

Fixation-related potentials reveal that confusing program code elicits a late frontal positivity

Annabelle Bergum^{1,+}, Anna-Maria Maurer^{1,+*}, Norman Peitek¹, Regine Bader², Axel Mecklinger², Vera Demberg^{1,3}, Janet Siegmund⁴, and Sven Apel¹

¹Computer Science, Saarland University

²Psychology, Saarland University

³Language Science and Technology, Saarland University

⁴Computer Science, University of Technology Chemnitz

*maureran@cs.uni-saarland.de

+These authors contributed equally to this work.

ABSTRACT

As software pervades more and more areas of our professional and personal lives, there is an ever-increasing need to maintain software and for programmers to efficiently write and understand program code. In the first study of its kind, we analyze *fixation-related potentials* (FRPs) to explore the online processing of program code patterns that are confusing to programmers, but not to the computer (so-called *atoms of confusion*), and their underlying neurocognitive mechanisms in an ecologically valid setting. Relative to clean counterparts in program code without an atom of confusion, confusing code elicits a late frontal positivity of about 400 to 700 ms after first looking at the atom of confusion. This frontal positivity resembles an event-related potential (ERP) component found during natural language processing that is elicited by unexpected but plausible words in sentence context. Thus, we suggest that the brain engages similar neurocognitive mechanisms in response to unexpected and informative inputs in program code and in natural language. In both domains, these inputs update a comprehender's situation model, which is essential for information extraction from a quickly unfolding input. Our results have far-reaching implications for programming and pave the way for interdisciplinary collaborations between software engineering and psycholinguistics.

1 Introduction

Software forms an integral aspect of our everyday life. Thus, it is highly relevant to understand the processes that underlie skilled programming, including the creation and maintenance of software. One central aspect of programming is reading and comprehending program code, which is crucial for various programming activities, such as implementing new functionalities or fixing program errors¹⁻⁵. During these activities, misconceptions about the code, such as confusion regarding its execution, purpose, or intended message, can lead to software failures with potentially fatal consequences^{4,6}.

The main goal of the present study is to explore the processing of confusing elements in program code and their underlying neurocognitive mechanisms in an ecologically valid setting. Recently, various studies have identified a set of small, recurring patterns⁷ prevalent in program code⁸⁻¹⁰, the so-called *atoms of confusion*. An *atom of confusion* is a syntactically correct code pattern that is unambiguous for the computer to execute, but, due to its unintuitive structure, is confusing to programmers, who encounter difficulties in understanding the code and might even misinterpret it^{7,11}.

In Figure 1, we contrast an illustrative example of a confusing code snippet (containing an atom of confusion) with a clean code snippet (containing functionally-equivalent code without the atom of confusion). The atom of confusion is located in Lines 7 and 8 of the confusing snippet displayed in Listing 1, in which the level of indentation of the `else` block subtly influences human understanding. The difficulty that programmers may encounter is that the indentation suggests the `else` block closes the first `if` statement, which leads to the incorrect result of `R = 7`. In the clean counterpart shown in Listing 2, the indentation aligns with the functionally correct second `if` statement. This eliminates the risk of confusion regarding the matching `if` statement, and the output of the variable `R` can be straightforwardly determined as 3, as several studies on atoms of

This work is licensed under a [Creative Commons "Attribution 4.0 International"](https://creativecommons.org/licenses/by/4.0/) license.



```

1 boolean V1 = true;
2 boolean V2 = false;
3 int R = 3;
4 if (V1)
5     if (V2)
6         R = R + 2;
7 else
8     R = R + 4;

```

Listing 1. Confusing code containing an atom of confusion (*indentation of the else block makes it difficult to match it to the correct if statement*)

```

1 boolean V1 = true;
2 boolean V2 = false;
3 int R = 3;
4 if (V1)
5     if (V2)
6         R = R + 2;
7 else
8     R = R + 4;

```

Listing 2. Equivalent clean code without an atom of confusion (*else block is unambiguously indented to make it easy to match it to the second if statement*)

Figure 1. A pair of corresponding code snippets with one containing an atom of confusion (left) and the other containing functionally equivalent code that is easy to understand (right).

confusion confirmed^{7,12,13}. It is important to note that the confusing and clean snippet conditions are both syntactically correct and semantically equivalent, but they differ in their risk that a human programmer may misunderstand their semantics due to the atom of confusion⁷. While such subtle differences may seem insignificant at first glance, they have caused severe bugs in the past, such as a mismatched indentation in Apple’s iOS operating system that rendered its secure network connections vulnerable^{14,15}. In general, several behavioral studies across different programming languages have shown that atoms of confusion require more time to process^{7,12}, are more effortful to comprehend^{16,17}, and lead to more errors^{7,13,18}.

Despite the insights into atoms of confusion obtained from behavioral experiments, the underlying mechanisms of the phenomenon remain unclear. One missing element is our limited understanding of which neurocognitive mechanisms are engaged in response to the confusing program code that is causing the difficulty. Identifying these mechanisms would allow us to gain a deeper understanding of the principles behind misconceptions in program code, with far-reaching implications for programming practice. Specifically, this line of research can provide empirical evidence for refining coding conventions, programming language design, and educational practices, ultimately leading to more secure and reliable software.

The premise of our work is that our understanding of how humans process natural language could be the key to interpreting the processing of atoms of confusion and discovering the underlying neurocognitive mechanisms. Indeed, parallels between programming languages and natural languages have been observed in several studies, hinting at an overlap in activated brain areas between natural language (e.g., Brodmann Areas 21, 44, and 47) and program comprehension^{19–23}. In general, considering analogies with what is already known about human language processing and the processing of program code might prove to be a very fruitful avenue.

The characteristics of program comprehension suggest that *event-related potentials* (ERPs) might be particularly well suited to study atoms of confusion. ERPs are online electrophysiological responses to linguistic or other eliciting events. They can track cognitive processes even when unfolding quickly, which is the case for both reading program code and natural language text²⁴. ERPs can provide a particularly informative way of exploring which neurocognitive processes underlie the processing of atoms of confusion, as there exists a rich literature in psycholinguistics that has related several ERP components to different cognitive processes that occur during language processing²⁵.

At present, it is unclear which aspect of the atoms of confusion causes the processing difficulties that have been observed for these elements. Is the difficulty (1) semantic in nature (i.e., the critical expression is unexpected and difficult to integrate with the meaning of the previous context, as in the sentence *I take coffee with cream and dog*)? In this case, the atom of confusion should elicit an enlarged N400 (i.e., a negative peak in the ERP approximately 400 ms after the eliciting event, similar to the N400 for the word *dog* in the aforementioned example^{26–29}).

Or should we rather think of atoms of confusion in terms of (2a) syntactic reanalysis similar to what is typically observed in garden path sentences, such as *The horse raced past the barn fell*, which, like an atom of confusion, in fact only has a single correct interpretation that is, however, difficult to work out for a human, as there might be uncertainty about the function of the verb *raced*. It can either be analyzed as a main verb of the sentence or alternatively as a past participle of a reduced relative clause *The horse (that was) raced past the barn fell* (which is similar to indentation in our code example, where there is uncertainty about the program execution). In this case, we should expect to see a late posterior positivity (P600)^{30,31}. Besides garden path sentences that involve syntactic reanalysis, the P600 is generally observed in sentences with syntactic errors, such as *The child throw the toy*, but also in sentences with syntactic ambiguities, as in *The woman persuaded to answer the door* relative to the syntactically unambiguous sentence *The woman struggled to prepare the meal*^{31,32}. In these cases, the P600 is typically interpreted as a correlate of a syntactic revision (e.g., passivizing the verb *persuaded* and attaching it to a reduced

relative clause), but it can also be evoked by semantic violations (see Paczynski and Kuperberg³³ for an example).

As a third alternative, one can think of atoms of confusion in terms of (2b) unexpected but plausible material, similar to materials in natural language that have been found to elicit a late frontal positivity in strongly constraining sentence contexts^{26,34}. In the sentence *he bought her a pearl necklace for her collection*, the word *collection* elicits a late frontal positivity, which is assumed to reflect the inhibition of the highly expected word *birthday* and the integration of the unexpected word *collection* in the current situation model – an overall meaning representation that is constructed during language comprehension³⁵. As outlined by Van Petten and Luka³⁶, both late components, the P600 and the late frontal positivity, are associated with unexpected linguistic inputs, but the main factor that distinguishes them is the plausibility of the input: The P600 is produced by highly implausible and anomalous words, whereas the frontal positivity is elicited by unexpected but plausible input.

Currently, it is unknown whether the confusion induced by atoms of confusion, as perceived by programmers, results (1) from the semantic integration of unexpected and challenging content (in which case we would expect to see an N400), (2) from syntactic ambiguity (see below), or (3) from something else entirely (producing an unexpected ERP signal). As for syntactic ambiguity, the confusion may arise from (2a) syntactic reanalysis (in which case we would expect to see a late posterior positivity) or from (2b) the integration of unexpected but plausible content into the situation model of the preceding context (in which case we would expect to see a late frontal positivity). Thus, we pose the following research question:

Which ERP components identified in psycholinguistic research are elicited when processing atoms of confusion?

The premise of our work is that building on natural-language studies would help us understand the sources and processing of such confusion, illuminating the intricacies of program comprehension. This, in turn, would be invaluable in better supporting programmers through dedicated programming languages and curricula. For example, if atoms of confusion evoke a perceived syntactic ambiguity, this could be avoided by better design choices made by programming language designers (e.g., by requiring explicit brackets around the bodies of *for*, *while*, and *if* statements); if they rather cause problems with semantic integration, this can be captured in dedicated training procedures within programming curricula and coding guidelines (e.g., defining clear and consistent rules that instruct programmers on the use of patterns). If we observe an entirely unexpected ERP signal, this could indicate that the confusion arising from the presence of an atom of confusion has no equivalent in natural-language comprehension, or that there is no unified cause for atoms of confusion. Thus, independent of the concrete results, it is highly relevant for the field of program comprehension to understand the electrophysiological signals and hence the underlying mechanisms of processing atoms of confusion. Looking at a broader picture, if we observe more and more analogies between processing program code and natural language text, this suggests that our premise of a methodological transfer between the two disciplines is worthwhile.

Since it is unclear which ERP component relates to processing confusing code, we employ a data-driven approach to detect differences between confusing and corresponding clean code without targeting a hypothesized component, as research on program comprehension and its underlying neurocognitive mechanisms is still in its infancy. Contrary to research on natural language comprehension, which builds on several decades of behavioral, neurophysiological, and brain imaging studies that have unraveled a variety of processing aspects and representational structures³⁷, studying program comprehension faces unique challenges, such as a non-linear reading flow (e.g., program statements are processed not in lexical order but along the control and data flow of possible program executions)³⁸⁻⁴¹. So far, only one study on program comprehension has investigated ERP responses to syntactic and semantic errors in program code: Using a token-wise presentation of code, Kuo and Prat⁴² identified P600 and N400 components for syntactically and semantically invalid program code, respectively. However, this setup of a token-wise presentation, while usually applied in reading studies with natural language materials⁴³, is too artificial and highly problematic for the analysis of program code: First, it enforces an unnatural linear reading flow. Second, with a token-wise presentation mode, each token is presented with a uniform duration, which allows the programmer only a cursory glance at the code but does not account for calculations that can strongly vary in complexity between tokens within a code snippet. Both aspects strongly compromise the construct and ecological validity of neurocognitive studies on program comprehension.

In an effort to increase the ecological validity of our study with regard to code presentation and code reading, we seek to present program code as a whole and allowed participants to read the code at their own pace. A major challenge in presenting a snippet of code all at once, however, is that the exact point in time at which the critical segment containing the atom of confusion is processed varies both between items and between participants and therefore cannot be assessed with conventional ERP averaging methods. That is, simply synchronizing electroencephalographical (EEG) activity with the onset of the presentation of the code snippet leads to a temporal misalignment of the ERP activity related to the processing of the atom within the trial period and, thus, diffuses the signal. This severely reduces statistical power and interpretability. To address this key challenge, we employ a more sophisticated and powerful approach based on *fixation-related potentials* (FRPs). In the FRP approach, event synchronization is not realized by external stimuli⁴⁴. Rather, eye movements and EEG are recorded simultaneously while the program code is presented as a whole, as illustrated in Figure 2a. The eye-tracking data are used to assess the participants' first eye fixations in the area of the atom of confusion (which is well-defined for every atom of confusion). This first fixation is then

used for aligning the EEG signals. This way, we are able to determine the exact point in time at which information from the atoms is processed and can use it for synchronizing the electrophysiological data.

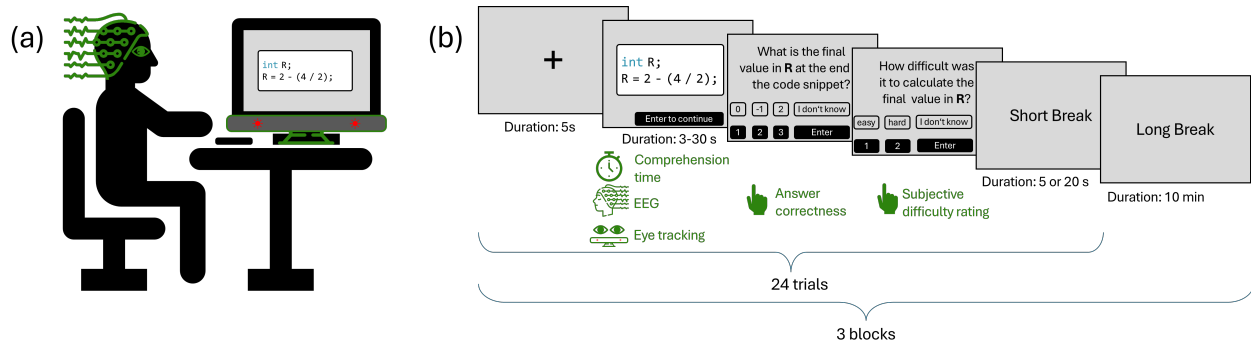


Figure 2. Overview of the experiment setup and design (EEG Icon by Freepik). Part (a) displays a symbolic image with the experiment setup of a participant with keyboard and EEG cap in front of an eye tracker attached to the bottom of the screen display. Part (b) visualizes the experiment design. The experiment included three blocks with 24 trials each: fixation cross (5 s to calibrate); snippet presentation (3–30 s, terminated by participant; the duration of snippet presentation is measured as comprehension time); answer correctness (the participant has to submit the output, i.e., the final value in variable R); subjective difficulty rating (the participant has to indicate their perceived difficulty of performing the task); short break (5 or 20 s) to reduce fatigue effects. After the block is finished, there is a long break of 10 min to reduce learning and fatigue effects.

To uncover which neurocognitive mechanisms are engaged in response to processing confusing program code, we conducted an FRP experiment involving 24 participants. We carefully adjusted the existing stimuli from previous work on atoms of confusion in code written in the Java programming language¹³ to optimize them for an FRP setup. In particular, each snippet was presented in a condition with and without an atom of confusion (across participants). Each pair of the snippet condition had a comparable number and length of code lines. Furthermore, to select suitable fixations for FRP analysis, we ensured that each snippet contained a single, clearly defined area of interest that included either the atom of confusion (*confusing snippets*) or an equivalent without an atom (*clean snippets*) (see Section 4 for details). The experiment consisted of three blocks, as illustrated in Figure 2b, in which we presented 24 snippets (half confusing and half clean). We asked the participants to comprehend the snippets as quickly and accurately as possible in order to report the output. During the experiment, we collected comprehension time, answer correctness, and subjective difficulty ratings to verify that our confusing snippets were generally more difficult to process on the behavioral level than clean snippets. As illustrated in Figure 2a, we collected EEG and eye-tracking data. We used this data to calculate FRPs by mapping fixations to the areas of interest within the code snippets. We took special care to identify and correct eye movements and muscular artifacts by means of independent component analysis (see Section 4, for further details). We synchronized the fixations with the EEG data stream to calculate FRPs. On the obtained FRPs, we tested for significant amplitude differences between confusing and clean code using pairwise cluster permutation tests.

2 Results

We report on the first-ever FRP analysis of program comprehension by measuring and pinpointing the processing of confusing elements and their neural underpinnings during program comprehension in an ecologically valid setting. Our results show that confusing code elicits a larger frontal positivity between 390 and 660 ms after first fixating the relevant area of the atom of confusion compared to fixating clean code (without an atom). We relate the identified component to the late frontal positivity occurring in natural language processing for unexpected but plausible input, which elicits an update of the situation model.

2.1 Behavioral results

The behavioral data comprised comprehension time, answer correctness, and subjective difficulty ratings, where the last two were mapped to a binary categorical scale (the submitted answer was correct or the submitted rating was easy). The descriptive trial statistics are illustrated in the left part of Table 1.

In line with prior work, our behavioral results based on general linear mixed-effects regression (GLMER) models show that the confusing snippets took longer to comprehend, had fewer correct answers, and were rated less frequently as easy to understand, as displayed in the right part of Table 1. The full analysis results, including all fixed and random effects, are provided in the supplementary information, Appendix C.

Table 1. Overview of the behavioral data and results of general linear mixed-effect regression (GLMER) models: The first two columns show the descriptive statistics on trial-level, where answer correctness and subjective difficulty rating were mapped to a binary scale. The other columns present an overview of the regression model results. For comprehension time, an LMER with Gaussian error probabilities was applied on logarithmic data; for correct answers and subjective difficulty rating, we applied a binomial GLMER. All models used random intercepts and slopes for condition under participants and snippet number. They also include block number and the item order within the block as fixed predictors. The full regression model results can be found in supplementary information, Appendix C. One of 1,728 trials was excluded due to hardware malfunction.

Trial count	Trial statistics		GLMER results for condition = confusing		
	Clean	864	Estimate	Statistic	Probability
	Confusing	863	$\beta \pm SD$	t/z value	p value
Comprehension time	Clean	11.2 \pm 7.4 s			
	Confusing	12.2 \pm 7.9 s	0.076 \pm 0.023	$t = 3.27$	0.001
Answer correctness	Clean	88%			
	Confusing	70%	-1.572 \pm 0.370	$z = -4.257$	< 0.001
Subjective difficulty rating	Clean	86%			
	Confusing	82%	-0.789 \pm 0.398	$z = -1.983$	0.047

2.2 Fixation-related potentials analysis

In a first step, we performed an ERP analysis of the EEG time-locked to the onset of the snippets with the goal of testing whether, without the eye-tracking information, ERPs differ between confusing and clean code. Consistent with our assumption that ERPs cannot reveal meaningful differences when presenting a code snippet as a whole, we found no significant ERP differences between the two conditions (cf. supplementary information, Appendix B).

Our primary goal was to determine the neurocognitive processes triggered by the presence of atoms of confusion in program code and to relate those processes to ERP components typically identified during natural language comprehension. In fact, there is a large delay and large variance in the point in time when participants first fixate on the critical region after the onset of the snippet: the first fixation on the critical region happens, on average, 3,772 \pm 3,367 ms after stimulus onset in the clean condition, and 3,246 \pm 3,084 ms in the confusing condition. The stimulus-onset-aligned ERPs, hence, cannot capture the first point in time at which information from the critical region is presumably processed. Fixation data show that on the first pass, participants fixated the critical region for 649 \pm 594 ms in the clean condition and 805 \pm 943 ms in the confusing condition. Therefore, we compared FRPs time-locked to the first fixation in the area of the atom of confusion and its counterpart without atom, spanning an epoch from 300 ms preceding the start of the fixation until 1,000 ms thereafter. For some trials, this epoch could not be extracted due to either low quality in the fixation data, no fixation in the area of interest, artifacts in the EEG signal, or a too short time period until the end of the snippet viewing phase. We excluded these trials from the subsequent analysis (341 of 1,728). Then, we calculated the participants' FRP averages for both conditions across all trials. The grand average FRP data across all participants for 9 electrodes at frontal, central, and parietal recording sites are presented in Figure 3a.

As illustrated in Figure 3, confusing code elicits a sustained frontal positivity relative to clean code. This positivity was most pronounced between 390 and 660 ms at frontal recording sites (i.e., F3, Fz, F4). At posterior recording sites, the positive amplitude difference between confusing and clean code was attenuated in this time interval. Figure 3b shows the difference in amplitude for all electrodes in the form of topographical distributions in consecutive 50 ms time intervals between 400 and 700 ms. Positive differences between both conditions were also present beyond 800 ms after fixation onset at some left frontal recording sites (e.g., F7 and F3), though they were less consistent. We statistically analyzed these effects using a cluster-based paired permutation t-test, which identifies significant clusters of adjacent data points (defined temporally and spatially).

The results from the permutation test with the alternative hypothesis *confusing* > *clean* reveal one significant cluster ($p = 0.031$) in the time interval from 390 to 660 ms (cf. Figure 3c). Although almost all electrodes are considered to be part of the significant cluster (at least, at one time point), the largest temporal extension of significant amplitude differences can be observed at midline frontal and fronto-central recording sites (i.e., Fz, FCz). Figure 3c displays the significant cluster across time and electrodes (surrounding black lines) over the amplitude difference in the grand average.

These results show that the late frontal positivity is significantly greater for confusing code than for clean code, indicating that both conditions are processed differently starting at around 390 ms after the onset of the first fixation.

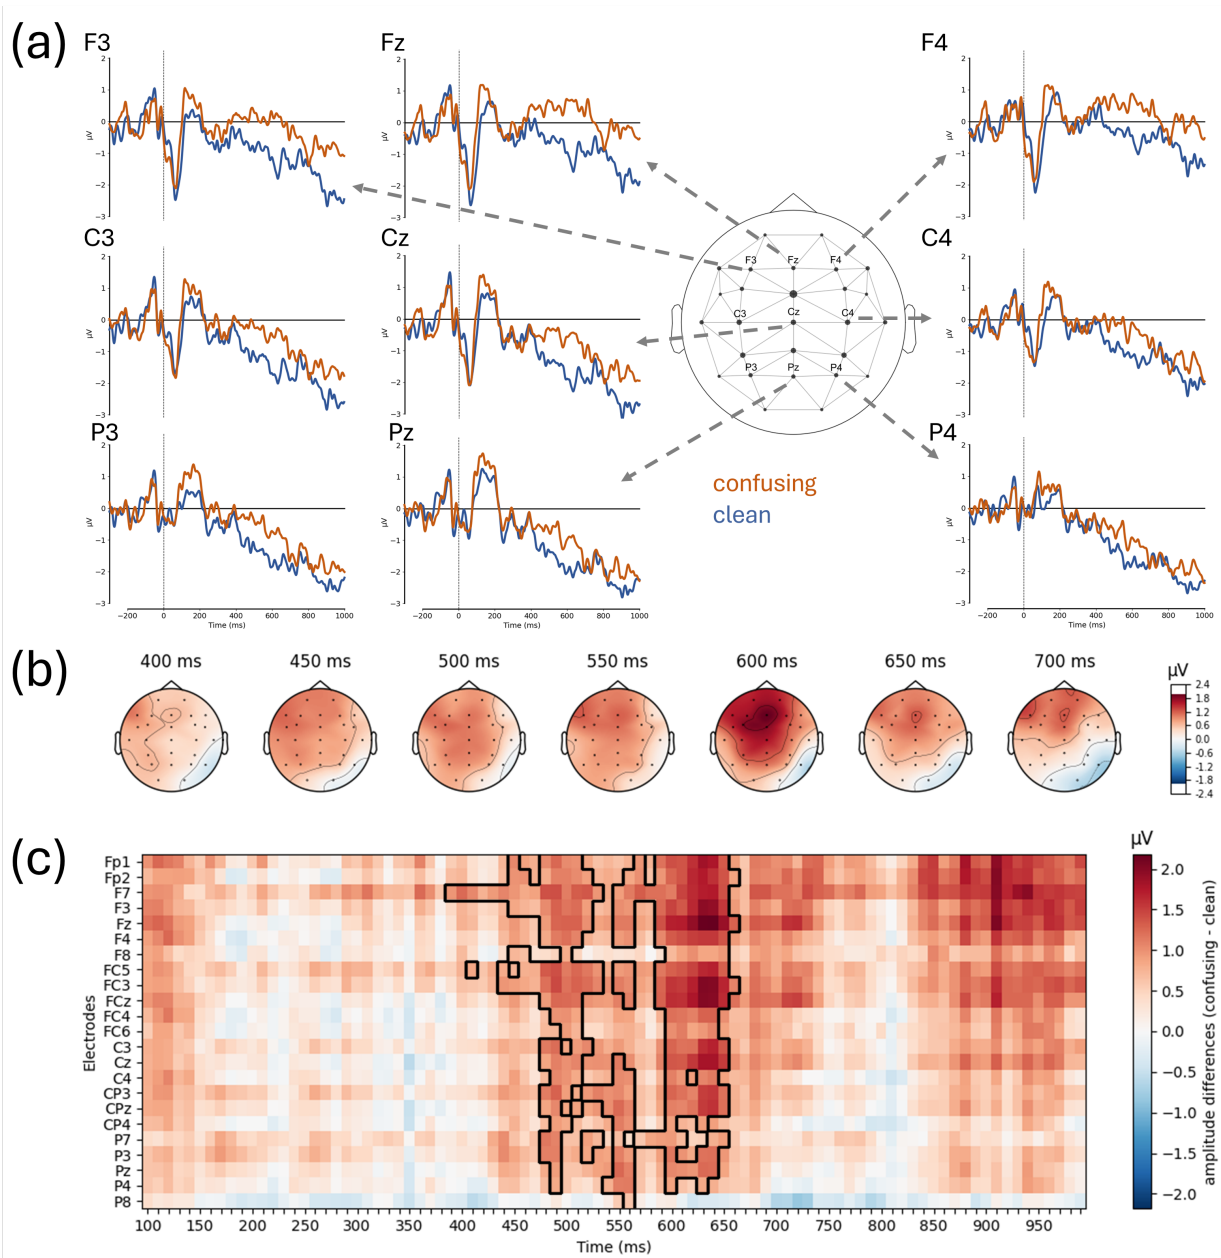


Figure 3. Results of the FRP analysis. Part (a) shows the FRP elicited by confusing (orange) and clean (blue) program code at nine scalp electrodes. The zero time points denote the onset of the fixation and positive voltages are plotted upwards. Part (b) portrays the topographic distribution of the amplitude difference between confusing and clean code in consecutive 50 ms time intervals. The onset of the time interval is indicated above each map. As apparent from the figure, the frontal slow wave is most pronounced at frontal and left frontal recording sites. Part (c) shows the amplitude differences between confusing and clean program code for all electrodes under investigation and all time points between 100 and 1,000 ms after fixation onset. The significant cluster is surrounded by black lines. Note that the late amplitude differences at frontal recordings did not reach the cluster-based significance threshold.

3 Discussion

The primary goal of this study was to apply an FRP approach – an established technique from psycholinguistics – to achieve a better understanding of the neurocognitive processes that occur when reading confusing program code (i.e., code that contains atoms of confusion) by exploring parallels in terms of ERP components between the processing of program code and natural

language processing in an ecologically valid experimental setting. Therefore, we investigated whether atoms of confusion trigger electrophysiological responses, which turned out to be similar to those engaged by unexpected but plausible inputs during language comprehension. In particular, our FRP analysis revealed a *late positivity with a frontal scalp topography*. It onsets around 400 ms after the first eye fixation in the area containing the atom of confusion and was significantly larger for confusing code compared to clean code snippets. Notably, this component resembles other late frontal positivities found in natural language processing studies, in which the late frontal positivity is elicited by unexpected words, often but not necessarily in highly constraining contexts⁴⁵⁻⁴⁷. As outlined in Section 1, the late frontal positivity can be dissociated from the P600, a late ERP component, that is elicited by words that are perceived as highly implausible during online comprehension^{27,48}.

So what does the late frontal positivity reveal about the processes involved in processing an atom of confusion? In the literature on natural language comprehension, the predominant account for a late frontal positivity holds that it reflects expectancy mismatch related processing^{45,46} or late processing *costs* due to the perceived difference between a prediction and the bottom-up input in highly constraining sentences^{49,50}. However, a range of studies has also observed that unexpected words within weakly constraining sentences can elicit a late frontal positivity^{26,51-55}. Brothers et al.^{34,56} proposed that the essential properties of a stimulus that lead to a late frontal positivity might be the presence of an unexpected but plausible word, which causes a substantial update of the situation model. More specifically, Kuperberg et al.⁵⁷ suggest that late frontal positivities arise in situations where there is initial uncertainty about the situation model, which triggers the retrieval of information from long-term memory, enabling the comprehender to accumulate confidence in their ability to successfully update their situation model. In the case of high contextual constraint and unexpected lexical material, this occurs because the unexpected word first reduces confidence in the situation model, triggers retrieval from long-term memory, and finally leads to a high confidence decision that the situation model can be successfully updated⁵⁷. In the case of a low-constraint context and unexpected but plausible lexical input, retrieval of relevant event-level information from long-term memory allows the comprehender to accumulate confidence about how to update their situation model. Kuperberg et al.⁵⁷ point out that the late frontal positivity effect in language processing seems to depend on the task, how engaged an individual is in deeply understanding the situation, and whether they can achieve confidence about the situation model based on the information they retrieved from long-term memory.

These ideas align well with a plausible mechanism for processing the atoms of confusion: Atoms of confusion are infrequent in code¹¹ and hence generally unexpected, but they are neither implausible nor incorrect, as they reflect syntactically correct code patterns that can be unambiguously executed by the computer. However, when a comprehender encounters an atom of confusion, they are likely to be unsure at first about its effect on program execution and the variables they need to keep track of for the task (i.e., they are unsure about how to update their situation model), which is in line with observations that relate confusing code patterns to increased cognitive load¹⁶. As a result, comprehenders access their long-term memory to remember exactly how to interpret the atom of confusion. If this retrieval is successful, the comprehender can confidently update their situation model. In the case of code without an atom, the situation model will also be updated, but there is no point of uncertainty followed by confidence on how to do that. This account also generates new predictions to be tested in future studies: We would expect that the late frontal positivity should only be present in trials where the comprehender could successfully resolve the ambiguity. It should therefore not appear in situations in which the atom of confusion contains a syntactically invalid code element, for which we would expect a P600 with a posterior scalp topography⁴². Also, we do not expect a late frontal positivity to appear when the comprehender remains uncertain about the semantics of the atom or when they are given a different task that does not require updating the situation model.

In summary, we presented, to our knowledge, the first-ever study that employs an FRP approach to explore the processing of confusing elements in program code. The results suggest that the brain engages highly similar mechanisms for processing unexpected but plausible elements in program code and natural language comprehension. This implies that studies on program comprehension can rely on established ERP components in natural language processing. The FRP approach offers new research perspectives with high ecological validity, focusing on the temporal characteristics of online program code comprehension. These results are significant for software engineering research, as they provide a foundation for a better understanding of code comprehension. Based on this common methodological ground, methods and findings from program comprehension research may be transferred back to research on natural language comprehension. The findings of our study, in particular, and this line of research, in general, also have far-reaching implications for the practice of software engineering. By shedding light on how programmers' brains process and respond to confusing code, we can enable programmers to develop and maintain more secure and reliable software. Specifically, our line of research supports programming practice by providing empirical evidence for coding conventions that avoid confusion-inducing code patterns. Moreover, programming language designers can leverage these insights to design consistent language constructs that minimize the potential for confusion and resulting severe mistakes. Finally, educators can incorporate these findings into their teaching practices, ultimately educating future programmers to apply more effective software development practices.

4 Methods

4.1 Participants

To estimate the number of participants required to find the hypothesized effect, we performed a power analysis with $power = 0.8$, $\alpha = 0.05$, and $d_z = 0.701$ (based on the reported $F(1, 23) = 11.8$) for the late frontal positivity detected by Thornhill and Peterson²⁶ for high-constraint unrelated sentence endings. This analysis estimated a minimum of 18 participants to correctly measure an existing effect (two-sided t-test). Due to various uncertainties in our setup and for the counterbalancing of the three version blocks across participants, we increased our sample to 24 participants. While this number might appear low, it is in line with ERP research in general^{44,46,58,59} and specifically with existing studies in the domain of software engineering^{60,61}. Further studies using larger samples shall be performed to generalize our results to a larger population, with an explicit focus on expert developers.

We recruited participants at our host institution (Saarland University) via e-mail lists, flyers, and (online) bulletin boards. We required participants to have at least an undergraduate degree in computer science or a related field. We further asked for basic experience with an object-oriented programming language (e.g., Java or C++), as the selected snippets use minimal language-specific operations. We only invited right-handed participants.

Our study was approved by the ethics committee of the Faculty of Human and Business Sciences at Saarland University. The experiment was conducted according to the relevant guidelines and regulations provided by the ethical review board. All participants provided their informed consent.

Overall, we invited 26 participants to our experiment. To ensure high data quality, we excluded two participants from our analysis due to hardware malfunctions during the experiment or frequent excessive artifacts in the EEG data (see Section 4.5.4 and 4.5.5). 9 out of the 24 included participants were female, and 15 were male, with a total age range from 21 to 45 years, and 14 participants being younger than 26. In Table 2, we provide an overview of our included participants' demographics and their programming experience.

Table 2. Overview of the programming experience and activity of our participants

Programming Experience		#	%	Frequency of programming activity		#	%
<1	Year	1	4.1	Multiple times a year	2	8.3	
1–3	Years	7	29.2	Multiple times a month	5	20.8	
3–5	Years	7	29.2	1–2 times a week	4	16.6	
5–10	Years	8	33.3	3–5 times a week	9	37.5	
>10	Years	1	4.1	Every day	4	16.6	

4.2 Experimental variables

The independent variable of our experiment is the snippet type, which has two conditions: either *confusing* or *clean*. As dependent variables, we measured the comprehension time, answer correctness, subjective difficulty rating, eye movements via an eye tracker, and brain activity via an EEG system for each snippet (see Figure 2).

4.3 Stimuli

The code snippets in our study are based on the snippet pairs from the study by Langhout and Aniche¹³, which we adjusted to reduce fixations and confounding neurocognitive processes in the FRP analysis. In a pilot study, we rated these snippet pairs considering three criteria regarding the snippet structure, the area of the atom of confusion, and behavioral response (cf. supplementary information, Appendix A). Specifically, we required that within each snippet pair, the snippets have a comparable structure to reduce the influence of confounding factors. This was determined by the number and length of the code lines. Furthermore, to facilitate the selection of suitable fixations for FRP analysis, each snippet in a pair must contain a single clear area of interest that either includes the atom of confusion (for confusing snippets) or a clarified equivalent (for clean snippets). Finally, we tested whether the snippets were of similar complexity and whether they were, in principle, understandable to students of our target population (familiarity with the syntax of Java) in a pilot study. To this end, we set a threshold of 50% answer correctness or 50% subjective difficulty rating as easy for either condition. Snippets that showed a clear violation of the expected behavior by lower correctness or fewer easy ratings were excluded from the stimulus materials.

We used 144 snippets (72 snippet pairs), each pair consisting of one confusing and one clean snippet, which are semantically equivalent and differ only in the presence of an atom of confusion. We presented only one condition of each snippet pair to each participant (72 snippets), thereby avoiding learning effects. In total, each participant viewed an equal number of snippets for both conditions (36 snippets each).

While the combined eye-tracking and EEG approach offers a higher ecological validity than using a token-wise presentation which is common in ERP studies with linguistic events, it is important to acknowledge that constraints inherent to eye tracking and EEG remain. These constraints limit the length and, consequently, the ecological validity of the presented code snippets. Nevertheless, because atoms of confusion consist of small, localized patterns, they can be meaningfully embedded within code even under these length limitations.

4.4 Procedure

At the beginning of the experiment, all participants provided their informed consent and completed a set of questionnaires, including a basic demographic questionnaire^{38,62}. To assess programming experience, we inquired about the years and recent frequency of programming activities, as well as all programming languages encountered and recently used. This provides us with a broad estimate of their capabilities and background⁶³. Additionally, participants filled out the Edinburgh Handedness Inventory⁶⁴ to ensure right-handedness.

After setting up the EEG (see Section 4.5.4), the participants were seated in front of the eye tracker attached to the screen. Then, we explained the purpose and setup of the study with introductory slides. To familiarize the participants with the task, they performed a brief training session consisting of one trial.

The experiment was divided into three consecutive blocks with 24 trials each. Between each block, participants took a break of ten minutes to reduce learning and fatigue effects. We pseudo-randomized the presentation order with the constraints that no more than three snippets of one condition appeared in a row and that at least two other snippets appeared between two snippets containing the same atom of confusion. Each trial started with the presentation of a fixation cross for 5 s. Next, the code snippet was presented (see Figure 2). The participants' task was to determine the output of variable R (as in **Result**) in each snippet as quickly and accurately as possible. Participants had up to 30 s to comprehend each snippet, but they could proceed earlier after at least 3 s by hitting the *enter* key. Then, we presented four response options: One option was the correct value, two were incorrect values that would follow typical erroneous thought processes associated with the respective atom (only one for two-choice answer options), and one *I don't know* option. The last option was intended to prevent participants from guessing. Finally, we asked participants whether the code snippet was easy or difficult to comprehend, also allowing the *I don't know* option.

After the experiment, we conducted a semi-structured interview that included questions on the participants' views about the experiment and each snippet. Finally, we paid the compensation of 12 per hour, calculated on a quarter-hourly basis.

4.5 Recording and preprocessing

4.5.1 Behavioral data preprocessing

After extracting the behavioral data, we checked whether each participant was generally able to perform the tasks based on their overall answer correctness, which we compared to the estimated correctness based on chance (i.e., 38% due to 21 snippets providing two and 51 snippets providing three answer options). All participants achieved a correctness of at least 55% (79.33% \pm 8.46%). Thus, no participants were excluded on the basis of task performance. One trial had to be excluded from the behavioral analysis due to a hardware malfunction, leaving us with 1,727 trials for the behavioral analysis. We transformed the subjective difficulty rating into a binary variable (1 if *easy*, 0 for *hard*, *I don't know*, or no answer), since we can assume that every participant who did not immediately answer the question with ease has experienced at least some difficulty.

4.5.2 Eye-movement recording

We recorded the eye movements using a Tobii Pro Spectrum eye tracker with a sampling rate of 1,200 Hz and positioned participants using the Tobii Eye Tracker Manager to ensure an optimal head position. Each block began with calibration and validation, checking all validation points to ensure high accuracy and precision. If there was a clear offset in the validation points, if some points did not receive any data, or if the data for a point were scattered across the screen, we repositioned the participant and restarted with calibration and validation.

4.5.3 Eye-movement preprocessing

The eye-tracking data, specifically the fixations, are the basis for our FRP analysis to reliably detect potential confusion at the correct moment. Missing a fixation or selecting an incorrect fixation would substantially affect the quality of our FRP analysis. Therefore, only trials with high data quality, including data correction, were considered in our eye-movement analysis.

To this end, we separated the raw eye-tracking data of all 1,727 trials based on the start and end of the snippet presentations. While the quality was overall high, for some trials, the eye tracker failed to identify the gaze position (i.e., resulting in missing values in the data). Because this poses a risk of missing a relevant fixation in the area of interest, we first excluded trials in which more than 25% of the snippet presentation consisted of missing values. We also removed trials that contained one or more intervals of missing values exceeding one second. Based on this, 34 trials were excluded, and 1,693 trials were left for further analyses.

As a next step, to reliably detect the correct fixation in the area of the atom of confusion, we manually checked the quality of all fixations. To this end, we first detected fixations and saccades in the raw eye-tracking data with I2MC, which is a noise-robust algorithm⁶⁵ with a minimal accepted fixation duration of 200 ms. Then, we manually reviewed each trial for fixations that are clear outliers (i.e., single fixations that could not be associated with a code line). 202 trials (out of 1,693) clearly contained outlier fixations. For these trials, we eliminated the outlier fixations (344) from the fixation data.

Another risk to data quality arises from fixations with an offset on the y-axis, because this can lead to assigning fixations to incorrect code lines. To reduce the risk, we used two automated fixation correction algorithms described by Carr et al.⁶⁶, which are suitable for code reading (i.e., no assumption of linear line-by-line reading), namely *cluster* and *stretch*. We fine-tuned these algorithms, because they both tend to place fixations in as many code lines as possible, although participants may ignore code lines that do not hold information relevant for program execution and variable values (i.e., lines containing only brackets { or }). Specifically, we derived variants that ignore such code lines in favor of other line matches. For each trial, we generated plots to compare the results of the four fixation corrections against the original fixation positioning and, for each fixation correction, to assess the precision on the code lines and suitability for detecting fixations within or near the area of interest later on. Two members of the research team manually reviewed the automated fixation corrections independently of each other. Upon a disagreement (275 out of 1,693 trials), the differences were discussed in detail, on 34 occasions with a third team member to reach consensus, resulting in an almost perfect inter-rater reliability of 81% determined by Cohen's Kappa⁶⁷. During these reviews, we detected that, for a few participants, there was also a noticeable offset on the x-axis. Since this offset was usually constant within a block, we corrected it with a constant offset.

We excluded trials from further analysis when we were unable to confidently align the fixations to the code (192 out of 1,693). This was the case when fixations were either entirely off the visualized code and could not be mapped to their correct position, or when single fixations could not be confidently assigned to the correct line. A visualization of all original fixations contrasted with our correction is part of the replication package.

For each of the remaining 1,501 trials, we selected the first fixation within the area of interest as the time point at which the atom of confusion was processed. 68 trials for which we did not receive a result using this strategy (i.e., the area of interest may not be fixated) were excluded from the FRP analysis. In total, we excluded 294 trials in the eye-movement preprocessing, which resulted in 1,433 trials usable for the FRP analysis.

4.5.4 EEG recording

We recorded continuous EEG from 28 scalp electrodes (Fp1, Fp2, F7, F3, Fz, F4, F8, FC5, FC3, FCz, FC4, FC6, T7, C3, Cz, C4, T8, CP3, CPz, CP4, P7, P3, Pz, P4, P8, O1, O2, and A2) placed in a cap according to the International 10–20 system⁶⁸ using BrainVision Recorder (Brain Products). Electrode AFz served as the ground electrode. An additional four electrodes were placed to the left of the left eye as well as above, beneath, and to the right of the right eye to monitor eye movements and blinks (electrooculography; EOG). Two more electrodes were placed over the left and right mastoids. Electrodes were referenced to the electrode placed over the left mastoid during recording. EEG signals were amplified with a bandpass filter of 0.016 to 250 Hz using a 16 bit BrainAmp DC amplifier. We treated all electrode locations with gel and cotton swabs to achieve an impedance below 5 k Ω for EEG and below 10 k Ω for EOG electrodes. Continuous analog-to-digital conversion of the EEG and stimulus trigger codes was performed at a sampling frequency of 500 Hz.

4.5.5 EEG preprocessing

We manually marked excessive artifacts caused by high tension or muscle movement, such as clenching teeth, using BrainVision Analyzer. Then, we applied two finite-impulse-response filters, consisting of a bandpass filter of 0.05 to 30 Hz of order 4 and a notch filter at 50 Hz. Next, we employed an independent component analysis (ICA) using the classical biased restricted info-max algorithm implemented in Brain Vision Analyzer 2.1 on the filtered continuous EEG data (excluding marked segments) to extract components of the signal that are associated with artifacts caused by eye or muscle artifacts. Next, in the Inverse ICA, we set the weight of these components on the final signal to 0. As a last step, we re-referenced the EEG channels to the mean of both mastoid electrodes to avoid bias from using one hemisphere as the reference.

After this preparation, we synchronized the event triggers in the EEG files with the behavioral data. For two trials in the behavioral data, we were unable to assign the corresponding EEG annotations due to hardware malfunctions. Then, we cut each participant's continuous EEG file into individual trials (from the start of the fixation cross to the end of the snippet presentation), resulting in 1,725 trials usable for ERP analysis. For the FRP analysis, this leaves us with 1,432 (= 1,725 – 294 + 1) trials due to an overlap of one trial in the excluded trials within the EEG and the eye-tracking data preparation.

In summary, we obtained 1,727 trials fulfilling the requirements for the behavioral data analysis. Among these, 1,433 trials were suitable for eye-movement analysis, and 1,725 were suitable for ERP analysis. Merging these two sets left us with an overlap of 1,432 trials to be used in the FRP analysis.

4.6 Data analysis

4.6.1 Behavioral analysis

For each of the 1,727 trials resulting from the processing of only the behavioral data, we obtained measures of comprehension time, answer correctness, and the subjective difficulty rating.

We analyzed the behavioral data using linear mixed effects models. As a model family, we chose a Gaussian error model with logarithmic values for fixation duration. Answer correctness and subjective difficulty ratings were coded as binary variables (*correct* = 1, *incorrect* = 0 and *easy* = 1, *other* = 0, respectively) and analyzed using binomial mixed effects models. For each behavioral measure, the regression models included the fixed parameter condition to identify whether a trial was confusing or clean (Condition: *confusing* = 1, *clean* = 0), as well as the block number (BlockNo: 1–3) and the trial number within a block (ItemOrder: 1–24) along with their interactions to account for potential learning and habituation effects. Linear mixed effects models included random intercepts and random slopes for condition under participants and items; the fixed effects structure was reduced using backward model selection based on the Akaike information criterion and a significance threshold of $p < 0.05$.

4.6.2 FRP analysis

To calculate the FRPs, we truncated the EEG segment of each trial relative to the start of the fixation into the epochs from –300 to 1000 ms including a baseline from –300 to 0 ms. The offset of the fixation compared to stimulus onset leads to a further reduction of trials since fixations at the end of the trial might not have the entire epoch length available within the snippet view and thus cannot be included, leaving us with 1405 (= 1432 – 27) trials. Next, the epochs were baseline-corrected to reduce the influence of pre-stimulus activity. That is, for each electrode and each sample point within the epoch (–300 to 1000 ms), the average pre-stimulus voltage was subtracted from the EEG signal. For determining trials containing artifacts, we ignored all EOG electrodes as well as the reference electrodes. We excluded all epochs that violated one of the following voltage constraints on any of the electrodes, which indicates the presence of an artifact in the epoch: First, a voltage step between two consecutive samples must be lower than 30 μV . Second, the maximal voltage difference within 0.2 s should not exceed 100 μV . Last, the absolute amplitude after baseline correction must remain below $\pm 70 \mu\text{V}$ at any point in time within the epoch. We excluded 18 trials (1387 = 1405 – 18 remaining). As a result, on average, 29 ($SD = 5$; range: 16–35) trials per condition and participant are suitable for the FRP analysis. Finally, we averaged the epochs per participant, condition, and electrode.

For statistical analysis, we calculated the FRP (*confusing* – *clean*) difference wave per participant and electrode to use it as a dependent variable in a pairwise cluster permutation test⁶⁹. This test identifies possibly significant data points consisting of time points and electrodes in which the test statistics exceed the significance threshold ($p < 0.05$) of a single paired t-test, and clusters these if at least two spatially or temporally adjacent points reach the significance threshold. Spatial adjacency was defined by adjacent electrodes (e.g., Fp1 and Fp2, see Figure 3a) and temporal adjacency was defined by two consecutive time points. Then, the test evaluates the significance of each detected cluster by a permutation test, which has the advantage of inherently controlling for multiple testing⁷⁰. To maximize the power in the permutation test, we followed the suggestions of Luck et al.²⁴: The epochs were down-sampled to 100 Hz, and electrodes T7, T8, O1, and O2 were not included in the analysis since we did not expect any significant effects at these electrode sites.

Additionally, we reduced the analysis interval to 100 to 1,000 ms. We performed two pairwise cluster permutation tests with this input for the difference being both a significantly smaller (*confusing*_{clean}) and a significantly higher (*confusing*_{clean}) amplitude than 0 μV and considered all resulting clusters with an alpha level of $p < 0.05$ as significant. For each of the permutation tests, we performed the maximum number of permutations (16,777,216)⁷⁰.

Acknowledgements

We would like to thank all the participants of the study who made this research possible, as well as our research assistants, Youssef Abdelsalam, Yannick Lehmen, and Marcellius William Suntoro, who helped with the technical setup.

Funding

This work has been supported by the European Union as part of the ERC Advanced Grant *Brains On Code* (101052182) and by the German Research Foundation as part of the Transregional Collaborative Research Center 248 *Foundations of Perspicuous Software Systems* (389792660). Axel Mecklinger and Regine Bader were funded by the German Research Foundation as part of the Collaborative Research Center 1102 Information Density and Linguistic Encoding (232722074), project A6.

Competing interests

There are no competing interests.

Data availability

We provide all relevant data in line with open data principles under a CC-BY license, respecting our local privacy laws (GDPR). Specifically, the datasets collected during the experiment and generated for the analysis are long-term archived in the Zenodo repository *Dataset for “Fixation-related potentials reveal that confusing program code elicits a late frontal positivity”*, <https://doi.org/10.5281/zenodo.14229848>.

Materials availability

This experiment used code available in Python (v3.11.5) as well as R (v4.3.2) for linear mixed effects models. Additionally, we used other open-source (PsychoPy, v2021.2.3) and commercial applications (Tobii Eye-Tracker Manager (v2.6.0), BrainVision Recorder (v1.20.0801), and BrainVision Analyzer (v2.3.0.8300)) to perform the experiment. We deposited the scripts and content in a GitHub Repository <https://github.com/brains-on-code/AoC-FRP-Code> and they are available for unrestricted open access under a CC-BY license.

References

1. LaToza, T., Venolia, G. & DeLine, R. "Maintaining mental models: A study of developer work habits". In Osterweil, L. J., Rombach, H. D. & Soffa, M. L. (eds.) *Proc Int Conf Softw Eng*, 492–501 (ACM, 2006).
2. Tiarks, R. What programmers really do: An observational study. *Softwaretechnik-Trends* **31**, 36–37 (2011).
3. Minelli, R., Mocci, A. & Lanza, M. "I know what you did last summer - an investigation of how developers spend their time". In Lucia, A. D., Bird, C. & Oliveto, R. (eds.) *Proceedings of the 2015 IEEE 23rd International Conference on Program Comprehension ICPC*, 25–35 (IEEE, 2015).
4. Meyer, A. N. *et al.* Detecting developers' task switches and types. *IEEE Transactions on Software Engineering* **48**, 225–240 (2022).
5. Meyer, A. N., Barr, E. T., Bird, C. & Zimmermann, T. Today was a good day: The daily life of software developers. *IEEE Transactions on Software Engineering* **47**, 863–880 (2021).
6. Meyer, A. N., Barton, L. E., Murphy, G. C., Zimmermann, T. & Fritz, T. The work life of developers: Activities, switches and perceived productivity. *IEEE Transactions on Software Engineering* **43**, 1178–1193 (2017).
7. Gopstein, D. *et al.* "Understanding misunderstandings in source code". In Bodden, E., Schäfer, W., van Deursen, A. & Zisman, A. (eds.) *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering ESEC/FSE*, 129–139 (ACM, 2017).
8. Gopstein, D., Zhou, H. H., Frankl, P. & Cappos, J. "Prevalence of confusing code in software projects". In Zaidman, A., Kamei, Y. & Hill, E. (eds.) *Proceedings of the 15th International Conference on Mining Software Repositories MSR* (2018).
9. Medeiros, F. *et al.* An investigation of misunderstanding code patterns in C open-source software projects. *Empir Softw Eng* **24**, 1693–1726 (2019).
10. Pinheiro, O., Rocha, L. & Viana, W. "How they relate and leave: Understanding atoms of confusion in open-source Java projects". In Moonen, L., Newman, C. D. & Gorla, A. (eds.) *2023 IEEE 23rd International Working Conference on Source Code Analysis and Manipulation SCAM*, 119–130 (IEEE, 2023).
11. Gopstein, D., Fayard, A.-L., Apel, S. & Cappos, J. "Thinking aloud about confusing code: A qualitative investigation of program comprehension and atoms of confusion". In Devanbu, P., Cohen, M. B. & Zimmermann, T. (eds.) *28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering ESEC/FSE*, 605–616 (Association for Computing Machinery, 2020).
12. de Oliveira, B. *et al.* "Atoms of confusion: The eyes do not lie". In Cavalcante, E., Dantas, F. & Batista, T. (eds.) *34th Brazilian Symposium on Software Engineering SBES*, 243–252 (ACM, 2020).
13. Langhout, C. & Aniche, M. "Atoms of confusion in Java". In O'Conner, L. (ed.) *29th IEEE/ACM International Conference on Program Comprehension ICPC*, 25–35 (2021).
14. Wheeler, D. The apple goto fail vulnerability: lessons learned. <https://dwheeler.com/essays/apple-goto-fail.html> (2014). Accessed: 2025-01-20.
15. Boyes, H., Norris, P., Bryant, I. & Watson, T. Trustworthy software: lessons from goto fail & heartbleed bugs. In *9th IET International Conference on System Safety and Cyber Security (2014)* (IET, 2014).

16. Yeh, M., Gopstein, D., Yan, Y. & Zhuang, Y. "Detecting and comparing brain activity in short program comprehension using EEG". In *2017 IEEE Frontiers in Education Conference FIE*, 1–5 (IEEE, 2017).
17. da Costa, J. A. S. *et al.* Seeing confusion through a new lens: On the impact of atoms of confusion on novices' code comprehension. *Empir Softw Eng* **28**, 81 (2023).
18. Yeh, M. K.-C., Yan, Y., Zhuang, Y. & DeLong, L. A. Identifying program confusion using electroencephalogram measurements. *Behaviour @AND@ Information Technology* **41**, 2528–2545 (2022).
19. Siegmund, J. *et al.* "Understanding understanding source code with functional magnetic resonance imaging". In Jalote, P., C., L. & van der Hoek, A. (eds.) *Proc Int Conf Softw Eng*, 378–389 (ACM, 2014).
20. Peitek, N. *et al.* A look into programmers' heads. *IEEE Transactions on Software Engineering* **46**, 442–462 (2020).
21. Durães, J., Madeira, H., Castelhana, J., Duarte, I. & Castelo-Branco, M. "WAP: Understanding the brain at software debugging". In Bradbury, J. (ed.) *27th IEEE International Symposium on Software Reliability Engineering ISSRE*, 87–92 (IEEE, 2016).
22. Castelhana, J. *et al.* The role of the insula in intuitive expert bug detection in computer code: An fMRI study. *Brain Imaging Behav* **13**, 623–637 (2019).
23. Floyd, B., Santander, T. & Weimer, W. "Decoding the representation of code in the brain: An fMRI study of code review and expertise". In Uchitel, S., Orso, A. & Robillard, M. P. (eds.) *Proc Int Conf Softw Eng*, 175–186 (IEEE, 2017).
24. Luck, S. *An introduction to the event-related potential technique* (MIT press, 2014).
25. Swaab, T. Y., Ledoux, K., Camblin, C. C. & Boudewyn, M. A. *The Oxford handbook of event-related potential components*, chap. "Language-related ERP components" (Oxford University Press, 2011).
26. Thornhill, D. E. & Petten, C. V. Lexical versus conceptual anticipation during sentence processing: Frontal positivity and N400 ERP components. *Int J Psychophysiol* **83**, 382–392 (2012).
27. Kuperberg, G. R., Brothers, T. & Wlotko, E. W. A tale of two positivities and the N400: Distinct neural signatures are evoked by confirmed and violated predictions at different levels of representation. *J Cogn Neurosci* **32**, 12–35 (2020).
28. Kutas, M. & Hillyard, S. A. Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science* **207**, 203–205 (1980).
29. Kasparian, K. & Steinhauer, K. Confusing similar words: ERP correlates of lexical-semantic processing in first language attrition and late second language acquisition. *Neuropsychologia* **93**, 200–217 (2016).
30. Núñez-Peña, M. & Honrubia-Serrano, M. P600 related to rule violation in an arithmetic task. *Brain Res Cogn Brain Res* **18**, 130–141 (2004).
31. Osterhout, L. & Holcomb, P. J. Event-related brain potentials elicited by syntactic anomaly. *J Mem Lang* **31**, 785–806 (1992).
32. Friederici, A. D., Gunter, T. C., Hahne, A. & Mauth, K. The relative timing of syntactic and semantic processes in sentence comprehension. *Neuroreport* **15**, 165–169 (2004).
33. Paczynski, M. & Kuperberg, G. R. Multiple influences of semantic memory on sentence processing: Distinct effects of semantic relatedness on violations of real-world event/state knowledge and animacy selection restrictions. *J Mem Lang* **67**, 426–448 (2012).
34. Brothers, T., Wlotko, E. W., Warnke, L. & Kuperberg, G. R. Going the extra mile: Effects of discourse context on two late positivities during language comprehension. *Neurobiol Lang (Camb)* **1**, 135–160 (2020).
35. Kuperberg, G. R. Separate streams or probabilistic inference? What the N400 can tell us about the comprehension of events. *Lang Cogn Neurosci* **31**, 602–616 (2016).
36. Van Petten, C. & Luka, B. J. Prediction during language comprehension: Benefits, costs, and ERP components. *Int J Psychophysiol* **83**, 176–190 (2012).
37. Kuperberg, G. R. *et al.* Multimodal neuroimaging evidence for looser lexico-semantic networks in schizophrenia: Evidence from masked indirect semantic priming. *Neuropsychologia* **124**, 337–349 (2019).
38. Peitek, N., Siegmund, J. & Apel, S. "What drives the reading order of programmers? An eye tracking study", 342–353 (ACM, 2020).
39. Aljehane, S., Sharif, B. & Maletic, J. "Determining differences in reading behavior between experts and novices by investigating eye movement on source code constructs during a bug fixing task". In Bulling, A. *et al.* (eds.) *Proc Eye Track Res Appl Symp*, 1–6 (ACM, 2021).

40. Busjahn, T. *et al.* "Eye movements in code reading: Relaxing the linear order". In Lucia, A. D., Bird, C. & Oliveto, R. (eds.) *Proceedings of the 2015 IEEE 23rd International Conference on Program Comprehension ICPC*, 255–265 (IEEE, 2015).
41. Peitek, N. *et al.* "Correlates of programmer efficacy and their link to experience: A combined EEG and eye-tracking study". In Roychoudhury, A., Cadar, C. & Kim, M. (eds.) *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 120–131 (2022).
42. Kuo, C.-H. & Prat, C. Computer programmers show distinct, expertise-dependent brain responses to violations in form and meaning when reading code. *Sci Rep* **14** (2024).
43. Wen, Y., Mirault, J. & Grainger, J. Fast syntax in the brain: Electrophysiological evidence from the rapid parallel visual presentation paradigm (RPVP). *J Exp Psychol Learn Mem Cogn* **47**, 99 (2021).
44. Hutzler, F. *et al.* Welcome to the real world: Validating fixation-related brain potentials for ecologically valid settings. *Brain Res* **1172**, 124–129 (2007).
45. Federmeier, K. D. Thinking ahead: The role and roots of prediction in language comprehension. *Psychophysiology* **44**, 491–505 (2007).
46. Hölzje, G. & Mecklinger, A. Benefits and costs of predictive processing: How sentential constraint and word expectedness affect memory formation. *Brain Res* **1788**, 147942 (2022).
47. Hubbard, R. J. & Federmeier, K. D. The impact of linguistic prediction violations on downstream recognition memory and sentence recall. *J Cogn Neurosci* **36**, 1–23 (2024).
48. Friederici, A. D. Towards a neural basis of auditory sentence processing. *Trends Cogn Sci* **6**, 78–84 (2002).
49. Kutas, M. In the company of other words: Electrophysiological evidence for single-word and sentence context effects. *Lang Cogn Process* **8**, 533–572 (1993).
50. Ness, T. & Meltzer-Asscher, A. Lexical inhibition due to failed prediction: Behavioral evidence and ERP correlates. *J Exp Psychol Learn Mem Cogn* **44**, 1269 (2018).
51. Chow, W.-Y., Smith, C., Lau, E. & Phillips, C. A "bag-of-arguments" mechanism for initial verb predictions. *Lang Cogn Neurosci* **31**, 577–596 (2016).
52. Freunberger, D. & Roehm, D. Semantic prediction in language comprehension: evidence from brain potentials. *Lang Cogn Neurosci* **31**, 1193–1205 (2016).
53. Zirnstein, M., van Hell, J. G. & Kroll, J. F. Cognitive control ability mediates prediction costs in monolinguals and bilinguals. *Cognition* **176**, 87–106 (2018).
54. Ng, S., Payne, B. R., Steen, A. A., Stine-Morrow, E. A. & Federmeier, K. D. Use of contextual information and prediction by struggling adult readers: Evidence from reading times and event-related potentials. *Sci Stud Read* **21**, 359–375 (2017).
55. Hubbard, R. J., Rommers, J., Jacobs, C. L. & Federmeier, K. D. Downstream behavioral and electrophysiological consequences of word prediction on recognition memory. *Front Hum Neurosci* **13**, 291 (2019).
56. Brothers, T., Morgan, E., Yacovone, A. & Kuperberg, G. Multiple predictions during language comprehension: Friends, foes, or indifferent companions? *Cognition* **241**, 105602 (2023).
57. Kuperberg, G. R., E., A. & T., B. The P600 does not reflect error correction: A confidence accumulation account of late posterior and frontal positivities. *Manuscript under review* (under review).
58. Boudewyn, M. A., Luck, S. J., Farrens, J. L. & Kappenman, E. S. How many trials does it take to get a significant erp effect? it depends. *Psychophysiology* **55**, e13049 (2018).
59. Clayson, P. E., Carbine, K. A., Baldwin, S. A. & Larson, M. J. Methodological reporting behavior, sample sizes, and statistical power in studies of event-related potentials: Barriers to reproducibility and replicability. *Psychophysiology* **56**, e13437 (2019).
60. Kitchenham, B. A. & Madeyski, L. Recommendations for analysing and meta-analysing small sample size software engineering experiments. *Empir. Softw. Eng.* **29**, 137 (2024).
61. Peitek, N., Apel, S., Parnin, C., Brechmann, A. & Siegmund, J. Program comprehension and code complexity metrics: An fmri study. *2021 IEEE/ACM 43rd Int. Conf. on Softw. Eng. (ICSE)* 524–536 (2021).
62. Zhuang, Y., Yan, Y., DeLong, L. A. & Yeh, M. K. Do developer perceptions have borders? Comparing C code responses across continents. *Softw Qual J* **32**, 431–457 (2024).

63. Siegmund, J., Kästner, C., Liebig, J., Apel, S. & Hanenberg, S. Measuring and modeling programming experience. *Empir Softw Eng* **19**, 1299–1334 (2014).
64. Oldfield, R. C. The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia* **9**, 97–113 (1971).
65. Hessels, R., Niehorster, D., Kemner, C. & Hooge, I. "Noise-robust fixation detection in eye movement data: Identification by two-means clustering (I2MC)". *Behav Res Methods* **49**, 1802–1823 (2017).
66. Carr, J., Pescuma, V., Furlan, M., Ktori, M. & Crepaldi, D. Algorithms for the automated correction of vertical drift in eye-tracking data. *Behav Res Methods* **54**, 287–310 (2022).
67. Landis, J. R. & Koch, G. G. The measurement of observer agreement for categorical data. *Biometrics* **33**, 159–174 (1977).
68. Jasper, H. Report of the committee on methods of clinical examination in electroencephalography. *Electroencephalogr Clin Neurophysiol* **10**, 370–375 (1958).
69. Candia-Rivera, D. & Valenza, G. Cluster permutation analysis for EEG series based on non-parametric Wilcoxon–Mann–Whitney statistical tests. *SoftwareX* **19**, 101170 (2022).
70. Groppe, D. M., Urbach, T. P. & Kutas, M. Mass univariate analysis of event-related brain potentials/fields II: Simulation studies. *Psychophysiology* **48**, 1726–1737 (2011).

Author contribution

Annabelle Bergum, Anna-Maria Maurer, Norman Peitek, Regine Bader, Axel Mecklinger, Janet Siegmund, and Sven Apel developed the concept and the general research idea. Annabelle Bergum, Anna-Maria Maurer, Norman Peitek, Regine Bader, Axel Mecklinger, and Sven Apel devised the study design and formulated the method. Anna-Maria Maurer conducted the study. Anna-Maria Maurer, Norman Peitek, and Vera Demberg performed the data analysis. All authors contributed to the interpretation, writing, editing, and reviewing. Anna-Maria Maurer and Annabelle Bergum prepared the visualizations. Axel Mecklinger, Regine Bader, Sven Apel, Vera Demberg, and Janet Siegmund provided the funding and supervised the project.