



**University of  
Zurich**<sup>UZH</sup>

**Zurich Open Repository and  
Archive**

University of Zurich  
University Library  
Strickhofstrasse 39  
CH-8057 Zurich  
[www.zora.uzh.ch](http://www.zora.uzh.ch)

---

Year: 2022

---

## **First come first served: The impact of file position on code review**

Fregnan, Enrico ; Braz, Larissa ; D'Ambros, Marco ; Çalıklı, Gül ; Bacchelli, Alberto

DOI: <https://doi.org/10.1145/3540250.3549177>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-232022>

Conference or Workshop Item

Published Version

The following work is licensed under a Publisher License.

Originally published at:

Fregnan, Enrico; Braz, Larissa; D'Ambros, Marco; Çalıklı, Gül; Bacchelli, Alberto (2022). First come first served: The impact of file position on code review. In: ESEC/FSE '22: 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Singapore Singapore, 14 December 2022 - 18 December 2022. ACM Digital Library, 483-494.

DOI: <https://doi.org/10.1145/3540250.3549177>

# First Come First Served: The Impact of File Position on Code Review

Enrico Fregnan, Larissa Braz  
{fregnan,larissa}@ifi.uzh.ch  
University of Zurich  
Switzerland

Gül Çalikli  
handangul.calikli@glasgow.ac.uk  
University of Glasgow  
Scotland

Marco D'Ambros  
marco.dambros@usi.ch  
CodeLounge at Software Institute  
Università della Svizzera Italiana, Switzerland

Alberto Bacchelli  
bacchelli@ifi.uzh.ch  
University of Zurich  
Switzerland

## ABSTRACT

The most popular code review tools (*e.g.*, Gerrit and GitHub) present the files to review sorted in alphabetical order. Could this choice or, more generally, the relative position in which a file is presented bias the outcome of code reviews? We investigate this hypothesis by triangulating complementary evidence in a two-step study.

First, we observe developers' code review activity. We analyze the review comments pertaining to 219,476 Pull Requests (PRs) from 138 popular Java projects on GitHub. We found files shown earlier in a PR to receive more comments than files shown later, also when controlling for possible confounding factors: *e.g.*, the presence of discussion threads or the lines added in a file. Second, we measure the impact of file position on defect finding in code review. Recruiting 106 participants, we conduct an online controlled experiment in which we measure participants' performance in detecting two unrelated defects seeded into two different files. Participants are assigned to one of two treatments in which the position of the defective files is switched. For one type of defect, participants are not affected by its file's position; for the other, they have 64% lower odds to identify it when its file is last as opposed to first. Overall, our findings provide evidence that the relative position in which files are presented has an impact on code reviews' outcome; we discuss these results and implications for tool design and code review.

**Data and Materials:** <https://doi.org/10.5281/zenodo.6901285>

## KEYWORDS

Code Review, Controlled Experiment, Cognitive Bias

### ACM Reference Format:

Enrico Fregnan, Larissa Braz, Marco D'Ambros, Gül Çalikli, and Alberto Bacchelli. 2022. First Come First Served: The Impact of File Position on Code Review. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '22)*, November 14–18, 2022, Singapore, Singapore. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3540250.3549177>

---

ESEC/FSE '22, November 14–18, 2022, Singapore, Singapore

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '22)*, November 14–18, 2022, Singapore, Singapore, <https://doi.org/10.1145/3540250.3549177>.

## 1 INTRODUCTION

Code review is a popular software engineering practice where developers manually inspect the code written by a peer [7, 39]. Code review aims to find defects [11], improve software quality [3, 12], and transfer knowledge [7, 47]. Over the years, code review has evolved from a formal strictly-regulated process [22] into a less strict practice. Contemporary code reviewing is informal, asynchronous, change-based, and supported by tools [9, 10, 45, 47].

The tools used to conduct code reviews share many similarities [11]. In particular, the vast majority of tools (including the popular Gerrit [27] and GitHub [28]) present the changes to review as a list/sequence of diff hunks [25] grouped by the file they belong to. Tools sort these files alphabetically, therefore the changes to a file named `org/Controller.java` are always presented before those to a file named `org/Model.java`. Could this choice or, more generally, the relative position in which a file is presented influence the outcome of code review?

This hypothesis seems to be supported by at least two factors. First, most developers tend to start their reviews in the order presented by the review tool [12]. Second, code review is a cognitively demanding task [8] whose outcome might be influenced by cognitive factors [41, 50] also related to the position of the file. For example, developers may be influenced by *attention decrement* (a decrease in attention when exposed to a list of elements [6]) or may deplete their *working memory capacity* (the memory for short-term storage during ongoing tasks [60]) near the end of longer reviews.

In this paper, we set to investigate this hypothesis. We do this by triangulating complementary evidence in a two-step study.

In the first step, we focus on the relation between file position and reviewers' activity. We collect and analyze 219,476 Pull Requests (PRs) from 138 GitHub open-source Java projects and investigate whether the position in which a file is presented in a PR is associated with the number of comments the file receives. In fact, the number of comments can be used to approximate reviewers' effectiveness and activity [45]. We find a significant correlation between the file's position and the number of comments this file receives: Files shown later in a PR receive fewer comments than files shown earlier.

In the second step, we focus on the influence of file position on defect finding in code review. We design and conduct a controlled experiment with 106 developers who have to review code in which we seeded two unrelated bugs (a Corner Case defect and a Missing Break defect) into two different files. By creating two treatments

that switch the position of the defective files (first Corner Case defect and last Missing Break defect, or vice versa), we measure the influence of file position. While we see no effect for the Missing Break defect, we find that developers have 64% lower odds to identify the Corner Case defect when the file containing it is displayed last as opposed to first. To further confirm our findings, we look at the files displayed on the participants' screen during the review task. We detect a statistically significant difference between the time participants spend viewing the first file and the last one.

Overall, our findings suggest that the relative position in which files are presented in a code review has an impact on the code review's outcome. This result has important implications for code review practice and for tool design. For instance, code review tool designers may consider rethinking alphabetically ordering files in favor of a more principled approach (e.g., showing first the most problematic files to review). Alternatively, code change authors may consider clearly signaling developers where to start their review, for example through the use of self-comments or in the change description, when they implemented more challenging parts.

## 2 BACKGROUND AND RELATED WORK

In this section, we present relevant literature on cognitive aspects in modern code review. Then, we illustrate the psychological concepts that might constitute the underlying causes of the effect of file position on code review. Finally, we report possible competing arguments against the hypothesis that file position plays a role in reviewing.

### 2.1 Cognitive Aspects in Code Review

Code review is a collaborative process where human factors play a crucial role [7]. Previous studies conducted at companies such as Microsoft [7] and Google [47] revealed how code review could foster knowledge transfer among developers in a team and improve shared code ownership. However, code review is not only a collaborative process, but it is also a cognitively demanding task for the single reviewer. In particular, a vast amount of research focused on reducing developers' cognitive load to improve reviewers' performance [11, 13, 30]. For instance, Baum et al. [13] conducted a controlled experiment to investigate how developers' code review performance relates to their cognitive load and working memory capacity. Their findings highlighted how working memory is associated with finding delocalized defects while only weakly associated with the ability to detect other defect types.

Another group of studies addresses developers' cognitive biases that might affect code review outcome [20, 35, 50, 54]. Huang et al. investigated how cognitive biases relate to code review process in a controlled experiment using medical imaging and eye-tracking. Chattopadhyay et al. synthesized helpful practices to reduce the effect of cognitive biases on software development activities, including code review. In their two-part field study, the authors identified the manifestation of various cognitive biases during code reviews (e.g., representativeness, availability, anchoring, confirmation bias). Spadini et al. [50] investigated the effect of *priming* on reviewers' ability to detect bugs due to their proneness to availability bias. In the controlled experiment, a review comment was used to prime the participants to look for defects of the same type. The

authors initially hypothesized that the participants would overlook defects of other types present in the code due to their proneness to availability bias. However, the experiment results show that the presence of a review comment identifying a bug does not prime reviewers towards looking for defects of similar type overlooking other defect types. On the contrary, such a review comment acts as a reminder for bugs developers usually do not look for during their daily practices.

### 2.2 Psychological Factors Related to Position

Based on the findings of Baum et al. [12], which suggest that developers start a code review from the first file in the change-set, the following psychological factors seem plausible explanations behind the potential effect of files' position on code review.

**Attention decrement.** Attention decrement is defined as “*the tendency for people to pay less attention to stimuli coming later in a sequential occurrence or presentation and thus to remember them less well*” [6]. To give a practical example, a student given a list of terms to memorize is likely to have more difficulties committing to memory those at the end. Hendrick et al. [32] identified attention decrement as the most feasible explanation of why people tend to remember the first information they see the most.

During code review, developers might be subjected to this phenomenon, slowly decreasing their attention the more the review continues. As the majority of reviewers begin their review from the first file in a change-set [12], this decrease in attention is more likely to impact the last files in a review change-set. This would sustain our hypothesis that the position of the files in code review matters.

**Working memory.** Working memory is defined as the part of human memory needed for short-term storage during ongoing cognitive processes [13, 60]. However, although the working memory capacity varies from individual to individual, its amount is still limited [21].

In software engineering, the study of Bergersen and Gustafsson investigated the effect of working memory on programming performance [14]. Their findings showed how working memory influences performance (albeit mediated by the developers' programming knowledge). Baum et al. studied the effect of working memory on code review performance [13]. Their experiment showed a correlation between working memory and the reviewers' effectiveness in finding delocalized defects. Reviewers might deplete their working memory looking at the first files in a review change-set, leading to an exhaustion of this resource that might negatively influence their code review performance on the last files they inspect.

### 2.3 Related Competing Arguments

Although previous evidence [12] suggests that developers start a code review from the first file in the change-set, some factors might affect the order in which developers review files during code review. Given the findings of some eye-tracking studies (e.g., [2, 19, 46]), it seems plausible to think that some files in a code change might attract reviewers' attention more due to their content (e.g., files that contain more call terms, control flow terms) regardless of their position. For instance, Abid et al. [2] found that developers visit call terms more often than non-call terms and spend the longest time

reading call terms. The next most visited and read locations were control flow terms and signatures. Furthermore, in an earlier study Rodeghero et al. [46] found that developers consider method signatures as the most important section of the code followed by call terms and the control flow terms. Moreover, token length and frequency in a source code might also influence developers' attention. The results Al Madi et al. [4] obtained indicate that participants fixate longer low-frequency tokens and tokens containing more characters.

Furthermore, in an eye-tracking study, Busjahn et al. [19] showed that people read source code in a more non-linear fashion than reading natural language. However, the findings by Busjahn et al. only explain how developers read code in a single code file: the authors showed the participants Java classes and pseudocodes ranging from a few lines of code to an entire screen full of text. Therefore, whether developers navigate files in a code change in a non-linear fashion during code review and how this is related to file position need further investigation.

In the presence of these arguments in the literature competing with the hypothesis that file location may influence code review outcome, our study aims to investigate this topic in-depth from complementary angles.

### 3 METHODOLOGY

In this study, we aim to understand whether the relative position in which a file is shown for review influences the code review process.

#### 3.1 Research Questions

We structured our investigation in two steps, which seek to collect complementary evidence. First, we investigate the relationship between file position and reviewers' activity. We do so by mining data from PRs belonging to a vast set of open source projects and answer the following research question:

**RQ<sub>1</sub>.** To what extent is the relative position of a file in a review associated with the amount of comments the file receives?

Second, we focus on the influence of file position on defect finding in code review. We do so by performing an online controlled experiment with software developers who have to conduct a review of code and answer the following question:

**RQ<sub>2</sub>.** What is the effect of the relative position of a defective file in a review on bug detection?

#### 3.2 RQ<sub>1</sub> - File Position and Review Activity

**Subject Projects.** For our analysis, we select 138 open-source projects from GitHub, focusing on Java projects with a star-count above 1,000. As previously reported [16, 17], the number of stars of a project can be effectively used as an indication of the projects' popularity and health. We focus on Java as (1) it is a widely popular programming language [56] and (2) focusing on one programming language might reduce potential bias introduced by different review practices of projects based on other programming languages. Moreover, we investigate large projects to reduce potential bias caused by project-specific review policies or characteristics.

**Data Collection.** Using PyGitHub<sup>1</sup> a Python library to access the GitHub REST API, for each PR, we extract the position of each file and the number of comments it receives. Moreover, we collect other factors that can influence the number of comments a file receives (*i.e.*, confounding factors). We consider: (1) The number of lines added and (2) deleted in a file because larger changes may require more comments; (3) whether the file is a test, because these tend to receive less comments [49] and to be ordered last alphabetically; and (4) the number of commenters, because more participants in the review of a PR might lead to more comments.

**Data Filtering.** As the presence of a bot might introduce bias in our analysis, we exclude bot comments from our analysis. To do so, we flag all users containing the word "bot" as bots and create a list of commonly used bots in GitHub (extracted from relevant gray literature [55, 59]) and remove them. The presence of discussion threads (*i.e.*, a series of comments where the reviewers and the author discuss solutions or improvements to the code) might act as a confounding factor in our analysis: A discussion thread can facilitate developers' engagement in adding comments to the thread regardless of the file's position. Therefore, we focus our investigation only on the first comment of a thread, disregarding subsequent ones. Moreover, the number of files in a Pull Request might influence the comments' distribution. Therefore, we group the PRs according to the number of files they contain to conduct initial analyses (*e.g.*, see Figure 4 for PRs with five files).

**Data Analysis.** To analyze the impact of the aforementioned factors on the number of comments in a file (dependent variable), we build a *Hurdle model* [24]: A statistical model that specifies two processes, one for zero counts and another for positive counts. We choose this model as it handles excess zeros in the dependent variables, in fact, a vast number of files in the collected PRs receive zero comments. To model the positive counts, we employ a *negative binomial* distribution. Furthermore, we check the multi-collinearity across the variables in the model computing their Variance Inflation Factor (VIF) and removing those with a VIF higher than five.

#### 3.3 RQ<sub>2</sub> - File Position and Defect Finding

Our controlled experiment is organized as an online experiment in which participants are asked to complete a review of a code change involving five files. Among these files, we seed two unrelated software defects<sup>2</sup> into two different files. We randomly assign participants to one of the following two treatments: (1) one file with a bug is presented first and the other file with the other bug is presented last (*i.e.*, fifth), or (2) the order of the files (hence the bugs) is reversed (*i.e.*, the fifth file is now first and vice versa). We measure the specific bugs detected by each participant, as well as how long each file is visible on the participants' screen.

In the following, we provide more details on the *experimental objects* (*i.e.*, code change to review and seeded defects) and the *experimental design* (*i.e.*, online platform and experiment flow). We conclude by describing the pilots we conducted, how we recruited participants, and how we analyzed the collected experimental data.

<sup>1</sup>PyGitHub: <https://github.com/PyGithub/PyGithub>

<sup>2</sup>The participants are not informed about the presence of these defects.

### 3.3.1 Experimental Objects

**Code change.** For the review task, we create a code change that satisfies the following requirements: (1) not belonging to any existing project (to avoid introducing bias caused by participants' previous knowledge of the review change-set); (2) written in Java, one of the most popular programming languages [56]; (3) self-contained (to minimize the previous knowledge participants need to understand the change fully); and (4) close to an actual code review scenario (to increase realism).

We create a code change spanning five files as a trade-off between small change-sets (e.g., three files) and more complex reviews. For smaller code changes (e.g., two files), we expect a smaller effect of a file's position. In contrast, even though larger code changes may present a stronger effect, more complex reviews require a longer review time, thus likely leading to more participants abandoning the experiment.

All files have similar sizes (ranging from 44 to 63 lines). In particular, we ensure that the two files containing the defects are similarly sized (46 lines in the Corner Case defect file and 44 lines in the Missing Break defect) to minimize the potential bias introduced by the file length.

We devise the code change to minimize links (e.g., method calls) between non-adjacent files. In particular, we checked that no connection existed between the two defective files (i.e., first and last). Our goal is to increase the control on how participants navigate the change, thus making the file position effect clearer, if any.

Finally, to re-position the files for the two different treatments, we rename them. This ensures that the reviewers are not influenced by a tool behavior that is unexpectedly different than what they are used to in common review platforms (where files are displayed in alphabetical order).

**Seeded Defects.** In our experiment, we investigate if the given treatment influences participants' ability to find bugs. In the review task, we seed the following two unrelated functional defects:

- **Corner Case (CC):** A corner case condition in an `if` statement is left unchecked, thus not respecting the documentation (Figure 1). We select this bug because it represents an issue developers typically check for [18, 50] and frequently exists in practice [36].
- **Missing Break (MB):** A missing break statement in a `switch` construct makes the execution incorrectly fall through the default case (Figure 2). We select this defect because it is reported as a common Java mistake by relevant gray literature [1, 58], also resonating with the infamous Apple *goto fail* [15].

We use two defects to make the review task as realistic as possible by preventing participants from focusing exclusively on one bug.

### 3.3.2 Experimental Design

**Experiment Platform.** We implement our experiment as an online platform. To this aim, we employ a publicly available tool, CRExperiment [48], already successfully used in previous code review experiments [18, 50]. CRExperiment allows participants to review code displaying a review change-set in two-pane diffs, as done in popular code review tools: e.g., Gerrit [27] or GitHub [28]. Moreover, the User Interface (UI) of CRExperiment adopts similar design choices (e.g., the color scheme) as the one of, for instance,

```

7      /**
8       * Returns the discount rate based on the membership level of the
          customer.
9       * Customers at level 1 do not receive any discount.
10      * Customers at level 2 to 4 receive a 10% discount.
11      * Customers from level 5 included receive a 25% discount.
12      * @param membershipLevel - the level of membership of the customer
13      * @return the discount rate applied to the customer
14      */
15      public double getSaleDiscountRate(int membershipLevel){
16          double discountRate = 0;
17          if(membershipLevel > 2 && membershipLevel < 5) {
18              discountRate = 0.1;
19          }
20          if(membershipLevel >= 5) {
21              discountRate = 0.25;
22          }
23          return discountRate;
24      }

```

**CC: Corner Case defect:** Here the if statement is missing a check for the condition where `customer.membershipLevel == 2`. According to the Javadoc of the function, customers with membership level equal to 2 should receive a 10% discount

Figure 1: Corner Case defect (CC) used in our experiment.

```

26      switch (destinationAddress.getCountry()) {
27          case "USA":
28              shippingCost = shippingCost * 1.2;
29              break;
30          case "Canada":
31              shippingCost = shippingCost * 1.18;
32              break;
33          case "Mexico":
34              shippingCost = shippingCost * 1.35;
35              break;
36          case "UK":
37              shippingCost = shippingCost * 1.27;
38          default:
39              shippingCost = shippingCost * 2;
40      }
41      return shippingCost;

```

**MB: Missing Break defect:** Here a break statement is missing. In this way, when the country is UK, the execution will fall through the default case and a wrong tax of  $1.27 * 2$  will be applied.

Figure 2: Missing Break defect (MB) used in our experiment.

Gerrit. This mitigates possible bias with the users not being familiar with the experiment platform.

CRExperiment also logs participants' interactions (e.g., scrolling) and the time participants spend in each phase of the study (e.g., in the review tasks). To collect information on which files in the change-set the participants focus on during code review, we extended the base experiment platform to record the files visible on the participants' screen during the review. For each file, the experiment tool records if this was on the participants' screen (displayed), partially displayed, or not displayed at a given time. We store all the collected data anonymously.

**Experiment flow.** We organize the flow of our experiment as depicted in Figure 3 and as described in the following:

- (i) **Welcome page.** On the experiment's landing page, we provide participants with information on the type of study and our data handling policy. Moreover, we ask for their consent to collect and use their data before proceeding in the experiment.
- (ii) **Code review task.** Participants are asked to perform a code review task. Before starting, we instruct participants to perform a code review as they would typically do in their normal practice. Participants are randomly assigned to one of two possible treatments:

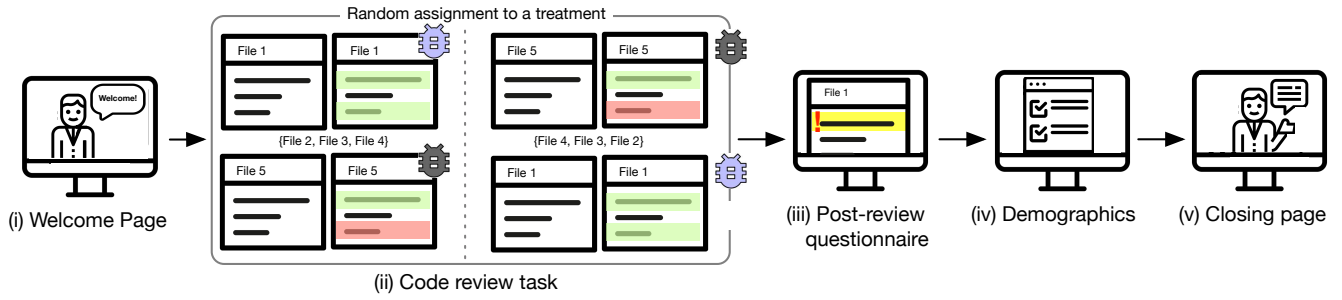


Figure 3: Design and flow of the online experiment.

- $CC_f$ - $MB_l$  - The file containing the Corner Case defect is shown *first* in the code change, while the one containing the Missing Break defect is shown *last*.
- $MB_f$ - $CC_l$  - The file containing the Missing Break defect is shown *first*, while the one with the Corner Case defect is shown *last*.

To avoid bias, we do not inform participants about the functional defects, the treatments, and the recorded metrics before the review task. At the end of the task, we ask participants whether and for how long they were interrupted during the review.

**(iii) Post-review questionnaire.** After participants declare their review task is completed, they are shown the code change again with the location and explanation of the defects (Figure 1 and Figure 2). We ask participants whether they found each of the two defects and the cause that, in their opinion, influenced their result. In a subsequent page (as not to influence their previous answers), we ask participants if they think the position of the file containing the bugs might have impacted their results.

**(iv) Demographics.** Afterward, we collect information about participants' gender, education, occupation, and experience with Java and code review. We analyze this information (1) to guarantee that the groups of participants assigned to different treatments are homogeneous, (2) as confounding factors in the analysis of the experiment results, and (3) to describe the population represented in our results.

**(v) Closing page.** On the final page of the experiment, we disclose the goal of our study and ask participants for any final remark. We also collect participants' consent to share their answers in an anonymized, yet publicly available, research dataset.

### 3.3.3 Pilot Runs

Before publicly releasing the online experiment, we conduct a pilot study to (1) verify the absence of technical issues with experiment settings (e.g., the online platform) and instructions, (2) check the goodness of the devised review task and seeded defects, and (3) improve the experiment based on the participants' feedback.

We conduct four iterations of our pilot study with a total amount of 15 participants. After each iteration, the authors discuss the results and feedback obtained and improve the experiment design accordingly. After the last iteration, the experiment is deemed ready to be released since the participants detected no significant issues.

We excluded the data gathered from the participants in the pilot study from the analysis of the experiment results.

### 3.3.4 Data Analysis

To analyze the data collected in our experiment, we perform a *Chi-square* test to assess the correlation between files position and participants' detection of the defects. We evaluate the strength of the correlation using the *phi-coefficient*. Then, we build two logistic regression models considering as dependent variables (1)  $CC_{Found}$  (whether the participants found the Corner Case defect), and (2)  $MB_{Found}$  (whether the participants found the Missing Break defect), respectively.

To identify if a participant found a defect, we follow the guidelines of previous studies [13, 30, 31]. We manually inspect the remarks left by the participants and classify a defect as found when a remark (1) is placed in proximity of the defect and (2) clearly describes the defect in the code. The first author performs this inspection, and a second author checks later the first author's inspection results.

To build our models, we follow the following procedure to ensure the goodness of our results: (1) We use the VARCLUS procedure to identify and remove variables from the model with Spearman's correlation higher than 0.5. (2) We compute the Variance Inflation Factors (VIF) among the variables to ensure none had a VIF value above 7. If a variable with a higher VIF is detected, we remove it from the model. (3) Finally, we build the models adding each variable one by one and checking that the coefficient remained stable, showing little interference among the considered variables.

We aim to evaluate if the position of the defect (in the first or last file) impacts participants' ability to find it. For this reason, we include the variable *position* among the independent variables in our model. We also consider the potential effect of other confounding factors, such as the participants' experience with Java and code review. Furthermore, we compute the time participants spend on the review task and include this variable in the model. Table 1 presents all the variables included in our logistic regression models.

### 3.3.5 Number of required participants.


We performed an *a priori* power analysis to compute the minimum sample size needed for our experiment. We used the G-Power software [23] and employed a *two-tail* test (with a manual distribution) having *odds ratio* = 1.5,  $\alpha$  = 0.05, *Power* = 0.8, and  $R^2$  = 0.3. The results of this analysis showed our experiment needed a minimum of 92 valid participants.

### 3.3.6 Participants' recruitment

We disseminated an invite to conduct the online experiment through the authors' professional network and social media accounts (e.g.,

**Table 1: Variables used in the logistic regression models.**

Metric	Description
<i>Dependent variables</i>	
CCFound	The participant found the Corner Case defect.
MBFound	The participant found the Missing Break defect.
<i>Independent variables</i>	
Position	The relative position of the bug (first or last) in the code change to review.
<i>Control variables</i>	
CRDuration	The time spent on the code review task.
DevExp	The participant’s development experience in a professional setting.
JavaExp	The participant’s experience with Java.
CRExp	The participant’s experience in code review.
OftenProg	How often the participant programmed in the three months prior to the experiment.
OftenCR	How often the participant reviewed code in the three months prior to the experiment.
Interruptions	How often and for how long the participant was interrupted during the review task.

**Figure 4: Distribution of comments in a PR with five files.**

LinkedIn), as well as on developers’ web forums and communication channels (e.g., Reddit). To prevent any bias in the results, we do not disclose the experiment’s aim to participants in the invite. To encourage developers’ participation, we commit to donate 5 USD to a charity for each valid participant in the experiment.

## 4 RESULTS

We present the results of our investigation by research question.

### 4.1 RQ<sub>1</sub> - File Position and Review Activity

Our first research question seeks to investigate the relationship between relative file position and reviewers’ activity, focusing on the comments left during code review. We analyze a total of 219,476 pull requests pertaining to 138 popular Java projects on GitHub.

**Table 2: Hurdle models for PR summary.**

	Positive count		Zero count	
	Estim.	Sig.	Estim.	Sig.
Intercept	-8.982	**	-3.571	***
Position	-0.187	***	-0.343	***
Lines added	0.012	***	0.007	***
Lines deleted	0.003	***	-0.001	***
isTest	-0.433	***	-0.707	***
N. commenters	0.485	***	1.142	***

Sig. codes: \*\*\*  $p < 0.001$ ; \*\*  $p < 0.01$ ; \*  $p < 0.05$

Among these, 26,685 pull requests received at least one review comment (which we computed excluding comments left by bots).

As the vast majority of PRs (81.06%) contain between one and ten files, we focused our investigation on this subset of PRs. We excluded from our analysis PRs containing only one file as the effect of the position in such cases is not relevant.

Code review on GitHub is an iterative process. When a developer uploads a commit, this is reviewed by fellow developers who may ask the commit author to perform some changes. These changes are addressed in a subsequent commit and the process continues iteratively until the code is accepted. To both consider and exclude the effect of the review process itself, we analyzed the data concerning two different moments in the history of pull requests:

- **Pull Request summary** : We consider the code change as it appears at the end of the review process and all the comments that the code has received during the entire process.
- **First commit** : We consider only the code as initially submitted in the pull request and the comments it receives. Thus, we exclude changes and comments induced by the review process.

#### 4.1.1 Pull Request summary

Our initial investigation showed that the number of comments on the files in a PR is higher in the first files and progressively decreases towards the last files. For instance, Figure 4 shows the distribution of the comments in PRs with five files. A similar pattern was observed also for all the other groups of PRs, e.g., with six files. We found similar results also when we restrict our investigation to the PRs that contain only Java files.<sup>3</sup>

To confirm and quantify the correlation between the relative position of a file in a pull request and the number of comments it receives, we performed a *Spearman correlation test* (a non-parametric test that measures the strength of the association between ordinal or continuous variables [53]). We obtained a p-value of  $< 2.2 \cdot 10^{-16}$  and a  $\rho = -0.203$ , thus confirming a statistically significant negative relation between these two variables.

Subsequently, we performed regression modeling to verify this association also when controlling for the other factors that may influence the number of review comments on a file.

Table 2 shows the Hurdle model predicting the number of comments in a file. We included in the model PRs containing only Java files to remove potential bias introduced by multiple file formats. Moreover, we excluded PRs having less than three files, since in these cases the effect of the file position is likely to be negligible.

<sup>3</sup>Our online appendix [26] provides graphs for all the cases we mention.

**Table 3: Hurdle models for PR’s first commit.**

	Positive count		Zero count	
	Estim.	Sig.	Estim.	Sig.
Intercept	-10.34	.	-4.200	***
Position	-0.183	**	-0.221	***
Lines added	0.019	***	0.005	***
Lines deleted	0.004	***	$-4.15 \cdot 10^{-4}$	*
isTest	-0.353	***	-0.667	***
N. commenters	0.607	***	1.685	***

Sig. codes: \*\*\*  $p < 0.001$ ; \*\*  $p < 0.01$ ; \*  $p < 0.05$ ; .  $p < 0.1$

To remove outliers, we limited our analysis to files that have fewer than 1,000 lines added. Finally, we normalized each file’s position over the number of files in the PR to allow for comparison among PRs with a different number of files. The model is summarized in Table 2 and shows how the position of a file is negatively correlated with the number of comments it receives, confirming how the last files in a review change-set tend to receive less review comments from developers, also when controlling for other factors.

Table 2 also shows that the other independent variables included in the model are statistically significant predictors of the number of comments in a file. The sign of the estimates show that the direction of the relationship follows the expected direction. For instance, a larger number of changed lines (added or deleted) is related to more reviewers’ comments. On the contrary, the expected number of comments decreases when the file is a test (this result is inline with previous research’s findings [49]). Finally, the number of commenters is positively correlated with the number of comments.

#### 4.1.2 First commit

By restricting our analysis to only the first commits of our set of Pull Requests, we noticed a similar pattern as the one reported for the PR summary. The first files in a commit receive more comments compared to the last ones. We analyzed commits with number of files ranging from two to ten, noticing similar patterns to the PR summary (available in the replication package [26]).

As for the PR summary, we corroborated our findings by restricting our analysis to commits containing only Java files to mitigate potential bias caused by multiple types of files. This verification, however, did not reveal any significant difference from the general case confirming our observations that a higher number of comments concentrates on the first files in a commit.

We performed a *Spearman correlation test* to verify the presence of a correlation between the number of review comments a file receives and the file’s position. The test achieved a  $p$ -value of  $< 2.2 \cdot 10^{-16}$  and a  $\rho = -0.152$ , showing how the two variables are correlated (although this correlation is weaker).

To build the Hurdle model in this scenario, we followed the same steps as the one reported for the Pull request summary. Our results (reported in Table 3) confirm how later files in a review change-set tend to receive a lower number of comments. Concerning the other independent variables included in our model, we achieved results similar to the ones presented for the PR summary (Section 4.1.1).

**Finding 1.** *The number of comments across the files in the analyzed pull requests is associated with the relative position of each file. Specifically, the later a file is presented in a pull request, the fewer the comments it receives.*

## 4.2 RQ<sub>2</sub> - File Position and Defect Finding

Encouraged by the findings for RQ<sub>1</sub>, we devised a controlled experiment to test this hypothesis with complementary evidence. Our second research question seeks to investigate the effect of files’ position on developers’ review effectiveness.

A total of 372 participants accessed the online experiment. Of these, we considered only the participants who completed all the steps. Furthermore, we removed participants who left no remarks and spent less than ten minutes doing the review, leaving us with 106 participants.

The vast majority of the participants (84.9%) have at least a B.Sc. degree (mostly in Computer Science). Overall, 72 participants reported to be software developers. Moreover, 84 participants self-described as male, four as female, two as non-binary, and 16 preferred not to disclose. Figure 5 reports participants’ experience and practice levels.

In total, 56 participants were in treatment  $CC_f-MB_l$ , while 50 were in treatment  $MB_f-CC_l$ . We compared the experience and practices (e.g., Java experience) of the participants assigned to the two groups and found no statistically significant difference.

### 4.2.1 Defect Finding

Overall, 52 participants found the Corner Case defect, while 50 detected the Missing Break defect. Table 4 reports the number of participants who found no defect, only one defect, or both defects in the review task, by treatment. Table 5 reports how many participants identified each type of bug, by treatment.

**Table 4: Participants who found (1) no defects, (2) only the Corner Case defect (CC), (3) only the Missing Break defect (MB), and (4) both defects.**

	No defect	CC	MB	CC & MB	Total
$CC_f-MB_l$	13	19	9	15	56
$MB_f-CC_l$	17	7	15	11	50
<b>Total</b>	30	26	24	26	106

**Table 5: Participants who found/not found each bug divided by treatment: Corner Case defect (CC) first or Missing Break defect (MB) first.**

	Corner Case		Missing Break	
	Found	Not found	Found	Not found
$CC_f-MB_l$	34	22	24	32
$MB_f-CC_l$	18	32	26	24
	p-value: <b>0.011</b>		p-value: 0.346	
	phi-coefficient: 0.247		phi-coefficient: 0.091	
	odds ratio: 2.75		odds ratio: 1.44	






Figure 5: Participants' experience and practice.

These results show how participants found more the first bug shown in the review task compared to the bug shown last. The Corner Case defect was found 34 times when the file containing it was shown first and only 18 when the file was shown last. A similar, albeit weaker, trend seems to appear for the Missing Break defect.

To verify the presence of a relation between the position of a bug and its detection, we performed a *Chi-square* test. In the case of the Corner Case defect, this test achieved  $\chi^2(df = 1, N = 106) = 6.45$ ,  $p$ -value = 0.011, rejecting the null hypothesis that no relationship exists between files' position and the detection of the Corner Case defect. However, in the case of Missing Break defect our test achieved  $\chi^2(df = 1, N = 106) = 0.88$ ,  $p$ -value = 0.346. Therefore, for this defect we could not draw any conclusion on the fact that its position influenced developers' ability to find it.

For the Corner Case defect, we also computed the phi-coefficient to measure the strength of the association obtaining a value of 0.247, which is close to a moderate positive correlation [29].

Table 6: Logistic regression models for RQ2.

	Dep. Var. = CC Found			Dep. Var. = MB Found		
	Estim.	S.E.	Sig.	Estim.	S.E.	Sig.
Intercept	-1.113	1.570		-1.573	1.536	
Position	-0.960	0.409	*	0.388	0.418	
CRDur.	-3 10 <sup>-4</sup>	0.009		0.005	0.010	
DevExp	-0.024	0.225		0.052	0.227	
JavaExp	0.120	0.189		0.251	0.191	
CRExp	0.045	0.228		0.326	0.236	
OftenProg	0.171	0.323		0.130	0.312	
OftenCR	0.024	0.220		-0.547	0.235	*
Interrupt.	0.005	0.156		-0.056	0.158	

Sig. codes: \*  $p < 0.05$


To investigate the association between the detection of a defect, its position in the code change, and other possible confounding factors (reported in Table 1), we built two logistic regression models, whose results are shown in Table 6. These findings are aligned with the results of the *Chi-square* test: The position of the Corner Case defect is statistically significant correlated with developers' ability to find it ( $p$ -value < 0.05).

**Finding 2.** The relative position of the file containing the Corner Case defect influences the likelihood of participants in detecting the defect. Participants are less likely to find the defect when its file is presented last. This effect is not significant for the Missing Break defect.

At the end of the code review task we asked participants if they thought the position of the defect had an influence on their ability to find it. The vast majority of them replied negatively (72.6% for the Corner Case defect and 76.4% for the Missing Break defect). However, our results showed how, on the contrary, the position of the bug had an effect on participants' ability to find it. This shows a mismatch between what developers *think* affects the code review outcome and which factors *actually* play a role in it.

#### 4.2.2 Visible Files

In the experiment, we measured the time that each file was visible on the participants' screen as a way to investigate further the effect of files' position on code review. Figure 6 reports the time (in seconds) that participants in the two treatments spent with each file visible on the screen. We removed from this analysis participants who declared to have been interrupted for more than ten minutes during the review task, as their data may bias this analysis.



**Figure 6: Time (in seconds) participants visualized each file. To improve the clarity, we limited the size of the Y axis.**

On average, participants regardless of their treatment ( $CC_f$ - $MB_l$  or  $MB_f$ - $CC_l$ ) spent more time looking at the first file compared to the last one. On average, participants spent 8.58 minutes with the first file displayed as opposed to 6.09 for the last file.

To verify whether the difference between the time participants visualized the first and last file is statistically significant, we performed a *Mann-Whitney U test*. For the whole participants (without dividing them by the treatment), our test obtained a p-value of 0.002, therefore showing that the two variables are indeed different. Dividing the participants by treatment, our test achieved a p-value of 0.022 for the treatment  $CC_f$ - $MB_l$ , and 0.031 for treatment  $MB_f$ - $CC_l$ . These results confirm those for the general population.

**Finding 3.** During the review task, participants spent significantly more time with the first file being visible compared to the last file.

#### 4.2.3 Robustness Testing

To challenge the validity of our findings and strengthen their goodness, we employed *robustness testing* [42].

**Participants' groups are not homogeneous.** Participants' experience in programming and reviewing might impact their performance during code review and, therefore, influence the results. To verify that the participants assigned to the two treatments have similar characteristics, we performed a *Chi-square test of homogeneity* on the variables measuring participants' experience and practice (e.g., participants' experience with Java). We obtained p-values above 0.1, therefore, not revealing significant differences between the two samples of participants. Moreover, we compared the time spent on the code review task of participants belonging to the two treatments. To this aim, we performed a *Mann-Whitney U test* (p-value of 0.453), which did not detect any difference between the two groups.

#### One defect might influence participants in finding the other.

Finding one defect might have biased participants towards finding the other one. For instance, participants who found one bug might have stopped looking for other defects assuming they found the only issue in the code. Another possibility is that one defect can give unintentional indications to participants to identify the other one. However, the results reported in Table 4 counter the existence of such bias: the number of participants who found only one bug is similar to the one of those who found both bugs. We further performed a *Chi-square* test to verify if finding one defect led participants to also find the second one. Our test obtained a p-value of 0.888, thus suggesting this bias not to affect the experiment.

**The defects are too easy/difficult.** The selected defects might have been too easy, or on the contrary, too hard to identify for the participants. To mitigate this possible bias we (1) selected defects among common Java mistakes based on relevant literature [1, 36, 51] and discussed them among the authors, and (2) verified our choice through a pilot study. Despite these measures, participants in the experiment might have still faced difficulties finding these defects (or identified them too easily). However, the results reported in Table 4 show that the number of participants who found no defects, only one defect, or both defects is almost evenly distributed, showing that no defect was too trivial or hard to find.

**A low number of participants.** A number of participants too low might bias the significance of our findings. The power analysis we run (Section 3.3) showed that we needed at least 92 participants. Our experiment exceeded this number: It was conducted by a total of 106 valid participants.

## 5 THREATS TO VALIDITY

**Construct validity.** The code changes selected for the review task might have introduced bias in the experiment. To mitigate this threat, the first and second authors prepared the code changes and selected the defects contained in them. The other authors also checked the goodness of the code review task.

The scenario represented in our experiment might differ from a real-world review. To reduce this threat, we adopted the following measures: we (1) created a code change as similar as possible to real ones (e.g., including documentation), (2) seeded defects that commonly happen in real-world scenarios [1, 50], and (3) used a review interface similar to widely used code review tools.

**Internal validity.** During code review, developers, instead of writing multiple similar comments, might leave only one comment requiring changes in all similar subsequent instances of an issue. Such comments might introduce bias in the results of our investigation in RQ<sub>1</sub>. For this reason, one of the authors manually inspected a randomly selected subset of 100 PRs, looking for such comments. Only eight PRs contained instances of these comments. Considering this proportion, we computed a statistically significant sample size with confidence level 95% and margin of error 5%, obtaining a sample size of 54, which confirmed that 100 PRs are sufficient to establish the proportion of this factor in our results.

In RQ<sub>2</sub>, before analyzing the results, we manually inspected the participants' logs. We removed all participants that did not take the review task seriously: Participants who spent less than ten minutes on the review and left no remarks. Since our experiment relied on an

online platform, the participants might have completed the review task with different set-ups and in different environments. However, this reflects the various conditions in which developers work in real-world scenarios. To mitigate the threats that interruptions might constitute for the validity of our study, we asked participants to state how long they were interrupted during the code review task and included their answers in our statistical analyses.

**External validity.** In RQ<sub>1</sub>, the selection of the projects might have introduced bias in our investigation. To reduce potential bias caused by the choice of a specific project, we considered a vast set of projects from different domains, all with star-count above 1,000. However, we cannot exclude the possibility that the choice of different projects might lead to different results. Further studies are needed to verify the generalizability of our findings.

Participants in our experiment possess a vast set of different backgrounds, yet our sample is not representative of all developers. Moreover, the specific defects we chose might have influenced the experiment results. Although these defects have been reported in the literature as common defects [1, 36, 50], they have specific characteristics that do not generalize to all other types of defects.

Furthermore, the size of the code change contained in the review task might have influenced the observed results. We carefully chose the review code change size to be a trade-off between a too small review (where the effect of the files' positions is likely not to matter) and a too-large task, which would have resulted in a higher participants drop-out rate. Although the results of RQ<sub>1</sub> suggest that this effect should be present also in the code change with a higher number of files, further studies are needed to corroborate our findings in different scenarios.

In our experiment, we considered only code written in Java. This was done to keep the necessary number of participants (Section 3.3) within an achievable range. In fact, considering multiple programming languages would have required to significantly extend the number of participants due to the following reasons. First, although languages such as Java, C#, and Python can be used to write code in a similar way, in reality they tend to follow different coding conventions and idioms [5, 44, 52]. Therefore, using multiple languages poses the issue of either not following the language's idioms (thus making the code snippet less "natural") or having snippets substantially different across languages. Second, past research has provided evidence on different programmers' behavior when comprehending code written in different languages [57]. Third, work on programming languages learning shows that novices face different difficulty levels with different programming languages [40]. This might be caused by the different cognitive load that each language poses. For these reasons, we cannot exclude that different programming languages might lead to different results. Further studies can be devised and conducted to investigate whether and how the results of our experiment change considering different programming languages.

## 6 DISCUSSION

In this section, we discuss the main implications of our findings.

**The importance of being first.** Our investigation provides evidence that the position of a file influences developers' review effectiveness: Developers have 64% lower odds to find a defect when it is

in the file shown last. Currently, code review tools (e.g., GitHub or Phabricator) display files in alphabetical order. Our results suggest that more principled ways of presenting the files in a code review can be used to support reviewers' effectiveness.

For example, tools could display first those files that are more critical. Previous work focused on identifying salient classes in a code change (i.e., classes that subsequently cause the modification of other classes in the change-set) [34]. Such classes can be used as the point from which developers should start a revision. On the contrary, another possibility would be to focus on identifying *arid files* and placing them last in a code review, taking inspiration from the approach developed at Google to identify *arid lines* [43]. Such an approach would allow developers to prioritize the review of more relevant files.

Other studies proposed an ordering approach to re-organize files in a code change to support reviewers' preferences [12, 13]. This approach focuses on how files to review should be grouped, suggesting to show files sequentially sharing a link (e.g., method call - method declaration). However, this ordering theory does not consider the absolute position of the files (i.e., which ones should be shown first). The ordering theory can be expanded to take into account our findings. Moreover, previous work [12] reported how developers adopt different tactics when reviewing code (e.g., starting from newly added files or files perceived as easier to review). For this reason, future code review tools might include a way to let reviewers customize the order of the files to review to fit different reviewers' needs and counter the effect of the relative position of the files.

Finally, if changing the alphabetical order of the files in a code review tool is too unnatural for reviewers, one could consider using warnings to point reviewers towards the file from which they should start the review. To this aim, new solutions could be devised taking inspiration from previous studies on program-analysis tools [33, 38], where the authors investigated how warnings should be displayed to be welcomed by developers. These warnings could, for instance, make developers aware of biases during code review (e.g., as created by the relative file position). Future studies can be designed and carried out to compare the effectiveness of using warnings as opposed to changing the position of files in a code review tool.

**Not all defects are the same.** In our experiment, participants' effectiveness in detecting the Corner Case defect was influenced by the relative position of the defective file. However, this was not the case for the Missing Break defect.

If we consider the nature of these bugs, we notice that they may require a different effort to be found, even though they were found with the same probability overall (as reported in Table 5). On the one hand, the Corner Case defect requires the reviewer to read and understand the documentation, then comprehend the code and identify a mismatch in the corresponding implementation. As one participant explained when asked why they did not find the Corner Case defect: "I didn't read the JavaDoc carefully enough, and there was not anything obviously wrong with the code." On the other hand, the Missing Break defect can be recognized as a pattern, even if one does not fully understand the semantic of the code. One participant who found it explained: "I am accustomed to looking

for break statements on every switch expression because I have missed them in my own code in the past.”

Therefore, the underlying reason for the difference in files’ position effect could be that these two bugs require a different cognitive effort to be found. These bugs could impose a different load on participants’ working memory, making, in turn, the effect of the defect position more or less prominent. We did not investigate this further, but studies can investigate this hypothesis and determine the role of cognitive effort in defect finding during review.

**The hidden psychological factors of code review.** At the end of our experiment, we showed the participants the defects we seeded in the code. When asked whether the position of the bugs might have had an impact on their ability to detect them, a vast amount of the developers replied negatively. A total of 77 (72.6%) participants rejected the idea that files’ position influenced their ability to find for the Corner Case defect. However, the results of our experiment found evidence of the contrary. The impact of file position on reviewers was also corroborated by the first part of our study when associating file position and review activity.

This mismatch between participants’ perceptions and actual results contributes to the discourse around the importance and the impact of cognitive biases that may affect developers during different software engineering activities [41]. After all, code review is a human effort and as humans we are influenced by biases and “hidden” psychological factors of which we are not naturally aware.

**Beyond code review.** Our results provided initial evidence on the effect of files’ position on developers’ activity and effectiveness in code review. However, this phenomenon might also affect other aspects of software engineering.

Previous studies investigated factors affecting developers’ adoption of program-analysis tools, showing that the way in which warnings are displayed is critical [33, 38]. Among the many factors identified to improve how warnings are displayed (e.g., warning should be well motivated [37]), our results may indicate that the position of the warning can play a role. Similarly to files, the position of the warning might affect, for instance, the attention that developers pay to them. Future studies can evaluate whether the effect of warnings’ position influences developers. This effect might be particularly significant, for instance, for those complex warnings that require developers’ effort to be understood and solved.

## 7 CONCLUSION

In this study, we investigated whether the relative position of a file has an impact on code review. To do so, we devised a two-step investigation to collect complementary evidence: We (1) collected and analyzed data from 219,476 Pull Requests (PRs) belonging to 138 open-source Java projects and (2) conducted an online controlled experiment with 106 participants.

In the first step of our investigation, we focused on reviewers’ activity, investigating the relationship between the relative position of file in a PR and the number of comments it receives. Having found evidence of a significant correlation between file position and review activity, we moved to the second step of our study to collect different evidence. We devised an online experiment where participants had to perform review a code change in which we two seeded defects (Corner Case and Missing Break). Participants were

randomly assigned to one of two possible treatments:  $CC_f-MB_l$  (where the file with the Corner Case defect was shown first, while the one with the Missing Break defect was shown last) and  $MB_f-CC_l$  (where the order of the files was reversed).

Our experiment found that the position of the file containing the Corner Case defect influences the likelihood of finding this defect. Specifically, 34 participants found the Corner Case defect when in the first file, while only 18 found it in the last. This effect was not significant for the less cognitive demanding Missing Break defect. Overall, our study provides evidence that the position in which files are presented during code review has an impact on code review’s outcome. This finding has implications for code review practices and review tool design, and may suggest that a similar effect may be present in other software engineering contexts.

## ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their thoughtful and important comments, which helped improving our paper. Fregnan and Bacchelli gratefully acknowledge the support of the Swiss National Science Foundation through the SNSF Projects No. PP00P2\_170529 and 200021\_197227. D’Ambros gratefully acknowledges the financial support of the Swiss National Science Foundation through the NRP-77 project 187353.

## REFERENCES

- [1] 2015. *Buggy Java Code: The Top 10 Most Common Mistakes That Java Developers Make*. Retrieved Feb 28, 2022 from <https://www.toptal.com/java/top-10-most-common-java-development-mistakes>
- [2] N. Abid, B. Sharif, N. Dragan, H. Alrasheed, and J. Maletic. 2019. Developer Reading Behavior While Summarizing Java Methods: Size and Context Matters. In *Proceedings of the International Conference on Software Engineering*. 384–395.
- [3] A. Frank Ackerman, Lynne S. Buchwald, and Frank H. Lewski. 1989. Software inspections: an effective verification process. *IEEE Software* 6, 3 (1989), 31–36.
- [4] N. Al Madi, C. Peterson, B. Sharif, and J. Maletic. 2021. From Novice to Expert: Analysis of Token Level Effects in a Longitudinal Eye Tracking Study. In *Proceedings of the International Conference on Program Comprehension*. 172–183.
- [5] Carol V Alexandru, José J Merchante, Sebastiano Panichella, Sebastian Proksch, Harald C Gall, and Gregorio Robles. 2018. On the usage of pythonic idioms. In *Proceedings of the 2018 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*. 1–11.
- [6] American Psychological Association. 2022. *APA Dictionary of Psychology*. Retrieved Feb 01, 2022 from <https://dictionary.apa.org/attention-decrement>
- [7] Alberto Bacchelli and Christian Bird. 2013. Expectations, outcomes, and challenges of modern code review. In *2013 35th International Conference on Software Engineering (ICSE)*. IEEE, 712–721. <https://doi.org/10.1109/ICSE.2013.6606617>
- [8] T. Baum. 2019. *Cognitive-support code review tools : improved efficiency of change-based code review by guiding and assisting reviewers*. Ph.D. Dissertation. Universität Hannover.
- [9] T. Baum, H. Leßmann, and K. Schneider. 2017. The choice of code review process: A survey on the state of the practice. In *Proceedings of the International Conference on Product-Focused Software Process Improvement*. Springer, 111–127.
- [10] T. Baum, O. Liskin, K. Niklas, and K. Schneider. 2016. A Faceted Classification Scheme for Change-Based Industrial Code Review Processes. In *Proceedings of the International Conference on Software Quality, Reliability and Security (QRS)*. 74–85.
- [11] T. Baum and K. Schneider. 2016. On the need for a new generation of code review tools. In *Proceedings of the International Conference on Product-Focused Software Process Improvement*. 301–308.
- [12] Tobias Baum, Kurt Schneider, and Alberto Bacchelli. 2017. On the optimal order of reading source code changes for review. In *2017 IEEE international conference on software maintenance and evolution (ICSME)*. IEEE, 329–340. <https://doi.org/10.1109/ICSME.2017.28>
- [13] Tobias Baum, Kurt Schneider, and Alberto Bacchelli. 2019. Associating working memory capacity and code change ordering with code review performance. *Empirical Software Engineering* 24, 4 (2019), 1762–1798. <https://doi.org/10.1007/s10664-018-9676-8>
- [14] G. Bergersen and J. Gustafsson. 2011. Programming skill, knowledge, and working memory among professional software developers from an investment theory

- perspective. *Journal of Individual Differences* (2011).
- [15] M. Bland. 2014. Finding more than one worm in the apple. *Commun. ACM* 57, 7 (2014), 58–64.
  - [16] K. Blincoe, J. Sheoran, S. Goggins, E. Petakovic, and D. Damian. 2016. Understanding the popular users: Following, affiliation influence and leadership on GitHub. *Information and Software Technology* 70 (2016), 30–39.
  - [17] H. Borges, A. Hora, and M. Valente. 2016. Understanding the factors that impact the popularity of GitHub repositories. In *Proceedings of the International Conference on Software Maintenance and Evolution*. 334–344.
  - [18] Larissa Braz, Enrico Fregnan, Gül Çalikli, and Alberto Bacchelli. 2021. Why Don't Developers Detect Improper Input Validation? ; DROP TABLE Papers ; --. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 499–511. <https://doi.org/10.1109/ICSE43902.2021.00054>
  - [19] T. Busjahn, R. Bednarik, A. Begel, M. Crosby, J. Paterson, C. Schulte, B. Sharif, and S. Tamm. 2015. Eye Movements in Code Reading: Relaxing the Linear Order. In *Proceedings of the International Conference on Program Comprehension*. 255–265.
  - [20] S. Chattopadhyay, N. Nelson, A. Au, N. Morales, C. Sanchez, R. Pandita, and A. Sarma. 2020. A Tale from the Trenches: Cognitive Biases and Software Development. In *Proceedings of the International Conference on Software Engineering*. 654–665.
  - [21] N. Cowan. 2010. The magical mystery four: How is working memory capacity limited, and why? *Current directions in psychological science* 19, 1 (2010), 51–57.
  - [22] M. Fagan. 1976. Design and code inspections to reduce errors in program development. *IBM Systems Journal* 15, 3 (1976), 182–211.
  - [23] F. Faul, E. Erdfelder, A. Lang, and A. Buchner. 2007. G\* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior research methods* 39, 2 (2007), 175–191.
  - [24] C. Feng. 2021. A comparison of zero-inflated and hurdle models for modeling zero-inflated count data. *Journal of Statistical Distributions and Applications* 8, 1 (2021), 1–19.
  - [25] Free Software Foundation. 2021. Comparing and Merging Files - Hunks. [https://www.gnu.org/software/diffutils/manual/html\\_node/Hunks.html](https://www.gnu.org/software/diffutils/manual/html_node/Hunks.html).
  - [26] Enrico Fregnan, Larissa Braz, Marco D'Ambros, Gül Çalikli, and Alberto Bacchelli. 2022. Artifacts Package - "First come first served: the impact of file position on code review". <https://zenodo.org/record/6901285>. <https://doi.org/10.5281/zenodo.6901285>
  - [27] Gerrit. 2021. Gerrit Code Review. <https://www.gerritcodereview.com>.
  - [28] GitHub. 2021. GitHub. <https://github.com>.
  - [29] S. Glen. 2016. *Phi Coefficient (Mean Square Contingency Coefficient)*. Retrieved Mar 11, 2022 from <https://www.statisticshowto.com/phi-coefficient-mean-square-contingency-coefficient>
  - [30] Pavlina Wurzel Gonçalves, Enrico Fregnan, Tobias Baum, Kurt Schneider, and Alberto Bacchelli. 2020. Do explicit review strategies improve code review performance?. In *Proceedings of the 17th international conference on mining software repositories*. 606–610. <https://doi.org/10.1145/3379597.3387509>
  - [31] Pavlina Wurzel Gonçalves, Enrico Fregnan, Tobias Baum, Kurt Schneider, and Alberto Bacchelli. 2022. Do explicit review strategies improve code review performance? Towards understanding the role of cognitive load. *Empirical Software Engineering* 27, 4 (2022), 1–46. <https://doi.org/10.1007/s10664-022-10123-8>
  - [32] C. Hendrick, A. F. Costantini, J. McGarry, and K. McBride. 1973. Attention decrement, temporal variation, and the primacy effect in impression formation. *Memory & cognition* 1, 2 (1973), 193–195.
  - [33] A. Henley, K. Muçlu, M. Christakis, S. Fleming, and C. Bird. 2018. Cfar: A tool to increase communication, productivity, and review quality in collaborative code reviews. In *Proceedings of the Conference on Human Factors in Computing Systems*. 1–13.
  - [34] Y. Huang, N. Jia, X. Chen, K. Hong, and Z. Zheng. 2018. Salient-class location: Help developers understand code change in code review. In *Proceedings of the Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 770–774.
  - [35] Y. Huang, K. Leach, Z. Sharafi, N. McKay, T. Santander, and W. Weimer. 2020. *Biases and Differences in Code Review Using Medical Imaging and Eye-Tracking: Genders, Humans, and Machines*. Association for Computing Machinery, 456–468.
  - [36] B. Jeng and E. Weyuker. 1994. A simplified domain-testing strategy. *Transactions on Software Engineering and Methodology* 3, 3 (1994), 254–270.
  - [37] B. Johnson, R. Pandita, J. Smith, D. Ford, S. Elder, E. Murphy-Hill, S. Heckman, and C. Sadowski. 2016. A cross-tool communication study on program analysis tool notifications. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 73–84.
  - [38] B. Johnson, Y. Song, E. Murphy-Hill, and R. Bowdidge. 2013. Why don't software developers use static analysis tools to find bugs?. In *Proceedings of the International Conference on Software Engineering*. 672–681.
  - [39] L. MacLeod, M. Greiler, M. Storey, C. Bird, and J. Czerwonka. 2017. Code reviewing in the trenches: Challenges and best practices. *IEEE Software* 35, 4 (2017), 34–42.
  - [40] L. Mannila, M. Peltomäki, and T. Salakoski. 2006. What about a simple language? Analyzing the difficulties in learning to program. *Computer science education* 16, 3 (2006), 211–227.
  - [41] R. Mohanani, I. Salman, B. Turhan, P. Rodríguez, and P. Ralph. 2018. Cognitive biases in software engineering: a systematic mapping study. *Transactions on Software Engineering* 46, 12 (2018), 1318–1339.
  - [42] E. Neumayer and T. Plümper. 2017. *Robustness tests for quantitative research*. Cambridge University Press.
  - [43] G. Petrović and M. Ivanković. 2018. State of mutation testing at google. In *Proceedings of the international conference on software engineering: Software engineering in practice*. 163–171.
  - [44] P. Phan-Udom, N. Wattanakul, T. Sakulniwat, C. Ragkhitwetsagul, T. Sunetnanta, M. Choetkiertikul, and R. Kula. 2020. Teddy: automatic recommendation of pythonic idiom usage for pull-based software projects. In *Proceedings of the International Conference on Software Maintenance and Evolution*. 806–809.
  - [45] P. Rigny and C. Bird. 2013. Convergent Contemporary Software Peer Review Practices. In *Proceedings of the Joint Meeting on Foundations of Software Engineering*. 202–212.
  - [46] P. Rodeghero, C. Liu, P. McBurney, and C. McMillan. 2015. An Eye-Tracking Study of Java Programmers and Application to Source Code Summarization. *Transactions on Software Engineering* 41, 11 (2015), 1038–1054.
  - [47] Caitlin Sadowski, Emma Söderberg, Luke Church, Michal Sipko, and Alberto Bacchelli. 2018. Modern code review: A case study at Google. In *Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Practice*. 181–190. <https://doi.org/10.1145/3183519.3183525>
  - [48] Davide Spadini. 2020. CRExperiment. <https://github.com/ishpard/CRExperiment>.
  - [49] Davide Spadini, Mauricio Aniche, Margaret-Anne Storey, Magiel Bruntink, and Alberto Bacchelli. 2018. When testing meets code review: Why and how developers review tests. In *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)*. IEEE, 677–687. <https://doi.org/10.1145/3180155.3180192>
  - [50] Davide Spadini, Gül Çalikli, and Alberto Bacchelli. 2020. Primers or reminders? The effects of existing review comments on code review. In *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*. IEEE, 1171–1182. <https://doi.org/10.1145/3377811.3380385>
  - [51] Davide Spadini, Fabio Palomba, Tobias Baum, Stefan Hanenberg, Magiel Bruntink, and Alberto Bacchelli. 2019. Test-driven code review: an empirical study. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 1061–1072. <https://doi.org/10.1109/ICSE.2019.00110>
  - [52] StackOverflow. 2014. *What does Pythonic mean?* Retrieved Jul 22, 2022 from <https://stackoverflow.com/questions/25011078/what-does-pythonic-mean>
  - [53] Laerd Statistics. 2018. *Spearman's Rank-Order Correlation*. Retrieved Mar 17, 2022 from <https://statistics.laerd.com/statistical-guides/spearman-rank-order-correlation-statistical-guide.php>
  - [54] P. Thongtanunam and A. Hassan. 2021. Review Dynamics and Their Impact on Software Quality. *Transactions on Software Engineering* 47, 12 (2021), 2698–2712. <https://doi.org/10.1109/TSE.2020.2964660>
  - [55] T. Thurium. 2015. *Beep Boop: 6 Bots To Better Your Open Source Project*. Retrieved Mar 8, 2022 from <https://www.twilio.com/blog/6-bots-better-open-source-project>
  - [56] TIOBE. 2022. *TIOBE-Index*. Retrieved Feb 15, 2022 from <https://www.tiobe.com/tiobe-index/>
  - [57] R. Turner, M. Falcone, B. Sharif, and A. Lazar. 2014. An eye-tracking study assessing the comprehension of C++ and Python source code. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. 231–234.
  - [58] Grand Canyon University. 2019. *5 Common Java Coding Mistakes to Avoid*. Retrieved Mar 15, 2022 from <https://www.gcu.edu/blog/engineering-technology/5-common-java-coding-mistakes-avoid>
  - [59] M. Wessel. 2021. *Awesome SE Bots*. Retrieved Mar 8, 2022 from <https://github.com/mairieli/awesome-se-bots>
  - [60] O. Wilhelm, A. Hildebrandt, and K. Oberauer. 2013. What is working memory capacity, and how can we measure it? *Frontiers in psychology* 4 (2013), 433.