Semantic Phasor Theory: A Wave-Based Model of Meaning for AI

Joseph. R. Carvalko, Jr.

Joseph R. Carvalko, Jr.
Independent Researcher,
Shelton, CT, USA
Carvalko@outlook.com
Jcarvalko@Quinnipiac.edu

ABSTRACT

Semantic Phasor Theory (SPT) proposes a wave-based model of meaning that complements traditional vector-space approaches used in contemporary large language models. Instead of treating words as static points, SPT represents meaning as a spectrum of semantic components, each with a magnitude and phase that determine how they align with context. This framework introduces wave-inspired mechanisms—superposition, interference, and phase rotation—to model ambiguity, sense selection, and contextual modulation in a more explicit and interpretable manner. By adapting tools from signal processing, such as filters and projectors, SPT provides a structured operator calculus for transforming and analyzing meaning. The aim is to offer a conceptually simple yet mathematically grounded explanation of how context shapes interpretation, suitable for both technical and non-technical audiences.

Keywords: semantic phasors, meaning representation, wave-based semantics, context modeling, signal processing analogy.

1. Why Today's AI Struggles with Meaning

Most current large language models (LLMs) like ChatGPT, Copilot, and Gemini turn words and sentences into long lists of numbers (token vectors) and then learn patterns in how these tokens follow each other in text. This works well for many tasks, but it also leads to familiar problems: such as fabricating ("hallucinating"), mixing up different meanings of the same word, losing track of context, or giving answers that sound confident but are logically shaky.[1]

We can think of today's models as treating meaning as points in a high dimensional multiverse space and compare them mostly by distance. They don't usually represent phase explicitly, so interference, alignment, and coherence are only implicit in the math. SPT treats meaning as if it were built from overlapping waves. Conventional transformer processing lacks a built-in way to talk about interference, alignment, or coherence of meanings as waves. This makes it hard for them to keep multiple roles of a word active at once like in this tricky example, "She broke his car and his heart.” Here, the single verb “broke” applies to two objects in different ways: literally to the car, and figuratively to his heart.

2. A Wave-Based View of Meaning

In engineering, phase tells us how a wave lines up with a reference wave. In SPT, we borrow this idea metaphorically: the ‘angle’ of a semantic phasor represents how a piece of meaning aligns with a chosen context—such as tone, viewpoint, or role in a sentence. It’s not a physical angle, but a conceptual tool for tracking alignment. SPT suggests each piece of meaning is like an arrow or hand of a clock that points to how a word, sentence, or paragraph is being used. The length of the arrow (magnitude) reflects how strong a semantic component is, while the angle (phase) reflects how that component aligns relative to a contextual reference, helping us model tone, viewpoint, or semantic role.[2][1]

But to complete the picture SPT does not look at just the one word but all the various tokens that support the word, because any word is a not defined by a single point in space, but part of a larger array of related concepts and works referred to as token vectors, where each word is composed of many concepts and is in reality like a musical chord made of many “notes” where we can think of each semantic eigenmode like a single note on a guitar string. A word is more like a chord: a combination of several notes played together. The overall meaning is a spectrum—a layered mix of these modes—so different parts of the meaning can grow, fade, or clash as context changes.[2][1]

To describe how meaning unfolds over time, SPT uses the metaphor of sound waves or water waves, which are easier to picture. As sentences and paragraphs move forward in something we might write, its semantic wave travels, overlaps, and interferes. Let me be clear these aren’t physical waves in air or water. Instead, we represent the unfolding of meaning *as if it were* a wave moving through a sentence: as words appear in sequence, their semantic phasors overlap and interfere, making some interpretations stronger (constructive interference) and others weaker (destructive interference). Its occurrence in amplitude, position and phase determines which interpretations rise like constructive waves and which get flattened out like destructive interference.[1]

3. What Is a Semantic Phasor?

In SPT, we model meaning using ideas from wave physics: superposition, interference, and phase sensitivity. Real language is more complicated than a simple wave in a tank, but this wave-based picture gives us a powerful way to describe how different shades of meaning can reinforce, compete, or flip depending on context. In this wave-based picture, three properties matter most:

- Superposition: multiple semantic components can exist at the same time, like overlapping waves on a lake.

- Interference: components can boost each other when aligned or cancel each other when out of step.

- Phase sensitivity: small changes in context can tilt which sense of a word comes to the front, the way shifting the timing of two speakers can change whether their voices blend or clash.

These are meant to mirror how humans use context to sharpen, soften, or redirect meaning from word to word and sentence to sentence.[2][1]

4. Tools for Shaping Meaning

In signal processing, whether it’s a sound recording or AM radio, a phasor encodes both how big a sinusoidal component is (amplitude) and how it shifts in time or angle i.e., the phase. The magnitude tells us the strength, and the phase tells us how it lines up with a reference signal.[1] In the full technical work, SPT shows that meaning can be modeled using the same kind of math used for AC signals. The structure of the equations mirrors signal processing, even though the ‘waves’ here are patterns of meaning, not physical vibrations.

Note that SPT borrows words like ‘frequency’ and ‘energy’ from signal processing, but here they’re abstract: a ‘semantic frequency’ is just one basic pattern of meaning, and ‘energy’ is how much that pattern contributes. The math looks like signal processing, even though the ‘signals’ live in meaning space, not in wires or air. Here, the magnitude of a semantic phasor says how strongly a semantic component shows up, and the phase says how that component lines up with a specific context such as tone, domain, or discourse role—like adjusting the angle of a loudspeaker so it catches a different part of a musical performance.[2][1]

Under this view, the meaning of a word or concept is a spectrum of phasors spread across semantic eigenmodes, like breaking a sound into pure tones using a frequency analysis. A sentence or utterance becomes a carefully structured superposition of these modes, where the pattern depends on both magnitudes and phases.[1][2]

To work with these meanings, the theory introduces a small toolkit of operators that act on semantic phasors much like standard operations in signal processing of any sort. You can think of this toolkit as a kind of “calculus of meaning” for rotating, filtering, projecting, and recombining semantic content:[2] Let’s look at vocabulary that describes this new way of analyzing communication based on not only a words simple definition, but what it means in the context of that is being used.

- Rotation operators: These change the phases of semantic phasors without changing their magnitudes. Intuitively, this model shifts in context—like changing stance, audience, or frame—while keeping the core content strength the same, similar to delaying a signal or shifting its carrier phase.[2]

- Semantic filters: These reshape how “energy” is spread across semantic modes, like low-pass, high-pass, or band-pass filters that reshape the spectrum of a sound through loudspeakers or any electronic signal for that matter. They can emphasize or dampen certain semantic components to support things like focusing on a topic or adapting language to a specific domain.[2]

- Projectors: These operators “pick out” components that line up with chosen semantic directions, behaving like correlation or matched-filter detection in signal processing. They provide a way to measure and extract a particular semantic pattern from a more complicated mix.[2]

Together, these operators form an operator calculus that mirrors classical linear systems theory in fields such as physics but now applied to semantic spaces instead of physical time-domain sounds, optical projections or electronic signals. The goal is to make the structure of meaning explicit in terms of magnitudes, phases, and spectral decompositions so that we can transform communicative content in a precise, mathematically controlled way.[1][2]

5. Why This Matters for AI Systems

When meanings are represented as semantic phasors, interference becomes a central mechanism: aligned phases reinforce each other, while opposed phases weaken or cancel. This offers a natural way to model ambiguity, contradiction, and sense selection, with context acting like a conductor who brings some instruments in phase and quiets others.[1][2]

Current transformer models can juggle multiple senses of a word, but they do so implicitly, buried in large real valued vectors. They don’t have an explicit phase-like structure that cleanly separates senses or makes interference visible. SPT proposes a representation where these effects are built into the math, making ambiguity, reinforcement, and contradiction easier to see and control.[1]

In the semantic phasor framework, different senses of a word live in different eigenmodes, much like different frequencies in a sound live in the string of a guitar. Context then selects or suppresses these senses by aligning or misaligning their phases, offering a structured alternative to purely pattern-based methods for choosing the right meaning.[1][2]

6. Applications and Future Possibilities

The SPT spectral picture suggests new ways to tackle familiar LLM transformer problems—such as hallucinations, semantic drift, and inconsistency—by viewing them as breakdowns in phase coherence in semantic space. Coherent, phase-aligned configurations correspond to internally consistent meanings, while strong destructive interference can signal contradictions or violations of contextual constraints.[1][2] From this foundation, SPT points toward several application areas:

- Semantic search: Projector-like operators can act as true semantic correlation tools, comparing both direction and phase patterns in semantic space, not just angle between static vectors. This could help distinguish queries with similar wording but different nuances, stance, or irony, where simple conventional cosine similarity is ambiguous.[2][1]

- Semantic memory and recall: Inspired by holography, the theory supports distributed, interference-based storage of concepts, where information is spread across many phasors so that partial or noisy cues can still retrieve relevant content. We use ‘semantic holography’ as an analogy: just as optical holograms store information in interference patterns spread across a surface, SPT imagines meaning stored in overlapping semantic phasors. The goal is a memory that’s more robust to paraphrase, even though it’s implemented in software, not in physical holograms.[1][2]

- Context-aware language interfaces: Phase rotation operators can capture changes in context—tone, audience, legal style—while keeping the overall “energy” of the core message. This gives a principled way to do style transfer, context-sensitive rephrasing, or jurisdiction-specific rewrites that preserve underlying meaning but adjust how it is expressed.[1][2]

- Cross-model and cross-agent alignment: Because meanings are spectral objects, independent systems can, in principle, align their semantic bases and exchange compact phasor codes instead of raw text or full model weights. This suggests paths to privacy-preserving knowledge transfer and modular AI systems that share structured meaning representations rather than copying parameters.[1]

- Interpretability and control: Viewing internal states through a spectral-operator lens gives “physical-style” tools—filters, modulators, projectors—for inspecting and reshaping what models encode. This can support interpretable debugging, controllable concept activation, and semantic “firewalls” that block access to restricted regions of semantic space.[2]

One especially important arena is legally and ethically constrained language use, where outputs must obey context-dependent rules like laws, regulations, and safety policies. Context-aware interfaces can attach a structured profile to each interaction that records jurisdiction, domain (such as medicine or finance), risk level, and applicable rules, so the same question can get different, locally compliant behavior in different places or sectors.[1]

On top of this, additional models can act as legal or ethics “judges,” reviewing candidate responses considering the active context and flagging potential violations such as unlawful advice, discrimination, or safety hazards. Because rules and internal policies change over time, the context layer can connect to up-to-date sources and update constraints without retraining the base semantic model, allowing ongoing policy alignment.[1]

In principle, semantic phasors could give more fine-grained control: by rotating, filtering, or projecting semantic components, so we can aim to steer content away from prohibited regions while preserving useful parts. Turning this into a real system would still require careful policy models and testing, but the spectral view offers a more precise ‘control panel’ than simple keyword filters. This makes it possible to selectively remove or reshape harmful or non-compliant elements—for example, stripping out verbal acts of incitement or bias, while preserving factual description—instead of relying on crude keyword blocking or overly broad refusals.[2][1]

7. Conclusion

Semantic Phasor Theory offers a simple but powerful way to think about meaning: not as fixed points, but as waves that shift and interact with context. By borrowing familiar ideas from signal processing, it gives us a clearer picture of how language carries nuance, ambiguity, and intent—and how future AI systems might handle meaning with more stability, transparency, and control.

For readers who want to explore the full technical details, the two main Semantic Phasor Theory papers are available online and describe both the unified spectral model of meaning and the formal operator calculus that underlies it.[2][1]

[1] Semantic Phasor Theory: A Unified Spectral Model of Meaning and ...
https://zenodo.org/records/18965947

[2] The Operator Calculus of Semantic Phasors: Formal Foundations ...
https://zenodo.org/records/19055628