

Are Humans and LLMs Confused by the Same Code? An Empirical Study on Fixation-Related Potentials and LLM Perplexity

Youssef Abdelsalam
Saarland University
Saarbrücken, Germany

Norman Peitek
Saarland University
Saarbrücken, Germany

Anna-Maria Maurer
Saarland University
Saarbrücken, Germany

Mariya Toneva
MPI for Software Systems
Saarbrücken, Germany

Sven Apel
Saarland University
Saarbrücken, Germany

Abstract

Already today, humans and programming assistants based on large language models (LLMs) collaborate in everyday programming tasks. Clearly, a misalignment between how LLMs and programmers comprehend code can lead to misunderstandings, inefficiencies, low code quality, and bugs. A key question in this space is whether humans and LLMs are confused by the same kind of code. This would not only guide our choices of integrating LLMs in software engineering workflows but also inform about possible improvements of LLMs. To this end, we conducted an empirical study comparing an LLM to human programmers comprehending clean and confusing code. We operationalized comprehension for the LLM by using *LLM perplexity*, and for human programmers using *neurophysiological responses* (in particular, EEG-based fixation-related potentials). We found that LLM perplexity spikes correlate, both in terms of location and amplitude, with human neurophysiological responses that indicate confusion. This result suggests that LLMs and humans are similarly confused about the code. Based on these findings, we devised a data-driven, LLM-based approach to identify regions of confusion in code that elicit confusion in human programmers.

CCS Concepts

• General and reference → Empirical studies; Metrics; • Computing methodologies → Natural language generation.

Keywords

Large Language Models, Program Comprehension, Atoms of Confusion, Perplexity, Fixation-Related Potentials

ACM Reference Format:

Youssef Abdelsalam, Norman Peitek, Anna-Maria Maurer, Mariya Toneva, and Sven Apel. 2026. Are Humans and LLMs Confused by the Same Code? An Empirical Study on Fixation-Related Potentials and LLM Perplexity. In *2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE '26)*, April 12–18, 2026, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3744916.3773219>



This work is licensed under a Creative Commons Attribution 4.0 International License. *ICSE '26, Rio de Janeiro, Brazil*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2025-3/26/04
<https://doi.org/10.1145/3744916.3773219>

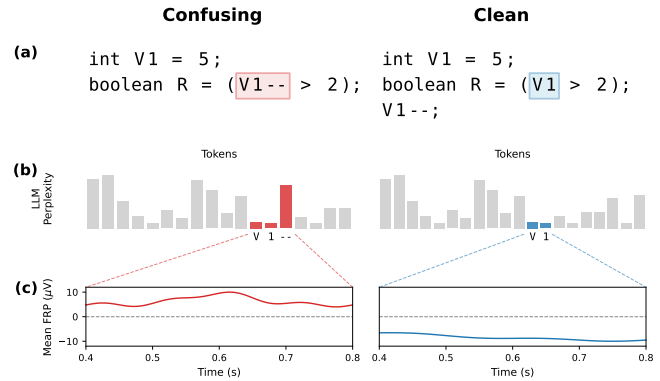


Figure 1: A comparison of **confusing and corresponding **clean** code (a) as well as effects on LLM perplexity (b) and neurophysiological responses (c): The **confusing** code snippet includes an atom of confusion (highlighted in red). The corresponding **clean** snippet avoids this atom by extracting the decrement-expression into a separate statement (a). The token-level perplexity from an LLM increases for the confusing code region (b). Corresponding EEG fixation-related potentials exhibit increased amplitude during the comprehension of the confusing snippet (c).**

1 Introduction

Large language models (LLMs) have rapidly advanced and demonstrated remarkable performance in many programming-related tasks [28], including code completion [31, 34], synthesis [54], and repair [27, 89]. Trained on massive corpora of data, LLMs are able to generate syntactically correct and semantically plausible code across all areas of software engineering, from education [24, 37, 62, 63] and research [11, 18, 75] to practice [23, 38]. Already today, LLM-based coding assistants and human programmers work in concert to produce substantial amounts of source code [9]. Clearly, a misalignment between LLM and human programmers in such a collaborative setting can lead to misunderstandings, inefficiencies, low code quality, and bugs [70, 79, 88]. To maximize the potential of LLMs, we are well advised to critically assess the alignment [1, 3, 53, 68] of how LLMs and humans comprehend source code.

The premise of our investigation is that understanding (the presence and extent of) this alignment would allow us to better understand possible risks and pitfalls of integrating LLMs into software engineering workflows and to improve the capabilities of

LLMs in collaborating with humans in programming-related tasks. Specifically, we are interested in the reaction of LLMs and humans to *confusing* code. Prior research has dedicated considerable attention to problematic or confusing code [14, 21], even offering validated stimuli that elicit distinct neurophysiological responses in humans [4, 86, 87]. Thus, confusing code provides an empirically grounded and well-studied setting to evaluate the alignment between human and LLM comprehension.

Research questions and methodology. Our investigation is guided by three research questions:

- RQ₁: How does confusing code affect LLM confusion?
- RQ₂: How does LLM confusion about code relate to human confusion about the same code?
- RQ₃: How well is a data-driven, LLM-based approach able to identify code regions that confuse humans?

For our investigation, we rely on a well-studied set of syntactically valid and semantically correct code idioms that are known to confuse human programmers, called *atoms of confusion* [21]. Figure 1(a) displays an example pair of an atom of confusion in a confusing and a corresponding clean (i.e., less confusing) code snippet. Atoms of confusion lead to longer reading times, higher error rates, and lower comprehension accuracy across programming languages [8, 21, 42]. Studies involving visual attention [8, 12] and neurophysiological measures using electroencephalography (EEG) [86, 87] further corroborate these results.

To analyze how humans comprehend code, we rely on EEG data from prior work. EEG has been a well-established method for measuring programmers' neurophysiological responses to source code [7, 33, 43, 46, 51, 58, 86]. In a study on atoms of confusion, Bergum et al. combined EEG with eye tracking to measure fixation-related potentials (FRPs). They found that confusing code reliably elicits a so-called *late frontal positivity* in the EEG signal of human programmers (see Figure 1(c)), which is a neurophysiological correlate of human confusion [41]. Therefore, we also measure human confusion using fixation-related potentials with the focus on the amplitude of late frontal positivity.

To analyze how LLMs comprehend code, we use the measure of *perplexity*. Perplexity reflects how unexpected a token is to the LLM given its context (see Figure 1(b)). In natural language processing and psycholinguistics, the unexpectedness of natural language text serves as a well-established proxy for human confusion and cognitive load [44], which has been confirmed also for source code [6, 26].

Combining the pieces, we have devised a methodology that (1) computes token- and snippet-level perplexity values with state-of-the-art LLMs, (2) extends this analysis to calculate perplexity at predefined areas of interest (AOIs), (3) aligns these with fixation-related potentials in EEG signals, and (4) detects perplexity peaks within code snippets to define new AOIs, extracts corresponding FRPs, and aligns these measures accordingly. We leveraged the dataset of Bergum et al. [4] for fixation-related potentials and evaluated several mainstream LLMs of varying sizes (see Section 4.2.1 for details on model selection). As the models yielded similar results, we report findings particularly for Qwen2.5-Coder-32B [30], which held the top Java pass@1 score on the Big Code Models Leaderboard at the time of our experiments.

Results. Our results demonstrate that the presence of atoms of confusion consistently triggers a spike in LLM perplexity, revealing that LLMs are also sensitive to these confusing code constructs, corroborating prior findings on human programmers. Moreover, when correlating the LLM perplexity values with the strength of human confusion measured by the amplitude of corresponding fixation-related potentials, we find a consistent, statistically significant positive relationship. The confusion of LLMs and humans is not only spatially aligned (based on tokens) but also in terms of their amplitude. Based on these findings, we have devised a data-driven, LLM-based approach to identify code regions of confusion that are associated with human confusion.

Overall, we find an alignment between LLM perplexity and human neurophysiological responses, which is an important step toward establishing LLMs as viable surrogate models of human program comprehension. This paves the way to deepen our understanding of how programmers comprehend code and supports the development of LLMs to be more closely aligned with human comprehension patterns. Our investigations suggest that LLMs may serve as a foundation for data-driven feedback methods that flag or revise cognitively demanding code based on LLM perplexity and facilitate the development of tools that automatically refactor confusing code or highlight areas where programmers may need additional information to comprehend the code.

Contributions. We make the following contributions:

- An investigation of the effects of confusing code (based on atoms of confusion) on LLM perplexity;
- A novel combined analysis of LLM perplexity and human confusion measured in terms of fixation-related potentials, showing a positive correlation between LLM perplexity and neurophysiological responses indicative of human confusion;
- A data-driven approach to detect potential regions of confusion using LLM perplexity, validated with EEG data to confirm the effect of the code region on human confusion; and
- A replication package with our data, the entire analysis pipeline, and additional details.¹

2 Background and Related Work

In the following, we provide the background and summarize prior work related to our study. We review three key concepts: (1) atoms of confusion, a set of small code patterns known to impair human comprehension despite being syntactically and semantically valid; (2) neurophysiological measures, particularly EEG and fixation-related potentials, which offer fine-grained insights into the neurophysiological responses underlying program comprehension and confusion; and (3) LLM perplexity, a measure of model uncertainty that has been linked to human difficulty in processing both natural language and source code. Together, these lines of work provide the theoretical and empirical grounding for our investigation into how LLMs and human programmers respond to confusing code.

¹<https://github.com/brains-on-code/llm-perplexity-AoC>

2.1 Atoms of Confusion

Over the past decade, atoms of confusion (AoCs) have emerged as a key concept for understanding the cognitive challenges programmers face when reading source code. These atoms are short, syntactically and semantically valid code snippets that often tend to cause confusion. First introduced by Gopstein et al., atoms of confusion were identified through behavioral studies in C and C++ that revealed certain patterns, such as misleading variable shadowing or misleading increment placement (see Figure 1a), that frequently led to human confusion during program comprehension tasks [21]. Subsequent work extended the study of atoms of confusion to other programming languages. Langhout and Aniche adapted the original atoms to Java and found that they led to significantly higher error rates and increased perceived difficulty [42]. Da Costa et al. identified similar effects in Python, showing that atoms of confusion increase the visual effort and comprehension difficulty [8].

Table 1: Illustrative **confusing code (atoms of confusion) and **clean** counterparts based on Bergum et al. [4]**

Atom of Confusion	Confusing Variant	Clean Variant
Arithmetic as Logic	<code>(V1 + 5 != 0);</code>	<code>(V1 != -5);</code>
Change of Literal Encoding	<code>12 & 3;</code>	<code>0b1100 & 0b0011;</code>
Constant Variables	<code>V1 = V2;</code>	<code>V1 = 5;</code>
Dead, Unreachable, Repeated	<code>R = 3; R = 2;</code>	<code>R = 2;</code>
Implicit Predicate	<code>!(V1 - V2 < 6)</code>	<code>(V1 - V2 > 5)</code>
Remove Indentation	<code>while (V1 > 0) V1--; R++;</code>	<code>while (V1 > 0) V1--; R++;</code>
Infix Operator Precedence	<code>2 - 4 / 2;</code>	<code>2 - (4 / 2);</code>
Omitted Curly Braces	<code>for (..) R++; R++;</code>	<code>for (..) { R++; } R++;</code>
Post-Increment/Decrement	<code>int R = 3 + V1++;</code>	<code>int R = 3 + V1; V1++;</code>
Pre-Increment/Decrement	<code>int R = ++V1 - 2;</code>	<code>++V1; int R = V1 - 2;</code>
Type Conversion	<code>(byte) V1;</code>	<code>(byte) (V1 % 256);</code>

Studies using mining software repository approaches have shown that atoms of confusion are not only prevalent in practice in large-scale projects, such as the Linux kernel, but are also frequently discussed or removed in later commits [22]. Moreover, mining-based analyses of Java repositories further linked atoms of confusion to bug-prone code and subsequent maintenance activity, highlighting their real-world relevance [52, 61].

Despite these consistent findings, atoms of confusion remain under-represented in programming guidelines [52]. A programmer survey has shown that, while a majority of programmers recognize certain patterns as confusing, they are often reluctant to refactor code that is syntactically correct and functionally complete [50]. A follow-up study using a think-aloud protocol suggests that behavioral accuracy may not always indicate true understanding: Some participants use external cues—such as test outputs—to guess correct results even when their mental models are incorrect [20].

2.2 Neurophysiological Correlates of Confusion

Detailed insights into program comprehension can be gained by analyzing the behavior of programmers on a neurophysiological level. *Electroencephalography* (EEG) is a non-invasive method for capturing neural activity with high temporal resolution, making it

particularly well-suited for investigating the time course of neurophysiological responses when humans are interpreting stimuli [35]. To reduce noise, it is established to average EEG responses from multiple trials time-locked to the start events, called *event-related potentials* (ERPs) [48]. For simple stimuli, ERPs are created time-locked onto stimulus onset, thus capturing the initial neurophysiological responses of participants regarding the stimuli. In contrast, for more complex stimuli, the setup can be extended with an eye tracker synchronized to the EEG, which allows the capture of neurophysiological responses relative to fixations in a region of interest, thereby calculating *fixation-related potentials* (FRPs) [32]. For both ERP and FRP, the EEG signal must be leveled to a baseline prior to aggregation, and commonly, the resulting ERP and FRP waves are compared in a contrast between multiple conditions in form of difference waves [48]. These difference waves capture the difference in neurophysiological activity at specific time points, and, to an extent, also scalp locations. This allows to observe which ERP components (i.e., established cognitive processes measured in the brain) are involved or show increased activity during the comprehension process. Multiple ERP components have been identified in prior research across many fields [15], including ones that specifically occur when processing natural language [76]. For example, one ERP component is the N400, a negative impact at about 400 ms after the event, which reacts to semantic incongruency [40, 41].

Recently, researchers have started to adapt these methods to software engineering settings. They have measured program comprehension using ERPs and FRPs and mapped the identified components to those during comprehending natural language [4, 39]. Kuo and Prat combined an ERP approach with token-wise presentation of code snippets to detect that programmers react to syntax and semantic violations in code with the same ERP components as when reading text in natural language [39]. Bergum et al. analyzed the neurophysiological response of programmers to atoms of confusion while determining the output of a code snippet using an FRP approach. They identified that these responses correspond to *late frontal positivity* [4]. This is a positive ERP component concentrated on the frontal part of the scalp and peaks at about 600 ms. It has been linked to unexpected but plausible sentence endings in natural language. They commonly occur in sentences that generate high expectations (i.e., high probability) for one specific sentence ending, but where the actual ending does not match that expectation (i.e., low probability) [5, 41, 77, 80]. For example, in the sentence *he bought her a pearl necklace for her ...,* people commonly expect *birthday*. When a less expected ending, such as *collection*, is presented, the neurophysiological response exhibits an increased amplitude of the late frontal positivity ERP component. Overall, these initial studies show fundamental similarities in the neurophysiological responses to natural language and source code. Complementary studies using functional magnetic resonance imaging (fMRI) provide corroborating evidence for this observation [57, 59, 60, 71, 72].

2.3 LLM Perplexity as a Measure of Confusion

Perplexity is an established metric for evaluating language models by quantifying their uncertainty in predicting sequences of tokens based on their probability [69, 83]. Concretely, a language model defines probability distributions over sequences, enabling it both

to generate plausible sentences and to evaluate the likelihood of existing ones [36]. A well-trained language model should assign higher probabilities to coherent and well-formed sequences, and lower probabilities to unexpected ones. Formally, given a token sequence $T = (t_1, t_2, \dots, t_n)$, the language model assigns a joint probability:

$$\Pr(T) = \prod_{i=1}^n \Pr(t_i \mid t_1, \dots, t_{i-1}).$$

Because this product decreases exponentially with sequence length regardless of sequence quality, a direct comparison of sequence probabilities across different token/sequence lengths is not meaningful [36]. To address this issue, the *geometric mean* of the token probabilities is used to normalize the sequence probability

$$\Pr_{\text{norm}}(T) = \Pr(T)^{1/n},$$

which corresponds to the average per-token likelihood. *Average Perplexity* is then defined as the inverse of this probability:

$$\text{PPL}(T) = \frac{1}{\Pr_{\text{norm}}(T)} = \left(\frac{1}{\Pr(T)} \right)^{1/n}.$$

This formulation highlights that perplexity quantifies the average uncertainty of the language model when predicting each token in the sequence, independent of sequence length [6].

In natural language, lower perplexity indicates more predictable constructs, whereas higher perplexity reflects greater uncertainty and difficulty in prediction [44, 73, 84]. These effects are well-established in psycholinguistics, where higher perplexity has been linked to slower reading times and increased cognitive load [44, 73].

In the context of source code, perplexity similarly reflects how “natural” or predictable the code is to the model [26]. Code exhibits strong statistical regularities, much like natural language, but with even greater predictability, as programmers tend to favor idiomatic and repetitive patterns [6]. Casalnuovo et al. experimentally tested whether higher perplexity in code actually relates to difficulty in humans understanding code. In a controlled user study, programmers were shown functionally identical code snippets with different perplexities. They found that the versions with higher perplexity led to slower response times and lower accuracy on comprehension questions [6]. Goodkind and Bicknell even showed that a token’s perplexity linearly predicts human reading time [19].

Building on these insights, recent research has further employed token-level perplexity to assess the predictability of individual tokens rather than entire code snippets and enable a fine-grained analysis [84, 85]. For instance, Xu et al. used this approach to distinguish between human-written and LLM-generated code, observing that code produced by LLMs exhibits more uniform perplexity [84].

These findings suggest that perplexity not only captures structural regularities in code, but also reacts to unexpected or complex tokens on a fine-grained level, making it a practical and cognitively relevant metric for code analysis. However, the use of token-level perplexity to directly model human neurophysiological responses during program comprehension remains largely unexplored.

2.4 Research Gap

Prior work has established that atoms of confusion negatively impact program comprehension across several programming languages. They occur frequently in software repositories and are associated with increased error rates, maintenance effort, and bug-prone code. It has been demonstrated that LLMs assign higher perplexity to certain code patterns, and that this perplexity correlates in cases with human difficulty during comprehension tasks. Recent studies using neurophysiological measures such as fixation-related potentials have shown that confusing code elicits distinct neurophysiological responses, particularly late frontal positivity, indicating increased mental effort during comprehension. Both LLM perplexity and late frontal positivity have one fundamental aspect in common: They are affected by what LLMs, respectively humans, perceive as a probable sequence based on the previously given context and training or experience.

Together, these lines of work suggest that there may be meaningful overlap between patterns that confuse humans, surprise language models, and trigger distinct neurophysiological responses. While atoms of confusion are known to induce confusion, and perplexity has been linked to comprehension difficulty in isolated studies, it is unclear whether LLMs assign higher perplexity to these known confusing constructs or whether LLM perplexity aligns with human neurophysiological responses during program comprehension. If such alignment exists, it raises the question of whether we can move beyond static, predefined patterns and use perplexity itself to automatically identify *regions of confusion* in code.

3 Research Questions

Based on this research gap, we structure our investigation around three main research questions, each exploring a different aspect of the relationship between confusing code constructs, LLM confusion, and human neurophysiological responses observed during program comprehension.

RQ1: How does confusing code affect LLM confusion? RQ₁ addresses whether LLM confusion reflects known cognitive challenges posed by confusing code. Here, we operationalize confusing code using the established atoms of confusion and LLM confusion to this code with perplexity. To establish perplexity as a computational measure of confusion, we compare the perplexity of code snippets containing atoms of confusion to their clean, functionally equivalent counterparts.

Based on the literature (cf. Section 2), we expect code snippets containing atoms of confusion to exhibit (1) higher perplexity particularly for the regions containing the atoms of confusion (as shown in Figure 1(b)), and (2) statistically significant differences in perplexity metrics (i.e., average and maximum) between confusing and clean versions of the code, especially within the regions that contain the atoms of confusion.

RQ2: How does LLM confusion about code relate to human confusion about the same code? Building on RQ₁, we investigate the extent to which LLM perplexity aligns with neurophysiological correlates of human confusion during program comprehension. To this end, we analyze the EEG recordings and eye-tracking data of Bergum et al. [4], in which they identified a late frontal positivity

using fixation-related potentials. This neurophysiological response, often associated with increased syntactic complexity and processing effort [41], provides a robust basis for evaluating whether code regions with high LLM perplexity correspond to increased cognitive demand in human programmers.

Specifically, we analyze whether the maximum token-level perplexity in atoms of confusion correlates with increased late frontal positivity in the FRP signal that is extracted time-locked to fixations positioned on this token sequence. We hypothesize that the LLM perplexity will positively correlate with the measured late frontal positivity, suggesting that LLM perplexity and human confusion are linked.

RQ₃: How well is a data-driven, LLM-based approach able to identify code regions that confuse humans? Finally, we explore whether perplexity can serve as a practical proxy for identifying confusing code regions by testing (1) whether regions of confusion derived from perplexity align with neurophysiological correlates of human confusion as measured by late frontal positivity, and (2) whether these detected regions of confusion result in similar correlations with EEG signals compared to known atoms of confusion.

Our third research question is concerned with the feasibility of using LLM perplexity directly to identify confusing regions in code, beyond the predefined set of atoms of confusion. For this purpose, we introduce a data-driven approach for extracting *regions of confusion* by locating peaks in token-level perplexity and aggregating them into semantically coherent sequences of tokens. These regions are then compared to the predefined atoms of confusion in terms of their alignment with FRP responses. We hypothesize that regions of confusion extracted through perplexity-based methods will yield correlations that are, at least, comparable, and possibly stronger, than those based on the predefined set of atoms of confusion—offering a scalable way to localize confusion in code.

4 Methodology

In this section, we present our research process, including the extraction and processing of the EEG data, the selection of language models, the perplexity analysis, and the approach to generate regions of confusion based on a data-driven approach.

4.1 EEG Data on Human Confusion

The dataset of Bergum et al. consists of EEG recordings collected from 24 participants, each completing 72 code comprehension trials, resulting in an initial total of 1 728 trials [4]. After preprocessing, artifact removal, and synchronization with eye-tracking data performed by Bergum et al., their dataset provides 1 432 high-quality samples usable as foundation for our analysis of RQ₂ and RQ₃. Each of these 1 432 trials was annotated as either *confusing* or *clean* based on whether the code snippet contained an atom of confusion.

4.1.1 EEG Preprocessing. To accurately retrieve EEG signals time-locked to fixations on confusing code regions, we first performed the epoching of the raw EEG data into time-locked segments around fixation events. For this, we used a slight modification of the algorithm used by Bergum et al. for fixation selection. Our algorithm identifies the first fixation that is closest to the area of interest (and not just the first fixation within the bounding box around

this defined area), which more accurately corresponds to the LLM receiving this token as input. The resulting epochs contained data from 27 EEG electrodes, spanning from -300 to 1 000 ms at a sampling rate of 500 Hz, and we used the mean EEG signal for -300–0 ms as baseline. Building on this fixation selection, we refined our analysis by focusing specifically on the frontal midline electrode (Fz) within a 400–800 ms window post-fixation, where the late frontal positivity identified by Bergum et al. was most pronounced [4]. For each atom, we followed the standard procedure of averaging the corresponding EEG epochs across participants to enhance the signal-to-noise ratio. We then extracted the mean amplitude for statistical comparison, as established for ERP analysis [49].

4.2 LLM-Based Perplexity Analysis

4.2.1 Language Model Selection. We based the selection of language models for our experiments on their performance on Java tasks from the MultiPL-E benchmark², a multilingual extension of the HumanEval benchmark designed to evaluate the functional correctness of synthesized programs from docstrings. Specifically, we relied on the Big Code Models Leaderboard³, which reports the average pass@1 scores for each model on a set of programming tasks across multiple languages. This metric reflects the proportion of generated code completions that are functionally correct on the first attempt. Because our EEG data stems from a study using Java code, we prioritized models with strong Java performance for perplexity evaluation. Accordingly, we selected Qwen2.5-Coder-32B⁴, which, at the time of writing, held the highest Java pass@1 score on the leaderboard at 65.49%. This choice ensured that our perplexity evaluations were based on a state-of-the-art model proficient in Java, thereby enhancing the validity of comparisons with neurophysiological responses to source code and improving the relevance of our findings to real-world program comprehension tasks. In addition, we selected two mainstream models, as well as three LLMs of different sizes and architectures with the best Java pass@1 scores. Thereby, we complement the range of model sizes and ensure the generalizability of our results across models (see Section 7.3). Table 2 summarizes the included models, their sizes, and leaderboard scores, if available.

Table 2: LLMs selected for perplexity measurement, including size and Java pass@1 leaderboard score, if available.

Model	Size	Java pass@1
Qwen2.5-Coder-32B	32B	65.49%
Yi-Coder 9B	9B	-
Artigenz-Coder-DS 6.7B	6.7B	56.84%
Deepseek-Coder 6.7B	6.7B	37.72%
Llama 3.2 3B	3B	-
Deepseek-Coder-1.3B	1.3B	27.16%

²<https://huggingface.co/datasets/nuprl/MultiPL-E/viewer/humaneval-java>

³<https://huggingface.co/spaces/bigcode/bigcode-models-leaderboard>

⁴<https://huggingface.co/Qwen/Qwen2.5-Coder-32B>, [30]

4.2.2 Token-Level Perplexity Calculation. To quantify model confusion at a fine-grained level, we compute per-token perplexity. This approach follows prior work that uses token-level perplexity to analyze code [84]. Our procedure for each code snippet is as follows:

- (1) Parsing and Tokenization. The snippet is first preprocessed to extract predefined atoms of confusion (based on Bergum et al. [4]) and is then tokenized using the model's tokenizer.
- (2) Token Probability Calculation. For each token position i , the model receives the previous sequence of tokens as context and predicts the probability distribution over the next token. Logits from the model's output are converted to probabilities via a softmax function, and the probability assigned to the actual next token t_i is extracted.
- (3) Per-Token Perplexity Computation. We then compute the per-token perplexity as the inverse of the predicted probability for all tokens in the snippet, producing a sequence of perplexity values.

4.2.3 Calculation of Perplexity for Snippets and Multiple Tokens. In addition to token-level analysis, we aggregate the perplexity values of multiple tokens to capture model confusion over the code snippets and over areas of interest (AOIs, i.e., manually defined atoms of confusion or data-driven regions of confusion). Using the standard sequence-level perplexity definition introduced in Section 2.3, we compute the following:

- Average perplexity: the standard sequence-level perplexity of the snippet or region, computed via the inverse geometric mean of token probabilities.
- Maximum perplexity: the highest perplexity observed for any token within the snippet or region.

We calculated these metrics separately for each code snippet and area of interest. They represent the language model's confusion and are used in subsequent analyses to evaluate whether regions associated with higher model confusion also correspond to higher correlates of human confusion (as indicated by EEG fixation-related potentials).

4.2.4 Statistical Analysis. For RQ₁, we first tested for normality with the Shapiro–Wilk test. Since the differences did not follow a normal distribution, we opted for the non-parametric Wilcoxon signed-rank test. We report p -values and the rank-based effect size $r = Z/\sqrt{N}$, where Z is the standardized test statistic and N is the number of paired observations. The four comparisons (snippet-level vs. AOI-level; average vs. maximum perplexity) are pre-specified complementary operationalizations of the same construct rather than post-hoc tests and are, therefore, interpreted independently without requiring multiple-testing correction [64, 67]. For RQ₂, we examined the relationship between model perplexity and neurophysiological correlates of confusion by applying a non-parametric, rank-based correlation metric, specifically Spearman's rank correlation coefficient (ρ). For RQ₃, we applied clustered bootstrapping for Spearman's correlation (10 000 replicates), resampling at the snippet level to control for non-independence of multiple detected regions of confusion per snippet. Our choice of 10 000 bootstrap replicates was guided by best practices, which recommend using thousands of replicates to ensure stable confidence intervals and

reduce Monte Carlo error [10, 16, 25]. While prior simulation studies on Spearman's correlation have used 2 000 bootstrap replicates when feasible [65], larger values offer improved precision and are increasingly adopted [17]. We report 95% confidence intervals to support robust statistical inference [29].

4.3 Data-Driven Detection of Regions of Confusion

For RQ₃, we explore whether language model perplexity can serve as a predictive signal for human confusion. To this end, we developed a data-driven method for identifying regions of confusion in code. Our method detects areas of locally-high perplexity and refines them through lexical, syntactic, and structural post-processing. The procedure consists of four main steps.

4.3.1 Peak Detection. Given a tokenized code snippet with token-level perplexity values, we identify local maxima using the prominence-based peak detection algorithm from `scipy`⁵. Prominence measures how much a peak stands out relative to its surroundings, ensuring that only meaningful perplexity peaks are detected as candidates. Through our testing, we set the prominence to 0.8, which reliably isolates sharp perplexity spikes for our dataset. In addition, we set the detection span to 1, such that each detected peak corresponds to a single high-perplexity token, which then forms the core of a candidate region of confusion.

4.3.2 Lexical and Syntactic Expansion. We first classify tokens using simple lexical rules to identify their categories. Based on this classification, we then expand each peak region syntactically by merging related tokens into meaningful units. This includes combining variable names with numeric suffixes (e.g., `V1`), or resolving operators with operands (e.g., `V1--`).

4.3.3 Merging Regions of Confusion. We merge adjacent or overlapping perplexity-based spans to form unified regions of confusion based on span proximity and overlap. This prevents fragmentation and captures compound expressions that may collectively contribute to confusion.

4.3.4 Categorization and Filtering. To better understand the nature of the regions flagged by our data-driven detection pipeline, we automatically annotate each identified merged region of confusion using a Java parser to identify the syntactic category (e.g., *Identifier*, *Punctuation*) of the covered code fragment. For each region of confusion, we extracted the highest node in the abstract syntax tree (AST) fully spanning the region and assigned it a label corresponding to the syntactic construct. We then grouped them into broader categories to facilitate interpretability. The mapping is based on typical program comprehension constructs. For example, both `if` and `for` constructs fall under *Control Flow*, while operators such as `++` and `==` are categorized as *Operator*.

We applied this categorization to all detected regions of confusion across the dataset. We filtered out the regions of confusion that were categorized as one of the following categories: *Type* and *Identifier* regions involved identifiers and type declarations that often appeared early in the snippet where the model lacked context, leading to artificially high perplexity that did not necessarily indicate

⁵version 1.14.1, [81]

confusion. Constructs of the category *Control Flow* typically span broad regions (e.g., entire conditional blocks), making it difficult to localize the specific source of model confusion. Constructs of the category *Punctuation* (e.g., semicolons), as the small size of the region likely leads to rare fixation. After filtering, the majority of retained regions of confusion fell into the *Literal*, *Program Structure*, *Expression*, and *Operator* categories.

We carefully chose these filtering criteria by manually inspecting the detected regions, aiming to retain only cognitively meaningful ones and exclude cases affected by missing context, sparse fixation data, or overlapping cognitive processes. All filtering steps were thoroughly documented to ensure transparency and reproducibility. While any filtering step carries a potential risk of bias, our approach tries to minimize such effects by focusing on regions with clear cognitive relevance. We provide all analyzed snippets, including their average and maximum token- and snippet-level perplexity values, detected regions of confusion, and included and excluded syntactic categories, as well as the entire script to reproduce this step in our replication package.

5 Results

In this section, we present the results regarding our research questions. We first compare perplexity between confusing and clean code snippets (RQ₁), then analyze correlations with EEG signals (RQ₂), and finally evaluate our data-driven detection of regions of confusion (RQ₃).

5.1 Atoms of Confusion & LLM Perplexity (RQ₁)

To address RQ₁, we analyzed whether atoms of confusion result in higher language model perplexity. For illustration, we begin with a qualitative comparison of a clean–confusing snippet pair as illustrated in Figure 1. The token-level perplexity values show a pronounced spike in the regions of the atoms of confusion compared to their clean counterpart, with both average and maximum perplexity increasing substantially (average: 1 → 9; maximum: 1 → 555). At the snippet level, we observe an increase in average perplexity (20 → 26), while the maximum (2 661) remains unchanged. These pronounced differences suggest that the presence of atoms of confusion contributes substantially to the model’s perplexity.

To evaluate whether this pattern generally holds, we systematically compared perplexity values between clean and confusing variants across all 72 snippet pairs using a Wilcoxon signed-rank test. As shown in Figure 2, both snippet-level metrics were statistically significantly higher in confusing snippets (avg.: $W = 941$, $p = 0.036$, $r = 0.25$; max.: $W = 97$, $p = 0.046$, $r = 0.39$).

For a finer-grained analysis, we narrowed our focus to the predefined atoms of confusion and their clean counterparts and evaluated them using the same perplexity metrics, but now at the AOI level (only the tokens related to the atoms and their counterparts). Both average and maximum perplexity were significantly higher in confusing snippets, with stronger effect sizes (avg.: $W = 686$, $p < 0.001$, $r = 0.42$; max.: $W = 638$, $p < 0.001$, $r = 0.45$). The difference in effect size between average and maximum perplexity at the AOI level is relatively small, likely due to the short AOI spans (typically 2–3 tokens) limiting the smoothing effect of averaging. Likewise, the small difference in maximum perplexity between snippet-level

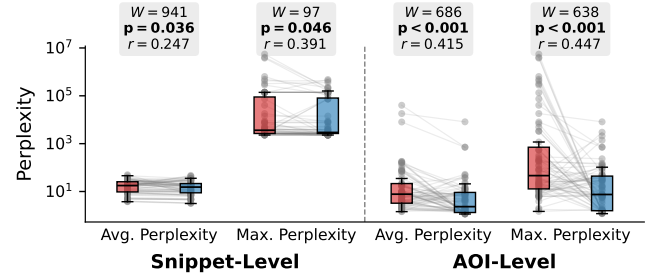


Figure 2: Comparing snippet-level and AOI-level perplexity metrics between confusing and clean variants with Wilcoxon signed-rank tests (all statistically significant).

and AOI-level is likely because the highest-perplexity token in the snippet often falls within the AOI. Still, maximum AOI-level perplexity offers the most targeted signal and achieves the highest effect size across all metrics making it the preferred measure for the remainder of our analyses.

RQ₁

We conclude that maximum AOI perplexity provides the most targeted signal of the model’s response to known confusing constructs. Consequently, we use maximum AOI perplexity in the remainder of our analyses.

5.2 LLM Perplexity & Neurophysiological Responses (RQ₂)

To address RQ₂, we analyzed the relationship between LLM perplexity and human neurophysiological responses in predefined AOIs in terms of fixation-related potentials, as established by Bergum et al. [4]. Each AOI corresponds to an atom of confusion in the confusing variant, or the corresponding region in the clean variant.

For each atom of confusion, we computed the maximum AOI-level perplexity and identified the longest first fixation inside the region to extract the corresponding EEG signal segment. Based on the preprocessed EEG data of Bergum et al., we computed the mean potential at the Fz electrode in the 400–800 ms post-fixation window, specifically capturing the previously identified late frontal positivity. Next, we computed a Spearman correlation between maximum AOI-level perplexity and mean FRP amplitude. Because the clean and confusing snippets are inherently different, we computed the correlations separately for the clean and confusing variants.

Overall, we found positive correlations with varying significance. For the confusing snippets, we observed a statistically significant moderate positive correlation (Spearman’s $\rho = 0.47$, $p < 0.001$). For the clean snippets, we observed a negligibly positive correlation, which is not statistically significant (Spearman’s $\rho = 0.05$, $p = 0.664$). We visualize these correlations in Figure 3.

RQ₂

For confusing code, the LLM perplexity significantly correlates with the neurophysiological response of human programmers associated with confusion.




Figure 3: Correlation between maximum AOI-level LLM perplexity and normalized mean FRP amplitude, separately for confusing (left) and clean (right) code variants.

5.3 Detecting Confusing Code (RQ₃)

To evaluate whether LLM perplexity can be used to detect confusing code regions, we implemented a data-driven method to identify such regions directly from token-level perplexity values. Specifically, we located local maxima in perplexity and aggregated them into syntactically coherent spans, which we refer to as *regions of confusion*. Depending on how many perplexity peaks occur in a snippet, the algorithm may identify multiple regions of confusion.

To validate this approach, we compared the perplexity-based regions of confusion to the annotated atoms of confusion introduced by [4]. A detected region was considered to overlap with an atom of confusion if it shared, at least, one token with a known confusing construct in the same code snippet.

The data-driven approach was able to detect regions of confusion in 64% of the manually labeled atoms of confusion, as well as regions of confusion that are possibly beyond the known atoms of confusion. Specifically, this approach detected a total of 227 distinct regions of confusion across the dataset: 115 regions in confusing snippets and 112 in clean snippets, many of which were novel and showed no overlap with manually labeled atoms of confusion. We visualize the overlap between our data-driven regions and the manually annotated atoms of confusion using an UpSet plot in Figure 4. Notably, the method detected 85 novel regions in confusing snippets and 69 in clean snippets, revealing potential regions of confusion not captured by manual annotation. The method largely overlapped with the manually labeled atoms of confusion in confusing snippets, indicating that a data-driven approach is well able to detect regions that actually contain confusion. However, 26 atoms of confusion were not detected, indicating that some of the confusing constructs could not be captured by our approach. Similarly, 45 out of 72 manually labeled clean atoms were not detected by the method, but 27 detected regions in clean snippets overlapped with manually labeled confusing atoms. These 27 overlaps could be considered false positives in a strict sense.

To better understand whether these newly detected perplexity-based regions of confusion are confusing to humans, we conducted a similar analysis to the one employed to answer RQ₂. Specifically, we extracted for all detected regions the closest fixation and computed the mean EEG amplitude in the 400–800 ms post-fixation window at the Fz electrode. Our clustered bootstrap Spearman analysis (resampling by snippet to account for multiple regions per snippet) shows a robust correlation with neurophysiological correlates of confusion in the confusing snippets ($\rho = 0.37$, 95% CI = [0.19, 0.53]). Note




Figure 4: UpSet plot showing the overlap between manually annotated atoms of confusion in clean and confusing snippets, and automatically detected regions (gray). Horizontal bars indicate the total number per category, while vertical bars represent intersections. Purple bars indicate overlaps.




Figure 5: Correlation between maximum AOI-level LLM perplexity and normalized mean FRP amplitude in automatically detected regions of confusion, for confusing (left) and clean (right) code variants. Markers distinguish overlap (o) with known atoms vs. non-overlap (x).

that, even in clean snippets, the data-driven approach detected regions align with FRP responses ($\rho = 0.31$, 95% CI = [0.10, 0.50]). We visualize the correlations of the perplexity in regions of confusion with the human response in Figure 5.

To understand whether the novel detected regions or manually labeled regions drive the robust correlation, we computed correlations separately for regions that did and did not overlap with previously annotated atoms of confusion. Overall, there are positive trends across all subgroups, but the stronger correlations appear in overlapping regions, both in clean and confusing snippets. Table 3 reports each Spearman rank correlation including 95% confidence intervals for all four subgroups.

RQ₃

The data-driven, perplexity-based method effectively detects confusing code regions, discovering nearly two-thirds of known confusing constructs in confusing variants. The detected regions of confusion show positive correlations with human confusion, even in (based on the literature) supposedly clean code snippets.

6 Discussion

In this section, we discuss our results and their implications by research question.

Table 3: Spearman rank correlations (ρ) between maximum AOI Perplexity and mean FRP response, with 95% confidence intervals computed via clustered bootstrapping (10 000 resamples) for sub-groups. * indicates significance ($0 \notin \text{CI}$).

Group	Spearman ρ	95% CI
Overlap with clean AoC	0.38	[0.10, 0.70]*
No overlap with clean AoC	0.24	[-0.05, 0.49]
Overlap with confusing AoC	0.33	[0.12, 0.52]*
No Overlap with confusing AoC	0.36	[-0.05, 0.71]

6.1 Atoms of Confusion & LLM Perplexity (RQ₁)

Our analysis reveals significant differences in perplexity measures between clean and confusing code variants at both snippet and AOI levels. However, maximum AOI-level perplexity provides the strongest and most robust effect ($r = 0.447$), likely because it is not influenced by noise within or outside the atom of confusion. As snippet length increases—such as when analyzing full functions or files—we expect both average and maximum perplexity at the snippet level to lose their effectiveness as reliable indicators of confusion. Likewise, we expect the average AOI-level to lose its effectiveness for confusing code constructs spanning more tokens. Visual inspection of token-level perplexity distributions further corroborates these quantitative findings, showing clear alignment between perplexity peaks and the location of known confusing code constructs (see Figure 1(b)).

Implications for Research. These results support the interpretation that LLM confusion, when triggered by atoms of confusion, is a localized phenomenon rather than a characteristic that spans entire snippets. Prioritizing such fine-grained, token-level perplexity metrics enables more accurate modeling and alignment between LLM and human confusion. Future work shall explore the variability of LLM perplexity responses across different categories of confusing constructs and investigate cases where perplexity peaks fall outside annotated AOIs, potentially reflecting deeper model confusion, ambiguous semantics, or gaps in existing atoms of confusion datasets. Additionally, alignment between perplexity peaks and annotated atoms of confusion may be further improved through prompting strategies that guide the model’s attention or through supervised fine-tuning on annotated atoms of confusion data.

Implications for Practitioners. From a practical standpoint, maximum AOI perplexity provides an interpretable and computationally efficient signal for identifying regions of potential confusion in code. Since perplexity can be obtained directly from pretrained language models without requiring additional instrumentation or manual labeling, it offers a scalable and lightweight approach for real-time confusion detection. Such signals have promising applications in developer tools and educational environments, where highlighting cognitively demanding code regions can assist in code review, debugging, and learning processes.

6.2 LLM Perplexity & Neurophysiological Responses (RQ₂)

Our results demonstrate a strong and statistically significant positive correlation between the maximum perplexity predicted by the language model and the amplitude of fixation-related potentials in response to confusing code variants (Spearman’s $\rho = 0.47$, $p < 0.001$). In contrast, clean code variants showed only a negligibly positive correlation ($\rho = 0.05$, $p = 0.664$). This pattern aligns well with our expectations: large spikes in perplexity—such as those observed in confusing code—should correspond to pronounced neurophysiological responses, whereas steady perplexity values—typically found in clean code—should elicit only weak responses because they do not cross the threshold required to trigger stronger neurophysiological activity. Importantly, these findings provide compelling evidence that LLM confusion is meaningfully aligned with human confusion during program comprehension, bridging computational measures and neurophysiological responses.

To the best of our knowledge, this is the first direct empirical validation linking LLM confusion and direct, neurophysiological correlates of human confusion in the domain of program comprehension. Thus, the results have far-reaching implications for both research and practice:

Implications for Research. From a research perspective, these findings contribute to an emerging line of research advocating for the use of neurophysiological measures to validate and interpret LLM behavior in software engineering contexts [74]. They highlight the prospects of incorporating neurophysiological signals into LLM evaluation and fine-tuning frameworks, potentially advancing model interpretability and alignment with human neurophysiological responses [2, 56, 78]. Furthermore, our results show that, when modeling human programmers, LLMs can serve as a viable surrogate model and a foundation for deriving and verifying hypotheses about programmer behavior when designing empirical experiments. In the future, this could even be extended by providing the approach with more information about the context, such as programmer experience (e.g., familiarity with certain constructs) or reading behavior (e.g., previously regarded tokens).

Implications for Practitioners. For practitioners, these findings provide evidence for the observation that LLMs are well able to collaborate with humans by being sensitive to human confusion. Additionally, they position perplexity as a reliable, lightweight measure for cognitive difficulty. Compared to behavioral metrics such as reading times or error rates, perplexity offers a scalable, token-level, and model-based signal that can be computed efficiently on large code bases. This opens avenues for developing human-aligned programming tools—such as adaptive integrated development environments and educational systems—that dynamically respond to perceived confusion, enhancing usability and learning efficiency.

6.3 Detecting Confusing Code (RQ₃)

To move beyond the limitations of manually annotated atoms of confusion, we devised a fully automatic method that identifies potentially confusing regions in source code by locating local maxima in token-level LLM perplexity. The method uncovered the majority of manually-labeled confusing snippets, while also detecting

additional regions of confusion in supposedly clean snippets (cf. Figure 4). It is important to note that the clean snippets are not guaranteed to be completely “free of confusion”, not even in the annotated area. Instead, they only proclaim to be less confusing than the corresponding atom of confusion. For example, bitwise operations, such as `0b0010 & 0b0100` (clean) versus `12 & 3` (confusing), may both plausibly generate confusion, explaining the model’s perplexity peaks in clean snippets. Thus, overlaps may reflect subtle or context-dependent confusion cues rather than outright errors in the automatic detection.

To assess the cognitive relevance of these data-driven regions of confusion, we aligned them with fixation-related potentials from EEG recordings. We found no significant correlations for clean snippets when relying solely on manually annotated clean atom of confusion counterparts in RQ₂. However, with the inclusion of the 85 additional regions detected by our method in clean snippets—alongside some overlapping snippets with annotated atoms—we now observe a moderate and significant correlation, suggesting that clean snippets can still contain subtle confusing constructs. This shift suggests that while the novel, non-overlapping regions alone do not yield a significant correlation (cf. Table 3), their combination with overlapping regions, especially after removing many from the original clean set, produces a more neurophysiologically grounded predictive signal.

Overall, the findings suggest that the effectiveness of our perplexity-based method lies not in novel detections alone, but in its ability to identify a targeted subset of regions in a data-driven way—some overlapping with atoms of confusion, others newly discovered—while excluding less informative ones. This selective refinement reveals confusing regions that align with neurophysiological signals and may be missed by manual annotation, further supporting perplexity as a meaningful and grounded predictor of confusion in source code, particularly when used to guide region selection in tandem with, rather than independently from, existing curated sets of atoms of confusion.

Implications for Research. Together with the results from RQ₂, we conclude that the relationship between perplexity and fixation-related potentials suggests that confusion is better conceptualized as a continuous spectrum, rather than a fixed set of discrete and equivalent constructs. These findings reinforce the idea that our perplexity-based approach captures varying levels of neurophysiological responses, enabling scalable and empirically validated detection of nuanced confusing regions. On a broader scale, this work provides a research avenue for a more human-aligned LLM working in concert with programmers. Future work shall investigate how LLMs might reflect other neurophysiological responses in their internal state. One example might be semantic implausibility, which elicits N400 responses in humans when reading natural language or code [39, 41], and is sometimes also associated with increased perplexity [47, 66, 82]. Incorporating automatically detected regions of confusion into research workflows—such as those used to study atoms of confusion or other confusing constructs—could support the development of hybrid frameworks that reduce reliance on manual curation through corpus studies and surveys.

Implications for Practitioners. This opens promising avenues for support tools that adapt dynamically to the programmer’s cognitive state by allowing data-driven feedback methods to flag or revise cognitively demanding code based on LLM perplexity. It also facilitates the development of tools that automatically refactor confusing code or highlight areas where programmers may need additional guidance. Future work shall refine detection thresholds to optimize sensitivity and specificity of the detection algorithm by varying prominence (e.g., to filter out regions with little to no confusion). Additionally, it might require deeper investigation of how larger context (i.e., bigger code snippets) influences the effect of atoms of confusion or creates further confusion regions in human programmers, whether LLMs also reflect these confusion patterns in their perplexity, and whether they might provide assistance in their detection. These findings should also be validated for a large and diverse population of programmers varying in experience and knowledge, and whether prior knowledge can be inserted into the model context to suit individual needs of programmers.

7 Threats to Validity

Our study combined the elements of confusing code, human thinking, LLMs, and their linkage, which all come with their own sets of threats to validity. Here, we discuss how we mitigated these threats.

7.1 Construct Validity

Our study relies on the constructs of “comprehending” code, which is “confusing” to the LLM and humans. Both are theoretical constructs that are inherently difficult to measure. Therefore, we relied on well-established operationalizations to minimize threats to construct validity. Specifically, we used the validated set of atoms of confusion [4, 21, 42] as a trigger for confusion. The used measure of human confusion, the fixation-related potential approach, is also a well-established indicator of cognitive processes in psychology and related fields [13, 32].

Similarly, we used perplexity as a commonly used, fast, and scalable measure of uncertainty for the LLM [69, 83], which we believe is the best comparable metric for LLM confusion. For our RQ₃, we detected regions of confusion based on peaks in the LLM perplexity and expanded it by the underlying code construct. This approach requires setting parameters for the peak detection and thresholds on how much to expand to the surrounding tokens. We minimized this threat to construct validity by using reasonable values and providing the full analysis script in the replication package.

Finally, we hypothesize in this paper that perplexity for LLMs and FRP-measured late frontal positivity for humans are related constructs. The data support this hypothesis, but further studies shall corroborate it.

7.2 Internal Validity

We view the diversity of the set of code snippets as the main threat to internal validity. Atoms of confusion cover several different constructs. We rely on the EEG data of Bergum et al. [4], which treat all snippets as equally confusing. However, as shown in Figure 4, there is a variety in the constructs and how well we can detect them with perplexity. It is possible that the different constructs, or

other confounding factors, affect the found link between human confusion and LLM confusion.

An additional threat to internal validity is the size of the dataset of Bergum et al. [4], on which we relied for our study. While providing a substantial contribution with their dataset and measuring the confusion across all atoms (with over 600 epochs per condition), the small sample size of up to 12 epochs for a single atom is challenging for our analysis. We minimized this threat by using only the data from electrode Fz (which showed the most prominent signal) and by averaging the EEG signal across the time period of 400–800 ms instead of performing a more detailed analysis over this time span.

7.3 External Validity

On the LLM side, we focused primarily on Qwen2.5-Coder-32B due to its strong performance on Java coding tasks. To minimize the potential that our findings are specific to this LLM model, we replicated our analysis with several other mainstream models of varying sizes (i.e., Deepseek-Coder 1.3B, Llama 3.2 3B, Artigenz-Coder-DS 6.7B, Deepseek-Coder 6.7B, Yi-Coder 9B). The results were largely consistent across models (Figure 6), suggesting that the observed relationships are not model-specific. Even a general-purpose model not fine-tuned on code, such as Llama 3.2, exhibited a comparable positive relationship between perplexity and mean FRP amplitude, albeit with a slightly weaker trend.

We conducted a sensitivity analysis of the *prominence* parameter used in the peak detection step (Figure 7). Varying the prominence (p) influences the number of detected regions (n) and reflects a trade-off between false positives (Type I errors) and false negatives (Type II errors), depending on how conservatively one wishes to detect peaks. The stable trends observed across different prominence thresholds indicate that our findings are robust to reasonable variations in this parameter. Detailed results for all models and parameter settings are available in our replication package.




Figure 6: Relationship between maximum AOI-level perplexity and mean FRP amplitude across multiple LLMs. The trends are consistent across models of different sizes.

Natural language studies have shown that there is a turning point at which larger models become less aligned with human natural language processing [45, 55]. Thus, a smaller coding model may show a closer link with how humans comprehend code. However, finding this turning point is beyond the scope of our study.

On the human side, we refer to the threats to validity of the used dataset of Bergum et al. [4]. Most critically, it is possible that the found effects of confusion via FRP are specific to the participant




Figure 7: Sensitivity analysis for the *prominence* parameter using Qwen2.5-Coder-32B. The relationship between maximum AOI-level perplexity and mean FRP amplitude remains stable across prominence thresholds, demonstrating robustness to parameter choice.

sample of intermediate programmers and may not generalize to programmers of all experience levels or backgrounds.

Finally, we used atoms of confusion as the best way to obtain a strong signal. It is possible our findings are specific to the confusion triggered by atoms of confusion and may not generalize to all confusing code constructs. Further studies shall investigate whether our results showing a link between LLM and human thinking hold up in other code, tasks, and contexts.

8 Conclusion

As humans increasingly collaborate with LLM-based programming assistants, it is critical to investigate whether and to what extent LLMs and humans align in this task. To this end, we conducted an empirical study examining an LLM comprehending clean and confusing code (in terms of LLM perplexity) as compared to human programmers comprehending the same code (in terms of neurophysiological responses). We found that LLM perplexity spikes correlate both in terms of location and amplitude with human neurophysiological responses that indicate confusion. This suggests that LLMs and humans are similarly confused about the same code. Thus, we have devised a data-driven, LLM-based approach to identify code regions of confusion that are associated with human confusion.

Our study revealed an alignment between LLM perplexity and human neurophysiological responses, which represents an important step toward establishing LLMs as surrogate models of human program comprehension. This opens new avenues for exploring the underlying cognitive processes of program comprehension and supports the development of LLMs to be more closely aligned with human comprehension patterns, such as other sources of confusion. We lay the groundwork for data-driven feedback methods that, using LLM perplexity, detect and revise cognitively demanding code. This approach can be optimized to generalize for larger code snippets and individual programmers, which facilitates the development of tools that automatically refactor confusing code or highlight areas in which programmers may need guidance.

Acknowledgments

This work has been supported by the European Union as part of ERC Advanced Grant “Brains On Code” (101052182).

References

- [1] Badr AlKhamissi, Greta Tuckute, Yingtian Tang, Taha Binhuraib, Antoine Bosse-lut, and Martin Schrimpf. 2025. From Language to Cognition: How LLMs Outgrow the Human Language Network. doi:10.48550/arXiv.2503.01830
- [2] Khai Loong Aw and Mariya Toneva. 2023. Training Language Models to Summarize Narratives Improves Brain Alignment. In *Int'l Conf. on Learning Representations (ICLR)*. 37 pages.
- [3] Shraddha Barke, Michael James, and Nadia Polikarpova. 2023. Grounded Copilot: How Programmers Interact with Code-Generating Models. *Proceedings of the Programming Languages* 7, OOPSLA1, Article 78 (2023), 85–111 pages.
- [4] Annabelle Bergum, Anna-Maria Maurer, Norman Peitek, Regine Bader, Axel Mecklinger, Vera Demberg, Janet Siegmund, and Sven Apel. 2025. Unexpected but Informative: What Fixation-Related Potentials Tell Us About the Processing of Confusing Program Code. doi:10.48550/arXiv.2412.10099
- [5] Trevor Brothers, Eddie Wlotko, Lena Warnke, and Gina Kuperberg. 2020. Going the Extra Mile: Effects of Discourse Context on Two Late Positives During Language Comprehension. *Neurobiology of Language* 1, 1 (2020), 135–160.
- [6] Casey Casalnuovo, Prem Devanbu, and Emily Morgan. 2020. Does Surprisal Predict Code Comprehension Difficulty?. In *Annual Meeting of the Cognitive Science Society*, Vol. 42. Cognitive Science Society.
- [7] Igor Crk and Timothy Kluthe. 2014. Toward Using Alpha and Theta Brain Waves to Quantify Programmer Expertise. In *Int'l Conf. Engineering in Medicine and Biology Society (EMBC)*. IEEE, 5373–5376.
- [8] José Aldo Silva da Costa, Rohit Gheyi, Fernando Castor, Pablo Roberto Fernandes de Oliveira, Márcio Ribeiro, and Balduino Fonseca. 2023. Seeing Confusion Through a New Lens: On the Impact of Atoms of Confusion on Novices' Code Comprehension. *Empirical Software Engineering* 28, 4 (2023), 81.
- [9] Simone Daniotti, Johannes Wachs, Xiangnan Feng, and Frank Neffke. 2025. Who Is Using AI to Code? Global Diffusion and Impact of Generative AI. doi:10.48550/arXiv.2506.08945
- [10] Russell Davidson and James MacKinnon. 2000. Bootstrap Tests: How Many Bootstraps? *Econometric Reviews* 19, 1 (2000), 55–68.
- [11] Vincenzo de Martino, Joel Castaño, Fabio Palomba, Xavier Franch, and Silverio Martínez-Fernández. 2024. A Framework for Using LLMs for Repository Mining Studies in Empirical Software Engineering. doi:10.48550/arXiv.2411.09974
- [12] Benedito de Oliveira, Márcio Ribeiro, José Aldo Silva da Costa, Rohit Gheyi, Guilherme Amaral, Rafael Maiani de Mello, Anderson Oliveira, Alessandro Garcia, Rodrigo Bonifácio, and Balduino Fonseca. 2020. Atoms of Confusion: The Eyes Do Not Lie. In *Brazilian Symposium on Software Engineering (SBES)*. ACM, 243–252.
- [13] Federica Degno and Simon Livesedge. 2020. Eye Movements and Fixation-Related Potentials in Reading: A Review. *Vision* 4 (2020), 11.
- [14] Edsger W. Dijkstra. 1968. Letters to the Editor: Go to Statement Considered Harmful. *Communications of the ACM* 11, 3 (1968), 147–148.
- [15] Thomas Donoghue and Bradley Voytek. 2024. Automated Meta-Analysis of the Event-Related Potential (ERP) Literature. *Scientific Reports* 12, 1867 (2024), 2045–2322. Issue 1.
- [16] Bradley Efron and Robert Tibshirani. 1998. An Introduction to the Bootstrap. *Psychometrika* 63, 1 (1998), 97–101.
- [17] Justin Esarey and Andrew Menger. 2019. Practical and Effective Approaches to Dealing With Clustered Data. *Political Science Research and Methods* 7, 3 (2019), 541–559.
- [18] Katia Romero Felizardo, Márcia Sampaio Lima, Anderson Deizepe, Tayana Uchôa Conte, and Igor Steinmacher. 2024. ChatGPT Application in Systematic Literature Reviews in Software Engineering: An Evaluation of Its Accuracy to Support the Selection Activity. In *Int'l Symposium on Empirical Software Engineering and Measurement (ESEM)*. ACM, 25–36.
- [19] Adam Goodkind and Klinton Bicknell. 2018. Predictive Power of Word Surprisal for Reading Times Is a Linear Function of Language Model Quality. In *Workshop on Cognitive Modeling and Computational Linguistics*. Association for Computational Linguistics, 10–18.
- [20] Dan Gopstein, Anne-Laure Fayard, Sven Apel, and Justin Cappos. 2020. Thinking Aloud About Confusing Code: A Qualitative Investigation of Program Comprehension and Atoms of Confusion. In *Proc. Europ. Software Engineering Conf./Foundations of Software Engineering (ESEC/FSE)*. ACM, 605–616.
- [21] Dan Gopstein, Jake Iannaccone, Yu Yan, Lois DeLong, Yanyan Zhuang, Martin K.-C. Yeh, and Justin Cappos. 2017. Understanding Misunderstandings in Source Code. In *Proc. Int'l Symposium Foundations of Software Engineering (FSE)*. ACM, 129–139.
- [22] Dan Gopstein, Hongwei Henry Zhou, Phyllis Frankl, and Justin Cappos. 2018. Prevalence of Confusing Code in Software Projects: Atoms of Confusion in the Wild. In *Proc. Int'l Conf. Mining Software Repositories (MSR)*. ACM, 281–291.
- [23] Huizi Hao, Kazi Amit Hasan, Hong Qin, Marcos Macedo, Yuan Tian, Steven Ding, and Ahmed Hassan. 2024. An Empirical Study on Developers' Shared Conversations with ChatGPT in Github Pull Requests and Issues. *Empirical Software Engineering* 29, 6 (2024), 150.
- [24] Arto Hellas, Juho Leinonen, Sami Sarsa, Charles Koutchme, Lilja Kujanpää, and Juha Sorva. 2023. Exploring the Responses of Large Language Models to Beginner Programmers' Help Requests. In *Conf. on International Computing Education Research (ICER)*. ACM, 93–105.
- [25] Tim Hesterberg. 2015. What Teachers Should Know About the Bootstrap: Resampling in the Undergraduate Statistics Curriculum. *The American Statistician* 69, 4 (2015), 371–386.
- [26] Abram Hindle, Earl Barr, Zhendong Su, Mark Gabel, and Premkumar Devanbu. 2012. On the Naturalness of Software. In *Proc. Int'l Conf. Software Engineering (ICSE)*. IEEE, 837–847.
- [27] Soneya Binta Hossain, Nan Jiang, Qiang Zhou, Xiaopeng Li, Wen-Hao Chiang, Yingjun Lyu, Hoan Nguyen, and Omer Tripp. 2024. A Deep Dive into Large Language Models for Automated Bug Localization and Repair. *Proceedings of the ACM on Software Engineering* 1 (2024), 66.
- [28] Xinyi Hou, Yanjie Zhao, Yue Liu, Zhou Yang, Kailong Wang, Li Li, Xiapu Luo, David Lo, John Grundy, and Haoyu Wang. 2024. Large Language Models for Software Engineering: A Systematic Literature Review. *Transactions on Software Engineering and Methodology (TOSEM)* 33, 8 (2024), 1–79.
- [29] Francis Huang L. 2018. Using Cluster Bootstrapping to Analyze Nested Data With a Few Clusters. *Educational and Psychological Measurement* 78, 2 (2018), 297–318.
- [30] Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Kai Dang, et al. 2024. Qwen2.5-Coder Technical Report. doi:10.48550/arXiv.2409.12186
- [31] Rasha Ahmad Husein, Hala Aburajouh, and Catagay Catal. 2025. Large Language Models for Code Completion: A Systematic Literature Review. *Computer Standards & Interfaces* 92 (2025), 103917.
- [32] Florian Hutzler, Mario Braun, Melissa L-H Võ, Verena Engl, Markus Hofmann, Michael Dambacher, Helmut Leder, and Arthur Jacobs. 2007. Welcome to the Real World: Validating Fixation-Related Brain Potentials for Ecologically Valid Settings. *Brain Research* 1172 (2007), 124–129.
- [33] Toyomi Ishida and Hidetake Uwano. 2019. Synchronized Analysis of Eye Movement and EEG During Program Comprehension. In *Int'l Workshop on Eye Movements in Programming (EMIP)*. IEEE, 26–32.
- [34] Maliheh Izadi, Roberta Gismondi, and Georgios Gousios. 2022. Codefill: Multi-Token Code Completion by Jointly Learning from Structure and Naming Sequences. In *Proc. Int'l Conf. Software Engineering (ICSE)*. ACM, 401–412.
- [35] Alice Jackson and Donald Bolger. 2014. The Neurophysiological Bases of EEG and EEG Measurement: A Review for the Rest of Us. *Psychophysiology* 51 (11 2014), 1061–1071. Issue 11.
- [36] Daniel Jurafsky and James Martin. 2025. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition with Language Models* (3rd ed.).
- [37] Majeed Kazemitabaar, Xinying Hou, Austin Henley, Barbara Jane Ericson, David Weintrop, and Tovi Grossman. 2024. How Novices Use LLM-based Code Generators to Solve CS1 Coding Tasks in a Self-Paced Learning Environment. In *Koli Calling Int'l Conf. on Computing Education Research*. ACM, Article 3, 12 pages.
- [38] Ranim Khojah, Mazen Mohamad, Philipp Leitner, and Francisco Gomes de Oliveira Neto. 2024. Beyond Code Generation: An Observational Study of ChatGPT Usage in Software Engineering Practice. doi:10.48550/arXiv.2404.14901
- [39] Chu-Hsuan Kuo and Chantel Prat. 2024. Computer Programmers Show Distinct, Expertise-Dependent Brain Responses to Violations in Form and Meaning When Reading Code. *Scientific Reports* 14 (2024), 5404.
- [40] Gina Kuperberg. 2016. Separate Streams or Probabilistic Inference? What the N400 Can Tell Us About the Comprehension of Events. *Language, Cognition and Neuroscience* 31, 5 (2016), 602–616.
- [41] Gina Kuperberg, Trevor Brothers, and Edward Wlotko. 2020. A Tale of Two Positives and the N400: Distinct Neural Signatures Are Evoked by Confirmed and Violated Predictions at Different Levels of Representation. *Cognitive Neuroscience* 32 (8 2020), 12–35. Issue 1.
- [42] Chris Langhout and Mauricio Finavaro Aniche. 2021. Atoms of Confusion in Java. In *Proc. Int'l Conf. Program Comprehension (ICPC)*. IEEE, 25–35.
- [43] Seolhwa Lee, Andrew Matteson, Danial Hooshyar, Songhyun Kim, Jaebum Jung, Gichun Nam, and Heuseok Lim. 2016. Comparing Programming Language Comprehension between Novice and Expert Programmers Using EEG Analysis. In *Int'l Conf. on Bioinformatics and Bioengineering*. IEEE, 350–355.
- [44] Roger Levy. 2008. Expectation-Based Syntactic Comprehension. *Cognition* 106, 3 (2008), 1126–1177.
- [45] Yi-Chien Lin, Hongao Zhu, and William Schuler. 2025. Vectors from Larger Language Models Predict Human Reading Time and fMRI Data More Poorly when Dimensionality Expansion is Controlled. doi:10.48550/arXiv.2505.12196
- [46] Yu-Tzu Lin, Yi-Zhi Liao, Xiao Hu, and Cheng-Chih Wu. 2021. EEG Activities During Program Comprehension: An Exploration of Cognition. *IEEE Access* 9 (2021), 120407–120421.
- [47] Adrielli Tina Lopes Rego, Joshua Snell, and Martijn Meeter. 2024. Language Models Outperform Cloze Predictability in a Cognitive Model of Reading. *PLOS Computational Biology* 20, 9 (2024), e1012117.
- [48] Steven Luck. 2014. *An Introduction to the Event-Related Potential Technique*. MIT Press.

- [49] Steven Luck. 2014. *An Introduction to the Event-Related Potential Technique*. MIT Press, Chapter 9: Quantifying ERP Amplitudes and Latencies.
- [50] Flávio Medeiros, Gabriel Lima, Guilherme Amaral, Sven Apel, Christian Kästner, Márcio Ribeiro, and Rohit Gheyi. 2019. An Investigation of Misunderstanding Code Patterns in C Open-Source Software Projects. *Empirical Software Engineering* 24, 4 (2019), 1693–1726.
- [51] Júlio Medeiros, Ricardo Couceiro, Gonalo Duarte, João Durães, João Castelhana, Catarina Duarte, Miguel Castelo-Branco, Henrique Madeira, Paulo de Carvalho, and César Teixeira. 2021. Can EEG Be Adopted as a Neuroscience Reference for Assessing Software Programmers' Cognitive Load? *Sensors* 21, 7 (2021), 2338.
- [52] Wendell Mendes, Oton Pinheiro, Emanuele Santos, Lincoln Rocha, and Windson Viana. 2022. Dazed and Confused: Studying the Prevalence of Atoms of Confusion in Long-Lived Java Libraries. In *Int'l Conf. on Software Maintenance and Evolution (ICSME)*. IEEE, 106–116.
- [53] Omer Moussa, Dietrich Klakow, and Mariya Toneva. 2023. Improving Semantic Understanding in Speech Language Models via Brain-tuning. In *Int'l Conf. on Learning Representations (ICLR)*. 59823–59849.
- [54] Nhan Nguyen and Sarah Nadi. 2022. An Empirical Evaluation of GitHub Copilot's Code Suggestions. In *Proc. Int'l Conf. Mining Software Repositories (MSR)*. IEEE, 1–5.
- [55] Byung-Doh Oh and William Schuler. 2023. Why Does Surprisal From Larger Transformer-Based Language Models Provide a Poorer Fit to Human Reading Times? *Transactions of the Association for Computational Linguistics* 11 (2023), 336–350.
- [56] Subba Reddy Oota, Emin elik, Fatma Deniz, and Mariya Toneva. 2024. Speech Language Models Lack Important Brain-Relevant Semantics. In *Annual Meeting of the Association for Computational Linguistics*. 8503–8528.
- [57] N. Peitek, S. Apel, C. Parnin, A. Brechmann, and J. Siegmund. 2021. Program Comprehension and Code Complexity Metrics: An fMRI Study. In *Proc. Int. Conf. Software Engineering (ICSE)*. IEEE, 524–536.
- [58] Norman Peitek, Annabelle Bergum, Maurice Rekrut, Jonas Mucke, Matthias Nadig, Chris Parnin, Janet Siegmund, and Sven Apel. 2022. Correlates of Programmer Efficacy and Their Link to Experience: A Combined EEG and Eye-Tracking Study. In *Proc. Europ. Software Engineering Conf./Foundations of Software Engineering (ESEC/FSE)*. ACM, 120–131.
- [59] N. Peitek, J. Siegmund, S. Apel, C. Kästner, C. Parnin, A. Bethmann, T. Leich, G. Saake, and A. Brechmann. 2020. A Look into Programmers' Heads. *IEEE Transactions on Software Engineering* 46, 4 (2020), 442–462.
- [60] N. Peitek, J. Siegmund, C. Parnin, S. Apel, J. Hofmeister, and A. Brechmann. 2018. Simultaneous Measurement of Program Comprehension with fMRI and Eye Tracking: A Case Study. In *Proc. Int. Conf. Empirical Software Engineering and Measurement (ESEM)*. ACM, 24:1–24:10.
- [61] Oton Pinheiro, Lincoln Rocha, and Windson Viana. 2023. How They Relate and Leave: Understanding Atoms of Confusion in Open-Source Java Projects. In *Int'l Working Conf. on Source Code Analysis and Manipulation (SCAM)*. IEEE, 119–130.
- [62] Nishat Raihan, Mohammed Latif Siddiq, Joanna Santos, and Marcos Zampieri. 2025. Large Language Models in Computer Science Education: A Systematic Literature Review. In *Technical Symposium on Computer Science Education*. ACM, 938–944.
- [63] Sanka Rasnayaka, Guanlin Wang, Ridwan Shariffdeen, and Ganesh Neelakanta Iyer. 2024. An Empirical Study on Usage and Perceptions of LLMs in a Software Engineering Project. In *Int'l Workshop on Large Language Models for Code (LLM4Code)*. ACM, 111–118.
- [64] Kenneth J. Rothman. 1990. No Adjustments Are Needed for Multiple Comparisons. *Epidemiology* 1 (1990), 43–46. <https://api.semanticscholar.org/CorpusID:5824330>
- [65] John Ruscio. 2008. Constructing Confidence Intervals for Spearman's Rank Correlation with Ordinal Data: A Simulation Study Comparing Analytic and Bootstrap Methods. *Modern Applied Statistical Methods* 7 (2008), 416–434.
- [66] Lavinia Salicchi and Yu-Yin Hsu. 2025. Not Every Metric Is Equal: Cognitive Models for Predicting N400 and P600 Components During Reading Comprehension. In *Int'l Conf. on Computational Linguistics*. 3648–3654.
- [67] D. J. Saville. 1990. Multiple Comparison Procedures: The Practical Solution. *The American Statistician* 44, 2 (1990), 174–180. doi:10.1080/00031305.1990.10475712
- [68] Agnia Sergeev, Olga Lvova, Sergey Titov, Anastasiia Serova, Farid Bagirov, Evgeniia Kirillova, and Timofey Bryksin. 2024. Reassessing Java Code Readability Models with a Human-Centered Approach. In *Proc. Int'l Conf. Software Engineering (ICSE)*. IEEE, 225–235.
- [69] Ola Shorinwa, Zhiting Mei, Justin Lidard, Allen Ren, and Anirudha Majumdar. 2025. A Survey on Uncertainty Quantification of Large Language Models: Taxonomy, Open Research Challenges, and Future Directions. doi:10.48550/arXiv.2412.05563
- [70] Mohammed Latif Siddiq, Lindsay Roney, Jiahao Zhang, and Joanna Cecilia Da Silva Santos. 2024. Quality Assessment of ChatGPT Generated Code and their Use by Developers. In *Proc. Int'l Conf. Mining Software Repositories (MSR)*. ACM, 152–156.
- [71] J. Siegmund, C. Kästner, S. Apel, C. Parnin, A. Bethmann, T. Leich, G. Saake, and A. Brechmann. 2014. Understanding Understanding Source Code with Functional Magnetic Resonance Imaging. In *Proc. Int. Conf. Software Engineering (ICSE)*. ACM, 378–389.
- [72] J. Siegmund, N. Peitek, C. Parnin, S. Apel, J. Hofmeister, C. Kästner, A. Begel, A. Bethmann, and A. Brechmann. 2017. Measuring Neural Efficiency of Program Comprehension. In *Proc. Europ. Software Engineering Conference and Int. Symp. Foundations of Software Engineering (ESEC/FSE)*. ACM, 140–150.
- [73] Nathaniel Smith and Roger Levy. 2013. The Effect of Word Predictability on Reading Time Is Logarithmic. *Cognition* 128, 3 (2013), 302–319.
- [74] Shashank Srikant, Ben Lipkin, Anna Ivanova, Evelina Fedorenko, and Una-May O'Reilly. 2022. Convergent Representations of Computer Programs in Human and Artificial Neural Networks. *Advances in Neural Information Processing Systems* 35 (2022), 18834–18849.
- [75] Igor Steinmacher, Jacob McAuley Penney, Katia Romero Felizardo, Alessandro Garcia, and Marco Gerosa. 2024. Can ChatGPT Emulate Humans in Software Engineering Surveys?. In *Int'l Symposium on Empirical Software Engineering and Measurement (ESEM)*. ACM, 414–419.
- [76] Tamara Swaab, Kerry Ledoux, Christine Camblin, and Megan Boudewyn. 2011. *The Oxford Handbook of Event-Related Potential Components*. Oxford University Press, Chapter Language-Related ERP Components, 397–439.
- [77] Dianne Thornhill and Cyma Van Petten. 2012. Lexical Versus Conceptual Anticipation During Sentence Processing: Frontal Positivity and N400 ERP Components. *Psychophysiology* 83 (2012), 382–392. Issue 3.
- [78] Mariya Toneva and Leila Wehbe. 2019. Interpreting and Improving Natural-Language Processing (In Machines) with Natural Language-Processing (In the Brain). *Advances in Neural Information Processing Systems* 32 (2019), 14954–14964.
- [79] Priyan Vaithilingam, Tianyi Zhang, and Elena Glassman. 2022. Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models. In *Conf. on Human Factors in Computing Systems (CHI): Extended Abstracts*. ACM, 332.
- [80] Cyma Van Petten and Barbara Luka. 2012. Prediction During Language Comprehension: Benefits, Costs, and ERP Components. *Psychophysiology* 83, 2 (2012), 176–190.
- [81] Pauli Virtanen, Ralf Gommers, Travis Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan van der Walt, Matthew Brett, Joshua Wilson, Jarrod Millman, Nikolay Mayorov, Andrew Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, Ilhan Polat, Yu Feng, Eric Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272.
- [82] Ethan Wilcox, Jon Gauthier, and Jennifer Hu. 2020. On the Predictive Power of Neural Language Models for Human Real-Time Comprehension Behavior. In *Annual Meeting of the Cognitive Science Society*.
- [83] Miao Xiong, Andrea Santilli, Michael Kirchhof, Adam Golinski, and Sinead Williamson. 2024. Efficient and Effective Uncertainty Quantification for LLMs. In *Workshop Neurips Safe Generative AI*. OpenReview.
- [84] Jinwei Xu, He Zhang, Yanjin Yang, Zeru Cheng, Jun Lyu, Bohan Liu, Xin Zhou, Lanxin Yang, Alberto Bacchelli, Yin Kia Chiam, and Thiam Kian Chiew. 2024. Investigating Efficacy of Perplexity in Detecting LLM-Generated Code. doi:10.48550/arXiv.2412.16525
- [85] Zhenyu Xu and Victor Sheng. 2024. Detecting AI-Generated Code Assignments Using Perplexity of Large Language Models. *AAAI Conf. on Artificial Intelligence* 38, 21 (2024), 23155–23162.
- [86] Martin Yeh, Dan Gopstein, Yu Yan, and Yanyan Zhuang. 2017. Detecting and Comparing Brain Activity in Short Program Comprehension Using EEG. In *Frontiers in Education Conf. (FIE)*. IEEE, 1–5.
- [87] Martin Yeh, Yu Yan, Yanyan Zhuang, and Lois Anne DeLong. 2022. Identifying Program Confusion Using Electroencephalogram Measurements. *Behaviour & Information Technology* 41, 12 (2022), 2528–2545.
- [88] Albert Ziegler, Eirini Kalliamvakou, X. Alice Li, Andrew Rice, Devon Rifkin, Shawn Simister, Ganesh Sittampalam, and Edward Aftandilian. 2022. Productivity Assessment of Neural Code Completion. In *Int'l Symposium on Machine Programming*. ACM, 21–29.
- [89] Fida Zubair, Maryam Al-Hitmi, and Catayatal. 2025. The Use of Large Language Models for Program Repair. *Computer Standards & Interfaces* 93 (2025), 103951.