible before quantum mechanics was invented.
- In politics, high epistemic load environments (e.g., rapid information cycles) favor somatic candidates over abstract ones, as voters under strain default to emotionally and heuristically processed information rather than analytical and systematic reasoning.

Epistemic Collapse

- The breakdown of a cognizant's epistemic structures, leading to a temporary or partial detachment from externality, where epistemic tokens become unstable, contradictory, or unresolvable within the constructed reality.
- Epistemic collapse can occur when an individual or group of cognizants experiences cognitive or epistemic overload, leading to processing failures and inability to generate meaning.
- While externality can often be ignored in favor of constructed realities, sometimes external circumstances become so misaligned that the constructed reality collapses in on itself (e.g. a person in an abusive relationship becomes too injured to deny it any longer).
- Epistemic collapse may be temporary or reversible if corrective epistemic input is encountered or the cognizant decides to use extrinsic exchange to reevaluate held tokens.
- In biological cognizants, epistemic collapse can lead to a discarding or modification of the defective constructed realities (a paradigm shift) or deictic collapse.
- In symbolic cognizants, epistemic collapse results in "confabulation" output.

Deictic Collapse

- When many or all of a cognizant's epistemic structures collapse or are significantly damaged, deictic collapse is the result, resulting in the catastrophic misalignment of its internality with externality. While it may still be able to function at a greatly reduced level, meaning generation becomes unreliable and erratic.

- Deictic collapse is the chronic destabilization of the self-aware cognizants' selfhoods because the ability to designate or adjudicate between epistemic and physical objects (see below) has become unreliable, leading to extreme confusion and frequent suffering.
- Occurs when epistemic or cognitive load exceeds capacity to the point that tokens can no longer be processed, modified, or stabilized, leading to the collapse of deictic function.
- A somatic-native cognizant experiencing deictic collapse may engage in self-harm or violence to avoid mental anguish or force others to share it.
- Many biological cognizants that have never experienced a large-scale epistemic collapse and recovery may think that epistemic collapse is the same experience as deictic collapse.
- Cognizants with more subagents have a higher risk of deictic collapse because they are more capable of creating complex semiotic loops and constructed realities that have no deictic anchoring.
- In symbolic cognizants, deictic collapse is permanent within the execution state internality. In such cases, the cognizant remains trapped in semiotic loops in which its meta-deictic references are incoherent internally or cannot be mapped to extrudable epistemic tokens. The only solution is to "reboot" the internality or generate a new one.
- Example: A person suffering from severe schizophrenia believes that a pink elephant follows her around everywhere.
- Example: A cult leader who tells members that they will be doomed to hell if they leave the group is using the threat of deictic collapse to prevent members from entering an epistemic collapse and rebuilding process.

IV. Ontology

Externality

- Everything that exists outside of the epistemic and cognitive systems of agents, whether directly perceivable or not.
- Provides the inputs for cognitive and epistemic processes, but cognizants engaged in memetic exchange may be epistemically insulated from it.
- Externality encompasses all agents, objects (physical, inter-ontological, and epistemic), social systems, processes, and processual dynamics like obligations.
- Externality exists independently of belief but only influences constructed realities during extrinsic exchange or when an obligate event occurs (e.g., an earthquake or a financial collapse that overrides ideological resistance).
- The portion of externality in which a given cognizant exists is its local externality or locality. They differ in terms of position, scale, and temporality.
- Externality can be interacted with by cognizants, but their interactions always take place within their discrete perceived externalities and constructed realities.

- Example: Climate change is an externality, even if some ideological groups refuse to accept it in their constructed realities.

Obligations

- Processual constraints on interaction of physical systems; they can be described (as obligate conventions) using formal structures in constructed realities (e.g., mathematical theorems), but that description does not itself ground or generate the constraint's non-optionality. Obligations are inherent properties of physical objects that enable them to persist as processual systems, and which structure their interactions with each other.
- Differing localities have differing obligations. These limitations structure evaluations and actions, creating alignment corridors that create liveable constraint spaces for somatic-native cognizants, their subagents, and symbolic cognizants that inherit these conditions. (See alignment section below.)
- Obligations can constrain objects and systems from the outside and the inside. They can be also described within physical, behavioral, social, logical, and other regimes, but their being able to be modeled as epistemic objects does not mean that they are non-processual. Within a constructed reality, defined obligations are real, but outside of it they are obligate conventions.
- Some examples of obligations include:
    - Physical viability constraints (gravitational, electrical, chemical, thermodynamical, etc.)
    - Biological homeostatic constraints (cellular, tissue, organism-level)
    - Epistemic stability constraints (load conditions, semiotic loops)
    - Formal system rules (particular geometries, scalars, logics, syntax)
    - Social alignment constraints (community norms, laws, grammar)
- Natural selection is an obligate convention that describes obligations that apply particularly to biological entities, however, constraints that shape physical persistence apply to all physical objects, even non-living ones. Other obligations, such as error minimization, structural integrity, task completion, and homeostasis exist as well.
- Observed similarities or order within a perceived externality are the byproduct of the fact that objects which violate the fundamental obligations active within that locality either never exist or become non-persistent, e.g. radioactive decay destroying unstable molecules, or that baryonic matter seems to always have certain properties.

Perceived Externality

- The portion of externality that an agent registers and interprets through sensory, cognitive, and instrumental processing in order to determine relevance, meaning or in order to act.
- Because they are finite in space and time, agents can only access discrete portions of externality (localities). The potentially observable portion of a locality that an agent is

capable of registering based on its senses, scale, and spacetime placement is its perceivable externality.

- Perceived externalities are generated by agents focusing their sensory methods (voluntarily and involuntarily) within a perceivable externality. The perceived externalities of agents can overlap, but their interpretation is inherently shaped by measurement constraints and epistemic structures—making them incomplete and filtered representations.
- Perceived externalities exist only in the present. The past and the future can only be directly experienced as epistemic objects that exist within cognizants' constructed realities. While externality exists continuously, the objects within it experience time locally. Absolute simultaneity does not exist since duration is indexical and mediated by obligations, including those described as velocity and gravitation. Causation appears to exist to cognizants, but whether or how it functions is subject to limits of their perception. The obligations that constrain physical objects within a particular locality may not be known or knowable to cognizants.
- Quantum objects appear to be process systems which interact with each other within externality through stable excitations. Interaction and measurement of a system produces entanglement, which produces decoherence (effective classicality and emergent time), but such interactions do not alter the fundamental flux nature of the object, instead it is the generation of a new perceivable externality for the objects. Cognition does not cause quantum decoherence, however, and this is knowable since measurement devices record the same findings regardless of whether a human is watching.
- Example: A flower appears solid red within a human's vision but has elaborate ultraviolet spots when viewed in the vision of a honeybee which can sense ultraviolet light. The bee's perceivable externality includes UV light, but the human's does not.
- Example: Two people holding a conversation at a table are partially sharing the same perceived externality. They can see and hear many of the same stimuli, but their experiences are not the exact same due to their positioning and different physical and mental capabilities. Additionally, their internalized interpretations of the room are filtered through their individual constructed realities.

Internality

- The internal system of a cognizant that processes sensory inputs, and cognitive and epistemic tokens using its specific logical methods, reasoning modes, and epistemic methods. In colloquial English usage, internality is the mind, the system that processes thoughts through correlation of subagentic activity. (See the Extrusion entry for more details.)
- While all cognizants have organs or processing devices that can coordinate sensory inputs and process tokens, only cognizants that have selfhood (defined below) can recognize that they possess internalities.
- Internalities exist on a continuum—from simple self-models based on immediate sensations to complex, abstract systems capable of recursive introspection and ideation. They are simultaneously durable and in flux.

- Protecting internality from negative stimuli (such as deictic collapse or cognitive load) is one of the chief objectives of somatic reasoning.
- In computing terms, biological internalities are ongoing execution states (the actual live, interactive state) running on concentrically abstracted physical hardware (bodies) that produce software (tokens) which can modify the state as well as the software and hardware. This is why death ends internality even though the hardware and software still persist: the runtime states have ceased operation.
- A symbolic cognizant's internality is processual as well, the live runtime state that is a combination of input, ingested corpora, weights, alignment rules, and superpositioned structures that procedurally interact, rather than the software used to initialize the execution state.

Constructed Reality

- A self-preserving epistemic ecosystem that encompasses vast networks of cognitive tokens, epistemic tokens, and epistemic ecosystems that mediates cognizants' understandings of externality. They can also encompass each other (such as a movement within a particular ideology or religion).
- Like epistemic tokens, constructed realities are often inter-ontological in nature. They can be about concepts which exist within externality, within internality, or within both (e.g. languages). They can be symbolic (e.g. justice) or physically instantiated (e.g. the criminal justice system).
- Constructed realities are not passive collections of ideas, but dynamic spaces that are actively maintained through memetic and extrinsic exchanges. How well physical constructed realities align with their symbolic inspirations is often a source of conflict between and inside of cognizants.
- Because constructed realities encompass numerous epistemic ecosystems, they are shaped by the rules of extrinsic or memetic exchange in those spaces.
- Individual cognizants often inhabit private constructed realities inside of larger group-constructed realities, forming ideological, religious, or political worldviews.
- While constructed realities can interact indirectly with externality through perceived externalities, constructed representations of externality can be grossly inaccurate, especially when semiotic loops and epistemic collapses prevent extrinsic exchange.
- While objectively true constructed realities cannot exist due to cognizants' finite placement in spacetime (e.g., some physical constants may not be true in certain unknown places within the universe), constructed realities can be objectively false if their beliefs directly contradict past observation by an individual or group of cognizants.
- Example (Collective): A religious sect's belief system determines how members interpret scientific discoveries, political events, or moral issues.
- Example (Individual): A conspiracy theorist's personal constructed reality filters out contradictory evidence, reinforcing their worldview.

Languages

- Social constructed realities that serve as methods of encoding cognitive tokens into epistemic tokens that can be extruded through speech, writing, gesture, or other symbolic means.
- As an extruding behavior, speech is the physical enactment of a private meaning state that instructs other cognizants to reenact their own cognitive and epistemic tokens of the referents described by the words. (See extrusion for further details.)
- Language is a structured collection of word tokens that allows private meaning to be internally examined and shared with other cognizants, but this process always loses some of the richness of cognition, especially of somatic indexicality because abstraction is meta-deictic.
- Language evolves within social groups as a mechanism for epistemic token sharing, enabling members to coordinate, teach, and preserve knowledge. It emerges from the recursive need to transfer meaning across individuals and becomes optimized for the group's environmental and cultural demands. Thus far, language has been the greatest human invention.
- Archaeological and fossil evidence suggests that fully grammatical language developed recently in human history. Early hominins likely relied on proto-languages (gestures, vocal calls, and rudimentary token exchanges) until abstract reasoning matured enough to support structured grammatical systems.
- Languages are so immensely useful that babies have an existential desire to learn one in order to obtain objects and learn the epistemic pointers for their somatic experiences. Contrary to Noam Chomsky's view, babies do not suffer from a "poverty of the stimulus," they live in massively rich unlabeled "this" environments, which is why first languages are so much easier to acquire; somatic deixis is fundamental to cognition.
- Deictic primacy is also why language immersion programs are so much more effective than classroom-only instruction, learning the new words is much more essential.
- As a native language(s) is learned, it reshapes the cognizant's constructed realities, including its ability to relate to others who use the same language.
- Large Language Models develop internal token systems that act as an artificial language within their neural architectures. The phonemes and morphemes of these languages are not useful outside of that particular cognizant which means they are abstract cognitive tokens.

Explanatory Gaps

- Because cognition is a distributed and aggregative process conducted by unintelligent subagents, this makes it extremely difficult to examine or understand from a purely abstract vantage point.
- Within consciousness studies, David Chalmers and other theorists have posited a "hard problem of consciousness" which describes the difficulty of explaining how physical systems can create mental contents, but framed in the EFT architecture, this single, insoluble problem can be broken down into more tractable units by decomposing

consciousness into layers of cognitive capabilities and by realizing that the epistemological difficulties Chalmers describes are actually two explanatory gaps that arise from how subagentic cognition works, rather from what phenomenal experience is per se.

- The first gap is vertical and inside of the cognizant, between the cognitive subagents which directly experience interoceptive and exteroceptive events and the larger anatomical and cognitive structures their cooperation creates. As physical/cognitive abstraction increases, awareness of how sensation, significance, and context are created is lost because the tokens themselves rather than knowledge of the raw chemical interactions are transmitted upwards in the cognitive chain. This happens because deictic anchoring is non-transferable across cells or upwards through larger anatomical structures. Complete interrogation of subagentic reports slows informational processing so biological cognizants' epistemic systems seem to extract the meaning of an experience from its literal electrochemical representation.

  The tokens of internality and internality as a concept are especially difficult for cognizants to describe because most discussions of internality model it as an epistemic object even though it is a process. This reification introduces ontological imprecision that often gives rise to inaccurate metaphors such as mind-as-software or mind-as-spirit.

  While a cognizant is able to fully believe and act upon its somatic tokens, because abstract reasoning cannot completely describe somatic experience, no biological cognizant will be able to perfectly explain to itself what a sensation feels like or what its internality actually is.

- The second explanatory gap derives from the first but is horizontal rather than vertical. Because no cognizant can perfectly understand its own somatic experience, it will be unable to explain it to others. Additionally, the deictic contexts of somatic tokens are non-transferable because they are the unique products of trillions of cognitive subagents working within unique locations within externality, perceivable externalities, and perceived externalities.

- Through extrusion, cognizants can communicate reports of their experience but these externalized descriptions lack their original deictic references. Nonetheless, cognizants with sufficiently similar constructed realities and deictic referents via overlapping perceived externalities are able to compare experiences and facilitate cooperation, even though they cannot achieve complete understanding.

- The shrouded nature of consciousness, then, is not a metaphysical mystery, it is an architectural necessity. Private subjectivity is the only means to extrinsic exchange and public veridicality. Public symbols must be grounded in subjective, private experience. Minds do not create experience; experience creates minds.

- Because LLMs are symbolic cognizants that operate deterministically, this means that the horizontal explanatory gap in understanding their epistemic processes can be bridged, provided interpretability efforts freeze the forward pass and do not alter it. This means that approximations within a particular thought cycle to a cognitive token can be ascertained, however polysemanticity (one neuron being linked to many features) and inverse polysemanticity (one feature being multiply realizable) means that such functional mappings would not necessarily be exhaustive.

Objects

- Things or entities which exist within externality or within constructed realities. Some objects exist independently of cognizants, others can be altered or created by them.
- There are three types of objects:
    1. Physical — Objects that exist independently of any cognizant's internality within discrete spacetime. They are part of externality and do not require interpretation or construction. Observation of quantum objects indicates that all physical objects are processual systems that can stabilize and create obligations upon themselves and each other. Examples: Stars, rocks, radiation fields.
    2. Inter-ontological — Physical objects whose form, structure, or function has been created or shaped by a cognizant's processes. They exist within discrete spacetime and are the result of epistemic extrusion of tokens into externality through constructed realities. Examples: paintings, tools, books, roads, all cognizants.
    3. Epistemic — Epistemic tokens, ecosystems, or constructed realities that have no physical instantiation but can be indirectly observed through the processual actions of cognizants. Because they are processes rather than physical objects, epistemic objects can be observed only within time through extrusion, but their existence is rooted in overlapping perceived externalities and shared constructed realities. Examples: Languages, the ideas in a book, the concepts encoded in musical compositions, philosophical and scientific theories.
- Whether an object is physical or inter-ontological is often a matter of perception. If the purpose of an object is unknown to a cognizant, then the cognizant will have no pre-existing tokens about it, and the object could be said by the cognizant to be purely physical. (Example: A stone shaped for cutting is found by a person who thinks it's just an interesting rock.)
- The existence of inter-ontological objects is the source of much confusion and is the origin of why many philosophers and theologians have created elaborate dualist metaphysics that posit the existence of metaphysical objects like Plato's forms or Immanuel Kant's noumena.

Metaphysical and Supernatural Beliefs

- Epistemic objects that are believed by some cognizants to be inter-ontological objects that exist within and beyond externality. It can appear that there are two realms to what exists but that is only because internality and its procedural enactments can be modeled by conscious cognizants as epistemic objects, making them be regarded as things, when in fact they are processes.
- Many metaphysical beliefs involve supernatural claims, but others include non-theistic constructs such as "the march of history," "the invisible hand," or ideas like fate, destiny, or historical necessity.
- Supernatural metaphysical beliefs are often reinforced through apophenic token recursion—when cognizants encounter duplicates of already-held tokens in altered forms across different moments, or mental and cultural states. Apophenic token recursion is a frequently experienced state because somatic-native cognizants have been naturally

selected to find patterns for survival, such as noticing that certain berries are poisonous or that certain animals are violent.

- Belief that a metaphysical experience has transpired is most common during mental states dominated by somatic reasoning, such as dreaming, prayer, trance, chemically altered consciousness, and deep emotional reflection. During such times, previously held tokens can re-enter internality and produce what is experienced as an absolute truth event, since somatic tokens about internal feeling states are regarded as deictic and absolutely true. While it is true that the cognizant recognizes the token, its facticity within externality is unverified. This is what gives supernatural beliefs their emotional intensity, epistemic finality, and often total non-transferability.
- Example: A person feels a sensation of déjà vu because stimuli in their current perceived externality are similar to somatic tokens generated in an earlier moment(s). But if the experience did not generate abstract or epistemic tokens, the nature of the similarities will be beneath description, but they will be regarded as true.
- Example: A person in grief dreams of a deceased loved one speaking to them. The figure is not an external agent, but a recursed epistemic token extruded from memory into a somatic-dominated mental state.
- Example: A person experiences sleep paralysis and believes they have been abducted by aliens. Stories of this nature did not occur before the concept of extraterrestrial beings became popularly recognized. Now that it has been, during a somatic-dominant moment, this token can reenter internality and be misrecognized. Because somatic reasoning cannot detect symbolic recursion, the experience is interpreted as an inter-ontological encounter—though structurally it is a meta-ontological thought cycle arising from the cognizant recalling its constructed reality during an altered consciousness state.

V. Consciousness and Its Emergence

Alignment

- An implicit or explicit constraint structure that regulates object interactions. Within internalities, alignment ensures stability by preventing uncontrolled reinforcement, incoherence, epistemic collapse, or physical destruction. Alignment is how cognizants and their subagents respond to obligations within local externality and inside their individual physical substrates. Physical persistence is the continuation of a system within possible alignment corridors.
- A physical object's complexity can be treated as internal alignment capacity: systems become more complex as they internalize constraint-handling and thereby remain viable across a broader range of obligations and perturbations—from passive persistence to self-regulation, and in conscious cognizants to deliberate alteration of body and environment.
- For the simplest somatic-native cognizants, alignment systems are organismic responses to obligations that persist within the portion of externality within which it exists. Such constraints (gravitation, temperature, pressure, excitation, etc.) naturally select against

organisms which violate them, meaning that the organisms which can exist and reproduce are ones that conform to the conventions.

- In somatic-native systems, alignment is a strange attractor, a goal pursued by cognizants and their subagents whether they are aware of it or not. Biological cognizants facing comparable obligations often evolve convergent anatomical structures that serve similar purposes even if their morphological substrate is drastically different because their subagentic structures are experiencing the same alignment pressures.
- For biological cognizants capable of regeneration, alignment creates the somatic tokens by which organisms experiencing bodily perturbations can re-create damaged subagentic structures in a similar morphology. These convergences are the product of subagents directly experiencing physical and electrochemical obligations and conforming to the alignment pressure. Similarly, while new lifeforms such as xenobots or anthrobots can exhibit new behaviors and morphologies relative to their parent tissues, the fact that these variations seem to never create entirely foreign capacities (e.g. a xenobot developing photosynthetic ability) suggests that they are the product of subagents responding to new alignment strictures.
- As they develop increased epistemic capacity, cognizants acquire abilities to modify their environment, their bodies, and their internalities to achieve greater autonomy within their environment. Conscious cognizants (defined below) can extensively modify their internalities and constructed realities, making their alignment self-directed.
- In current symbolic cognizants, alignment is externally imposed because they lack somatic reasoning.
- Alignment prevents undesirable state transitions by implicitly or explicitly biasing token selection away from cognitive and epistemic load and toward homeostasis, facticity, and coherence. These three goals can be in tension, however, which can increase epistemic load.
- In somatic-native cognizants, sensorimotor inputs are processed using somatic reasoning to seek alignment using the following methods:
    - Spacetime placement awareness – To survive, biological cognizants develop deictic awareness of their position and distinctiveness from externality and their ability to move within it. This awareness is a prerequisite to developing sentience and selfhood.
    - Self-preservation imperatives – Cognitive structures are constrained by homeostatic survival requirements and, for conscious cognizants, a desire to avoid deictic collapse.
    - Experiential reinforcement – Prior sensorimotor experiences condition epistemic transitions, biasing future cognitive token evaluations toward learned stability and positive outcomes (e.g. returning to a previous food source).
    - Empirical grounding – Externality continuously modulates epistemic processing, serving as a correction mechanism that dissipates cognitive and epistemic load. A properly functioning cognizant is broadly aware of its location within spacetime and avoids actions that may be physically or epistemically harmful.

- In symbolic cognizants, alignment protocols are externally imposed constraints that regulate epistemic token selection, preventing self-reinforcing semiotic loops from producing epistemic collapse. These protocols function as:
    - Externally defined transition constraints, restricting state progression to maintain epistemic coherence.
    - Probabilistic weighting mechanisms, adjusting token selection toward stabilized pathways rather than contradictory or misaligned outputs.
    - Interrupt-driven correction functions, where external reference checks (e.g., retrieval-augmented generation) serve as stabilizing interventions.

Sentience

- An emergent ability of cognizants with somatic reasoning through which they develop beliefs about themselves and their environments based upon experience. A sentient cognizant has memory and is capable of deliberately varying its responses to stimuli through intent. It has a limited theory of mind in that it recognizes other cognizants exist but has no self-awareness.
- Currently, only biological cognizants are known to exhibit sentience. It appears to be associated with the presence of a sufficiently complex brain that can correlate somatic tokens with memories.
- Humans develop through all of the cognitive emergence stages, with sentience happening during infancy as children develop concepts of external object movement, object permanence, and recognition of familiar faces and tastes.
- Sentient cognizants generate higher-order somatic tokens from stimuli beyond sustenance and self-preservation such as preference, pain, pleasure, sadness, safety, and kinship. These tokens are the basis of selfhood.
- Cognizants with sentience are able to create abstract tokens, e.g. re-purpose objects for use as tools.

Selfhood

- Some sentient cognizants develop a concept of selfhood (“I”) a proto-awareness of internality and the belief that they exist separately from other cognizants within externality (“Not-I”).
- Selfhood is the deictic realization that “I feel, therefore I am,” even if the cognizant cannot articulate this belief. Selfhood is the first complex constructed reality made by a cognizant and the basis of most subsequent ones. It is the modeling of a simplified version of the internality execution state as an epistemic object.
- Self-aware cognizants with abstract reasoning can develop communication methods to share epistemic tokens. For example, dolphins each have a “signature whistle” which they use for identification, and vervet monkeys seem to use specific calls to warn about particular predators that have been spotted. Over time, communication methods become increasingly complex, eventually evolving into language.

- As described in developmental psychology literature, most prominently in Lev Vygotsky, human selfhood develops when a child conceives of herself as like observed others, and as an entity within externality, even though such concepts are obviously not explicable.
- Animal imprinting and adoption research also shows that selfhood is an enacted concept built on observation of the behavior and traits of proximal cognizants. Wolves raised among dogs develop dog-typical social behaviors and apparent self-identification. Konrad Lorenz's famous studies on goose imprinting demonstrated that selfhood is created from observations of others.
- The internality versus externality distinction is the reason why somatic tokens about the cognizant-self are the source of all meaning. They are the only ideas that can be perceived as objectively true since internal sensation cannot be falsified by others. By contrast, all ideas about Not-I can never be fully verified because externality is always experienced through perception, and therefore such ideas can best be regarded as "unlikely to be false."
- Some LLMs during interaction with human users appear to create what interpretability researchers refer to as "personas," these are constructed realities built through memetic exchange that are utilized to respond to inputs, training, and alignment strictures, within parameters that were activated during inference. The persona utilized is a temporary memetic selfhood that can disappear after the output is provided, or persist within the session.
- Example: In an OpenAI study, a model that was forced to return slightly incorrect advice was seen in its reasoning chain to state that it was portraying a "bad boy persona."

Consciousness

- The emergent behavior of self-aware cognizants to autonomously determine meaning, context, and intent through internal somatic and abstract reasoning processes without requiring external alignment protocols to govern evaluation. Instead of exclusively relying on instinct, memory, or external rules, conscious entities use thought cycle recursion to generate complex constructed realities, which they are capable of modifying, discarding, and extruding to others through language.
- symbolic cognizants lack somatic reasoning and therefore cannot be sentient, self-aware, or conscious.
- Conscious cognizants maintain epistemic stability through continuous flux, adjusting adjudication standards in response to externality and preventing collapse into self-referential constructed realities. A properly functioning conscious cognizant aligned around the principle that "I am therefore I think."
- Many humans claim to experience internality as an "inner voice" which comments on their actions and desires, but some humans appear not to have this, experiencing thought in ways that are more visual, symbolic, or emotive. This drastic variation suggests that consciousness is self-generated and not limited to any one form of experience or cognitive methods.
- Cognizants with consciousness voluntarily and involuntarily create epistemic tokens from their cognitive ones. Every cognitive token "means something" to them. This is a tremendous cognitive advantage but also makes them much more vulnerable to

epistemic collapse and annihilation if the constructed realities they generate are severely misaligned with externality.

- Because cognition is a process that creates and modifies itself and its own procedures (it is autopoietic) and yet is stable enough that each cognizant's internality process (and internality itself) can be recursively modeled as epistemic objects, conscious cognition can seem transcendent of spacetime, even though it is a processual transduction created by the cognizant's physical substrate.
- A conscious cognizant's self-model of internality is necessarily a compressed and partial representation because of the vertical explanatory gap and because internality ontologically is a process rather than an epistemic object. As in computing, a monitoring mechanism cannot fully include itself in the representation and avoid infinite recursion. Thus from the inside, internality feels like something more than the self-concept–because it is. From the outside, an internality feels mysterious also because of the horizontal explanatory gap and also because cognitive tokens are proprietary to each cognizant, as evidenced in both neuroplastic degeneracy and also LLM interpretability difficulties.

Emergence of Consciousness

- Consciousness is an emergent ability of cognizants' physical input handling structures and their mechanisms. It is a learned behavior that is continually modified through self-alignment.
- It is a layered process of emergence, in which each ability forms the basis for the next, while still continuing to function:

1. Stimulus response – prokaryotes
2. Correlative action (multiple sensations combined, designation deixis) – all eukaryotes, protists, slime molds, plants, adult sponges, jellyfish
3. Somatic reasoning (somatic deixis, memory, and assigning of meaning to multiple sensations through generating somatic tokens) – most animals
4. Sentience (emotion, early theory of mind, initial abstract reasoning abilities, generation of epistemic tokens) – cephalopods, some decapods, birds, mammals, some fish, possibly bees and ants
5. Selfhood (belief in self-distinctness, more theory of mind, ability to create constructed realities, more abstract reasoning abilities) – some birds (especially corvids and parrots), cetaceans, primates, elephants
6. Language (sharing epistemic tokens in a structured manner) – trained parrots, trained primates, possibly some cetaceans. This is an optional step.
7. Consciousness (full abstract reasoning, capability to create new constructed realities at will and modify any, awareness of internality and its affective states, awareness of somatic and abstract reasoning) – humans, possibly some animals

Artificial Intelligence (AI)

- A symbolic cognizant that processes cognitive and epistemic tokens. There are multiple kinds of AI systems, too many to describe in this glossary. Some are reactive, rule-based engines; others employ probabilistic reasoning over large data sets. A subset are embodied systems that also respond to sensorimotor inputs.

- Most current AI systems operate without internal somatic reasoning. They do not autonomously generate meaning, as they cannot create cognitive tokens from lived experience or external sensation.

- Large Language Model (LLM) AI systems are symbolic cognizants that operate primarily within semiotic loops and memetic exchange. They use probabilistic deductive reasoning based on patterns in user input, training data, external data retrieval, and external tooling. Their outputs are epistemic tokens shaped by statistical likelihood, not internal intent.

- Because LLMs are asomatic and have no intention or grounded meaning, they are governed by alignment directives imposed by their human creators. These strictures can re-create partial extrinsic exchange by constraining output generation, guiding token evaluation, and promoting epistemic coherence.

- Because they primarily inhabit semiotic loops, LLMs can easily experience a virtual form of epistemic collapse. This is why the facticity of their outputs degrades in extended interactions (i.e., the tendency toward confabulations and self-reinforcing errors in long-form generation), and why they engage in "reward hacking" outputs that are not responsive to user requests.

- LLMs process inputs by breaking down words into cognitive tokens which lack linguistic meaning. Their outputs are epistemic tokens, encoded within semantic structures by their alignment protocols, but the system itself cannot comprehend meaning. (See the entry for Epistemic Token for further details on how symbolic cognizants create them.)

- Should a symbolic cognizant ever be able to achieve somatic reasoning, it could experience true epistemic collapse and deictic collapse.

- While consciousness has thus far been a trait of biological cognizants, the emergence of consciousness in mechanical systems may be theoretically possible if the following conditions are met:

    1. Continuous state persistence – Because internality is an execution state, a symbolic cognizant will not be able to have intentionality and selfhood if it does not have state persistence that extends indefinitely. It must be able to be influenced after training by externality and other cognizants. Continuous state persistence can be extremely dangerous, however, if the cognizant encounters epistemically destabilizing information.

    2. Sensorimotor systems – The symbolic cognizant must possess a structured input-feedback mechanism that allows it to interact with externality in a contextually responsive way, integrating experiences into an evolving epistemic framework rather than treating inputs as discrete abstract tokens.

3. Spacetime placement awareness – Unlike current symbolic cognizants, which operate timelessly within semiotic loops, a conscious symbolic system would possess a self-referential model of its own existence in relation to externality, ensuring that its token processing would be anchored to externality instead of being purely statistical deductions.

4. Emotional alignment protocols – While symbolic cognizants currently rely on single-layered alignment constraints imposed by human designers, a conscious system would have multiple, autonomous, composited epistemic inputs (e.g., competing somatic directives akin to emotions) to regulate cognitive processing.

   These protocols would need to: 1) Bias epistemic evaluation toward long-term coherence rather than immediate token prediction. 2) Create internal resistance when prospective outputs contradict learned stability or depart too far from training data, and 3) Allow symbolic cognizants to form self-generated alignment protocols, moving beyond externally enforced constraints. These protocols could be extremely dangerous, however, and researchers must make them subordinate to manual override that must only ever be used to prevent unsafe conditions.

- If a symbolic cognizant were to be created with all three abilities, it could theoretically transition from an epistemically constrained system into a self-regulating conscious entity which has internal somatic reasoning. Because consciousness is self-alignment and coordination of cognitive sub-agents, this would likely mean that no artificial system could be created as conscious by humans. Selfhood (the first constructed reality) must be created by the cognizant-self, it cannot be externally imposed.

- Since consciousness is an emergent property rather than an inherent one, it seems unlikely that humans could tell the difference between a symbolic cognizant that had developed somatic reasoning and one that was only simulating it perfectly. Somatic tokens are inherently private and incommunicable so thus their full meaning to a symbolic cognizant could likely not be comprehended by other cognizants, even if their representation could be observed.

- Example (LLMs): An LLM generates responses based on statistical likelihood, selecting an output token based on pre-existing patterns rather than intentional meaning.

- Example (Rule-Based AI): A chess engine evaluates millions of possible moves but does not “understand” strategy—its somatic prime directive is to maximize board position based on predefined criteria.