Between Fodor and Sellars -- A Middle Ground for Language-like Neural Representations

by

Hanzhe Dong

BA in Philosophy & BS in Media Science – Undergrad, Boston University, 2022

A Thesis

Submitted to The Graduate School of the
University of Missouri-St. Louis
in partial fulfillment of the requirements for the degree

Master of Arts
In

Philosophy

May 2024

<u>Advisory Committee</u>

Gualtiero Piccinini, Ph.D.
Chairperson

Lauren Olin, Ph.D.

Jon McGinnis, Ph.D.

Between Fodor and Sellars -- A Middle Ground for Language-like Neural Representations

Abstract

The recent resurgence of the language of thought (LOT) hypothesis has drawn much attention. The history of philosophy and cognitive science has provided us with various versions of LOT. From Sellars to Fodor, theorists have offered us considerations on the vehicles, content, and functions of such a representation format. However, it's been more than 50 years since Fodor's publication on LOT (1975), and the resurgence suggests that we need a modern iteration of LOT to fit with recent developments in the study on cognition. Therefore, in this paper, I argue that we need a middle ground between Fodor and Sellars's views on LOT – specifically, we should abandon Fodor's view on the vehicle and content of LOT, and we should also acknowledge that perception might involve LOT (contrary to what Sellars thought). What we should focus on is LOT representations' format's function, that at least some neural processes are analogous to linguistic structures, and how such a function is able to facilitate a modern situational view of cognitive representation. Having this account of LOT, we can make it a more versatile tool fitting with the other cognitive frameworks, such as the free energy principle and Bayesian inference.

Keywords: Language of Thought (LOT), Jerry Fodor, Wilfrid Sellars, Neural representations, Representation formats, Symbolic representation, Neural computation

1. Introduction

In the study of the mind, the Language of Thought (LOT) emerges as an important framework, offering a unique lens through which we can understand cognitive processes. LOT, famously advocated by Jerry Fodor (1975) and later refined by various scholars, posits that thinking is structured in a language-like symbolic system. It suggests that our cognitive processes are akin to linguistic episodes, where mental "words" form the building blocks of more complex mental "sentences". The essence of this hypothesis lies in the belief that each mental sentence derives its meaning from the meanings of its constituent parts and the way in which these parts are assembled.

Fodor, in his works of 1975 and 1987, advanced the idea that these mental representations are processed through mental computations, akin to a digital computer. This notion, deeply influenced by advancements in digital computation and the Turing machine model, implies that the brain functions in a manner similar to programmable processors with memory registers. However, as neuroscience has developed, evidence has emerged suggesting that the digital computation model might not fully encapsulate the complexities of mental computations (Piccinini, 2020, 2023). This revelation, however, does not entirely discredit the LOT but rather invites a reevaluation of its key features, such as symbolic representation, compositionality, and the rule-governed manipulation of symbols.

On the other hand, Wilfrid Sellars' perspective offers an alternative account of LOT, proposing that while certain neural processes parallel linguistic structures, they do not necessarily imply a digital computation process (Sellars, 1963, 1975). This divergence in viewpoints sets the stage for a reexamination of the LOT framework, as well as its application to understanding thinking and other cognitive processes like perception.

Recent contributions to the field, such as the work of Quilty-Dunn et al. (2022), have further elaborated on the LOT, outlining features like discrete constituents, role-filler independence, predicate-argument structure, logical operators, and inferential promiscuity. These features underscore the hypothesis' potential applicability in contemporary cognitive studies.

This paper seeks to navigate a "middle ground" between the Fodorian and Sellarsian interpretations of LOT. By dissecting the framework into three components – content, vehicle, and function – I aim to reconcile the differences between Fodor's digital computation model and Sellars' more permissive approach. I will argue that while the discrete constituents of LOT may appear digital, this does not necessitate a commitment to digital representations and computations. Instead, this paper posits that a holistic, situational framework, which considers the integrative and contextual processing of the brain, aligns more closely with the dynamic nature of neural representations.

Moreover, the paper will challenge the monolithic view of digital computation in the brain, advocating for a pluralistic approach that encompasses both digital and analog computational elements. This perspective aligns with recent neurolinguistic research (Pylkkänen, 2019) and clinical evidence from patients with brain damage, which collectively offer insights into how the brain processes logical operations and syntactic and semantic aspects of language.

Finally, the paper also addresses the scope of LOT, particularly its applicability to cognitive processes beyond thinking, such as perception. While Fodor posits that perception benefits from a LOT framework, Sellars provides a contrasting non-LOT-based account of perception. The recent resurgence of LOT (Quilty-Dunn et al., 2022), while credits LOT as a useful too, doesn't address the divergence of LOT. For example, Piccinini (2020) argues that the empirical evidence suggests that the neural computation is not entirely digital. Therefore, the divergence highlights the need for a more inclusive approach that both recognizes the linguistic nature of perceptual representations to some degree and address the controversy on the understanding of LOT.

2. Language of Thought – An Overview

In this section, I will provide an overview of the characteristics of LOT, in which I will outline Fodor's contribution to this framework and how Sellars' stance is different from his, distinguishing between a "strong" (Fodorian) and a "weak" (Sellarsian) version of the theory.

LOT proposes that thinking is structured in a language-like symbolic system. In other words, we use "mental language" (mentalese) to think, and the thinking process involves manipulating a language-like code consisting of mental representations or symbols. The reason we call it a "language" is that mental processes are similar to linguistic episodes: there are mental "words" that constitute mental "sentences," and the meaning of each mental "sentence" depends on the meanings of its constituents (mental words) and how those constituents are combined. Historically, it has been rooted in the works of early thinkers like Augustine, Boethius, Thomas Aquinas, and John Duns Scotus. A comprehensive and systematic approach to LOT was first developed by William of Ockham in his seminal work, "Summa Logicae" (circa 1323), where he analyzed the intricacies of the structure and meaning of Mentalese expressions. LOT enjoyed considerable popularity during the late medieval period, but its prominence waned in the sixteenth and seventeenth centuries, leading to a period where it was largely absent from significant mind-related theorizing up until the mid-twentieth century. The 1970s marked a significant resurgence of interest in LOT, predominantly triggered by Jerry Fodor's book "The Language of Thought" in 1975. This book was influential in rekindling interest in LOT, stimulating a wave of discussions, both in support of and against the hypothesis (Rescorla, 2019).

LOT is anchored in the Representational Theory of Mind, which argues that mental states (beliefs, desires, intentions, etc.) are propositional attitudes. These attitudes are relational, connecting an individual to a proposition with a truth value, and are closely linked to mental representations, which are the vehicles of these propositional attitudes.

Fodor, as a key proponent of LOT, emphasizes the role of mental representations in propositional attitudes (Fodor 1981, 1987). He suggests that these mental states inherit semantic properties, including truth conditions, from the mental representations they are related to. Fodor’s approach was predominantly abductive, arguing that the most effective scientific theories about psychological processes imply the existence of Mentalese, thus giving us solid grounds to believe in its existence. He introduces functionalism to analyze psychological relations. In his view, the psychological relation between an individual and a mental representation is characterized by functional roles, which are integral to understanding cognitive processes like reasoning and decision-making.

Moreover, LOT underscores the compositional nature of mental representations. As argued by Fodor (2008), mental representations also have the feature of compositionality, that is, the meaning of complex mental representations is determined by the meaning of their constituent parts and their arrangement, paralleling the structure of natural languages. Fodor extends this to include the logical structure within Mentalese, proposing that mental representations have logical forms similar to natural languages, encompassing elements like logical connectives and structured sentences.

LOT also intersects with the Computational Theory of Mind (CTM), which conceptualizes the mind as a computational system, processing mental representations similarly to Turing-style computations. Fodor thinks that the processing steps and vehicles of the thinking process are digital, and he argues that mental processes are executions of programs stored in memory registers similar to modern digital computers. Fodor (1975, 1981) integrates this by advocating a formal-syntactic conception of computation, wherein mental computations manipulate symbols based on their syntactic properties, contributing to the semantic coherence in cognition.

Here, we can roughly conclude that LOT (especially Fodorian) at least supports the idea that 1) thoughts are analogous to sentences in a language in terms of their structure, syntax, and semantics, and 2) thinking is computational.

What is more controversial is Fodor's argument on the content and vehicle of language-like representation formats. For example, neuroscience seems to offer empirical evidence against the thesis that mental computations are entirely digital (Piccinini, 2020, Ch. 13, 2023).

On the other hand, Sellars' version of LOT seems to be more compatible with modern concerns. He argued that psychological concepts are akin to theoretical concepts in the sciences, suggesting a naturalistic view of mental states. His view that our concepts of intentionality are derived from concepts pertaining to linguistic behavior indicates a complex relation between thought and language. He suggests that thoughts are postulated to explain complex patterns of reasonable behavior, indicating a belief in representational states but with a significant emphasis on their functional role in a larger behavioral and linguistic context (Sellars, 1997, 1957).

In terms of mental representations, Sellars viewed mentalistic concepts as compatible with naturalism, like scientific concepts. This perspective would likely align with some aspects of the LOT, particularly the idea that mental processes can be understood in representational terms. However, Sellars' emphasis on the normative aspects and the functional role of mental states suggests a broader view than a straightforward internalist or computational perspective on thought. While acknowledging that at least some cognitive processes are linguistic, Sellars argued that for public language, linguistic utterances are typically the product of internal thinking activity, which they express, and that language is prior to thought in the order of knowing. This would imply that while he might see a role for a 'mental language', it is not primary but rather a secondary, expressive mechanism for pre-existing cognitive processes. Yet, this would show that at least some cognitive processes are linguistic, similar to Fodor (Sellars, 1997, 1975, 1981).

In light of Sellars' idea on the relationship between thinking and language, we can also see that he has a different opinion on the scope of LOT. Sellars' treatment of sensory states as theoretical constructs introduced to explain perceptual phenomena suggests a complex view of mental representation. He acknowledges that non-linguistic thought exists, such as visual perception and imagination, indicating a recognition of diverse forms of mental representation beyond just a language-like structure (Sellars, 1975,1981).

In short, by endorsing the analogy between thinking and language (Sellars, 1963), Sellars offers another reasonable account of LOT. He suggests that certain neural processes parallel linguistic structures, without necessarily implying that mental computations are digital (Sellars, 1975). On the scope of LOT, Sellars has a non-LOT-based account of perception (Reider, 2012; deVries, 2021). Thus, while both Fodorian and Sellarsian versions of LOT agree that thinking is analogous to language to some extent, they disagree on the type of computation (whether it's digital or not) and the possible application of LOT (whether it can be helpful in accounting for other cognitive processes such as perception)

Quilty-Dunn et al. offer a new summary of the features of the language of thought hypothesis to show its applicability in modern studies of the mind (Quilty-Dunn et al., 2022):

1) Discrete constituents: Mental representations are made up of distinct elements that correspond to individuals and their separable characteristics.

2) Role-filler independence: The architecture of LOT combines components in a way that there is no interdependence between the constituents that fill the syntactic roles and the roles themselves.

3) Predicate-argument structure: Predications are applied to an argument to yield a truth-evaluable structure.

4) Logical operators: LOT architecture also uses logical symbols.

5) Inferential promiscuity: LOT representations are usable for inference in a way that is automatic and independent of natural language.

In the next section, I will argue for a "middle ground" between Fodorian and Sellarsian LOT with respect to the features of LOT that Quilty-Dunn et al. offer.

3. A Middle Ground

As summarized above, though theorists arguing for LOT agree that some neural processes parallel linguistic structures, they differ on many nuances of the implications of LOT. For example, as described in the last section, there is a debate on the types of computation that LOT performs. As modern sciences develop, what seems to be reasonable 50 years ago can be obsolete now. Yet, some core ideas of LOT are still useful for understanding mental phenomena and cognitive processes, as Quilty-Dunn et al. (2022) suggest. So, we should come up with a modern iteration of LOT that fits the needs of modern science (Coelho Mollo & Vernazzani, 2023). Here, I will argue for a middle ground between Fodor and Sellars's versions of LOT to meet such a need. I will dissect the argument into three parts: the content (the general ontology of LOT), vehicle (types of computation), and function (essential features of LOT supported by neuroscientific evidence) of language-like representation format.

3.1 Content of LOT

Content-wise, I want to argue for a shift in the “symbolic” (atomistic) and “situational” (holistic) cognitive paradigm (Vilarroya, 2017).

As Villarroya (2017) suggests, the “symbolic” framework refers to the traditional idea from the Representational Theory of Mind that focuses on how a single element of the environment is encoded in the brain, and the element is represented by patterns of neural activity that respond to such external stimuli, information, or states. This also echoes the Fodorian LOT that stresses symbolic representation and the manipulation of such symbols. In other words, the Fodorian LOT stands on the idea that our minds represent the environment in a symbolic way, and the representations, and symbols, have a mapping (1-to-1) relationship with the environment such that each symbol represents an attribute of an external object. I will call this framework atomistic as it presumes certain neural processes are specific for a certain function.

In contrast, the “situational” framework focuses on how the brain processes information in the context of the specific situation in which the individual is involved. I will call this framework holistic as it encompasses the situations where one pattern of neural activity might respond to various external stimuli. Studies on place cells have given us a good reason to conclude that this framework is needed.

Place cells, primarily found in the hippocampus, are known for their role in encoding spatial representation. However, research (Schiller et al., 2015) shows that place cells are sensitive to a wider range of cues, both spatial and non-spatial. This includes responses to odors, sounds, tactile inputs, timing, rewards, and more. This sensitivity suggests that place cells are involved in processing a richer array of information than just spatial location, supporting a situational rather than a strictly spatial or representational framework.

So, place cells change their activity patterns in response to changes in context, indicating their role in encoding more than just spatial information. The context here is understood as a coherent set of features occurring within a particular environment, including environmental, cognitive, and motivational factors. This aligns with the situational framework, where the hippocampus is seen as processing complex situations rather than just mapping the physical environment.

Other studies (Zucker and Ranganath, 2015) have shown that hippocampal networks exhibit sequential organization in their neuronal firing patterns, suggesting a role in encoding sequences of events. This capability for sequential processing and prediction further supports the idea that the hippocampus is involved in managing and interpreting complex, situational information, rather than merely representing discrete environmental stimuli. Here, the traditional representational model, which assumes that neurons in the hippocampus encode simple spatial representations (like a map), is challenged by the evidence of complex firing patterns in response to a variety of stimuli and contexts.

More generally, from the perspective of tuning curves, we observe a notable complexity in deciphering the relationship between neural representation and external stimuli (Brette, 2021). Tuning curves refer to the chart that shows how a neuron's response (typically recorded in the number of action potentials or spikes per unit time) varies with a continuous stimulus characteristic like orientation, wavelength, or frequency. When a neuron shows its highest response to a specific stimulus, it is considered "tuned" to that stimulus. The curve's breadth at half its maximum response on either side of its peak reveals the range of tuning of a neuron to a particular stimulus characteristic. In the context of the auditory system, this graph is used to assess how selectively a neuron responds to different frequencies. For instance, in auditory nerve fiber recordings, the threshold is often identified as a consistent increase in firing rate in response to a single-frequency sound.

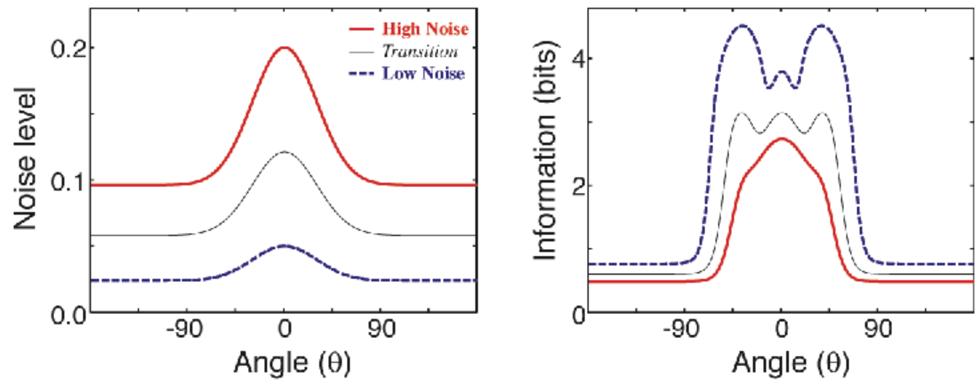

(Butts&Goldman, 2006)

A neuron's response to a specific stimulus, like the color encoded by its firing rate, is not absolute but context-dependent. A change in an external factor, such as light intensity, alters the neuron's tuning curve, thereby affecting its response to the same stimulus. This demonstrates that interpreting a neuron's firing rate to infer a specific property, like color, is fraught with uncertainty and ambiguity. It's not just the firing rate, but also the context or additional dimensions, that determine the interpretation of the neural signal.

This intricacy shows that the traditional atomistic presumption of neural representations, the simplistic, one-to-one correspondence as Fodor suggests, can have trouble staying compatible with the evidence from cognitive science studies. In turn, this highlights why a holistic approach, considering the integrative and contextual processing of the brain, aligns more accurately with the dynamic nature of neural representations. The above examples should show that the situational framework should be a better fit with neuroscience studies on the brain. So, this would be the first reason to argue against a Fodorian LOT in terms of the idea of symbolic one-to-one representation.

3.2 Vehicle of LOT

Vehicle-wise, it has always been controversial what kind of computation is performed in the brain. The idea of Computational Theory of Cognition (CTC) began with McCulloch and Pitts (1943), who proposed neural activity as computation similar to Turing's concept. This computation was thought to be digital but was countered with the notion of it being analog. While classicists view cognitive computation as digital, connectionists and computational neuroscientists see it as neural network activity, which might not be purely digital. In the understanding of LOT, Fodor argues that the neural computation should be digital while some neural scientists argue that at least some kind of information, such as magnitude (Peacocke, 2019), should be analog. So, the monolithic view from Fodorian LOT should be abandoned, and I also see a possible ground for LOT to accommodate both kinds of computation.

The concept of digital computation, as applied by Fodor in the context of LOT, is analogous to the operations of a digital computer. This is again related to the Representational Theory of Mind. Under the framework of Fodorian LOT, these mental representations are manipulated through mental processes that are akin to digital computations in computers. This includes the manipulation of symbols based on their syntactic properties, without regard for their actual content, akin to how a computer processes data. In Fodor's view, mental processes are like executions of programs, where mental representations are processed in a step-by-step, discrete, and rule-governed manner, similar to how a digital computer operates. This implies a clear, structured, and binary manner of processing information. There is neuroscientific evidence that supports this view. For example, neurons operate on an all-or-none principle, like digital signals (Pareti, 2007). When a neuron fires, it sends an action potential (a spike) down its axon, which is a discrete event. In digital systems, information is processed in discrete units (bits), which are either in an 'on' (1) or 'off' (0) state. The All-or-None principle mirrors this binary operation, where a neuron's firing (action potential) is analogous to the 'on' state, and its resting state to the 'off' state. This resemblance supports Fodor's view of mental processes as digital-like operations. Moreover, Fodor argues that mental processes involve the manipulation of internal symbols. In the context of the All-or-None principle, the action potentials of neurons can be seen as the physical basis for these symbols. Just as digital computers use binary codes to represent and manipulate complex information, the brain could use the firing patterns of neurons to represent and process symbolic mental information. In short, the digital kind of LOT makes sense as there is existing evidence supporting the idea of a digital computing mind.

In contrast, analog computation in the context of neural processing refers to a more continuous, less discrete, and often non-symbolic form of information processing. Unlike digital computation's emphasis on discrete symbols and step-by-step manipulation, analog computation involves processing that is more akin to continuous variables and can represent and process information in a way that mirrors the continuous nature of the real world. In analog models, information is processed in a more fluid, holistic manner, often relying on the strength of connections and the flow of activations rather than on discrete symbolic manipulation. Unlike action potentials, some neural signals are graded. These are changes in membrane potential that vary in magnitude and are not all-or-none. This graded response can be analog, as it represents a continuum of values, much like analog signals in electronics. As argued by Peacock (2019), certain types of information, like sensory data (e.g., visual and auditory stimuli), are encoded in the brain in a continuous manner. For instance, the encoding of visual information in the retina and processing in the visual cortex involves analog-like computations to handle the continuous nature of visual scenes.

Here, we can see that there are both the digital and analog aspects of neural computations, and it remains to be a debate on the ontological characterization of the kinds of computation that the brain performs.

Mayley (2023) proposes one way to re-think about the characterization of neural computation, arguing that the common belief that analog computation is solely about continuity is incorrect. The distinction between analog and digital computation is rooted in the types of representation they involve. Analog computation essentially uses analog representation, while digital computation is based on digital representation. For analog representation, Maley characterizes it based on the Lewis-Maley account, which argues for a model where analog representations can exist both as continuous and discrete forms. The core of this account is the principle of monotonic covariation, which suggests a consistent and predictable correspondence between changes in a representational property (such as voltage, rotation, or physical displacement) and the represented quantity. This relationship is governed by a homomorphic function, ensuring that while the representation need not be an exact mirror of the represented quantity, it must maintain a systematic and structured correlation. Additionally, the Lewis-Maley account introduces the concept of a resolution factor in discrete systems, recognizing that minor variations or 'jitter' do not necessarily impact the represented quantity, thus distinguishing significant representational changes from trivial fluctuations. By expanding the scope of analog representation to include both continuity and discreteness, the Lewis-Maley account offers a realistic framework that aligns closely with the complexities of computational and cognitive processes. Thus, analog computation is described as the mechanistic manipulation of analog representations. Maley's approach aligns with general computational theories but emphasizes the role of representation types.

With Maley's account, let's consider Rate Coding: it is the principle that the information a neuron conveys is in the frequency of its firing. It suggests that while the action potential itself is a digital-like event, the frequency (rate) at which these events occur can represent a continuum of values, much like an analog signal (Enoka & Duchateau, 2017). Thus, Rate Coding incorporates both digital and analog aspects. The digital aspect is that each individual spike in a neuron is a discrete, all-or-none event, aligning with digital processing, and the analog aspect is that the rate of these spikes over time can represent a range of values, not just binary states. This is akin to analog signals, where the signal can vary continuously and represent a spectrum of values. In the context of LOT, this suggests a more nuanced view where neural representations (the representational aspects of neuronal signals. those are, primarily, the frequency and sometimes timing of spikes) might not be strictly digital.

On the other hand, Piccinini (2020) reconstructs the ontological debate by arguing that the existing evidence showing the duality of both digital and analog computation should lead us to understand neural computation as sui generis. As mentioned above, neural signals, particularly spikes, are not continuous like analog signals. Their most crucial aspect is their all-or-none nature, contrasting with the continuous variation required in analog computation. On the other hand, while spikes may have a discrete appearance, they lack the precise and deterministic character of digital events. The inherent variability and probabilistic nature of spike occurrences challenge their classification as digits. Furthermore, the physiological significance of spikes often does not correspond to individual spikes but rather to patterns of activity, such as firing rates, which do not conform to digital computation's demand for discrete, countable events. Piccinini extends his argument by addressing the concept of spike trains as strings of digits. He contends that for spike trains to be considered strings, there would need to be an unambiguous way to determine the start and end of a string and to assign spikes to specific positions within it. However, due to the stochastic nature of neuronal firing and the often-noisy background activity, delineating precise spike train boundaries becomes an impractical task. Moreover, the functional significance in neural processing frequently lies in the collective patterns formed by multiple spike trains rather than in any individual spike train.

Piccinini (2020) therefore argues that neural computation is sui generis. He points out that the current understanding of neural mechanisms and processes is rooted in sophisticated mathematical models that do not rely on digital or analog computational theories. These models, reflective of the intricate dynamics of the brain, encompass aspects like spiking activity, graded changes, and network interactions, which cannot be accurately captured by traditional computational paradigms. These realistic mathematical models, which are empirically testable and grounded in experimental neuroscience, have shown that neural processes cannot be fully explained by digital computability theory or computer design. The implication of this argument suggests a shift in perspective, where cognitive scientists must align their theories more closely with the unique nature of neural computation, rather than forcing them into the existing frameworks of digital or analog computation. This alignment, Piccinini posits, should lead to the development of new mathematical tools and theories specifically tailored to the distinctive properties of neural computation.

While Maley and Piccinini approach the debate on the vehicle of neural computation from distinct paradigms and methodologies, their work converges in a critical challenge to the monolithic digital view of Fodorian LOT. Maley, through his examination of analog representation, broadens the perspective on computation to include both continuous and discrete forms. This perspective not only accommodates a more nuanced understanding of neural computation but also aligns with evidences in neuroscience, such as Rate Coding, which embody a blend of digital and analog elements. On the other hand, Piccinini's argument for neural computation as sui generis directly confronts the simplifications inherent in the digital-only approach of Fodorian LOT. He underscores the unique characteristics of neural processes that elude strict classification as either digital or analog, advocating for a paradigm shift towards a new computational framework specifically designed for neural phenomena. Together, Maley's expansion of the analog realm and Piccinini's sui generis classification of neural computation offer robust evidence against the restrictive digital bias of Fodorian LOT. Their contributions urge a rethinking in cognitive science, advocating for a model that accommodates the intricate, multifaceted nature of neural computation, rather than confining it to a digital framework.

3.3 Function of LOT

Function-wise, there are at least three essential features of LOT that need to be addressed before we can call this kind of representation format “language-like.” The three features are discrete constituents, predicate-argument structures, and logical operators. I will explain each sequentially.

3.3.1 Discrete Constituents

First, let’s start with discrete constituents. As characterized earlier, it refers to the idea that representations are made up of distinct elements that correspond to individuals and their separable characteristics. This echoes the traditional characterization of LOT that it has systematicity of thought, which enables “mental languages” to have compositionality. For example, MARY, LOVE, and JOHN can be combined into "MARY LOVES JOHN," and they can also be combined as "JOHN LOVES MARY." In this way, each constituent maintains its identity independently of its role. The “discrete” part seems to imply a sense of digital computation as it involves discrete elements used to perform computation. However, if we limit the role of discrete constituents to the scope of digital computation, we are quickly entangled with the debate on the types of computation that our brains are performing. This would also face the earlier criticism about the vehicle of Fodorian LOT. So, we should also adopt a pluralistic understanding of discrete constituents. By doing this, we can avoid controversies on the types of neural computation. Though features like Discrete Constituents sound digital, this does not necessarily entail a commitment to digital representations and computations. Rather, we should keep it simple: when LOT theorists mention discrete constituents, they refer to separable, distinct elements that mental representations are composed of. So, the "constituents" can be both digital and analog. What matters is that they allow for flexible and combinatorial mental operations of these constituents, irrespective of their specific informational content or format. This would also benefit from the earlier discussion on the vehicle of LOT as it would embrace a more explanatory powerful and versatile account of neural computation.

3.3.2 Predicate-Argument Structure

Now, let's move to the predicate-argument structure. This linguistic structure is based on the functions of lexical items. The functions determine the roles to be played by the other words in the sentence. A predicate specifies a relationship between objects or a state that characterizes an object. For example, [SEE (BOY, DOG)] means that "the boy sees the dog." Arguments refer to real-world objects about which something is predicated. So, predictions are applied to an argument to yield a truth-evaluable structure.

The study of the neural basis of predicate-argument structure is an ongoing topic, and there is neuroscientific evidence to support its existence and computational mechanism. As early as 2003, Hurford argued that the neural basis of predicate-argument structure in language can be traced back to fundamental brain processes used in both linguistic and nonlinguistic tasks. Specifically, he highlights how the brain integrates sensory information through distinct neural pathways for locating objects (the "where" pathways) and identifying their attributes (the "what" pathways), paralleling the logical structure of PREDICATE(x) in language. This dual processing involves deictic variables, which are mental pointers to objects in the environment, and these variables are integrated with semantic information about the objects, akin to the predicates in language. The ability to mentally represent scenes and objects, a capability found in primates and foundational for practical tasks, is proposed as a pre-adaptive basis for the development of complex linguistic functions, including the description of scenes and events. This evolutionary perspective implies that the neural underpinnings of predicate-argument structures in language have their roots in more basic sensory and perceptual systems of the brain, which are shared across primates and are integral to interaction with the physical world.

In more recent years, this study has progressed with the emerging need for natural language processing (NLP) and its implications for artificial intelligence. In the realm of cognitive neuroscience and computational linguistics, recent explorations into the neural basis of language processing present compelling parallels with the architectural design of Predictive Language Models (PLMs). These parallels, particularly in the context of predicate-argument structures (PASs), offer profound insights into the fundamental underpinnings of language processing both in the human brain and artificial intelligence systems.

Pylkkänen (2019) provides an insightful exploration into the neurobiological explanation of how humans process and understand language, specifically focusing on the interplay of syntax and semantics.

One of the key findings discussed in his paper is the involvement of specific brain regions in language comprehension and production. The left anterior temporal lobe (LATL) and the ventromedial prefrontal cortex (vmPFC) are identified as crucial areas in this process. Pylkkänen highlights that these regions are engaged in both the comprehension and production of language and are active in processing both spoken and signed languages, which suggests a shared neurobiological basis for different forms of language. The LATL, in particular, is emphasized as playing a significant role in the rapid combination of concepts during language processing. This region shows increased activity when processing two-word phrases, suggesting that it is sensitive to the conceptual integration (which shares the feature of a predicate-argument structure of words rather than their syntactic arrangement. On the other hand, vmPFC appears to be involved in a later stage of language comprehension and may contribute to integrating the composed meanings into a broader context, such as social cognition and episodic memory.

Moreover, the paper discusses the way that our brains automatically and instinctively compose meanings from individual words, a process that involves complex neural activities. It shows how we discern different meanings from similar sentence structures based on syntax, using examples like “Sally baked the black beans” versus “Sally baked the beans black.” These examples illustrate how syntax guides our interpretation of sentences, influencing whether we perceive 'black' as a descriptor of the beans or as a result of the baking process. The brain's capacity to handle these structures in both language comprehension and production mirrors the LOT's proposition of a mental syntax, providing a neurobiological basis for the theory.

On the side of computational linguistics, Conia and Navigli (2022) advanced the research on pre-trained language models (PLMs), a compositional language model that leverages Predicate-Argument Structures (PASs) to enhance the understanding of complex language structures, particularly in the context of natural language processing (NLP). The model is designed to learn word representations and their compositional functions using contexts based on both bag-of-words and dependencies, differing from traditional word-sequence-based models.

Traditional NLP models often struggle with long-range dependencies in language, such as verb-object and subject-verb-object relations. The PAS-based model addresses this by composing arguments into predicates using category information from the PASs. By using predicate-argument structures, the model can capture the syntactic roles of words and their semantic relationships simultaneously.

The shift from individual word representations to phrase representations is significant in understanding language processing. This model's ability to learn meaningful representations for various dependencies (like adjective-noun, noun-noun, and verb-object) demonstrates an advanced understanding of how different elements in a sentence interact to convey meaning. This is particularly evident in the model's performance on phrase similarity tasks, where it excels in capturing the semantic essence of phrases.

The model's performance on semantic similarity tasks (like adjective-noun, noun-noun, and verb-object phrase similarities) indicates a significant advancement in how machines understand and process language. By achieving state-of-the-art results in these areas, particularly with a smaller training corpus and without pre-trained word vectors, the model demonstrates its efficiency and effectiveness in language comprehension. The success of these models in capturing the essence of complex language structures and their meanings aligns with the LOT's view of mental processes as inherently structured and rule-governed.

By putting both studies together, we can see that the convergence of findings from neurolinguistic research and the operational mechanisms of PLMs not only enriches our understanding of linguistic cognition but also paves the way for philosophical inquiries into the modern understanding of LOT.

The neurolinguistic research led by Pylkkänen (2019) shows the neurobiological roots of language comprehension and production, identifying key regions such as the left anterior temporal lobe (LATL) and the ventromedial prefrontal cortex (vmPFC). These areas are instrumental in semantic composition, akin to the components in PLMs that integrate and interpret the meanings of words in various contexts. This resemblance suggests a shared cognitive architecture between human neural processing and algorithmic models in understanding language. Furthermore, the role of the LATL in conceptual combination mirrors how PLMs utilize contextual information to ascertain meanings, underscoring the pivotal role of context in both human and artificial language processing.

The other study (Conia and Navigli, 2022) shifts the focus to a compositional language model that utilizes PASs, reflecting a similar approach to language processing as observed in human neural mechanisms, specifically within areas like Broca's and Wernicke's. The model's emphasis on integrating arguments into predicates echoes the brain's natural propensity to organize language elements hierarchically and relationally. In human cognition, Broca's area is implicated in syntactic processing, while Wernicke's area handles semantic understanding. This functional distinction finds a parallel in PLMs, where complex language structures and dependencies are processed in a manner akin to the brain's handling of syntax and semantics.

The alignment between the processing strategies of PLMs and the neural activity observed in language comprehension and production suggests a fundamental role for PASs in mental representations of language. The ability of the human brain to interpret different meanings from similar sentence structures based on syntax and context, as highlighted in the first paper, is mirrored in the capability of PLMs, as demonstrated in the second paper, to represent and comprehend complex language constructs.

The synthesis of neurolinguistic findings and advancements in computational linguistics provides compelling evidence in support of LOT, especially in terms of predicate-argument structures. The studies by Pylkkänen and Conia & Navigli collectively underscore the neural and computational mechanisms that parallel LOT's conceptual framework.

3.3.3 Logic Operators[1]

With previous discussions of discrete constituents and predicate-argument structure, we can see there is a key element of LOT that makes both possible, that is, logical operators. The general idea of LOT is that thoughts are structured much like sentences in a language, composed of elements akin to words and governed by rules analogous to grammar. So, logical operators (like AND, OR, NOT, IF…THEN) are essential for the structure and rules of Mentalese. They allow for the formation of complex thoughts by combining simpler elements. Just as grammar in natural language enables the construction of meaningful sentences from words, logical operators in Mentalese enable the construction of coherent and logically structured thoughts. Therefore, logical operators are crucial for both the systematicity and productivity of LOT, as they allow for the combination and recombination of mental symbols in structured and meaningful ways. On the other hand, logical operators are essential to complex reasoning processes. They allow for the formation of conditional

[1] As Jon McGinnis suggests during the thesis defense (2024) of this paper, there is a place to further discover which kind of logic (Stoic or Aristotelian) can be applied to neuroscientific studies.

statements, hypotheses, deductions, and inductions. When logical operators are used with predicate-argument structures, they help form complex propositions. These complex propositions can express relationships between different events, conditions, or states of affairs. For example, using "or" in "Alice reads a book or Bob watches a movie" suggests an alternative between two scenarios.

Having a modern iteration of LOT, we need to have a neuroscientific basis for the logical operation in the brain. Luckily, modern science has provided us with a lot evidence on this issue. As studied, logical reasoning relies on a varied network of brain regions, which is task-dependent. This is demonstrated through functional neuroimaging, brain lesion, and behavioral studies (Goel et al., 2017; Pamplona et al., 2015). For instance, different neural systems contribute to semantic bias and conflict detection, and the neural bases of logical syllogisms are influenced by the emotional context of tasks. While the neural network of the brain is heterogeneous, certain areas like the Inferior Parietal Lobule (IPL) are repeatedly shown to be crucial in various aspects of reasoning (Arora et al., 2015). Enhanced activity in the parietal cortex, for example, may compensate for reasoning deficits in sub-clinically depressed participants. While these studies indicate a complex interplay of different brain areas in processing logical reasoning, they do not explicitly show that the brain uses representations of logical operators. This evidence, though indirect, is crucial as it sets the foundation for understanding the brain's capacity for complex logical processes. Here, enhanced activities in regions like the Inferior Parietal Lobule (IPL) during reasoning tasks highlight the brain's engagement in sophisticated cognitive functions, which could underpin the use of logical operators.

Cognitive neuroscience research has inspired new computational theories of how individuals segment perceptual information into event representations. For example, the research by Khemlani et al. (2015) illustrates how neuroscientific findings can inform a novel computational theory of event segmentation, which is essential in logical reasoning and aligns with the LOT's perspective on the computational nature of the brain.

The direct evidence that bridges this theoretical gap comes from the study (Salto et al., 2021) examining the differences in brain activity between logically valid and invalid deductions. Specifically, this study reveals distinct patterns of neural activity in the medial left prefrontal cortex and the anterior cingulate cortex (ACC) when participants engage in logically valid reasoning tasks. The distinction in brain activity, particularly the reduced intensity and delayed response in valid conditions, aligns with the idea of neural representation (the encoding of information through specific patterns of neuronal signals, including the frequency and timing of spikes). When engaging in tasks that require the manipulation of logical operators, the brain appears to exhibit a more measured and structured pattern of activity. This observation implies that the representation of logical operators involves a deliberate arrangement and sequencing of neuronal firing patterns. The recursive and computational nature of these processes, highlighted in the study, supports the implication of logical operators in LOT that the brain operates with representations of logical operators.

These findings provide direct neuroscientific support for the claim that the brain does indeed use representations of logical operators, as proposed in the Fodorian version of LOT and expanded upon by Quilty-Dunn et al (2022). This evidence substantiates the hypothesis that the brain's cognitive processing involves a 'Mentalese,' where logical operators are explicitly represented and manipulated within a structured, language-like system.

So far, I have examined the functional account of LOT through features like discrete constituents, predicate-argument structures, and logical operators, showing the ways in which our cognitive processes parallel linguistic structures and fit with modern scientific studies. We have seen how discrete constituents provide the fundamental elements for representation, allowing for a flexible and dynamic composition of ideas. The predicate-argument structure, supported by neurolinguistic research and advances in computational linguistics, reinforces the notion that our cognitive processes are deeply analogous to linguistic syntax and semantics. Furthermore, the essential role of logical operators in facilitating complex thought formation and reasoning, supported by cognitive neuroscience findings, solidifies the conceptual framework of LOT as not only structurally but also functionally akin to a language.

4. Scope of LOT - Perception

As we transition to discussing the scope of LOT, particularly its relevance to perception, it becomes evident that the language-like structure of cognition may also provide insights into how we perceive and interpret the world around us. The same principles of systematicity, compositionality, and logical structuring that govern thought could potentially illuminate the neural mechanisms underlying perceptual processes. Thus, extending the framework of LOT to encompass perception offers a promising avenue for a more unified and comprehensive understanding of the mind's capabilities, bridging the gap between internal cognitive mechanisms and our interaction with the external world.

First, let’s start with the border between perception and cognition. As mentioned in the second section, LOT argues and deals with the propositional and conceptual representations used in cognition. Fodor also argues that the psychological relation between an individual and a mental representation is characterized by functional roles, which are integral to understanding cognitive processes like reasoning and decision-making. Sellars also seems to agree with Fodor that LOT can be used to explain cognitive processes to some extent. Yet, Sellars disagrees with Fodor by arguing that LOT is too limited to help us understand perception. Sellars’ argument makes perception and cognition two distinct areas to understand the brain. Indeed, traditionally, perception and cognition have been viewed as distinctly separate processes – perception being the initial sensory processing of the environment, and cognition being the higher-level processing involving thought, reasoning, and memory. However, as recent studies suggest (Tacca, 2011; Clarke and Beck, 2023), the perception-cognition border might be thinner than Sellars expected.

As Clarke and Beck (2023) argued, from an eliminativism perspective, which rejects the traditional scientific ontology that separates perception and cognition, controversial cases and theories like predictive processing suggest deep interactions between perception and cognition that challenge their separation. They further examine five approaches to defining this border: phenomenology (the subjective experience of mental states), revisability (the ability to change mental states through will or evidence), modularity (the idea that perceptual processes operate independently from cognitive processes), format (the structure of mental representations, like pictorial or sentence-like), and stimulus-dependence (the reliance on sensory input). Each approach faces challenges and disagreements, such as the interaction between perception and cognition, the role of unconscious perception, and the influence of cognition on perception. So, what we can learn from the study is the diversity and complexity of understanding the perception-cognition border, and multiple approaches might be needed for a comprehensive understanding, which raises the question of whether a clear-cut border of perception-cognition even exists.

In exploring the intersection of the LOT and perception, Quilty-Dunn et al. (2022) assert that given the need for interaction between perceptual and cognitive systems, elements of perception might exhibit LOT-like characteristics. This perspective aligns with Tacca’s study, which investigates the intricate relationship between visual perceptual representations and cognitive processes. Tacca's research (2011) reveals parallels in the structural and content characteristics of these two domains. An important point in Tacca's argument is that the spatial structure of visual object representation demonstrates systematicity. Systematicity, typically a feature of propositional cognitive representations, suggests that visual representations are not random or unorganized sensory inputs but instead possess a structured, rule-governed framework akin to language. This alignment implies that such representations can be manipulated and reconfigured in ways similar to linguistic constructs. Consequently, Tacca posits that if visual representations exhibit systematicity, they could be regarded as rudimentary conceptual representations. This hypothesis challenges traditional distinctions between sensory perception and higher cognitive functions and underscores the deep interconnectivity of cognitive processes and perceptual information. Tacca’s insights bolster the concept that perception and cognition are not separate but are profoundly integrated, each influencing and shaping the other.

In the early discussion of the neural basis for predicate-argument structure, we can also see that Hurford (2003) presents a compelling argument that links the neural underpinnings of language, specifically predicate-argument structures, to fundamental brain processes engaged in both linguistic and nonlinguistic tasks. He points out the brain's ability to integrate sensory information via distinct neural pathways – those for object localization (the 'where' pathways) and for identifying attributes (the 'what' pathways). This dual processing mirrors the logical structure found in linguistic expressions like PREDICATE(x), suggesting an inherent parallelism between language and perception. The concept of deictic variables in Hurford's theory, which is akin to mental pointers to objects in the environment, aligns closely with the notion of variables in language. These variables, when integrated with semantic information about the objects, function similarly to predicates in language.

Moreover, Hurford (2003) proposes that the capacity for the mental representation of scenes and objects, a skill evident in primates and crucial for practical tasks, serves as a preadaptive foundation for the evolution of complex linguistic functions, including scene and event description. This evolutionary perspective reinforces the idea that the neural mechanisms supporting predicate-argument structures in language are deeply rooted in the sensory and perceptual systems of the brain, systems that are not only shared across primates but also integral to their interaction with the physical world.

This understanding is particularly illuminating in the context of LOT. It underscores the hypothesis's relevance and applicability in explaining perception, as it suggests that the cognitive structures used in language have their origins in these fundamental perceptual processes. Thus, Hurford's insights not only provide a neurobiological basis for the LOT but also expand its scope, highlighting how language-like structures within our brain may govern our perception of the world.

So far, I have presented some reasons showing the complexity and dynamism at the intersection between cognitive and perceptual processes, challenging traditional binary views. This helps us to move forward to understand the role of LOT in perception.

LOT's applicability to perception, especially through the lens of systematicity and structural properties in visual representations, as discussed in Tacca's study, illustrates those perceptual representations, with their systematic and predicate-argument structure, may not only mirror the characteristics of language-like cognitive representations but also serve as fundamental elements in the formation of conceptual thought.

However, there is still an ongoing discussion in philosophy and cognitive science regarding the boundaries and interplay between perception and cognition. By further digging into various approaches to defining this border and examining the implications of systematicity in visual representations, we need a more integrated and comprehensive approach to understanding the mind.

5. Conclusion

This paper has navigated the philosophical debate between Fodor and Sellars, seeking a reconciliatory middle ground for understanding LOT in the context of modern cognitive science. In the pursuit of this middle ground, I have dissected the LOT into its core components—content, vehicle, and function—and explored how each component contributes to our understanding of cognitive processes and perception.

Firstly, in terms of content, I advocated for a transition from the traditional Fodorian atomistic symbolic model to a more nuanced, situational paradigm. This shift aligns with recent neuroscientific evidence, highlighting the complex and contextual nature of neural representations. This perspective respects the complexity of environmental stimuli and their diverse neural responses, moving beyond the simplistic one-to-one correspondence of neural representation and object.

Regarding the vehicle of thought, I challenged Fodor's monolithic digital computation model, incorporating the dynamic interplay of both digital and analog processes in neural computation. This pluralistic approach is more reflective of the nuanced and multifaceted nature of brain functions as illuminated by recent scientific inquiries. It allows for a more flexible and inclusive interpretation of neural computation, recognizing the coexistence of discrete and continuous processes in mental representations.

Functionally, I discussed the essential features of LOT—discrete constituents, predicate-argument structures, and logical operators—demonstrating their alignment with linguistic structures and their compatibility with modern scientific studies. I showed how these features enable the flexible and combinatorial operations of thought, akin to linguistic processes. The discussion illuminated the intricate connections between neurolinguistic research, computational linguistics, and cognitive neuroscience, all converging to support the LOT's functional framework.

Finally, I extended the scope of LOT to encompass perception, challenging Sellars' view on the limitation of LOT in explaining perceptual processes. This extension, which bridges the often-assumed divide between perception and cognition, is grounded in the systematicity and structural properties observed in visual representations. By doing so, I suggest that perceptual processes might also exhibit LOT-like characteristics, contrary to Sellars' more restrictive stance. This approach integrates insights from various approaches to the perception-cognition border, advocating for a more holistic understanding of cognitive processes that seamlessly incorporate both perceptual and conceptual elements.

In conclusion, the exploration between Fodor and Sellars' perspectives on LOT reveals the potential for a middle ground that accommodates the evolving understandings in cognitive science. This middle ground honors the foundational principles of LOT while integrating contemporary scientific insights, thus offering a more versatile and realistic framework for understanding the complexity of cognitive processes. By embracing this integrated approach, we reach a situational view of LOT and can make LOT a more versatile tool fitting with the other cognitive frameworks, such as the free energy principle and Bayesian reference. This choice of understanding LOT enables us to emphasize the structural and functional aspects of the brain's internal model of cognition while maintaining a broader perspective that encompasses both perception and other cognitive processes such as thinking and acting.

Bibliography

Arora, A., Weiss, B., Schurz, M., Aichhorn, M., Wieshofer, R. C., & Perner, J. (2015). Left inferior-parietal lobe activity in perspective tasks: identity statements. Frontiers in Human Neuroscience, 9, 360. https://doi.org/10.3389/fnhum.2015.00360

Butts, D.A., & Goldman, M.S. (2006). Tuning curves, neuronal variability, and sensory coding. PLoS Biology, 4.

Clarke, S., & Beck, J. (2023). Border disputes: Recent debates along the perception–Cognition border. Philosophy Compass, 18(8), e12936. https://doi.org/10.1111/phc3.12936

Coelho Mollo, D., & Vernazzani, A. (2023). The formats of cognitive representation: A computational account. Philosophy of Science.

deVries, W. (2021). Wilfrid Sellars. The Stanford Encyclopedia of Philosophy. https://plato.stanford.edu/archives/fall2021/entries/sellars/

Enoka, R. M., & Duchateau, J. (2017). Rate coding and the control of muscle force. Cold Spring Harbor Perspectives in Medicine, 7(10), a029702. https://doi.org/10.1101/cshperspect.a029702

Fodor, J. A. (1975). The Language of Thought. New York: Thomas Y. Crowell.

Fodor, J. A. (1981). Representations. Cambridge, MA: MIT Press.

Fodor, J. A. (1987). Psychosemantics. Cambridge, MA: MIT Press.

Fodor, J. A. (2008). LOT 2: The Language of Thought Revisited. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199548774.001.0001

Goel, V., Navarrete, G., Noveck, I. A., & Prado, J. (2017). Editorial: The reasoning brain: The interplay between cognitive neuroscience and theories of reasoning. Frontiers in Human Neuroscience, 10, 673. https://doi.org/10.3389/fnhum.2016.00673

Hurford, J. R. (2003). The neural basis of predicate-argument structure. Behavioral and Brain Sciences, 26(3), 261–283. https://doi.org/10.1017/S0140525X03000074

Khemlani, S. S., Harrison, A. M., & Trafton, J. G. (2015). Episodes, events, and models. Frontiers in Human Neuroscience, 9, 590. https://doi.org/10.3389/fnhum.2015.00590

Knauff, M. (2007). How our brains reason logically. Topoi, 26(1), 19-36.

Marras, A. (1973). Sellars on thought and language. Noûs, 7(2), 152-163.

Meltzer-Asscher, A., Mack, J. E., Barbieri, E., & Thompson, C. K. (2015). How the brain processes different dimensions of argument structure complexity: Evidence from fMRI. Brain and Language, 65. https://doi.org/10.1016/j.bandl.2014.12.005

Ockham, W. of. (c. 1323 [1957]). Summa Logicae. In P. Boehner (Ed. and Trans.), Philosophical Writings, A Selection. London: Nelson.

Pamplona, G. S. P., Santos Neto, G. S., Rosset, S. R. E., Rogers, B. P., & Salmon, C. E. G. (2015). Analyzing the association between functional connectivity of the brain and intellectual performance. Frontiers in Human Neuroscience, 9, 61. https://doi.org/10.3389/fnhum.2015.00061

Pareti, G. (2007). The "all-or-none" law in skeletal muscle and nerve fibres. Archives Italiennes de Biologie, 145(1), 39–54.

Peacocke, C. (2019). The Primacy of Metaphysics. Oxford University Press.

Piccinini, G. (2020). Neurocognitive Mechanisms: Explaining Biological Cognition. Oxford: Oxford University Press.

Piccinini, G. (2023). Why Neuroscience Refutes the Language of Thought. Retrieved from https://philosophyofbrains.com/2023/01/02/why-neuroscience-refutes-the-language-of-thought.aspx

Pylkkänen, L. (2019). The neural basis of combinatory syntax and semantics. Science. https://doi.org/aax0050

Quilty-Dunn, J., et al. (2022). The best game in town: The re-emergence of the language of thought hypothesis across the cognitive sciences. Behavioral and Brain Sciences, 1-55. https://doi.org/10.1017/S0140525X22002849

Rescorla, M. (2019). The Language of Thought Hypothesis. The Stanford Encyclopedia of Philosophy. https://plato.stanford.edu/archives/sum2019/entries/language-thought/

Romain, B. (2021). What is a neural representation? Mathematical Consciousness Science. https://www.youtube.com/watch?v=OAmB5SOS2LQ

Salto, F., Requena, C., Álvarez-Merino, P. et al. Brain electrical traits of logical validity. Sci Rep 11, 7982 (2021). https://doi.org/10.1038/s41598-021-87191-1

Schiller, D., Eichenbaum, H., Buffalo, E. A., Davachi, L., Foster, D. J., Leutgeb, S., et al. (2015). Memory and space: Towards an understanding of the cognitive map. Journal of Neuroscience, 35, 13904–13911. https://doi.org/10.1523/JNEUROSCI.2618-15.2015

Sellars, W. (1957). Intentionality and the mental. In Minnesota Studies in the Philosophy of Science, vol. II, H. Feigl, M. Scriven, & G. Maxwell (Eds.), Minneapolis: University of Minnesota Press, 507–39.

Sellars, W. (1963). Science, Perception and Reality. London: Routledge & Kegan Paul.

Sellars, W. (1975). The Structure of Knowledge. In Action, Knowledge, and Reality: Studies in Honor of Wilfrid Sellars, H.-N. Castañeda (Ed.), Indianapolis: Bobbs-Merrill, 295–347.

Sellars, W. (1981). Mental Events. Philosophical Studies, 81, 1981: 325–45.

Sellars, W. (1997). Empiricism and the Philosophy of Mind. Cambridge, Mass.: Harvard University Press. Edited by Richard Rorty & Robert Brandom.

Simone Conia & Roberto Navigli. (2022). Probing for predicate argument structures in pretrained language models. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (pp. 4622–4632). Dublin, Ireland: Association for Computational Linguistics.

Tacca, M. C. (2011). Commonalities between perception and cognition. Frontiers in Psychology, 2, 358. https://doi.org/10.3389/fpsyg.2011.00358

Vilarroya, O. (2017). Neural representation: A survey-based analysis of the notion. Frontiers in Psychology, 8, 278840. https://doi.org/10.3389/fpsyg.2017.01458

Zucker, H. R., & Ranganath, C. (2015). Navigating the human hippocampus without a GPS. Hippocampus, 25, 697–703. https://doi.org/10.1002/hipo.22447