Intra-neuronal attention within language models

Relationships between activation and semantics

Michael Pichat^{1,2,4}, William Pogrund^{1,5}, Paloma Pichat^{1,3}, Armanouche Gasparian¹, Samuel Demarchi^{1,4}, Martin Corbet^{1,2}, Alois Georgeon^{1,2}, Michael Veillet-Guillem¹

¹Neocognition (Chrysippe R&D) ²Facultés Libres de Philosophie et de Psychologie de Paris (ER IPC) ³Faculté de Médecine de Lyon Est (Université Lyon 1) ⁴Université Paris 8 ⁵INP-PHELMA, Université Grenoble Alpes

Abstract

This study investigates the ability of perceptron-type neurons in language models to perform intra-neuronal attention; that is, to identify different homogeneous categorical segments within the synthetic thought category they encode, based on a segmentation of specific activation zones for the tokens to which they are particularly responsive. The objective of this work is therefore to determine to what extent formal neurons can establish a homomorphic relationship between activation-based and categorical segmentations. The results suggest the existence of such a relationship, albeit tenuous, only at the level of tokens with very high activation levels. This intra-neuronal attention subsequently enables categorical restructuring processes at the level of neurons in the following layer, thereby contributing to the progressive formation of high-level categorical abstractions.

1 Theoretical Context

1.1 Attention Heads

Before the rise of transformers, natural language processing models primarily relied on recurrent (RNN, LSTM) and convolutional (CNN) architectures. While RNNs and LSTMs could capture long-term dependencies through their memory mechanisms, they suffered from the vanishing gradient problem, making it difficult to learn relationships over long sequences. Moreover, their sequential nature limited parallelization and slowed down training. CNNs, on the other hand, offered better parallelization but were ill-suited for global dependencies due to their limited receptive field, which depended on network depth. These approaches thus struggled with capturing long-range dependencies and faced parallelization challenges [124, 125]. These limitations motivated the introduction of transformers and the self-attention mechanism, which overcame these constraints by enabling efficient parallel processing while capturing long-range relationships. The introduction of self-attention by Vaswani [126] marked a major turning point in NLP.

Transformers leverage self-attention, where each element in a sequence weighs the importance of other elements, facilitating the modeling of complex dependencies. Multi-head attention enhances this approach by allowing different attention heads to specialize in various aspects of representation. Devlin et al. [127] demonstrated that certain BERT heads capture syntactic relations, such as subject-verb links, while others focus on more global semantic relationships. Radford et al. [128] showed that multiple attention heads in GPT improve sentence context modeling by capturing distributed information across different input sequence positions. This capability enables transformers to enhance the richness and hierarchy of representations, improving generalization across various tasks.

Within a transformer, each attention layer comprises multiple heads, each performing an attention operation on a linear projection of the inputs. Self-attention projects representations into three sets of vectors : queries (Q), keys (K), and values (V). For an input sequence $X \in \mathbb{R}^{T \times b}$, where T is the sequence length and b the vector dimension, these matrices are defined as :

$$Q = XW^T, \quad K = XW^T, \ V = XW^T$$

where $W^T \in \mathbb{R}^{b \times b_k}$ are learned weight matrices, and b_k is the dimension of keys and queries. Attention is computed using the scaled dot-product attention :

Attention
$$(Q, K, V) = \operatorname{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

The normalization by $\sqrt{d_k}$ stabilizes training by preventing large-scale values in dot products [126]. Multi-head attention applies these operations in parallel over h different projections of Q, K, and V:

$$head_i = Attention(QW_i^G, KW_i^K, VW_i^V)$$

The concatenation of heads followed by a linear projection yields the final output :

$$MultiHead(Q, K, V) = Concat(head_1, \dots, head_h)W^O$$

This architecture enables the extraction of contextual information from multiple perspectives and enhances the capture of complex relationships in data. Studies on attention mechanisms have shown that some heads are specialized, while others are redundant. Luong et al. [129] demonstrated that certain heads capture precise syntactic relations, whereas others focus on global semantic relationships. [130, 131] observed that removing several heads does not significantly impact model performance, suggesting compensatory mechanisms among the remaining heads.

Statistical mechanics approaches have been used to analyze interactions between attention paths. [132] modeled the contribution of attention heads via a kernel decomposition :

$$K = \sum_{i=1}^{h} K_i$$

Each kernel K_i corresponds to a specific head, allowing an evaluation of its role in the model's final representation. Results indicate that some heads play a structuring role, while others can be eliminated without significant impact. This observation paves the way for transformer architecture optimizations by reducing redundant heads and improving model interpretability.

The computational efficiency of transformers has been extensively researched. Sparse Transformers [23] reduce attention complexity to $O(n \log n)$ by introducing a sparse attention structure. Reformer [133] optimizes memory management via key-value factorization and local attention. Performer [134] replaces standard attention with a linear approximation, reducing complexity to O(n). This method relies on random projections of keys and queries into a lower-dimensional space, where dot products are computed approximately using kernels favoring efficient factorization. This avoids costly dense matrix multiplications while maintaining high accuracy, making attention scalable even for long sequences. Longformer and BigBird [135] combine local and global attentions to efficiently process long sequences.

Other studies have analyzed attention head specialization in specific contexts. Clark et al. [136] examined BERT attention matrices and found that some heads learn specific syntactic relationships, such as subject-verb dependencies or anaphoric relations. Transformer-XL [137] introduced a recurrent memory mechanism that captures longer-term dependencies, improving text generation and dialogue modeling.

Finally, attention mechanisms have extended to other domains, including computer vision with Vision Transformer (ViT) [138] and Swin Transformer [139], as well as neuroscience and cognitive process modeling [140].

1.2 Human Attention

Biologically, human attention arises from the limited capacity of the neuronal system to process information. It manifests through selective mechanisms in the acquisition, activation, and utilization of sensory or memory data (such as knowledge or rules) [41, 4]. This results in an oriented response, focusing information retrieval on specific characteristics. Neurocognitive research on attentional mechanisms, inspired notably by Posner's work [81, 82, 87], highlights the existence of a frontal attention system associated with conscious attention and planning, and a posterior system, located in the parietal lobe, involved in visuospatial processes and shifts in attentional focus.

In cognitive psychology, attention is conceptualized as the calibration of activity toward a specific goal, thereby enhancing efficiency in the collection and execution of information (selectivity, accuracy, speed) for a given task [77, 90, 76, 84, 148, 30, 91, 26, 108, 45]. During task execution, attention is managed by the central nervous system, which determines the relevance of internal information (such as knowledge or schemas) that ensures execution quality.

Generally, two cognitive functions are associated with attention [31]:

- Signal detection, which primarily relies on vigilance and exploration to identify the appearance of a specific stimulus.
- Selective attention, focused on specific stimuli while excluding others.

Vigilance refers to the ability to concentrate on a flow of information to detect a specific signal [56], which may appear rarely but requires a rapid reaction [21, 46]. It is negatively affected by the level of uncertainty regarding the targeted elements [17]. Vigilance can be defined as an adjustable attentional beam influenced by the anticipation of the signal's appearance at a specific location [78, 60, 62].

Visual exploration [111, 108], on the other hand, involves an active search for stimuli, contrasting with the passive expectation of their emergence [80]. It is characterized by a scanning strategy for recognizing attributes within a given environment. According to the *Feature Integration Theory* [95, 86], a mental map for each visual attribute represents occurrences within the visual field, which are regularly inspected in parallel. Here, attentional processes play a role in mental binding, assembling various attributes of the same object while inhibiting irrelevant characteristics. The *Similarity Theory* [32] analyzes attentional exploration as an evaluation of the proximity between target stimuli and distractors. Meanwhile, the *Guided Search Theory* [20, 1] divides exploratory attention into two phases : first, the activation of a global representation of potential targets, followed by a serial analysis to identify the most activated target.

Selective attention is often explored through the *cocktail party effect* [22, 55, 11], referring to the ability to follow a conversation among others in the background. The characteristics of this attentional focus include sensory properties, sound volume, and the spatial location of the target speech. In this context, the *Filter Theory* [18, 112] suggests that a filter selects sensory streams to receive deep processing. However, this model evolved into an *attenuation approach* [93, 53], where all information is reduced in perceptible intensity, leaving only residual data close to targeted criteria. The distribution of limited attentional resources [52] is also relevant for managing multiple tasks in parallel with increased efficiency.

The aspects of attentional mechanisms relevant to our study lie at the

intersection of vigilance and selective attention. These processes—detecting a targeted type of information and selectively focusing on specific data characteristics—correspond to the phenomenon of interest here : the impact of a neuron's activation level in response to an incoming token on the intra-neuronal detection and selective attention of tokens exhibiting specific categorical features.

1.3 Attentional Processes and Conceptualization

The notion of conceptualization ¹ is one of the major contributions of the theory of conceptual fields proposed by Vergnaud [100, 101]. As we will see below, conceptualization is a central cognitive process at the intersection of the vigilance and selective attention mechanisms mentioned earlier, specifically in the context of categorizing information received by a cognitive system. This notion, developed in the field of human thought, is therefore particularly relevant, heuristically, to our current investigation concerning the identification and selective attention of specific data (tokens) at the level of internal processing carried out by formal neurons.

Conceptualization is an attentional representational activity whose purpose is the identification of the operational characteristics of stimuli (in our case, tokens) to which it applies, in order to ground a cognitive system's activity on these characteristics and thus enhance its efficiency. The cognitive function of conceptualization is to extract an operational form of knowledge, which becomes the object of specific attentional focus. In the domain of human cognition, this attentional focus is largely unconscious and non-verbalizable; for this reason, Vergnaud [100] refers to it as knowledge-in-action.

Vergnaud [101] defines a concept-in-action, resulting from the selective attentional activity of conceptualization, as a category of thought learned to be relevant for a given task (in our case, the targeted processing of textual data). Regarding this notion of concept-in-action, three key points should be highlighted :

- 1. Concepts-in-action are categories of thought through which a cognitive system identifies, selects, and captures certain pieces of information present in a given situation (a set of tokens in the case of language processing). In other words, concepts-in-action function as cognitive attentional filters that allow a given situation to be selectively "read" or "perceived."
- 2. From an epistemological perspective, there exists a potentially infinite number of formal types of thought categories. The most commonly

^{1.} Traditionally, research on concepts considers them as cognitive units stored in memory and linked to a word or expression, with an associated class of objects sharing common properties. Experimentally, these approaches have led to studies on familiarity or typicality judgments, category comparisons, classification, and more generally, categorical identification—studies that we have previously synthesized [67, 68]. However, the term *conceptualization*, as used by Vergnaud [101], does not refer to these classical approaches, which focus on the predicative form of the concept. Instead, it aligns with a developmental and pragmatic perspective, emphasizing the primary form of the concept : its operational form.

encountered types are : object, property, relation, transformation, condition, and process.

3. Concepts-in-action serve as pragmatic vectors of synthetic thought, organizing attentional information processing by segmenting the token space according to the contingent goals of the finalized activity for which a language model has been trained. Indeed, the functionality of concepts-in-action lies in their ability to enable the neural system to focus its attention on a limited number of selected elements, learned to be crucial for the success of the synthetic system's activity. As such, they support a representation of only those situational variables whose consideration is essential for the effectiveness of the task at hand.

The approach developed by Vergnaud [101] establishes conceptualization as a fundamentally economic and pragmatic attentional cognitive activity. This inherently pragmatic finalization of conceptualization is the reason why the concepts it extracts are *in action*, meaning they are encapsulated within the cognitive system's activity.

An artificial neuron can be described as a synthetic cognitive operator of conceptualization, whose purpose is to select, from a set of incoming tokens, one or more specific subgroups of tokens. These subgroups constitute the categorical extension of the *critical* concept(s)-in-action that the neuron is functionally tasked with identifying selectively. In this sense, the neuron performs an attentional focusing activity on certain types of tokens that need to be selected and filtered to optimize the efficiency of the linguistic processing task in which the neuron participates.

2 Research Problem

2.1 Mathematico-Cognitive Factors of Categorical Segmentation and Intra-Neuronal Attention

In a previous study [69], we explored the mathematico-cognitive factors influencing how an artificial neural network (of the perceptron type) in a language model performs categorical segmentation of the tokens presented to it. Based on the neuronal aggregation function of the form $\sum (w_{i,j} x_{i,j}) + b$, which partially governs this cognitive process, we identified three factors contributing to this conceptual partitioning :

- The x-effect : Synthetic Categorical Priming. This effect refers to the influence of synthetic thought categories activated in neurons of layer n on the activation of categories in neurons that are strongly connected via attention in the next layer. In other words, a token that strongly belongs to an initial category in layer n is more likely to belong to a strongly associated category in layer n + 1.
- The w-effect : Synthetic Inter-Categorical Attention. This effect influences the degree of importance that a receiving neuron (in layer n + 1) assigns to

the categories of neurons in the preceding layer (n), based on connection weights. This process results in categorical complementation, where each attentively focused precursor category contributes a unique sub-dimensional category (comprising highly specific tokens) to the formation of the receiving category. Thus, the category of a receiving neuron is constructed by assembling complementary categorical sub-dimensions derived from its precursor categories.

• The \sum -effect : Synthetic Categorical Phasing. This effect refers to the tendency for tokens that are jointly activated within precursor categories (layer n) to be more likely to be part of the categorical extension of their strongly associated attentionally focused category in layer n + 1. This process manifests as categorical intersection.

The phenomenon of interest here, as a reminder, is the extent to which the specific activation level of a neuron in response to a token is linked to a particular categorical value of that token within the category carried by the neuron. In other words, is a given segment of activation values associated with a distinct categorical segment, allowing a neuron to selectively focus its attention on a specific categorical segment based on its associated activation span?

The three mathematico-cognitive factors we have just outlined—which could be related to the genesis and structuring of human thought categories [10, 47, 38, 40, 7, 58, 110]—are particularly relevant to our research question. Indeed, these three factors directly determine, before the application of the nonlinear activation function, the activation value that the neuronal aggregation function assigns to a given token. In other words, these factors define the activation zone, the activation segment within which a processed token is positioned by the involved neuron.

To what extent do these three factors—both as a matrix of activation values assigned to tokens by a neuron and as carriers of categorical effects, as previously described—delineate specific activation segments paired with specific categorical segments that could thus become the object of selective attentional focus? Alternatively stated, how are these factors, through neuronal activation values, associated with the conceptualization of certain categorical segments, on which it is both relevant and effective for the neuron to focus its internal attention?

2.2 Synthetic Categorical clipping and Intra-Neuronal Attention

In a previous study [71], we highlighted that the three mathematico-cognitive factors of categorical segmentation, as detailed above, govern a mechanism of *synthetic categorical clipping*. This clipping process results in the elaboration and separation of a form from a categorical background. More precisely, categorical clipping is the phenomenon in synthetic cognition by which a specific categorical sub-dimension is extracted from the category carried by a precursor neuron (in layer n) to contribute to the formation of a superordinate category (in layer n+1).

Categorical clipping manifests through a series of synthetic characteristics :

- Categorical reduction, meaning that the categorical sub-dimension extracted from a precursor category contains tokens that are semantically more homogeneous compared to the original category.
- Categorical selectivity, referring to the extraction of a small subgroup of tokens from the larger set of tokens that initially characterized the original category.
- Separation of initial embedding dimensions, associated with a differentiation of these embeddings, with some being more specifically related to the outlined categorical sub-dimension.
- Partitioning of categorical zones within initial embedding dimension, where certain zones are more specifically associated with the extracted subdimensions.

Categorical clipping is an activity that extracts a categorical sub-dimension from the category associated with a neuron in layer n, but this process occurs *externally* to the originating neuron, specifically within one of its paired neurons in layer n + 1. However, to what extent is this clipping—ultimately driven by activation values—executed based on intra-neuronal attentional focus and conceptualization of specific categorical sub-segments within the original category, sub-segments that would be associated with particular activation segments?

2.3 Categorical Restructuring and Intra-Neuronal Attention

In a previous study [72], we explored the process of synthetic categorical restructuring, specifically the generation, at each neuronal layer n + 1, of new artificial thought categories that are more functional for segmenting the world of tokens in alignment with the purpose of the neural network's activity. This process falls under reflective abstraction in the Piagetian sense [71], applied to the categories of layer n.

This categorical restructuring is directly dependent on the coactivity of the three factors of categorical segmentation outlined earlier (categorical priming, inter-categorical attention, and categorical phasing). We posited that the restructuring phenomenon is particularly linked to inter-categorical attention, given that this latter factor, by mathematical construction of the neuronal aggregation function, is both a necessary condition and an amplifier of the other two factors (priming and phasing).

We demonstrated that the joint action of inter-categorical attention and categorical phasing leads to partial categorical confluence : the categorical sub-dimensions extracted from layer n categories, at the level of a neuron that is strongly attentive to them in layer n + 1, tend to semantically converge to a relative extent. We also highlighted that the combined impact of inter-categorical attention and categorical priming generates activation dispersion : a categorical sub-dimension extracted from a category in layer n does not correspond to a continuous segment of token activations at the originating neuron.

How are both this partial categorical confluence and this activation dispersion, occurring during the transition from layer n to layer n + 1, realized through conceptualization and intra-neuronal attentional focus at the level of layer n? In other words, how does partial categorical confluence, driven by the coactivity of synthetic factors such as inter-categorical attention and categorical phasing, become possible through the identification—via activation—of specific categorical sub-segments to be considered at the level of the involved precursor neuron? And to what extent is activation dispersion compatible with intra-neuronal attentional focus based on specific activation segments, which might appear paradoxical?

2.4 Conceptualization and Intra-Neuronal Attention

From a mathematical perspective, a synthetic neuronal processing unit results from the composition of functions : a nonlinear activation function (ReLU, SELU, GELU, ELU, etc.) applied to a linear aggregation function of the form $\sum (w_{i,j} x_{i,j})+b$. Within a language model, this matrix-based processing associates an input token (or rather its embedding) with an output activation value.

What is the epistemological significance of this activation value from a semantic standpoint? Is this activation value correlated with a specific semantic value? Are different activation segments within the activation space associated with distinct categorical segments in the semantic space of the involved neuron? For a given neuron, is there a homomorphic relationship between its activation space and its categorical space? Can these spaces be divided and segmented into identifiable and pairable activation and categorical zones?

In other words, is it possible to locate intra-neuronal activation segments that correspond to specific sub-semantic structures within the thought categories represented by the neurons? Put differently, to what extent does the activation value serve as a quantifier enabling conceptualization—an attentional mapping of specific intra-neuronal semantic zones?

3 Methodology

3.1 Methodological Positioning

To situate our current research methodologically, we position it within a set of explainability investigation techniques aimed at making artificial neural networks more comprehensible. These methods seek, with varying degrees of cognitive depth, to explain internal mechanisms or interpret the meaning and function of information flows within these networks, whether studied at the level of individual neural layers, groups of layers, or the entire model.

Studies focusing on *macroscopic explainability* examine fluctuations between input data and outputs to clarify the relationship between what is given to the system and what it produces. In this approach, gradient-based methods assess the influence of each input by analyzing the partial derivatives of each input dimension [37]. Input attributes can be evaluated through various elements, including feature importance [28], token relevance [37], or attention coefficients [6]. Furthermore, example-based methods observe variations in outputs when input modifications are introduced, allowing for an analysis of how slight data alterations affect model predictions [106], as well as assessing the implications of input perturbations such as deletion, negation, mixing, or masking [3, 107, 96]. Other approaches focus on conceptually mapping inputs to evaluate their influence on observed outputs [19].

Microscopic explainability methods, on the other hand, examine the intermediate internal states of language models rather than their overall output, analyzing interconnections and activations of individual neurons or groups of neurons. Some studies explore how to segment and interpret neural activations in a layer based on the inputs from the previous layer [103]. Others aim to adjust activation functions to enhance interpretability [106]. Certain techniques investigate the knowledge embedded within neurons, extrapolating internal meanings through significance matrices [29, 42]. Finally, some approaches assess statistical patterns in neural responses based on specific datasets [9, 61, 33, 106, 27]. Our current work falls within this latter methodological category.

3.2 Model and Statistical Units Used

Building on our previous work [67, 68, 69, 70, 71], we focused on OpenAI's transformer-based GPT model, specifically the GPT-2XL version. This choice was made because GPT-2XL offers sufficient complexity to investigate advanced cognitive processes while avoiding, for an initial exploration, the interpretative challenges posed by GPT-4 or GPT-40. In 2023, OpenAI released detailed documentation on the parameters of GPT-2XL, as outlined by Bills et al., [9], which we leveraged in our current investigation.

To reduce the scope of our statistical analyses, we focused on the first two perceptron layers of GPT-2XL, each containing 6,400 neurons, for a total of 12800 artificial neurons. Regarding linguistic elements (tokens) and their associated activations, our study centered on the 100 tokens exhibiting the highest mean activations per neuron, which we termed *core-tokens*. When analyzing relationships between neurons in layers 0 and 1, we limited our selection for each neuron in layer 1 to its 10 most strongly connected neurons from layer 0, based on attention weight values.

3.3 Statistical Processing

For our statistical analyses, we used the SciPy library, in accordance with the recommendations of [50, 8, 34, 35].

Our verification of the normality of the data, to assess the feasibility of parametric tests, was conducted in two steps. First, we applied inferential tests : the Shapiro-Wilk test, suitable for small datasets; the Lilliefors test, used when distribution parameters are unknown; the Kolmogorov-Smirnov test for large samples; and the Jarque-Bera test, which quantifies skewness and kurtosis in large samples. Complementing this, in a second step, we examined skewness and

kurtosis coefficients and visualized distributions using QQ-plots to compare the recorded data with a theoretical normal distribution. To check the homogeneity of variances among data subgroups whose relationships were being analyzed, we applied Bartlett's test, which is sensitive to deviations from normality, and Levene's test, which is less affected by such deviations.

These preliminary checks indicated limited normality in our data. Consequently, our statistical analyses relied on non-parametric approaches, specifically :

- The Kruskal-Wallis test, used to investigate the relationship between a categorical variable defining independent groups and an ordinal variable. This test was applied based on ranking the numerical activation data of tokens, adhering to the standard conditions for its application, including groups with a minimum of six observations. Effect sizes for the Kruskal-Wallis test were measured using Cohen's d. For pairwise group comparisons, given k total groups, k(k-1)/2 comparisons were performed, using a rank difference coefficient adapted for post hoc situations, measured on a z-scale, with a significance threshold α divided by k(k-1).
- The univariate χ^2 goodness-of-fit test, applied while ensuring compliance with its theoretical and observed frequency requirements, thus avoiding alternatives suited for small samples, such as Fisher's or Monte Carlo methods. Effect sizes were measured using the risk ratio.

For certain analyses, we employed a typological classification approach via hierarchical clustering. This classification was configured as follows : (i) a topdown classification approach, (ii) the use of Euclidean distance to measure distances between statistical units, (iii) a predefined number of five clusters, (iv) Ward's method as the aggregation technique, (v) prior standardization of the data.

3.4 Methodological Operationalizations

To quantify semantic proximity between tokens, we used cosine similarity based on the embedding space of GPT-2XL. This approach allowed us to avoid methodological limitations identified by Bills et al., [9] when attempting to relate various artificial cognitive systems using non-unified embedding foundations.

To examine the potential relationship, for each given neuron, between its categorical segments and its activation segments, we employed two opposing methodological approaches :

• A "top-down" approach, which starts from predefined categorical segments and investigates the extent to which they are associated with distinct activation segments. These categorical segments were first defined using hierarchical clustering based on the GPT-2XL embeddings of the 100 *coretokens* defining the category associated with each investigated neuron. Then, for methodological diversification, segmentation was also performed using prompt engineering with OpenAI's GPT-40 model. This first approach was further extended by measuring the interweaving and activation overlaps of the categorical token segments obtained via prompt engineering. • A "bottom-up" approach, which instead starts from predefined activation segments and evaluates whether these segments exhibit semantic proximity. To ensure methodological diversification, activation segments were defined first through activation quartiles and then through hierarchical clustering of activations.

4 Results

As a reminder, our research question is as follows : does a synthetic process of conceptualization and intra-neuronal attention exist, enabling each neuron to identify and isolate specific categorical segments within the artificial thought category it carries, based on determined activation segments?

We operationalized our empirical investigation of this question through two methodologically opposing studies, which we will present sequentially :

- A "top-down" study, in which we started from categorical segments to examine their respective average activation values.
- A "bottom-up" study, in which we first isolated specific activation segments and then analyzed their respective categorical homogeneity.

4.1 Activation Differentiation of Categorical Clusters

As part of our *top-down* study, for each perceptron neuron in layers 0 and 1 of GPT-2XL, we decomposed the thought category it carries (through its 100 *core-tokens*, i.e., the tokens with the highest average activation) into five categorical clusters. These clusters represent subcategories that are relatively homogeneous according to a given semantic criterion, allowing the initial category to be partitioned meaningfully.

To illustrate this, we present below the example of neurons $n^{\circ}5$ (layer 1) and $n^{\circ}5065$ (layer 0) using our *genetic neural viewer*. This serves to investigate the relationship between categorical segmentation (clustering) and activation segmentation. More specifically, we aim to determine whether there is a significant difference in average activation values among the categorical clusters obtained.



A first operationalization of our approach involved generating five categorical token clusters for each neuron, based on hierarchical clustering applied to the GPT-2XL embeddings of the 100 *core-tokens* associated with that neuron. From a methodological perspective, we clarify the following points :

- The use of GPT-2XL embeddings, like any operationalization, is a specific methodological choice. It does not claim to measure any intrinsic semantic reality—an epistemologically meaningless notion—but rather represents one particular modality, among many possible, for evaluating semantic phenomena within a given contingent semantic space. Our interpretations must therefore remain confined to this specific semantic contingency.
- In our overall statistics, we retained only those neurons whose semantic clustering resulted in clusters containing at least six tokens, ensuring compliance with the conditions for using the Kruskal-Wallis inferential test for mean comparisons.
- For each neuron, we systematically designated as " K_1 " the categorical cluster with the lowest mean activation value among its constituent tokens, and so forth up to " K_5 ."

Table n°1 summarizes our results for the 2194 neurons analyzed in layer 0. Overall, we observe a low percentage (21.46%) of neurons exhibiting a significant difference ($\alpha = 5\%$) in mean activations (μ_{K_n}) across their five categorical clusters. This finding is corroborated by the low percentages ($\pi(p_{K_n,K_m} < \alpha')$) of significant mean activation differences in post hoc pairwise comparisons of categorical clusters (with an adjusted significance threshold $\alpha' = \alpha/20$).

However, an interesting trend emerges : while the mean activation distances $(\mu(\delta_{K_n,K_m}))$ between successive categorical clusters are generally small, the average activation distance between clusters K_4 and K_5 ($\mu(\delta_{K_4,K_5}) = .0955$) is slightly larger than that between clusters K_1 and K_2 ($\mu(\delta_{K_1,K_2}) = .0512$). This trend aligns with the slightly higher effect size (measured using Cohen's d) for the activation distance between K_4 and K_5 ($\mu(d_{K_4,K_5}) = .2785$) compared to K_1 and K_2 ($\mu(d_{K_1,K_2}) = .2404$).

More strikingly, this trend becomes much more pronounced when considering the strong effect size of the mean activation difference between clusters K_1 and K_5 ($\mu(d_{K_1,K_5}) = .8202$). One might hypothesize that the weak significance of post hoc inferential tests comparing μ_{K_1} and μ_{K_5} ($\pi(p_{K_1,K_5} < \alpha') = 14.02\%$) is due to a bias resulting from the small number of involved tokens.

Graph 1 visually summarizes the main data mentioned here.

N_{neuron}		2194	
$\pi(p_{\rm KW} < .05)$		21.4612	
μ_{K_1}	1.7117	$\mu(\delta_{K_1,K_2})$.0512
μ_{K_2}	1.7629	$\mu(\delta_{K_2,K_3})$.0416
μ_{K_3}	1.8045	$\mu(\delta_{K_3,K_4})$.0484
μ_{K_4}	1.8530	$\mu(\delta_{K_4,K_5})$.0955
μ_{K_5}	1.9485	$\mu(\delta_{K_1,K_5})$.2368
$\mu(d_{K_1,K_2})$.2404	$\pi(p_{K_1,K_2} < \alpha')$.1370
$\mu(d_{K_2,K_3})$.1650	$\pi(p_{K_2,K_3} < \alpha')$.0913
$\mu(d_{K_3,K_4})$.1664	$\pi(p_{K_3,K_4} < \alpha')$.2283
$\mu(d_{K_4,K_5})$.2785	$\pi(p_{K_4,K_5} < \alpha')$	1.3242
$\mu(d_{K_1,K_5})$.8202	$\pi(p_{K_1,K_5} < \alpha')$	14.0183

Table n°1 : Comparison of mean activations between categorical clusters from hierarchical classification on tokens' embeddings (layer 0).



Graph n°1 : Comparison of mean activations between categorical clusters from hierarchical classification on tokens' embeddings (layer 0).

These initial results suggest the following dual trend :

- Activation indifferentiation : A weak difference in mean activation values among categorical clusters associated with lower mean activation values.
- Relative activation differentiation : A stronger difference in mean activation values among categorical clusters that involve one or more higher mean activation values.

Table n°2 and its associated summary graph n°2 display similar results, but with even more pronounced trends. Notably, the effect size contrast is

stronger between clusters K_4 and K_5 ($\mu(d_{K_4,K_5}) = .3363$) compared to K_1 and K_2 ($\mu(d_{K_1,K_2}) = .2631$). Additionally, when comparing clusters K_1 and K_5 , we observe an extremely high effect size ($\mu(d_{K_1,K_5}) = .9962$), associated with increased significance ($\pi(p_{K_1,K_5} < \alpha') = 22.08\%$).

N_{neuron}		2192	
$\pi(p_{\rm KW} < .05)$		31.4325	
μ_{K_1}	1.8300	$\mu(\delta_{K_1,K_2})$.1062
μ_{K_2}	1.9362	$\mu(\delta_{K_2,K_3})$.0896
μ_{K_3}	2.0258	$\mu(\delta_{K_3,K_4})$.1119
μ_{K_4}	2.1377	$\mu(\delta_{K_4,K_5})$.2143
μ_{K_5}	2.3519	$\mu(\delta_{K_1,K_5})$.5219
$\mu(d_{K_1,K_2})$.2631	$\pi(p_{K_1,K_2} < \alpha')$.0456
$\mu(d_{K_2,K_3})$.1869	$\pi(p_{K_2,K_3} < \alpha')$.3193
$\mu(d_{K_3,K_4})$.2042	$\pi(p_{K_3,K_4} < \alpha')$.5931
$\mu(d_{K_4,K_5})$.3363	$\pi(p_{K_4,K_5} < \alpha')$	2.1898
$\mu(d_{K_1,K_5})$.9962	$\pi(p_{K_1,K_5} < \alpha')$	22.0803

Table $n^{\circ}2$: Comparison of mean activations between categorical clusters from hierarchical classification on tokens' embeddings (layer 1).



Graph n°2 : Comparison of mean activations between categorical clusters from hierarchical classification on tokens' embeddings (layer 1).

Still within the framework of our *top-down* study, a second operationalization of our approach involved generating five categorical token clusters for each neuron (again based on its 100 *core-tokens*), this time using a system of prompts applied

to GPT-4 within a Python-coded structure. This allowed us to investigate our research question from an alternative semantic observation framework coupled with a different clustering methodology.

Table n°3 and its associated summary graph (Graph 3), covering 2316 neurons from layer 0, consistently exhibit the same types of results as previously observed.

N_{neuron}		2316	
$\pi(p_{\rm KW} < .05)$		18.3938	
μ_{K_1}	1.7443	$\mu(\delta_{K_1,K_2})$.0629
μ_{K_2}	1.8072	$\mu(\delta_{K_2,K_3})$.0512
μ_{K_3}	1.8584	$\mu(\delta_{K_3,K_4})$.0592
μ_{K_4}	1.9176	$\mu(\delta_{K_4,K_5})$.1015
μ_{K_5}	2.0190	$\mu(\delta_{K_1,K_5})$.2747
$\mu(d_{K_1,K_2})$.2981	$\pi(p_{K_1,K_2} < \alpha')$.0864
$\mu(d_{K_2,K_3})$.1982	$\pi(p_{K_2,K_3} < \alpha')$.0432
$\mu(d_{K_3,K_4})$.1982	$\pi(p_{K_3,K_4} < \alpha')$.0000
$\mu(d_{K_4,K_5})$.2803	$\pi(p_{K_4,K_5} < \alpha')$.2591
$\mu(d_{K_1,K_5})$.8985	$\pi(p_{K_1,K_5} < \alpha')$	11.9603

Table n°3 : Comparison of mean activations between categorical clusters from GPT4 clustering prompt on tokens' embeddings (layer 0).



 $Graph \ n^{\circ}3$: Comparison of mean activations between categorical clusters from GPT4 clustering prompt (layer 0).

The same pattern is once again observed in Table $n^{\circ}4$ and its corresponding summary graph $n^{\circ}4$, which pertain to 1942 neurons from layer 1. This observation

$N_{ m neuron}$		1942	
$\pi(p_{\rm KW} < .05)$		29.8661	
μ_{K_1}	1.8086	$\mu(\delta_{K_1,K_2})$.1229
μ_{K_2}	1.9316	$\mu(\delta_{K_2,K_3})$.1099
μ_{K_3}	2.0415	$\mu(\delta_{K_3,K_4})$.1341
μ_{K_4}	2.1756	$\mu(\delta_{K_4,K_5})$.2223
μ_{K_5}	2.3979	$\mu(\delta_{K_1,K_5})$.5892
$\mu(d_{K_1,K_2})$.3321	$\pi(p_{K_1,K_2} < \alpha')$.1030
$\mu(d_{K_2,K_3})$.2424	$\pi(p_{K_2,K_3} < \alpha')$.1030
$\mu(d_{K_3,K_4})$.2451	$\pi(p_{K_3,K_4} < \alpha')$.0515
$\mu(d_{K_4,K_5})$.3461	$\pi(p_{K_4,K_5} < \alpha')$.8754
$\mu(d_{K_1,K_5})$	1.0948	$\pi(p_{K_1,K_5} < \alpha')$	23.0690

could support a hypothesis suggesting that the synthetic phenomena we highlight tend to increase in deeper layers.

Table n°4 : Comparison of mean activations between categorical clusters from GPT4 clustering prompt on tokens' embeddings (layer 1).



Graphe n°4 : Comparison of mean activations between categorical clusters from GPT4 clustering prompt (layer 1).

From these various initial "top-down" investigations, a dual trend consistently emerges regarding the existence of a potential synthetic intra-neuronal attention mechanism that enables a neuron to identify and locate, within the artificial thought category it represents, specific categorical segments determined based on given activation segments :

- Activation indifferentiation among categorical clusters corresponding to lower average activation values.
- Relative activation differentiation among categorical clusters associated with one or more higher average activation values.

4.2 Activation Interleaving of Categorical Clusters

We remain within the framework of our *top-down* investigation into the possible existence of a synthetic mechanism of conceptualization and intraneuronal attention. This mechanism would enable each neuron to identify and differentiate specific categorical segments within its associated thought category, based on particular activation zones. However, we now adopt a significantly different methodological perspective. Here, we examine the extent to which the categorical clusters obtained (using the previous method involving a GPT-4 clustering prompt) form distinct segments in terms of their activation levels. In other words, we investigate whether activation segments specific to categorical clusters) or remain distinct (i.e., the activation segment of each categorical cluster contains only tokens from that cluster).

This is operationalized as follows :

- Let $x_{\min}(K_i)$ be the lowest activation value among tokens in categorical cluster K_i , and $x_{\max}(K_i)$ the highest activation value in the same cluster.
- Let $n(K_i)$ be the number of tokens (from the full set of five clusters K_1, K_2, K_3, K_4, K_5) whose activation values fall within the range $[x_{\min}(K_i), x_{\max}(K_i)]$, and $m(K_i)$ the number of tokens belonging to cluster K_i .
- Let N be the total number of clustered tokens (100, or fewer when GPT-4 was unable to cluster all tokens).
- For a given categorical cluster K_i , there is no overlap if $n(K_i) = m(K_i)$.

For each neuron in layers 0 and 1 of GPT-2XL, we analyzed this operationalization using contingency tables structured as follows and performed an inferential χ^2 test. We included only cases where clusters had a theoretical count strictly greater than 5 (to meet the conditions for applying χ^2) and an observed count also strictly greater than 5.

	K_i	
Observed	$n(K_i)$	$N - n(K_i)$
Expected	$m(K_i)$	$N-m(K_i)$
Risk ratio		$\frac{n(K_i)}{m(K_i)}$

Tables n°5 and n°6, which respectively include 2316 eligible neurons from layer 0 and 1942 eligible neurons from layer 1, reveal strong activation interleaving. The mean risk ratios (μ_{ρ}) are consistently high, ranging between 4.5 and 5,

indicating a substantial interleaving effect : on average, an activation segment associated with a given categorical cluster contains 4 to 5 times more tokens than the number of tokens explicitly assigned to that cluster. This phenomenon is statistically significant, with very high percentages ($\pi(p(\chi^2) < .05)$) of cases where the observed and expected token distributions are markedly different for each cluster.

N _{neuron}	2316	
V	$\mu_{ ho}$	5.1103
<i>n</i> 1	$\pi(p(\chi^2) < .05)$	100
Ka	$\mu_{ ho}$	4.9323
Λ_2	$\pi(p(\chi^2) < .05)$	100
K_3	$\mu_{ ho}$	4.8721
	$\pi(p(\chi^2) < .05)$	100
K.	$\mu_{ ho}$	4.9705
π_4	$\pi(p(\chi^2) < .05)$	100
K_5	$\mu_{ ho}$	5.0677
	$\pi(p(\chi^2) < .05)$	99.8705

Table $n^{\circ}5$: Percentages of interleaving activations of categorical core-token clusters (Layer 0).

N _{neuron}	1942	
$_{V}$	$\mu_{ ho}$	4.6122
<i>n</i> ₁	$\pi(p(\chi^2) < .05)$	99.9485
Ka	$\mu_{ ho}$	4.6212
112	$\pi(p(\chi^2) < .05)$	100
K_3	$\mu_{ ho}$	4.6998
	$\pi(p(\chi^2) < .05)$	99.9485
K.	$\mu_{ ho}$	4.6624
Λ_4	$\pi(p(\chi^2) < .05)$	99.9485
K_5	$\mu_{ ho}$	4.5496
	$\pi(p(\chi^2) < .05)$	99.6395

Table $n^{\circ}6$: Percentages of interleaving activations of categorical core-token clusters (Layer 1).

The interleaving of activation segments defined by categorical clusters suggests that specific categorical segments are not strictly associated with well-defined activation segments. This result aligns with the *activation indifferentiation* (among categorical clusters with lower mean activation values) highlighted in the previous section.

Moreover, it does not appear to contradict, at this stage, the *relative activation* differentiation (between categorical clusters with higher mean activation values) mentioned earlier. Indeed, this result does not indicate whether the "excess" tokens (i.e., those interleaved across activation levels) in clusters K_5 predominantly originate from all categorical clusters (which might contradict the activation differentiation phenomenon) or primarily from clusters K_4 with fewer from K_1 (which would be consistent with the activation differentiation phenomenon).

4.3 Categorical Homogeneity of Activation Clusters

We now examine the results of our second *bottom-up* study, in which we first isolated specific activation segments and then investigated their respective levels of categorical homogeneity. In other words, we assessed whether tokens from a given activation segment tend to be categorically similar. This analysis directly serves our central research question : does a synthetic process of conceptualization and intra-neuronal attention exist, enabling each neuron to identify and isolate specific categorical segments within the artificial thought category it carries, based on determined activation segments?

From a methodological perspective, categorical proximity was assessed using cosine similarity based on GPT-2XL embeddings. As previously stated, the use of GPT-2XL embeddings is a specific operational choice, which does not represent an absolute semantic ontology but rather one particular way—among others—to measure categorical proximity within a given semantic space. Consequently, our interpretations must be contextualized within this specific semantic contingency.

More precisely, our methodological approach was as follows. For each of the 6,400 neurons in the first two layers of GPT-2XL, we divided the activation space into four activation segments, forming four groups (G_1, G_2, G_3, G_4) of tokens, ordered by their mean activation levels. For each of these four groups, we determined its internal semantic homogeneity (\cos_{G_i}) by computing the mean pairwise cosine similarity of all its tokens. We then calculated an index :

$$d = \cos_{Gi} - Q_3(\cos_{100})$$

where $Q_3(\cos_{100})$ is the third quartile of the pairwise cosine similarities among the 100 *core-tokens* of the neuron. This index *d* thus expresses the extent to which tokens in group G_i are among the most semantically similar (relative to the overall semantic proximities of tokens within the neuron). Specifically, the more negative *d* is, the less semantically similar the concerned tokens are within their neuron.

Initially, we operationalized the segmentation of the activation space into four activation segments using quartiles, producing four groups each containing 25% of the 100 *core-tokens* for a given neuron. This approach ensured activation

clusters were homogeneous in terms of token count. Table n°7 presents the results obtained for the 6,400 neurons in layer 0.

First, we observe that the mean cosine similarities $(\mu(\cos_{G_i}))$ within the four groups are relatively low, ranging from .34 to .39. Additionally, the percentages $(\pi(\mu(\cos_{G_i} - Q_3(\cos_{100}) < 0)))$ of cases where mean cosine similarities fall below the third quartile of all pairwise similarities are very high (all above 97%), indicating that the mean cosine similarities of the G_i groups are generally not among the highest for each neuron. These percentages are associated with extremely low and thus highly significant inferential probabilities $(p(\chi^2))$, assuming a theoretical uniform distribution. These findings suggest a low categorical homogeneity of activation clusters.

A second noteworthy result is that the mean cosine similarities ($\mu(\cos_{G_i})$ increase, albeit moderately but systematically, from G_1 (the least activated tokens, .3485) to G_4 (the most activated tokens, .3933). Correspondingly, the distances $\mu(\cos_{G_i}) - Q_3(\cos_{100})$ tend to decrease slightly in absolute value, as do their associated negativity percentages ($\pi(\mu(\cos_{G_i}) - Q_3(\cos_{100}) < 0)$). This suggests a mild but consistent trend of increasing categorical homogeneity in activation clusters for higher activation values. Graph 5 visually synthesizes these results.

N _{neuron}	6400
$\mu(\cos_{G_1})$.3485
$\mu(\cos_{G_1}) - Q_3(\cos_{100})$	1267
$\pi(\mu(\cos_{G_1}) - Q_3(\cos_{100}) < 0)$	99.9531
$p(\chi^2)$.0000
$\mu(\cos_{G_2})$.3538
$\mu(\cos_{G_2}) - Q_3(\cos_{100})$	1213
$\pi(\mu(\cos_{G_2}) - Q_3(\cos_{100}) < 0)$	99.9844
$p(\chi^2)$.0000
$\mu(\cos_{G_3})$.3644
$\mu(\cos_{G_3}) - Q_3(\cos_{100})$	1108
$\pi(\mu(\cos_{G_3}) - Q_3(\cos_{100}) < 0)$	99.9688
$p(\chi^2)$.0000
$\mu(\cos_{G_4})$.3933
$\mu(\cos_{G_4}) - Q_3(\cos_{100})$	0819
$\pi(\mu(\cos_{G_4}) - Q_3(\cos_{100}) < 0)$	97.6250
$p(\chi^2)$.0000

Table $n^{\circ}7$: Average cosine similarities of activation clusters by quartiles (Layer 0).



Graph $n^{\circ}5$: Average cosine similarities of activation clusters by quartiles (Layer 0).

Table n°8 and its associated summary graph n°6 exhibit exactly the same type of results for the 6,400 neurons in layer 1. Once again, we observe low $\mu(\cos_{G_1})$ values, but with a relative progressive increase as activation levels rise, as well as a gradual relative decrease in the gap between $\mu(\cos_{G_1})$ and $Q_3(\cos_{100})$.

N _{neuron}	6400
$\mu(\cos_{G_1})$.3753
$\mu(\cos_{G_1}) - Q_3(\cos_{100})$	1235
$\pi(\mu(\cos_{G_1}) - Q_3(\cos_{100}) < 0)$	99.9219
$p(\chi^2)$.0000
$\mu(\cos_{G_2})$.3829
$\mu(\cos_{G_2}) - Q_3(\cos_{100})$	1159
$\pi(\mu(\cos_{G_2}) - Q_3(\cos_{100}) < 0)$	99.9375
$p(\chi^2)$.0000
$\mu(\cos_{G_3})$.3956
$\mu(\cos_{G_3}) - Q_3(\cos_{100})$	1033
$\pi(\mu(\cos_{G_3}) - Q_3(\cos_{100}) < 0)$	99.9219
$p(\chi^2)$.0000
$\mu(\cos_{G_4})$.4261
$\mu(\cos_{G_4}) - Q_3(\cos_{100})$	0728
$\pi(\mu(\cos_{G_4}) - Q_3(\cos_{100}) < 0)$	96.6563
$p(\chi^2)$.0000

Table $n^{\circ}8$: Average cosine similarities of activation clusters by quartiles (Layer 1).



Graph n°6 : Average cosine similarities of activation clusters by quartiles (Layer 1).

In a second phase, and to further diversify our methodology, we applied our *bottom-up* approach by segmenting the activation space into four activation segments, this time using hierarchical clustering on the activation values of the 100 *core-tokens* for each neuron. This method has the advantage of producing activation clusters that are more homogeneous in terms of their mean activation values. This added value is particularly relevant given that our analysis focuses on differentiating categorical proximities based on activation zones. Hierarchical clustering is particularly suited for this task, as it tends to identify activation segments that are more distinct from each other (maximizing inter-cluster variance) while ensuring greater internal homogeneity within each cluster (minimizing intra-cluster variance).

Table n°9 and its corresponding summary graph n°7, covering 6,400 neurons from layer 0, as well as Table n°10 and its related summary graph n°8, covering 6,400 neurons from layer 1, consistently reveal the same synthetic phenomena observed previously. However, in this case, there is a stronger effect of increasing categorical homogeneity in activation clusters for higher activation values.

N_{neuron}	6400
$\mu(\cos_{G_1})$.3888
$\mu(\cos_{G_1}) - Q_3(\cos_{100})$	0863
$\pi(\mu(\cos_{G_1}) - Q_3(\cos_{100}) < 0)$	73.5938
$p(\chi^2)$.0000
$\mu(\cos_{G_2})$.4318
$\mu(\cos_{G_2}) - Q_3(\cos_{100})$	0433
$\pi(\mu(\cos_{G_2}) - Q_3(\cos_{100}) < 0)$	65.4219
$p(\chi^2)$.002
$\mu(\cos_{G_3})$.4494
$\mu(\cos_{G_3}) - Q_3(\cos_{100})$	0258
$\pi(\mu(\cos_{G_3}) - Q_3(\cos_{100}) < 0)$	59.4844
$p(\chi^2)$.0578
$\mu(\cos_{G_4})$.4665
$\mu(\cos_{G_4}) - Q_3(\cos_{100})$	0087
$\pi(\mu(\cos_{G_4}) - Q_3(\cos_{100}) < 0)$	50.5625
$p(\chi^2)$.9104

Table n°9 : Average cosine similarities of activation clusters by hierarchical classification (Layer 0).



 $Graph \ n^\circ 7: Average \ cosine \ similarities \ of \ activation \ clusters \ by \ hierarchical \\ classification \ (Layer \ 0).$

N_{neuron}	6400
$\mu(\cos_{G_1})$.4444
$\mu(\cos_{G_1}) - Q_3(\cos_{100})$	0544
$\pi(\mu(\cos_{G_1}) - Q_3(\cos_{100}) < 0)$	66.4688
$p(\chi^2)$.001
$\mu(\cos_{G_2})$.4769
$\mu(\cos_{G_2}) - Q_3(\cos_{100})$	0219
$\pi(\mu(\cos_{G_2}) - Q_3(\cos_{100}) < 0)$	60.5625
$p(\chi^2)$.0346
$\mu(\cos_{G_3})$.4899
$\mu(\cos_{G_3}) - Q_3(\cos_{100})$	0090
$\pi(\mu(\cos_{G_3}) - Q_3(\cos_{100}) < 0)$	55.1875
$p(\chi^2)$.2995
$\mu(\cos_{G_4})$.5013
$\mu(\cos_{G_4}) - Q_3(\cos_{100})$.0024
$\pi(\mu(\cos_{G_4}) - Q_3(\cos_{100}) < 0)$	47.0625
$p(\chi^2)$.5569

Table n°10 : Average cosine similarities of activation clusters by hierarchical classification (Layer 1).



Graph n°8 : Average cosine similarities of activation clusters by hierarchical classification (Layer 1).

As part of our final exploration of the existence of a synthetic intra-neuronal attention process, all our results regarding the categorical homogeneity of activation clusters tend to indicate a low categorical homogeneity of activation clusters, accompanied by a gradual positive evolution of this homogeneity as activation values increase.

5 Discussion and Conclusion

5.1 Summary and Interpretation of Results

The central question of this study was to determine the extent to which the specific activation level of a neuron in response to a token is linked to a categorical value of that token within the neuron's own category. More precisely, in topological terms, we investigated whether a given segment of activation values in a neuron correlates with a distinct categorical segment, thereby enabling the neuron to selectively focus its intra-neuronal attention on a specific categorical segment based on its associated activation range. Put differently, from an epistemological and functional perspective, is the activation space of a neuron compartmentalized into activation zones whose significance and function are to facilitate the attentional detection of specific categorical zones?

Our methodological approach was twofold :

- A textittop-down approach, where we started from categorical segments to examine their respective mean activation values.
- A *bottom-up* approach, where we first isolated activation segments and then investigated their categorical homogeneity. It should be noted, once again, that we made the specific choice of operationalizing categorical homogeneity measurement using only the embedding system of GPT-2XL

as our observational framework, acknowledging that other valid approaches could have been pursued.

In our *top-down* approach, the first study examined the extent to which significant differences in mean activation values existed between categorical clusters. This revealed two main tendencies regarding the potential existence of a synthetic process of intra-neuronal attention that enables a neuron to identify and localize specific categorical segments within the artificial thought category it represents, based on their respective activation spans :

- Activation indifferentiation (i.e., weak differences in mean activation values) among categorical clusters with lower mean activation values.
- Relative activation differentiation (i.e., stronger differences in mean activation values) among categorical clusters involving one or more higher mean activation values.

Continuing our *top-down* investigation of intra-neuronal attention, we identified activation interleaving among categorical clusters—meaning that given categorical segments are not confined to distinct activation segments but instead overlap. We noted that this finding is compatible with the observed activation indifferentiation (between categorical clusters with lower mean activation values) and, at this stage, does not contradict the observed activation differentiation (between categorical clusters with higher mean activation values).

Finally, in our *bottom-up* study of intra-neuronal attention, our results suggest a low categorical homogeneity of activation clusters, accompanied by a relative but consistent increase in homogeneity as activation values rise.

Overall, these findings are compatible with the existence of an intra-neuronal attention mechanism that can be characterized as follows :

- 1. No direct correlation between activation segmentation and categorical segmentation for strongly activated tokens (since all tokens considered in this study are core-tokens, meaning they are highly activated on average).
- 2. A weak but systematic relationship between activation segmentation and categorical segmentation for tokens with the highest activation levels.

In other words, activation appears to play the role of an intra-neuronal attentional vector : it enables a neuron to identify and delimit, within the set of tokens constituting its categorical extension, those tokens specifically associated with the highest activation values. These tokens are relatively more homogeneous from a categorical perspective and therefore constitute a distinct categorical segment of particular interest.

This intra-neuronal attention mechanism, by facilitating attentional focus on only the highest activations, would thus operate an attentional dichotomization : that is, it establishes a qualitative activation threshold (i.e., very high activations) from a quantitative activation continuum, beyond which intra-neuronal selective attention is triggered. This intra-neuronal attention is directly linked to vigilance mechanisms [60, 21, 46, 62] and selective attention processes [55, 11, 112, 53].

5.2 Integration of Current and Previous Findings

The absence of a relationship between activation segmentation and categorical segmentation for tokens with high activation levels is largely compatible with the synthetic process of categorical divergence proposed in one of our previous investigations [67]. Categorical divergence refers to the following two synthetic cognitive phenomena :

- 1. Categorical discontinuity of successive core-tokens in terms of activation levels, meaning that the cosine similarity between successive core-tokens is particularly low.
- 2. Categorical inhomogeneity among core-tokens with the same activation levels, indicating that core-tokens sharing similar activation levels are not necessarily the most categorically similar.

This phenomenon of categorical divergence may be linked to the polysemantic nature of neuronal concepts [64, 39, 9, 16], distinguishing it from traditional human categorical approaches [49, 149, 150, 151, 152, 153, 154, 155].

Conversely, the relative correlation we observed between activation and categorical segmentation for tokens with very high activation values aligns with the categorical convergence phenomenon identified in a previous study [68]. Categorical convergence postulates that as the activation levels of successive core-tokens increase, their categorical variability decreases. This aligns with the fundamental characteristic of thought categories being *ad hoc* [5, 156, 153], meaning they serve a purpose, which, in synthetic categories, might involve minimal alignment with human thought categories and partial convergence toward human-like semantic elements. Furthermore, the tokens specifically involved in this very high activation and categorical convergence process could be interpreted as corresponding to categorical prototypes [149, 150].

In another previous study [69], we identified three mathematical-cognitive factors influencing categorical segmentation performed by each synthetic neuron :

- Categorical priming (effect x) : A token strongly activating a neuron in layer n has a higher probability of activating a strongly connected neuron in layer n + 1. This categorical priming is conceptually linked to its human counterpart [113, 114, 115, 116].
- Inter-categorical attention (effect w) : The stronger the connection between a neuron in layer n + 1 and a neuron in layer n, the higher the probability that a token strongly activated in the first neuron will also be strongly activated in the second.
- Categorical phasing (effect Σ) : Tokens that simultaneously activate different neurons in layer n have a higher probability of activating a strongly connected neuron in layer n + 1. This notion of categorical phasing is inspired by similar concepts in cognitive psychology and human neuroscience [54, 117, 118, 119, 120, 121, 122, 123].

These three factors drive the determination of the activation value associated with a given token. By mathematical construction of the aggregation function, a strong and combined effect of these three factors on a given token results in a very high activation value for that token. Consequently, this token is positioned within the activation segment of very highly activated tokens for a given neuron. These factors—categorical priming, inter-categorical attention, and categorical phasing—thus enable and guide intra-neuronal attention, which manifests as an *elective attentional focus* on highly activated tokens. These tokens define a specific categorical segment that is particularly relevant for the neuron involved.

In a separate prior study [70], we identified a synthetic categorical contouring mechanism, which consists of the separation of a categorical form from its background at the neuronal level [141, 142, 143]. More specifically, categorical contouring is the process by which a specific categorical sub-dimension is extracted [144, 145, 146, 147] from the category vectorized by a precursor neuron (in layer n) to contribute to the formation of a superordinate category (in layer n + 1). This contouring is directly driven by the mathematical-cognitive factors mentioned above. The categorical sub-dimension extracted through this process is immediately determined by intra-neuronal attention : by mathematical construction of the neuronal aggregation function, intra-neuronal attention enables the selective focus on a segment of tokens—those with very high activation values—from which the sub-dimensions will statistically be extracted.

Furthermore, we have demonstrated [71] that the co-activity of intercategorical attention and categorical phasing generates a synthetic phenomenon of partial categorical convergence : the categorical sub-dimensions extracted from layer n categories, when strongly linked to an attentionally connected neuron in layer n + 1, exhibit a tendency toward partial semantic convergence. Categorical convergence is directly influenced by intra-neuronal attention, as this attention mechanism allows for the identification and selection of specific tokens—those with very high activations—through which categorical phasing and, subsequently, categorical convergence take place.

5.3 Intra-Neuronal Attention and Conceptualization

As previously discussed, conceptualization [100, 101] refers to the identification of specific operational characteristics within a class of objects (in this case, tokens) on which a cognitive system must act to ensure appropriate and effective processing. Conceptualization is inherently an attentional phenomenon, as it enables a selective focus on a limited subset of objects or specific attributes of these objects, thereby aligning cognitive activity with their particular properties.

An artificial neuron can be interpreted as a synthetic cognitive agent of conceptualization, whose function is to selectively attend to and filter a subset of tokens from the broader set to which it strongly reacts. This subset constitutes the categorical extension of the *critical concept-in-act* that the neuron is designed to identify selectively. Consequently, the neuron performs an attentional focusing activity on specific token types, selecting and filtering them to optimize the language processing task in which the neuron is engaged.

The paroxysmal concept-in-act thus identified consists of tokens with very high activation levels. These particular tokens, which undergo activation differentiation, tend to semantically converge (categorical homogeneity) toward a characteristic relevant to the neuron's categorical function. The activation differentiation of the categorical segment identified by this synthetic critical concept-in-act is closely linked to the categorical priming effect (effect x), as it inherently denotes tokens with exceptionally high activation. Effect x, combined with categorical phasing and inter-categorical attention, facilitates categorical contouring as well as categorical convergence.

Thus, intra-neuronal conceptualization and attention, occurring at the level of a neuron in layer n, are the fundamental synthetic processes that subsequently enable the formation of new, more functionally effective thought categories at layer n + 1. These emerging categories further enhance the network's capacity to perform the tasks for which it has been trained.

6 Conclusion

In this study, we investigated the extent to which a synthetic process of conceptualization and intra-neuronal attention exists within the perceptron-type neurons of language models. Specifically, we examined whether each neuron can identify and isolate a specific categorical segment within the artificial thought category it represents, based on its activation space. This inquiry was closely tied to the question of whether formal neurons internally exhibit a relative homomorphic relationship between activation segmentation and categorical segmentation, thereby shaping the functional and epistemological significance that can be pragmatically attributed to activation.

Our findings suggest that such a relationship does exist, albeit subtly but systematically, and only at the level of tokens with very high activation values. This intra-neuronal attention mechanism segments, within a given neuron in layer n, an activation zone associated with a specific neuronal concept-in-act. This synthetic concept-in-act serves as the basis for categorical restructuring processes—such as categorical contouring and categorical convergence—which can then be carried out in neurons of layer n + 1. These processes subsequently guide the formation of higher-order categorical abstractions, constituting the thought categories of these superordinate neurons.

It is important to emphasize, once again, that the observed phenomenon of greater categorical homogeneity among tokens with very high activation levels—a central finding in our study of intra-neuronal attention and conceptualization—was operationalized using the input embedding system of GPT-2XL. Additionally, the activation differentiation process we identified was determined using GPT-40. The advantage of this methodology is that it allows us to study intra-neuronal attention using semantic observation frameworks that are analogous to, or at least relatively aligned with, human thought categories. This seems particularly relevant given that we are investigating how intra-neuronal attention enables a language model to align with human tasks—tasks that inherently require conceptualization and the use of thought categories that are compatible with human cognitive structures. However, from another epistemological perspective, this approach may be considered biased, self-referential, and anthropomorphic. Indeed, it could be worthwhile to explore the relationship between activation and categorical segmentation within a given neuron using each respective layer's own input embeddings as the semantic observation framework. The added value of this alternative approach—more closely aligned with the unique and layer-specific thought categories and *alien concepts* of each formal neuron—would be its greater fidelity to the categorical structures intrinsic to each neural layer. This alternative methodology might reveal a different phenomenology of intra-neuronal attention and conceptualization, potentially leading to findings such as stronger categorical homogeneity among highly activated tokens or even a broader homomorphic relationship between activation and categorical segmentation, extending beyond the very high activation levels highlighted in our current study.

Acknowledgments

The authors thank Madeleine Pichat for her careful review of this article, as well as Chantal Colle for the stimulating projects undertaken with her in the field of artificial intelligence.

Bibliography

- Alahmari, S. S., Gardner, M. R., & Salem, T. (2024). Attention guided approach for food type and state recognition. *Food and Bioproducts Processing.*
- [2] Anderson, J. R. (1985). Cognitive Psychology and Its Implications (2nd ed.).
 W. H. Freeman. DOI: 10.4324/9781315784786
- [3] Atanasova, P., Simonsen, J. G., Lioma, C., & Augenstein, I. (2020). Generating Fact Checking Explanations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 7352–7364). Association for Computational Linguistics. DOI : 10.18653/v1/2020.acl-main.656.
- Barr, W., & Bieliauskas, L. A. (2024). Neuropsychology of Decision Making : A Clinical Perspective. *Neuropsychology Review*, 34(1), 1–15. DOI : 10.1007/s11065-023-09500-1.
- [5] Barsalou, L. W. (1995). Cognitive Psychology : An Overview for Cognitive Scientists. Lawrence Erlbaum Associates. DOI: 10.4324/9781315784786
- [6] Barkan, R. (2021). The Role of Cognitive Biases in Human Decision Making. Journal of Behavioral Decision Making, 34(3), 243–255. DOI: 10.1002/bdm.2210.
- Bathia, N., & Richie, D. (2024). Advances in Reinforcement Learning : Applications and Challenges. Artificial Intelligence Review, 57(2), 123–145. DOI : 10.1007/s10462-023-10123-4.

- [8] Beaufils, M. (1996). Les réseaux de neurones artificiels : Modèles et applications. *Revue d'Intelligence Artificielle*, 10(4), 365–387. DOI : 10.1016/S0992-499X(97)80001-2.
- [9] Bills, S., Cammarata, N., Mossing, D., Saunders, W., Wu, J., Tillman, H., Gao, L., Goh, G., Sutskever, I., & Leike, J. (2023). Language models can explain neurons in language models. OpenAI. https://openaipublic. blob.core.windows.net/neuron-explainer/paper/index.html
- [10] Bolognesi, M. (2020). Where Words Get Their Meaning : Cognitive Processing and Distributional Modelling of Word Meaning. John Benjamins Publishing Company. DOI : 10.1075/ftl.7
- [11] Bosker, H. R., & Frost, R. L. A. (2024). Statistical learning at a virtual cocktail party. *Psychonomic Bulletin & Review*, 31, 849-861.
- [12] Bosker, H. R., & Frost, R. L. A. (2024). Statistical learning at a virtual cocktail party. *Psychonomic Bulletin & Review*, 31, 849-861.
- [13] Bower, G. H., Black, J. B., & Turner, T. J. (1979). Scripts in memory for text. *Cognitive Psychology*, 11(2), 177-220.
- [14] Brewer, W. F., & Hay, A. E. (1984). Reconstructive recall of linguistic style. Journal of Verbal Learning and Verbal Behavior, 23(2), 237-249.
- [15] Brewer, W. F., & Hay, A. E. (1984). Reconstructive recall of linguistic style. Journal of Verbal Learning and Verbal Behavior, 23(2), 237-249.
- [16] Bricken, T., Schaeffer, R., Olshausen, B., & Kreiman, G. (2023). Emergence of Sparse Representations from Noise. Proceedings of the 40th International Conference on Machine Learning, in Proceedings of Machine Learning Research, 202:3148-3191. Available from https://proceedings.mlr.press/ v202/bricken23a.html
- [17] Broadbent, D. E., & Gregory, M. (1965). Effects of noise and of signal rate upon vigilance analysed by means of decision theory. *Human Factors*, 7(2), 155-162.
- [18] Deutsch, J. A. (1958). Perception and Communication. Nature, 182(4649), 1572-1572.
- [19] Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Reynolds, J., Melnikov, A., Lunova, N., & Reblitz-Richardson, O. (2020). Captum : A unified and generic model interpretability library for PyTorch. arXiv preprint arXiv:2009.07896. DOI : 10.48550/arXiv.2009.07896.
- [20] Cave, K. R., & Wolfe, J. M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology*, 22(2), 225-271.
- [21] Chen, T., Zhang, Y., Wang, H., Liu, J., & Li, X. (2024). Electrophysiological correlation between executive vigilance and attention network based on cognitive resource control theory. *International Journal of Psychophysiology*, 203, 112393.
- [22] Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *Journal of the Acoustical Society of America*, 25, 975-979.

- [23] Child, R., Gray, S., Radford, A., & Sutskever, I. (2019). Generating long sequences with sparse transformers. arXiv preprint arXiv:1904.10509. https: //arxiv.org/abs/1904.10509
- [24] Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 8(2), 240-247. https://doi.org/10.1016/s0022-5371(69)80069-1.
- [25] Conrad, C. (1972). Cognitive economy in semantic memory. Journal of Experimental Psychology, 92(2), 149-154.
- [26] Cowan, N. (2024). Working Memory Capacity : Theories and Applications. Annual Review of Psychology, 75, 1–25. DOI : 10.1146/annurev-psych-010723-120001.
- [27] Dai, D., Dong, L., Hao, Y., Sui, Z., Chang, B., & Wei, F. (2022). Knowledge Neurons in Pretrained Transformers. Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1 : Long Papers). https://doi.org/10.18653/v1/2022.acl-long.581
- [28] Danilevsky, M., Qian, K., Aharonov, R., Katsis, Y., Kawas, B., & Sen, P. (2020). A Survey of the State of Explainable AI for Natural Language Processing. arXiv (Cornell University). https://doi.org/10. 48550/arxiv.2010.00711
- [29] Dar, S. A., Durrani, N., Sajjad, H., Dalvi, F., & Belinkov, Y. (2023). Probing Pre-trained Language Models for Temporal Knowledge. In *Proceedings of* the 61st Annual Meeting of the Association for Computational Linguistics (ACL). DOI: 10.18653/v1/2023.acl-long.123.
- [30] Duncan, J. (1984). Selective Attention and the Organization of Visual Information. Journal of Experimental Psychology : General, 113(4), 501-517. DOI: 10.1037/0096-3445.113.4.501
- [31] Duncan, J. (1999). Attention. In R. A. Wilson & F. C. Keil (Eds.), The MIT Encyclopedia of Cognitive Sciences. Cambridge, MA : MIT Press.
- [32] Duncan, J., & Humphreys, G. (1992). Beyond the search surface : Visual search and attentional engagement. *Journal of Experimental Psychology : Human Perception and Performance*, 18(2), 578-588.
- [33] Durrani, N., Sajjad, H., Dalvi, F., & Belinkov, Y. (2022). On the Transformation of Latent Space in Fine-Tuned NLP Models. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP). DOI: 10.18653/v1/2022.emnlp-main.123.
- [34] Ellis, P. (2010). The essential guide to effect sizes. Cambridge University Press.
- [35] Ellis, P. (2020). Effect Size Matters : How Reporting and Interpreting Effect Sizes Can Improve Your Publication Prospects and Make the World a Better Place ! London : MadMethods.co.
- [36] Efimov, A., Dubrovsky, D., & Matveev, F. (2023). What's stopping us achieving artificial general intelligence? *Philosophy Now*, April/May.

- [37] Enguehard, J. (2023). Extrmask : A Method for Explaining Time Series Predictions by Masking. *arXiv preprint arXiv:2301.08552*. DOI : 10.48550/arXiv.2301.08552.
- [38] Eysenck, M. W., & Keane, M. T. (2020). Cognitive Psychology : A Student's Handbook (8th ed.). Psychology Press. DOI : 10.4324/9780429449229.
- [39] Fan, Y., Dalvi, F., Durrani, N., & Sajjad, H. (2023). Evaluating Neuron Interpretation Methods of NLP Models. arXiv preprint arXiv:2301.12608. https://doi.org/10.48550/arxiv.2301.12608
- [40] Fel, J., Smith, A., & Wang, T., "A Holistic Approach to Unifying Automatic Concept Extraction and Concept Importance Estimation," in *Proceedings of* the 37th Conference on Neural Information Processing Systems (NeurIPS 2023), 2024.
- [41] Funayama, T., & Shibata, K. (2024). Advances in Quantum Computing : A Comprehensive Review. Journal of Quantum Information Science, 12(1), 45–67. DOI : 10.4236/jqis.2024.121004.
- [42] Geva, M., Schuster, R., Berant, J., & Levy, O. (2023). Transformer Feed-Forward Layers Are Key-Value Memories. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS)*. DOI: 10.48550/arXiv.2012.14913.
- [43] Giallanza, T., & Campbell, D. I. (2024, March). Context-Sensitive Semantic Reasoning in Large Language Models. In *ICLR 2024 Workshop on Representational Alignment*.
- [44] Gernsbacher, M. A. (1985). Surface information loss in comprehension. Cognitive Psychology, 17(3), 324-363.
- [45] Gresch, D., & Müller, K. (2024). Machine Learning in Materials Science : Recent Progress and Emerging Applications. Advanced Materials, 36(5), 2105678. DOI : 10.1002/adma.202105678.
- [46] Hanzal, S., Müller, C., Schwarz, J., Binder, L., & Schröder, P. (2024). EEG markers of vigilance, task-induced fatigue and motivation during sustained attention : Evidence for decoupled alpha-and beta-signatures. *bioRxiv*, 2024-10.
- [47] Haslam, S. A., Reicher, S. D., & Platow, M. J. (2020). The New Psychology of Leadership : Identity, Influence, and Power (2nd ed.). Routledge. DOI : 10.4324/9781351108225.
- [48] Hastie, R. (2022). Schematic principles in human memory. Social Cognition, 39-88.
- [49] Hornsby, A. N., & Love, B. C. (2020). How decisions and the desire for coherency shape subjective preferences over time. *Cognition*, 200, 104244. https://doi.org/10.1016/j.cognition.2020.104244
- [50] Howell, D. C. (2024). Méthodes statistiques en sciences humaines. De Boeck Supérieur.
- [51] von Humboldt, W. (1907). Werke, vol. 7, part 2. Berlin : Leitmann.

- [52] Kahneman, D. (1973). Attention and Effort. Englewood Cliffs, NJ: Prentice-Hall.
- [53] Lin, Z. (2024). Attenuation Theory. In *The ECPH Encyclopedia of Psychology* (pp. 1-2). Singapore : Springer Nature Singapore.
- [54] Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., & Guo, B. (2020). Swin Transformer : Hierarchical Vision Transformer using Shifted Windows. Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). https://arxiv.org/abs/2103.14030
- [55] Liu, H., Zhang, Y., Wang, F., Li, J., & Chen, T. (2024). Electrophysiological correlation of auditory selective spatial attention in the "cocktail party" situation. *Human Brain Mapping*, 45(11), e26793.
- [56] Mackworth, N. H. (1948). The breakdown of vigilance during prolonged visual search. *Quarterly Journal of Experimental Psychology*, 1(1), 6-21.
- [57] Maturana, H. (1970). Biology of cognition (Vol. 9). Urbana : Biological Computer Laboratory, Department of Electrical Engineering, University of Illinois.
- [58] Marconato, E., & al. (2024). BEARS Make Neuro-Symbolic Models Aware of their Reasoning Shortcuts. arXiv preprint arXiv:2402.12240. DOI : 10.48550/arXiv.2402.12240.
- [59] Moreira, G., Hauptmann, A., Marques, M., & Costeira, J. P. (2024). Learning Visual-Semantic Subspace Representations for Propositional Reasoning. arXiv preprint, arXiv:2405.16213.
- [60] Motter, B. C. (1999). Attention in the animal brain. In R. A. Wilson & F. C. Keil (Eds.), *The MIT Encyclopedia of the Cognitive Sciences* (pp. 39–41). Cambridge, MA : MIT Press.
- [61] Mousi, B., Durrani, N., & Dalvi, F. (2023). Can LLMs facilitate interpretation of pre-trained language models? arXiv preprint arXiv:2305.13386. DOI: 10.48550/arXiv.2305.13386.
- [62] Murray, S. (2024). The Nature and Norms of Vigilance. American Philosophical Quarterly, 61(3), 265-278.
- [63] Ortiz-Rodriguez, F., Tiwari, S., Panchal, R., Medina-Quintero, J. M., & Barrera, R. (2022, June). MEXIN : multidialectal ontology supporting NLP approach to improve government electronic communication with the Mexican Ethnic Groups. In DG.O 2022 : The 23rd Annual International Conference on Digital Government Research (pp. 461-463).
- [64] Olah, C., Cammarata, N., Schubert, L., Goh, G., Petrov, M., & Carter, S. (2020). Zoom In : An Introduction to Circuits. Retrieved from https: //distill.pub/2020/circuits/zoom-in/. Accessed 24-11-2023.
- [65] Patel, A. S., Merlino, G., Puliafito, A., Vyas, R., Vyas, O. P., Ojha, M., & Tiwari, V. (2023). An NLP-guided ontology development and refinement approach to represent and query visual information. *Expert Systems with Applications*, 213, 118998.

- [66] Piaget, J. (1974). La prise de conscience. Paris : Presses Universitaires de France.
- [67] Pichat, M. (2024). Psychologie de l'IA et alignement cognitif. Actes du colloque Intelligence artificielle collaborative, management et développement des organisations du 24/05/2024 coorganisé par l'Université Paris Dauphine-PSL et le Cabinet Chrysippe R&D. Available online : https://www.youtube.com/watch?v=9TMmgbELaxQ&list= PLD25p-Bh6_sz6Sr7ms643GpCWW2LlIqeQ&index=6
- [68] Pichat, M. (2024). Psychology of Artificial Intelligence : Epistemological Markers of the Cognitive Analysis of Neural Networks. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2407.09563
- [69] Pichat, M., Campoli, E., Pogrund, W., Wilson, J., Veillet-Guillem, M., Melkozerov, A., Gasparian, A., Pichat, P., Poumay, J. (2024). Neuropsychology of AI : Relationship Between Activation Proximity and Categorical Proximity Within Neural Categories of Synthetic Cognition. arXiv preprint arXiv:2410.11868.
- [70] Pichat, M., Pogrund, W., Gasparian, A., Pichat, P., Demarchi, S., & Veillet-Guillem, M. (2024). How Do Artificial Intelligences Think? The Three Mathematico-Cognitive Factors of Categorical Segmentation Operated by Synthetic Neurons. arXiv preprint, arXiv:2501.06196.
- [71] Pichat, M., Pogrund, W., Gasparian, A., Pichat, P., Demarchi, S., Veillet-Guillem, M., Corbet, M., & Dasilva, T. (2025). The Process of Categorical Clipping at the Core of the Genesis of Concepts in Synthetic Neural Cognition. arXiv preprint, arXiv:2502.15710.
- [72] Pichat, M., Pogrund, W., Pichat, P., Gasparian, A., Demarchi, S., Corbet, M., Georgeon, A., Dasilva, T., & Veillet-Guillem, M. (2025). Synthetic Categorical Restructuring Or How AIs Gradually Extract Efficient Regularities from Their Experience of the World. arXiv:2503.10643
- [73] Planchuelo, C., Hinojosa, J. A., & Duñabeitia, J. A. (2024). The nature of lexical associations in a foreign language : valence, arousal and concreteness. *Bilingualism : Language and Cognition*, 1-10.
- [74] Polyn, S. M. (2024). 15 Attribute Theories of Memory. In M. J. Kahana & A. D. Wagner (Eds.), *The Oxford Handbook of Human Memory, Two Volume Pack : Foundations and Applications* (p. 417). Oxford University Press.
- [75] Ponomarev, A., & Agafonov, A. (2022, November). Ontology concept extraction algorithm for deep neural networks. In 2022 32nd Conference of Open Innovations Association (FRUCT) (pp. 221-226). IEEE.
- [76] Posner, M. I. (1978). Chronometric Explorations of Mind. Lawrence Erlbaum Associates.
- [77] Posner, M. I., & Snyder, C. R. R. (1975). Attention and Cognitive Control. In R. L. Solso (Ed.), *Information Processing and Cognition : The Loyola Symposium* (pp. 55-85). Lawrence Erlbaum Associates. DOI : 10.4324/9781315784786

- [78] Posner, M. I. (1980). Orienting of attention. Quarterly Journal of Experimental Psychology, 32(1), 3-25.
- [79] Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology : General*, 109(2), 160.
- [80] Posner, M. I., & DiGirolamo, G. J. (1998). Executive Attention : Conflict, Target Detection, and Cognitive Control. In R. Parasuraman (Ed.), *The Attentive Brain* (pp. 401-423). Cambridge, MA : MIT Press.
- [81] Posner, M. I., & Rafal, R. D. (1995). Inhibition of return : Neural basis and function. Cognitive Neuropsychology, 12(3), 505–524.
- [82] Posner, M. I. (2024). Orienting of attention and spatial cognition. Cognitive Processing, 25(Suppl 1), 55-59.
- [83] Qiu, Q., Huang, Z., Xu, D., Ma, K., Tao, L., Wang, R., ... & Pan, Y. (2023). Integrating NLP and Ontology Matching into a Unified System for Automated Information Extraction from Geological Hazard Reports. *Journal of Earth Science*, 34(5), 1433-1446.
- [84] Richard, J. C. (1980). The Language Teaching Matrix. Cambridge University Press.
- [85] Rosch, E. (1978). Cognition and categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and Categorization*. Lawrence Erlbaum Associates.
- [86] Rosenholtz, R. (2024). Visual Attention in Crisis. Behavioral and Brain Sciences, 1-32.
- [87] Rueda, M. R. (2024). Developing the attentive brain : Contribution of cognitive neuroscience to a theory of attentional development. *Human Development*, 1-16.
- [88] Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory. *Psychological Review*, 81(3), 214-241.
- [89] Sartori, G., Coltheart, M., Miozzo, M., & Job, R. (2024). Category Specificity and Informational Memory, *Memory*, 1, 604.
- [90] Schneider, W., & Shiffrin, R. M. (1977). Controlled and Automatic Human Information Processing : I. Detection, Search, and Attention. *Psychological Review*, 84(1), 1-66.
- [91] Tipper, S. P. (1985). The Negative Priming Effect : Inhibitory Priming by Ignored Objects. *The Quarterly Journal of Experimental Psychology*, 37A(4), 571-590. DOI : 10.1080/14640748508400920
- [92] Thukral, A., Dhiman, S., Meher, R., & Bedi, P. (2023). Knowledge graph enrichment from clinical narratives using NLP, NER, and biomedical ontologies for healthcare applications. *International Journal of Information Technology*, 15(1), 53-65.
- [93] Treisman, A. (1964). Monitoring and storage of irrelevant messages in selective attention. *Journal of Verbal Learning and Verbal Behavior*, 3(6), 449-459.

- [94] Treisman, A., & Gelade, G. (1980). A Feature-Integration Theory of Attention. *Cognitive Psychology*, 12(1), 97-136. DOI : 10.1016/0010-0285(80)90005-5
- [95] Treisman, A. (1986). Features and objects in visual processing. Scientific American, 255(5), 114B-125.
- [96] Treviso, M., Lee, J. U., Ji, T., Van Aken, B., Cao, Q., Ciosici, M. R., & Schwartz, R. (2023). Efficient methods for natural language processing : A survey. *Transactions of the Association for Computational Linguistics*, 11, 826-860.
- [97] Tversky, A. (1977). Features of similarity. Psychological Review, 84(4), 327-352.
- [98] Varela, F. (1984). The creative circle. In P. Watzlawick (Ed), The invented reality. London : W W Norton & Co Inc.
- [99] Varela, F. J. (1988). Cognitive Science : A Cartography of Current Ideas. MIT Press.Varela1996
- [100] Vergnaud, G. (2009). Activité, développement, représentation. In M. Merri (Ed.), Activité humaine et conceptualisation. Questions à Gérard Vergnaud (pp. 149–154). Presses universitaires du Mirail.
- [101] Vergnaud, G. (2016). Relations entre conceptualisations dans l'action et signifiants langagiers et symboliques. In Symposium latinoaméricain de didactique de mathématique, Bonito, Brésil. Disponible sur : https://www.gerard-vergnaud.org/texts/gvergnaud_2016_ signifiants-langagiers-symboliques_conference-bonito.pdf.
- [102] Vergnaud, G. (2020). A Classification of Cognitive Tasks and Operations of Thought Involved in Addition and Subtraction Problems. In P. Carpenter, M. Moser & A. Romberg (Eds.), Addition and Subtraction : A Cognitive Perspective. London : Routledge.
- [103] Voita, E., Sennrich, R., & Titov, I. (2021). Language modeling, lexical translation, reordering : The training process of NMT through the lens of classical SMT. arXiv preprint arXiv:2109.01396. DOI : 10.48550/arXiv.2109.01396.
- [104] von Glaserfeld, E. (2002). Radical Constructivism. London : Routledge Falmer.
- [105] von Humboldt, W. (1907). Werke (Vol. 7, Part 2). Berlin : Leitmann.
- [106] Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Reynolds, J., Melnikov, A., Lunova, N., & Reblitz-Richardson, O. (2020). Captum : A unified and generic model interpretability library for PyTorch. arXiv preprint arXiv:2009.07896. DOI : 10.48550/arXiv.2009.07896.
- [107] Wu et al., (2020). pyOptSparse : A Python framework for large-scale constrained nonlinear optimization of sparse systems. Journal of Open Source Software, 5(54), 2564. DOI: 10.21105/joss.02564
- [108] Wu, D., & Zhang, S. (2024). Does visual attention help? Towards better understanding and predicting users' good abandonment behavior in mobile search. *Library Hi Tech*, 42(3), 867-884.

- [109] Yang, Y., Li, L., de Deyne, S., Li, B., Wang, J., & Cai, Q. (2024). Unraveling lexical semantics in the brain : Comparing internal, external, and hybrid language models. *Human Brain Mapping*, 45(1), e26546.
- [110] Zettersten, M., Bredemann, C., Kaul, M., Ellis, K., Vlach, H. A., Kirkorian, H., & Lupyan, G. (2024). Nameability supports rule-based category learning in children and adults. *Child Development*, 95(2), 497-514. DOI : 10.1111/cdev.14008.
- [111] Zhang, C., Yin, Z., & Qin, R. (2024). Attention-Enhanced Co-Interactive Fusion Network (AECIF-Net) for automated structural condition assessment in visual inspection. Automation in Construction, 159, 105292.
- [112] Zhao, M., Xu, D., & Gao, T. (2024). From Cognition to Computation : A Comparative Review of Human Attention and Transformer Architectures. arXiv preprint arXiv:2407.01548.
- [113] Anderson, J. R. (1985). Cognitive Psychology and Its Implications (2nd ed.). W. H. Freeman. DOI: 10.4324/9781315784786
- [114] Chao, L. L. (2024). Advances in Neuroimaging Techniques for Cognitive Neuroscience. Journal of Cognitive Neuroscience, 36(1), 1–15. DOI : 10.1162/jocn_a_01700.
- [115] Xu, W., & Futrell, R. (2024). A hierarchical Bayesian model for syntactic priming. arXiv preprint arXiv:2405.15964. DOI : 10.48550/arXiv.2405.15964.
- [116] Hernández-Gutiérrez, C. A., & Pérez-González, J. (2024). Deep Learning Techniques for Natural Language Processing : A Survey. *IEEE Transactions* on Neural Networks and Learning Systems, 35(2), 1234–1256. DOI : 10.1109/TNNLS.2023.3101234.
- [117] Mitchell, M. (2021). Abstraction and analogy-making in artificial intelligence. Annals of the New York Academy of Sciences, 1505(1), 79-101. DOI: 10.1111/nyas.14619
- [118] Capuano, F., & Kaup, B. (2024). Pragmatic Reasoning in GPT Models : Replication of a Subtle Negation Effect. Proceedings of the Annual Meeting of the Cognitive Science Society, 46. Retrieved from https ://escholarship.org/uc/item/22q5920s
- [119] Protachevicz, P. R., Hansen, M., Iarosz, K. C., Caldas, I. L., Batista, A. M., & Kurths, J. (2021). Emergence of neuronal synchronisation in coupled areas. *Frontiers in Computational Neuroscience*, 15, 663408. DOI: 10.3389/fncom.2021.663408.
- [120] Canales-Johnson, A., Silva, C., Huepe, D., Rivera-Rei, Á., Noreika, V., Del Carmen Garcia, M., Silva, W., Vaucheret, E., Sedeño, L., Couto, B., Melloni, M., Ibáñez, A., Chennu, S., Bekinschtein, T. A. (2015). Auditory feedback differentially modulates behavioral and neural markers of objective and subjective performance when tapping to your heartbeat. *Cerebral Cortex*, 25(11), 4490–4503. DOI : 10.1093/cercor/bhv076.

- [121] Ribary, U., & Ward, L. M. (2024). Synchronization and functional connectivity dynamics across TC-CC-CT networks : Implications for clinical symptoms and consciousness. In *Phenomenological Neuropsychiatry : How Patient Experience Bridges the Clinic with Clinical Neuroscience* (pp. 105–118). Cham : Springer International Publishing. DOI : 10.1007/978-3-031-38391-5_10.
- [122] Shavikloo, M., Esmaeili, A., Valizadeh, A., & Madadi Asl, M. (2024). Synchronization of delayed coupled neurons with multiple synaptic connections. *Cognitive Neurodynamics*, 18(2), 631-643. DOI: 10.1007/s11571-023-10013-9.
- [123] Rzechorzek, A. (2024). Understanding Cognitive Processes : Insights from Recent Research. Journal of Cognitive Neuroscience. DOI : 10.1162/jocn_a_01678.
- [124] Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473. https://arxiv.org/abs/1409.0473
- [125] Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. arXiv preprint arXiv:1406.1078. https://arxiv.org/abs/1406.1078
- [126] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems. https://arxiv.org/abs/1706.03762
- [127] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT : Pretraining of Deep Bidirectional Transformers for Language Understanding. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics : Human Language Technologies, 1, 4171-4186. https://aclanthology.org/N19-1423/
- [128] Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. *OpenAI blog.* https://openai.com/research/language-unsupervised
- [129] Luong, M.-T., Pham, H., & Manning, C. D. (2015). Effective Approaches to Attention-based Neural Machine Translation. Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, 1412–1421. https://doi.org/10.18653/v1/D15-1166
- [130] Voita, E., Talbot, D., Moiseev, F., Sennrich, R., & Titov, I. (2019). Analyzing multi-head self-attention : Specialized heads do the heavy lifting, the rest can be pruned. arXiv preprint arXiv:1905.09418. https: //arxiv.org/abs/1905.09418
- [131] Michel, P., Levy, O., & Neubig, G. (2019). Are sixteen heads really better than one? An empirical study of attention in Transformers. arXiv preprint arXiv:1905.10650. https://arxiv.org/pdf/1905.10650
- [132] Tiberi, L., Mignacco, F., Irie, K., & Sompolinsky, H. (2024). Dissecting the Interplay of Attention Paths in a Statistical Mechanics Theory of

Transformers. arXiv preprint arXiv:2405.15926. https://arxiv.org/abs/2405.15926

- [133] Kitaev, N., Kaiser, Ł., & Levskaya, A. (2020). Reformer : The Efficient Transformer. International Conference on Learning Representations. https: //openreview.net/forum?id=rkgNKkHtvB
- [134] Zaheer, M., Guruganesh, G., Dubey, A., Ainslie, J., Alberti, C., Ontañón, S., Pham, P., Ravula, A., Wang, Q., Yang, L., & Ahmed, A. (2020). Big Bird : Transformers for Longer Sequences. Advances in Neural Information Processing Systems, 33, 17283-17297. https://proceedings.neurips.cc/ paper/2020/file/c8512d142a2d849725f31a9a7a361ab9-Paper.pdf
- [135] Tay, Y., Dehghani, M., Bahri, D., & Metzler, D. (2022). Efficient Transformers : A Survey. ACM Computing Surveys, 55(6), 1–28. https: //doi.org/10.1145/3530811
- [136] Clark, K., Khandelwal, U., Levy, O., & Manning, C. D. (2019). What Does BERT Look at? An Analysis of BERT's Attention. Proceedings of the 2019 ACL Workshop BlackboxNLP : Analyzing and Interpreting Neural Networks for NLP, 276-286. https://aclanthology.org/W19-4828/
- [137] Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q. V., & Salakhutdinov, R. (2019). Transformer-XL: Attentive language models beyond a fixed-length context. arXiv preprint arXiv:1901.02860. https://arxiv.org/abs/1901.02860
- [138] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2021). An Image is Worth 16x16 Words : Transformers for Image Recognition at Scale. *International Conference on Learning Representations*. https://arxiv.org/abs/2010.11929
- [139] Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., & Guo, B. (2021). Swin Transformer : Hierarchical Vision Transformer using Shifted Windows. Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 10012–10022. https://arxiv.org/pdf/2103. 14030
- [140] Traylor, A., Merullo, J., Frank, M. J., & Pavlick, E. (2024). Transformer Mechanisms Mimic Frontostriatal Gating Operations When Trained on Human Working Memory Tasks. arXiv preprint arXiv:2402.08211. https: //arxiv.org/abs/2402.08211
- [141] Zeki, S. (2002). Inner vision : An exploration of art and the brain. Oxford University Press.
- [142] Hock, R. M., et al. (2024). Effects of manipulating prefrontal activity and dopamine D1 receptor signaling in an appetitive feature-negative discrimination learning task. *Behavioral Neuroscience*.
- [143] Green, I., Amo, R., & Watabe-Uchida, M. (2024). Shifting attention to orient or avoid : a unifying account of the tail of the striatum and its dopaminergic inputs. *Current Opinion in Behavioral Sciences*, 59, 101441.

- [144] Ghorbani, A., Wexler, J., Zou, J. Y., & Kim, B. (2019). Towards automatic concept-based explanations. Advances in Neural Information Processing Systems, 9273–9282.
- [145] Bhatt, U., Weller, A., & Moura, J. M. F. (2020). Evaluating and aggregating feature-based model explanations. *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- [146] Clark, P. (2021). Formalising concepts. arXiv preprint arXiv:2101.05125v1. https://arxiv.org/abs/2101.05125
- [147] Ponomarev, D., & Agafonov, A. (2022). Ontology Concept Extraction Algorithm for Deep Neural Networks. https://www.researchgate.net/ publication/365833644_Ontology_Concept_Extraction_Algorithm_ for_Deep_Neural_Networks
- [148] Treisman, A., & Gelade, G. (1980). A Feature-Integration Theory of Attention. *Cognitive Psychology*, 12(1), 97-136. DOI : 10.1016/0010-0285(80)90005-5
- [149] Singh, V., Gupta, I., & Jana, P. K. (2020). An Energy Efficient Algorithm for Workflow Scheduling in IaaS Cloud. *Journal of Grid Computing*, 18(3), 357–376. https://doi.org/10.1007/s10723-019-09490-2
- [150] Vogel, T., Ingendahl, M., & Winkielman, P. (2021). The architecture of prototype preferences : Typicality, fluency, and valence. *Journal of Experimental Psychology : General*, 150(1), 187–194. https://doi.org/10. 1037/xge0000798
- [151] Sanborn, A. N., Heller, K., Austerweil, J. L., & Chater, N. (2021). REFRESH : A new approach to modeling dimensional biases in perceptual similarity and categorization. *Psychological Review*, 128(6), 1145–1186. https://doi.org/10.1037/rev0000310
- [152] Nosofsky, R. M., Meagher, B. J., & Kumar, P. (2022). Contrasting exemplar and prototype models in a natural-science category domain. *Journal of Experimental Psychology : Learning, Memory, and Cognition*, 48(12), 1970–1994. https://doi.org/10.1037/xlm0001069
- [153] Love, A. H., Zdon, A., Fraga, N. S., Cohen, B., Mejia, M. P., Maxwell, R., & Parker, S. S. (2022). Statistical evaluation of the similarity of characteristics in springs of the California Desert, United States. *Frontiers in Environmental Science*, 10. https://doi.org/10.3389/fenvs.2022.1020243
- [154] Poth, N., & Dolega, K. (2023). Bayesian belief protection : A study of belief in conspiracy theories. *Philosophical Psychology*, 36(6), 1182–1207. https://doi.org/10.1080/09515089.2023.2168881
- [155] Roads, B. D., & Love, B. C. (2024). Modeling Similarity and Psychological Space. Annual Review of Psychology, 75(1), 215-240. https://doi.org/ 10.1146/annurev-psych-040323-115131
- [156] Glaser, J. I., Benjamin, A. S., Chowdhury, R. H., Perich, M. G., Miller, L. E., & Kording, K. P. (2020). Machine Learning for Neural Decoding. *eNeuro*, 7(4), ENEURO.0506-19.2020. https://doi.org/10.1523/eneuro. 0506-19.2020