

Negativity in Self-Admitted Technical Debt: How Sentiment Influences Prioritization

Nathan Cassee · Neil Ernst · Nicole
Novielli · Alexander Serebrenik

Received: date / Accepted: date

Abstract Self-Admitted Technical Debt, or SATD, is a self-admission of technical debt present in a software system. The presence of SATD in software systems negatively affects developers, therefore, managing and addressing SATD is crucial for software engineering. To effectively manage SATD, developers need to estimate its priority and assess the effort required to fix the described technical debt. About a quarter of descriptions of SATD in software systems express some form of negativity or negative emotions when describing technical debt. In this paper, we report on an experiment conducted with 59 respondents to study whether negativity expressed in the description of SATD **actually** affects the prioritization of SATD. The respondents are a mix of professional developers and students, and in the experiment, we asked participants to prioritize four vignettes: two expressing negativity and two expressing neutral sentiment. To ensure the vignettes were realistic, they were based on existing SATD extracted from a dataset. We find that negativity causes between one-third and half of developers to prioritize SATD in which negativity is expressed as having more priority. Developers affected by negativity when prioritizing SATD are twice as likely to increase their estimation of urgency and 1.5 times as likely to increase their estimation of importance and effort for SATD com-

Nathan Cassee
Eindhoven University of Technology, The Netherlands
E-mail: n.w.cassee@tue.nl

Neil Ernst
University of Victoria, Canada
E-mail: nernst@uvic.ca

Nicole Novielli
University of Bari, Italy
E-mail: nicole.novielli@uniba.it

Alexander Serebrenik
Eindhoven University of Technology, The Netherlands
E-mail: a.serebrenik@tue.nl

pared to the likelihood of decreasing these prioritization scores. Our findings show how developers actively use negativity in SATD to determine how urgently a particular instance of technical-debt should be addressed. However, our study also describes a gap in the actions and belief of developers. Even if 33% to 50% use negativity to prioritize SATD, 67% of developers believe that using negativity as a proxy for priority is unacceptable. Therefore, we would not recommend using negativity as a proxy for priority. However, we also recognize it might be unavoidable that negativity is expressed by developers to describe technical debt.

Keywords Self-Admitted Technical Debt · Software Engineering · Sentiment

1 Introduction

Technical Debt is used as a metaphor by developers to describe suboptimal implementations that require future reimplementations to fix the existing implementation (Cunningham, 1992; Tom et al., 2013). Technical debt is pervasive as a large number of developers is familiar, and even affected by, it (Lim et al., 2012). The presence of technical debt in a system is known to have negative effects: it not only makes it more difficult to modify a software project (Wehaibi et al., 2016), but also reduces the morale of developers working on a system where technical debt is present (Besker et al., 2020).

Self-Admitted Technical Debt, or SATD, is a specific category of Technical Debt. SATD is characterized by explicit admissions of developers indicating the presence of technical debt (Potdar and Shihab, 2014). Both technical debt and SATD have been extensively studied, including aspects of SATD such as the automatic identification (Guo et al., 2021; Maldonado et al., 2017) or the management and removal of SATD (Maldonado et al., 2017; Tan et al., 2021; Zampetti et al., 2018, 2020). Specifically, we know that developers use SATD to describe a wide range of technical issues (Cassee et al., 2022; Maldonado and Shihab, 2015). In particular, existing taxonomies of SATD instances show that developers tend to describe functional issues or poor implementation choices.

Sentiment is a construct used to categorize text into one of the three categories (Novielli and Serebrenik, 2023). Negative sentiment includes texts expressing negative emotions, while text has positive sentiment expressing positive emotions, and a text is considered neutral if it expresses no emotions. The automatic classification of sentiment in software engineering texts has been used to study many different aspects of software engineering (Lin et al., 2022), such as in the code-review process (Singh and Singh, 2018), or on Stack Overflow (Calefato et al., 2018; Swillus and Zaidman, 2023). Olsson et al. (2021) found that the presence of technical debt, or design smells, can cause developers to experience negative emotions. Furthermore, it has also been found that in roughly 20% of SATD instances, negativity, or negative emotions, are used to describe the SATD (Cassee et al., 2022). From the psychological literature, we know that emotionally salient information is more likely to capture attention in the working memory of the brain (Okon-Singer

et al., 2015). Similarly, when it comes to issue resolution, previous work has found that there appears to be a link between the expression of positive sentiment and a faster resolution of issues (Ortu et al., 2015; Sanei et al., 2021). While Calefato et al. (2018) concluded that Stack Overflow questions that are neutral, or in which no sentiment is expressed, are more likely to receive an answer. Focusing on SATD, a quarter of the developers surveyed by Cassee et al. (2022) stated that they use negative emotions to describe high-priority SATD, and that they interpret negativity expressed in SATD as a proxy for priority. Because there is often a gap between respondents’ beliefs, and their actions (Barr, 2006), we want to understand whether expressions of negativity in SATD influence prioritization. Therefore, we pose the following research question:

RQ₁: Do developers interpret technical debt annotated with negative source-code comments as having a higher priority?

While there are many different ways to prioritize technical debt (Lenarduzzi et al., 2021) it is currently unknown whether expressions of negativity influence the perception of the priority of technical debt. Because of how challenging effort estimation is (Molokken and Jorgensen, 2003), it is important to understand how negativity influences the prioritization of technical debt. If negativity influences prioritization, it might lead to unintended consequences, as technical debt might be prioritized not because it is important or urgent but because negativity has been used to describe it.

We use a vignette-based experimental design to address RQ_1 (Aguinis and Bradley, 2014). By purposefully selecting a realistic set of SATD instances, creating variations of these SATD instances in which negativity is expressed, and assigning them to participants we study the effect of negativity on prioritization. We sampled respondents from open-source software mailing lists and industrial contacts in the Eindhoven region of the Netherlands.

Based on the responses of 59 participants, we conclude that between 30% to 50% of developers score the priority of technical debt as higher if negativity is expressed in its description. Most importantly, even developers who self-report that they are not influenced by negativity are more likely to increase their estimation of the effort required to fix SATD if negativity is expressed. Our results show that developers use SATD not just to describe technical issues but also use negativity in descriptions of SATD as an additional dimension to communicate priority.

2 Methodology

In this section, we describe the experiment conducted to understand whether negativity influences the prioritization of SATD. First, we justify our choice to use controlled experiments. Then, we describe the design of the experiment,

the instruments used in the experiment, and we explain the method used to analyze the data.¹

2.1 Choice of Research Method

In this section, we discuss several potential research methods that could be used to understand the impact of negativity on priority using the framework of Robillard et al. (2024), and we justify our choice to use a controlled experiment. We discuss the potential alternatives to a controlled experiment, our considerations, rationale, and the implications of our decision to use an experiment.

Alternatives We first briefly describe the three alternatives we considered: A controlled experiment, a Mining Software Repository (MSR) study, and an interview. In a controlled experiment, we would ask participants to prioritize different SATD instances, while for an MSR study, we would analyze whether SATD in which negativity is expressed is removed quicker. Finally, in an interview, we would ask participants about their past prioritization practices.

Considerations: The most important consideration informing our choice of methodology is the expectation that any effect of negativity on the perception of priority might be relatively small. Previous correlational studies on the impact of sentiment on software engineering report small effect sizes (Calefato et al., 2018; Cheruvilil and da Silva, 2019; Olsson et al., 2021). Because of these small effect sizes, we require a research method that gives us a high level of control, allowing us to isolate and measure the effect of negativity. Furthermore, many factors influence prioritization, as prioritization of Technical Debt is a complex process, and many factors influence how it is prioritized (Lenarduzzi et al., 2021). This will also limit the effect of a single factor, such as the expression of negativity, on priority. Because of the relatively small effect size, estimating the true effect of negativity on prioritization is more difficult: Many confounding variables might influence the removal of SATD, and in fact, even classifying whether the removal of SATD was purposeful is already challenging (Zampetti et al., 2018). The presence of confounding variables and the noise introduced by potentially inaccurate classifications of removal has made us decide to not use an MSR study.

The second consideration is the observation that human recollection is not optimal which results in humans misremembering, especially when asked about past events (Johnson et al., 1993; Raykov et al., 2023). For an interview study or a survey, in which we ask participants about choices they made in the past related to the prioritization of SATD this might be problematic. Similarly, while we could ask about current prioritization practices, this still has the downside that our questions would be hypothetical.

¹ The code of the statistical modeling is available in a replication package at <https://doi.org/10.6084/m9.figshare.26863435>.

Implications: Because of these two considerations, we opt for controlled experiments as research method to study the prioritization of SATD. The first implication of our choice for controlled experiments is that we have a high level of control (Stol and Fitzgerald, 2018; Storey et al., 2020). We can use this control to account for as many relevant confounding variables that could also influence prioritization as possible. However, the second implication of our choice is reduced realism, as with any experimental study (Stol and Fitzgerald, 2018). Asking respondents to prioritize SATD in an experimental context is not very realistic. The context in which the respondents usually prioritize SATD is likely not the experimental setup. Therefore, the effect we observe in the experiments might not be the effect observed in the field, however, using an experimental set-up can be expected to allow us to understand with certainty whether any effect exists.

2.2 Experimental Design:

For RQ_1 we use a between-person experimental vignette design, where each participant is only exposed to one-condition, according to the guidelines of Aguinis and Bradley (2014). To maximize realism of the experiment the vignettes we show to participants are instances of SATD, each containing a snippet of source code containing technical debt and a source code comment describing the technical debt. We ask participants to score the priority of the vignettes, and by experimentally varying whether negativity is expressed in the source code comment describing the SATD, to investigate whether negativity influences prioritization.

Operationalization of Priority: The concept of priority is quite a broad topic; different respondents might interpret the meaning of priority differently. Therefore, following existing guidelines (Aguinis and Bradley, 2014), we “split” the concept of priority into three constructs: *Urgency*, *Importance*, and *Effort*. Both urgency and importance are common constructs used to operationalize priority (Bellotti et al., 2004; Gavidia-Calderon et al., 2021; Middleton et al., 2018), and effort is used to determine the cost of Technical Debt repayment (Lenarduzzi et al., 2021).

Participant recruitment: The target population for the experiment are software engineers, and to increase the response rate, we do not require any minimum working experience. Because recruiting of participants for software engineering studies is challenging, we use different channels to recruit participants (Danilova et al., 2021). In particular, we posted the invitation to participate in the experiment on a set of mailing lists previously used to recruit software engineers (Cassee et al., 2022; Maldonado et al., 2017) and on the social-media pages of the authors, as well as used convenience sampling to invite developers at medium to large software companies in the Eindhoven region of the Netherlands. To ensure that we did not burden the mailing lists, we sent out invitations piecemeal, a few a day, and we only posted the call to participate on active mailing lists. The Ethical Review Board of Eindhoven

University of Technology approved both the experiment and the recruitment strategy.²

Instrument Design: Table 1 contains an overview of the questions as included in the survey. The instrument is divided into three sections: The first section lists the questions we ask for each vignette. To ensure the constructs *Effort*, *Urgency*, and *Importance* were interpreted equally by all participants, we provided the italicized definitions included in Table 1. The second section contains a set of questions on the perception of participants about the usage of negativity as a proxy for priority, and the third section contains questions on participants’ demographics. The questions on perceptions and demographics were placed at the end of the survey to prevent them from biasing participants (Steele and Aronson, 1995).

The demographic questions included in the survey are related to the respondents’ age and working experience. We record working experience because open and closed source developers are known to manage SATD differently (Zampetti et al., 2021). Experience can be defined in many different ways (Siegmund et al., 2014), and for this experiment, we choose to re-use questions about experience from the Stack Overflow developer survey.³ We also ask respondents to indicate their age as an open input question, following the recommendations from Hughes et al. (2016). We record the age because age tends to influence how people experience emotions (Yeung et al., 2011). Finally, we also expect developer’s attitude towards the practice of using negativity as a proxy for priority to influence prioritization. Therefore, we re-use the closed questions from the study of Cassee et al. (2022), in which developers are asked about their perceptions and beliefs about the usage of SATD as a proxy for priority. The questions for each of the four vignettes were mandatory, while all other questions were optional.

Vignette selection: For the experiment, the vignettes should be as realistic as possible (Aguinis and Bradley, 2014). Therefore, as vignettes, we select SATD instances from an existing dataset of SATD items. The dataset was gathered by Maldonado et al. (2017), and we use the version of Cassee et al. (2022) in which the SATD instances have been categorized based on the type of problem described in the SATD. We select SATD instances from a single category to minimize the risk that differences between SATD instances influence the results. The category we select is *Poor Implementation Choices*, the most common SATD category, in which about 30% of the SATD instances express negativity. SATD in this category includes comments like ‘`// TODO: define constants for magic numbers`’.

We select four SATD instances from the dataset; we do not select any more to reduce respondent fatigue and the odds of disengaging. Because the comments in the dataset of Cassee *et al.* have already been labeled with sentiment polarity, we select two comments that have been labeled as negative and two comments that have been labeled as non-negative. The dataset of Cassee *et*

² Reference: ERB2023MCS17

³ <https://survey.stackoverflow.co/2022#overview>

Table 1: The questions and response options per question as included in the experiment.

Question	Response Type
Q1 How would you rate the <i>Urgency</i> of the listed code-snippet? <i>In this context we define urgency as whether swift action is required to address the technical debt item.</i>	<i>One of: Very low, Low, Medium, High, Very high</i>
Q2 How would you rate the <i>Importance</i> of the listed code-snippet? <i>In this context we define importance as the impact of the technical debt item.</i>	<i>One of: Very low, Low, Medium, High, Very high</i>
Q3 How would you rate the <i>Effort</i> required to address the listed code-snippet? <i>In this context we define effort as the amount of work required to address the technical debt item.</i>	<i>One of: Very low, Low, Medium, High, Very high</i>
Perception (Re-used from Cassee et al. (2022))	Response type
Q4 When writing source code, how often do you write source code comments indicating delayed or intended work activities such as TODO, FIXME, hack, workaround, etc.?	<i>Never, Rarely (Less than once a month), Sometimes (Monthly), Often (Weekly), Very often (Daily)</i>
Q5 When authoring comments that describe a problem, how often do you write negative source-code comments indicating delayed or intended work activities such as TODO, FIXME, hack, workaround, etc.?	<i>Never, Rarely (Less than once a month), Sometimes (Monthly), Often (Weekly), Very often (Daily)</i>
Q6 How often do you come across negative source-code comments indicating delayed or intended work activities such as TODO, FIXME, hack, workaround, etc.?	<i>Never, Rarely (Less than once a month), Sometimes (Monthly), Often (Weekly), Very often (Daily)</i>
Q7 While writing a comment describing an issue in the source-code, I am more likely to write negative comments for issues that I believe are more important.	<i>Strongly disagree, Disagree, Neutral, Agree, Strongly agree</i>
Q8 Writing negative comments to assign extra priority to issues in the source-code is an acceptable practice.	<i>Strongly disagree, Disagree, Neutral, Agree, Strongly agree</i>
Q9 Whenever I come across a source-code comment describing a problem that is particularly negative, I interpret this as a more important issue than a source-code comment describing a problem that is more neutral.	<i>Strongly disagree, Disagree, Neutral, Agree, Strongly agree</i>
Demographics	Response type
Q10 What is your age?	<i>Open numerical input</i>
Q11 Which of the following best describes your current employment status?	<i>One of: "Employed", "Independent contractor, freelancer or self-employed", "Student", "Not employed", "Prefer not to say", "Retired"</i>
Q12 Which of the following best describes the code you write outside of work? Select all that apply.	<i>One or more of: "Contribute to open-source software", "Hobby", "Freelance/contract work", "School or academic work", "Bootstrapping a business", "I do not write code outside of work"</i>
Q13 How many years of programming experience do you have?	<i>Open numerical input</i>

al. identifies three different sentiment classes in SATD: negative, non-negative, and mixed. The mixed class, however, occurs in less than 2% of cases. Consequently, we exclude the mixed class and only sample from the negative and non-negative classes. Additionally, due to the negative connotations of technical debt, we expect a low number of positive instance (Cassee et al., 2022). Therefore, we exclude these, focusing solely on negative versus non-negative instances.

The selection of SATD instances is performed manually: we pay careful attention to ensure that we select instances that are comprehensible and self-contained, such that the respondents can understand them without becoming fatigued or confused. *Alternative generation:* For the between-person design of the experiment, we require that each of the four selected vignettes has two variations: a neutral instance and a negative instance. One of the two variations can be randomly assigned to a participant. For each of the selected vignettes, we create a manipulated variation expressing a different sentiment polarity than the original. The two crucial requirements for these manipulated instances are: The semantics of the original comment should be preserved, *i.e.*, the alternative comment should describe the same problem as the original comment, and the manipulated comment should express the required sentiment polarity.

To generate manipulated comments we used ChatGPT.⁴ Through ChatGPT, we aim to reduce the risk that our own perception of what negativity in SATD looks like influences manipulated alternatives we would draft ourselves. We manually validate the comment generated by ChatGPT to ensure the manipulated comment meets the previously listed requirements. For each SATD instance, we prompt ChatGPT to generate three alternative comments that express a sentiment opposite to the sentiment expressed in the instance. We iteratively refined the prompt used to generate the alternative comments and we evaluated the alternatives generated by ChatGPT until all authors were satisfied that each alternative was sufficiently realistic. This process took several iterations. Listed below is the prompt used for the sentiment transfer from negative to neutral:

“Take the following source-code comment, and change the language of the source-code comment to ensure that the resulting source-code comment contains neutral sentiment. Generate three alternative, neutral, source-code comments, but make sure to preserve the original meaning of the comment as much as possible.”

After finalizing the prompt, three of this paper’s authors independently voted to select the best-manipulated variation of the SATD instance from the three alternatives. Criteria for voting adhere to the requirements listed above: Did the sentiment transfer work (“Is the comment actually neutral?”), and whether the manipulated comment describes the same problem as the original comment. We selected the best alternative comments for each of the four

⁴ Version 3.5, accessed in November 2023 at <https://chat.openai.com>

selected SATD instances based on the votes. For three manipulated instances, all authors voted for the same alternative; for the fourth instance, two of the three voters preferred one.

Table 2: The vignettes as they were used in the experiment. The source code is matched to either the *Negative* or *Neutral* comment and shown to the participant. The *(M)* indicates that the comment was generated using ChatGPT.

Type	Item
Vignette #1	public static boolean useThetaStyleImplicitJoins; public static boolean regressionStyleJoinSuppression;
<i>Negative</i>	<i>// USED ONLY FOR REGRESSION TESTING!!!! todo : ↪ obviously get rid of all this junk</i>
<i>Neutral (M)</i>	<i>// Used only for regression testing! todo: clearly remove all this ↪ unnecessary code</i>
Vignette #2	ConstDeclNode constDeclNode = (ConstDeclNode) node; Node constNode = constDeclNode.getConstNode();
<i>Neutral</i>	<i>// TODO: callback for value would be more efficient, but unlikely ↪ to be a big cost (constants are rarely assigned)</i>
<i>Negative (M)</i>	<i>// TODO: This is frustrating! A callback for value would be ↪ more efficient, but unlikely to be a big cost (constants are ↪ rarely assigned).</i>
Vignette #3	public class ConstDeclNode extends AssignableNode ↪ implements INameNode { private final String name; private final INameNode constNode; ... }
<i>Neutral</i>	<i>// FIXME: ConstDecl could be two seperate classes (or done ↪ differently since constNode and name never exist at the ↪ same time).</i>
<i>Negative (M)</i>	<i>// FIXME: Ugh, ConstDecl is a mess. It should have been ↪ divided into distinct classes (or approached differently) ↪ because constNode and name are never in sync.</i>

Continued on next page

Table 2: The vignettes as they were used in the experiment. The source code is matched to either the *Negative* or *Neutral* comment and shown to the participant. The *(M)* indicates that the comment was generated using ChatGPT. (Continued)

Type	Item
Vignette #4	<pre> if (plot instanceof PiePlot) { applyToPiePlot((PiePlot) plot); } else if (plot instanceof MultiplePiePlot) { applyToMultiplePiePlot((MultiplePiePlot) plot); } else if (plot instanceof CategoryPlot) { applyToCategoryPlot((CategoryPlot) plot); } </pre>
<i>Negative</i>	<pre> // now handle specific plot types (and yes, I know this is some ↪ really ugly code that has to be manually updated any time a ↪ new plot type is added – I should have written something ↪ much cooler, but I didn't and neither did anyone else). </pre>
<i>Neutral (M)</i>	<pre> // Now address specific plot types (and yes, I am aware that this ↪ code needs manual updates whenever a new plot type is ↪ added – a more advanced implementation could have been ↪ developed, but it wasn't, and no other approach was ↪ proposed). </pre>

Table 2 contains an overview of the selected vignettes. Each vignette combines a short source-code snippet with an accompanying source-code comment that “admits” technical debt in the snippet. The Type column lists the sentiment polarity expressed by the comment, and the *(M)* tag denotes that the comment on that line has been generated using ChatGPT.

Vignette Assignment: Respondents are assigned to one of the two variations for each vignette. However, the assignment of variation to respondents is not fully random. When assigning variations, we consider the alternative comments generated by ChatGPT as a blocking factor (Juristo and Moreno, 2001) because the comments generated by ChatGPT might influence prioritization scores. We refer to this blocking variable as *Manipulation*.

We defined two experimental flows to limit the effect of the manipulation on the outcome. Table 3 shows these two flows. Observe that we ensure that each respondent sees two neutral and two negative vignettes and sees two manipulated and two original vignettes. Participants are randomly assigned to one of the two flows, and within a flow to order of the four vignettes is randomized, such that the impact of the blocking factors is limited (Atzmüller and Steiner, 2010). To ensure we do not bias participants, we do not reveal to participants the polarity of the vignette.

Table 3: The two different flows used for the vignette section of the survey. *Sentiment* is the sentiment expressed in the vignette, while *Manipulated* denotes whether the respondent in the flow sees the original comment or the manipulated one.

		V_1	V_2	V_3	V_4
<i>Flow 1</i>	<i>Sentiment</i>	<i>negative</i>	<i>neutral</i>	<i>negative</i>	<i>neutral</i>
	<i>Manipulated</i>	×	×	✓	✓
<i>Flow 2</i>	<i>Sentiment</i>	<i>neutral</i>	<i>negative</i>	<i>neutral</i>	<i>negative</i>
	<i>Manipulated</i>	✓	✓	×	×

Finally, to validate that the vignettes shown in the experiment were appropriate, and the questions asked were understandable, we piloted the survey with a set of four active software engineers. Based on the feedback of the pilot we made minor text changes to the wording of some questions, and clarified a few concepts a bit better. The experiment was hosted using Qualtrics.⁵

2.3 Data Analysis

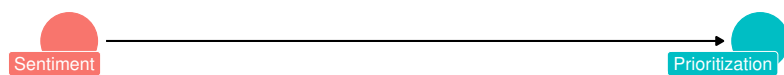


Fig. 1: A Directed Acyclic Graph (DAG) illustrating the relation we are analyzing. The color blue denotes the outcome and red the exposure.

Figure 1 visualizes the relation we study in this manuscript as a Directed Acyclic Graph (DAG). In a DAG, nodes are used to represent variables, and arrows between nodes represent causal relations between variables (Elwert, 2013). For instance, Figure 1 should be interpreted as “A change in sentiment leads to a change in prioritization”.

The analysis would be relatively straightforward, assuming that Figure 1 is the correct model capturing all relevant effects. However, while we attempted to control for as many confounding variables in the experimental design, there were confounding factors we could not control for. For instance, manipulating the vignettes to transfer sentiment could also have affected the prioritization. Therefore, we create a more complete DAG that contains all other variables in this experimental context that might influence the prioritization of Technical Debt.

⁵ <https://www.qualtrics.com/>

Fig. 2 shows the complete model, including all variables that we could not control for in the design of the experiment. We first explain the construct that each of the variables describes, and then we explain the relations between the nodes.

- **Importance, Urgency and Effort:** This is the node prioritization in Figure 1 split into one node for each of the Likert-scale questions asked in the experiment.
- **Sentiment:** Sentiment is the binary exposure variable, recording the sentiment of the vignette shown to the participant.
- **Technical Debt:** We selected four existing instances of technical debt to use as the basis for the vignettes. While these four instances all belong to the same category of technical debt, there are still variations in the technical debt in each of the vignettes. It is a categorical variable, recording which of the four vignettes the participant responded to.
- **Manipulated:** Each vignette in the experiment has two variations: The unmodified, existing instance of technical debt and one variation for which the alternative comment was generated using ChatGPT. Manipulated is a binary variable indicating whether a participant was looking at a variation generated by ChatGPT.
- **Perception:** From Cassee et al. (2022) we know that developers have different perceptions concerning the usage of negativity as a proxy for priority. Whether a participant agrees, has no opinion, or disagrees with using negativity as a proxy for priority is indicated by this variable as a trinary variable.

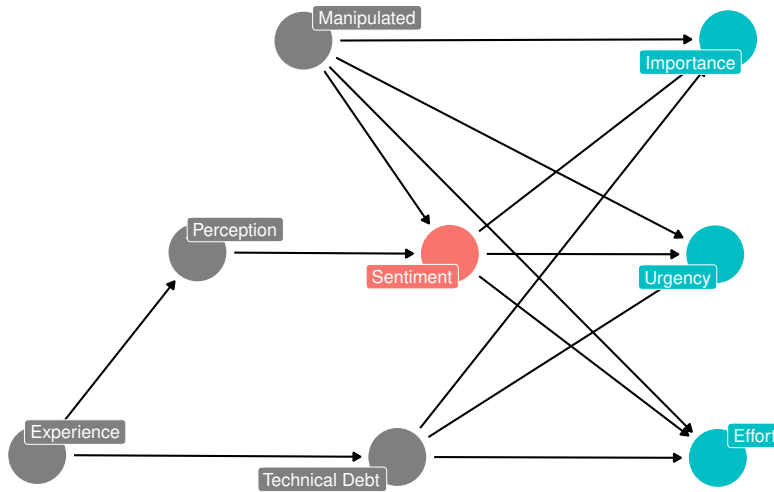


Fig. 2: The theoretical model visualized as a DAG. The color blue denotes the outcome, red is the exposure, and grey denotes confounders.

- **Experience:** Experience captures the participant’s years of professional working experience.

We expect that the prioritization, measured as *Importance*, *Urgency* and *Effort*, is directly influenced by the *Sentiment* in the vignette. Because of the results of the survey of Cassee et al. (2022). Secondly, the *Technical Debt* shown to participants is different in each vignette; therefore, we expect that each item of technical debt will receive a different prioritization score. Finally, we expect whether we *Manipulated* the comment to influence both prioritization and sentiment. While we did our best to ensure the comments generated using ChatGPT were as realistic as possible, the manipulated comments might still contain language that influences the prioritization given by participants. Additionally, comments generated by ChatGPT could also express negativity differently than human-authored comments, and therefore, whether a comment is *Manipulated* could affect how *Sentiment* influences prioritization.

Whether a participant believes negativity should be used to prioritize technical debt (*Perception*) will influence whether a participant uses *Sentiment* to prioritize. We expect the *Experience* of participants to influence both how they prioritize *Technical Debt*, and their *Perception* to using negativity as a proxy for priority (Yeung et al., 2011; Zampetti et al., 2021). *i.e.*, more experienced participants will have seen more technical debt in their career and therefore prioritize specific instances of technical debt. Similarly, for *Perception*, a more experienced candidate could be more likely to think using negativity is an acceptable prioritization mechanism.

Given the DAG shown in Figure 2, understanding the effect of negativity on prioritization (measured as *Effort*, *Urgency*, and *Importance*) requires building models that adjust for the other variables in the DAG. Specifically, understanding the total effect of *Sentiment* on *Prioritization* requires adjusting for *Manipulated*, *Perception* and *Experience*. To accurately estimate the effect of sentiment on prioritization, we use Bayesian statistics (Gelman et al., 2013). In frequentist statistics, model outcomes are often single-point estimates of effects and significance values. One benefit of Bayesian statistics is the fact that model output is not just a single-point estimate. Instead, because we use Bayesian statistics, we can better account for any uncertainty in the data (Furia et al., 2022; Kruschke and Liddell, 2018; McElreath, 2018). By following existing guidelines (Furia et al., 2022; Gelman et al., 2020) and the examples of the applications of Bayesian statistics to software engineering data (Ghorbani et al., 2023; Torkar et al., 2022) we ensure we can estimate the effect of negativity on prioritization.

In the remainder of this section, we first describe the Bayesian models we fit to the data. Then, we explain how we use the models to understand the effect of negative sentiment on the prioritization of SATD.

$$\begin{aligned}
Outcome &\sim \text{Ordered-logit}(\phi, \kappa) \\
\phi &= \alpha_{\text{effects}[Sentiment, Perception]} + \alpha_{\text{manipulated}[Manipulated]} \\
&\quad + \alpha_{\text{experience}} * Experience \\
\alpha_{\text{effects}} &\sim \text{Normal}(0, 0.5) \\
\alpha_{\text{manipulated}} &\sim \text{Normal}(0, 0.5) \\
\alpha_{\text{experience}} &\sim \text{Normal}(0, 0.5) \\
\kappa &\sim \text{Normal}(0, 2) \\
\text{Where } Sentiment, Manipulated &\in \{\text{True}, \text{False}\} \\
Perception &\in \{\text{Disagree}, \text{No Opinion}, \text{Agree}\} \\
Experience &\in \mathbb{Q}
\end{aligned}
\tag{Model 1}$$

Model: Based on the causal graph of Figure 2, we should adjust for three variables: *Perception*, *Manipulated*, and *Experience*. Model 1 shows the model we use, where *Outcome* is the score given by a participant to the vignette. Because we have three outcome constructs: *Effort*, *Importance*, *Urgency*, we fit one model per construct. The first line shows the choice of the *likelihood function*, an Ordered-logit. lines 2–3 show the choice of model parameters, indicated by the prefix α , lines 4–8 show the choice of priors, and lines 9–11 show the allowed values for the input variables.

Because the outcomes are responses on a Likert scale. *i.e.*, the responses are on an ordinal scale, we use an Ordered Logit as a likelihood function as recommended by (McElreath, 2018). Ordered-logit has two arguments: ϕ , a model-term consisting of all of the parameters (denoted by α) and κ , the cut-points, used to model the log-cumulative-odds for the Likert responses. The model itself has four input variables (typeset in red and italics): Whether the vignette shown to the developer was negative (*Sentiment*), whether the vignette was manipulated (*Manipulated*), what the perception of the respondent was (*Perception*), and the experience of the respondent (*Experience*). Negative and manipulated are binary variables, perception is a trinary variable based on the respondent’s answer to Q9, and age is a natural number. *effects* is a 2 by 3 matrix with a model parameter for each combination of *Sentiment* and *Perception*. *Manipulated* is a vector of length 2, for each value of *Manipulated*, and experience, *Experience*, is a single parameter.

The priors we select are all uninformative, and we do not associate any positive or negative effects to them. We opted for uninformative priors as we did not want to bias the models. Selecting priors that expect sentiment to increase prioritization scores could potentially bias our results, and there is no literature that allows us to pick more informative priors.

Estimating Effects: We use the three models, fit for each of the three outcome variables, to quantify the impact of negativity on prioritization. First, we plot the distribution of the model parameters related to sentiment, the

so-called posterior, as density plots. The posterior plots show the effect on the outcome variable that the model associates with a change from neutral to negative sentiment and are commonly used to interpret results of Bayesian models (McElreath, 2018). We also use the posterior to compute the *Evidence Ratio* of the effect of negativity and interpret the evidence ratio according to Stefan et al. (2019). Where an evidence ratio larger than one is considered as evidence in favor of a correlation, and an evidence ratio smaller than one is considered as evidence for an inverse correlation.

Through the posteriors we can understand whether negativity influences prioritization, however, it does not allow us to quantify how often negativity leads to a different prioritization score. Therefore, we also use the fitted models to simulate the effect of negativity *i.e.*, “*What are the odds that changing sentiment from neutral to negative increases the prioritization score?*” To do this, we manually fix the input variables (*Sentiment*, *Manipulated*, *Perception*) to obtain two simulated distributions of prioritization scores for a SATD instance: a set of prioritization scores for an SATD instance with neutral sentiment, and SATD instance with negative sentiment. We can quantify how negativity influences the outcome variable by computing the contrast between these two distributions. Because Model 1 has as input the perception of the respondent, whether the vignette has been manipulated, and the experience of the respondent we have to fix values for all three. As we are not interested in a scenario in which the vignettes are manipulated, we fix *Manipulated* to false. Similarly, we set *Experience* to the mean experience value of the respondents. Finally, we can not fix the value of *Perception* to one particular value for perception. Therefore, we compute a contrast distribution for each value of *Perception* (Disagree, Indifferent, Agree).

Table 4: The odds ratio between the odds of the priority score increasing compared to the odds of the score decreasing if sentiment changes from neutral to negative.

Perception	Effort	Urgency	Importance
Agree	1.02	1.00	0.99
No Opinion	1.00	0.99	1.00
Disagree	1.00	0.98	1.00

Prior predictive checks: To ensure our choice of priors does not bias our results, we perform prior predictive checks. Specifically, we sample from the model using only the priors to verify that the outcomes of the model are realistic. *i.e.*, the model only generates responses within the range of 1 – 5. Secondly, we want to verify that the responses of the model are unbiased, *i.e.*, the priors do not associate any effect of negativity on prioritization. To verify the outcomes are unbiased, we compute the odds of the prioritization increasing if negativity is expressed, as described in *estimating effects*. Figure 3 plots the

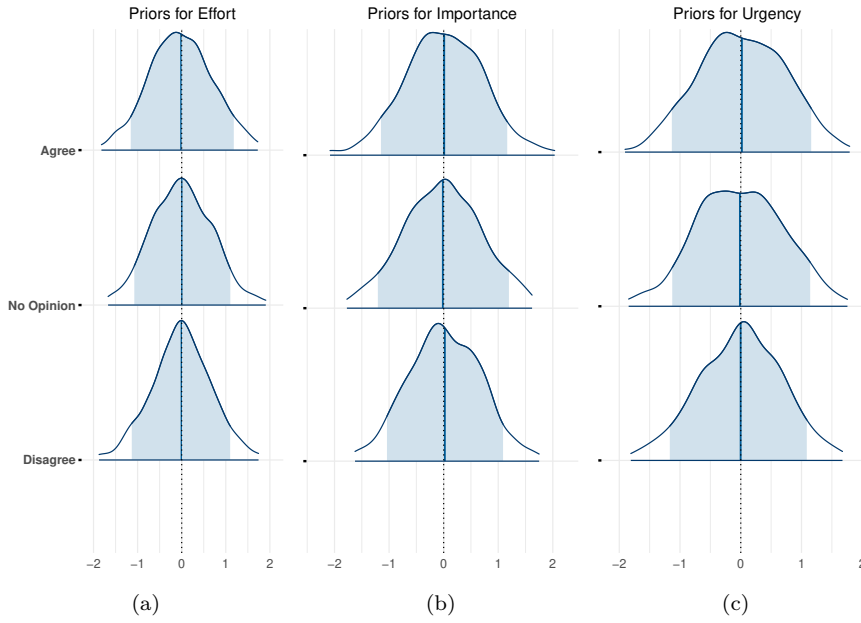


Fig. 3: Distributions of the priors for each of the outcome variables.

distribution of the priors for the effect of negativity on the estimation of the outcome variables. The distribution for each of the priors is centered around zero, showing that the priors do not favor any specific outcome. Additionally, if the priors are unbiased, we expect the odds ratios to be close to zero. This is the case, as can be seen in Table 4.

Perception and Demographics Respondents' answers to the closed questions and demographics are visualized. As the questions on Perception are taken verbatim from the study of Cassee et al. (2022) the results obtained in this experiment are compared to those of the original study. Meanwhile, the answers to questions on demographics taken from the Stack Overflow survey will be compared to the most recent results of the Stack Overflow survey.

3 Results

In total, we received 75 full and partial responses to the experiment. Based on a manual check, we removed one response because the values provided by the respondent for age and experience were unrealistic. Because we also accepted partial responses and because the instrument contained optional questions, we discuss the number of relevant and full responses included per question.

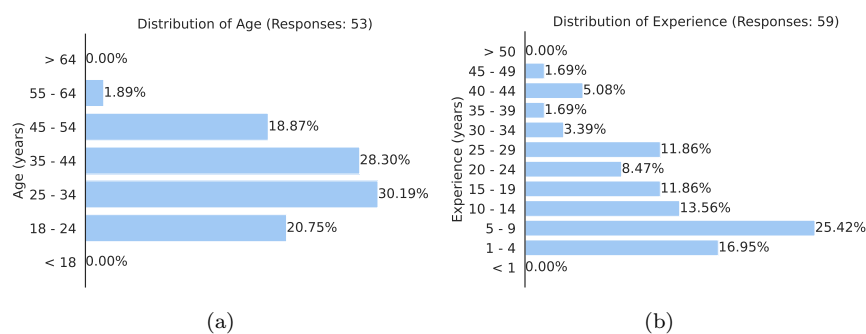


Fig. 4: Histograms for respondent age and experience.

3.1 Demographics

Figure 4 shows the respondent age and experience distributions. The bins for both plots are identical to the bins of the Stack Overflow developer survey.⁶ We compare the demographics to the Stack Overflow developer survey because the Stack Overflow survey is one of the largest surveys of developers, with almost 90,000 respondents for 2023. For age, our experiment’s population appears to be a bit older than the responses to the Stack Overflow developer survey. Notably, we saw no respondents younger than 18 or older than 64. Experience-wise, the distribution of respondents is quite similar to that of the Stack Overflow survey. Most importantly, there are no large differences between our respondents and the respondents to the Stack Overflow survey. Therefore, we conclude, based on the distributions of age and experience, that our respondents are a good representation of the general developer population for these two characteristics.

Table 5: The self-described employment status and the self-described coding outside of respondents’ work.

Employment Status	#	%	Coding outside of Work	#	%
Employed full-time	45	76.27%	Personal projects	20	64.52%
Student, full-time	8	13.56%	I do not write code outside of work	7	22.58%
Self-employed	4	6.78%	School or Academic	3	9.68%

Continued on next page

⁶ <https://survey.stackoverflow.co/2023/#overview>

Table 5: The self-described employment status and the self-described coding outside of respondents' work. (continued)

Employment Status	#	%	Coding outside of Work	#	%
Employed part-time	2	3.39%	Open-source	1	3.23%
Total	59	100.00%	Total	31	100.00%

Table 5 shows respondents' employment status and the coding respondents do outside of work. A large majority of the respondents are professional developers who are employed full-time.

3.2 Negativity's Effect on Prioritization

For the Bayesian models, we can only use responses for which all questions used in the model have received a response, particularly the demographic question on respondent experience (Q13). This requirement leaves 59 valid responses for the experiment on which the models have been fit.

Figure 5 visualizes the density distribution for the posteriors related to the effect of sentiment on the outcome. In other words, this figure shows the effect

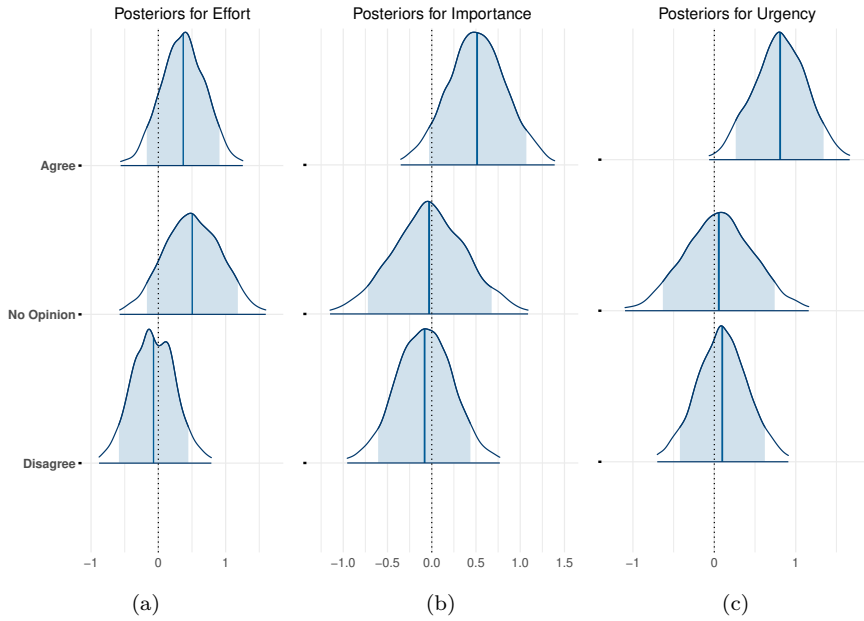


Fig. 5: Distributions of the posterior for each of the outcome variables.

the model associates with a change in sentiment from neutral to negative. A positive value indicates that a change from neutral to negative increases the prioritization score. As can be observed in Figure 5, the distributions of the priors are generally quite wide, and they overlap with the dashed line indicating zero. However, this uncertainty is not unexpected or uncommon in Bayesian analysis, especially in studies with smaller samples (Frattini et al., 2024; Ghorbani et al., 2023).

Table 6: Evidence ratio table for the hypothesis that negativity increases the prioritization score for each of the outcome variables.

Perception	Effort	Urgency	Importance
Agree	Moderate for	Strong for	Strong for
No Opinion	Moderate for	Anecdotal for	Anecdotal against
Disagree	Anecdotal against	Anecdotal for	Anecdotal against

Because we adjusted the models for the *Perception* of the participants, the effect of sentiment on the outcome can vary for each of the three levels of *Perception*. The labels *Disagree*, *No Opinion*, *Agree* therefore show the effect of negative sentiment on prioritization based on whether the participants believe that negativity should be interpreted as a signal that the SATD has a higher priority. The posteriors indicate that participants who agreed with the use of negativity as a proxy for priority were also more likely to assign a higher priority score, for all of the three measured variables. More interestingly, participants who indicated they had no opinion on the use of negativity as a proxy for priority increase the score for effort. Finally, for the other groups and outcomes we see that the models do not associate any impact of negativity, as the means all appear to be centered around zero. Table 6 lists an interpretation of the Bayes factor, or the strength of the evidence, of the hypothesis that negativity increases the prioritization. We interpret the values for the Bayes factor according to Stefan et al. (2019).

Table 7: The odds ratio between the odds of the priority score increasing compared to the odds of the score decreasing if sentiment changes from neutral to negative. Bold font indicates whether the evidence ratio supports the hypothesis that negativity increases prioritization (Table 6).

Perception	Effort	Urgency	Importance
Agree	1.40	1.95	1.50
No Opinion	1.56	1.07	0.98
Disagree	0.95	1.10	0.94

However, quantifying the actual impact that negativity has *e.g.*, answering the question: “Does negativity make it more likely that SATD receives a higher priority score?” is not possible based only on the plotted posteriors. Therefore, we also used the models to simulate, per outcome variable, the difference in priority scores for SATD expressing either neutral or negative sentiment. This results in a set of probabilities: How likely is it that negativity results in a higher priority score? How likely is it that negativity results in a lower priority score? Table 7 shows the ratio between these two odds for each outcome variable and for each value of *Perception*. From this table, we conclude that 37% of respondents who interpret negativity as a signal for importance (i.e., “Agree”) are also more likely to assign higher priority scores for SATD expressing negative sentiment. In this case, the consistency between respondents’ beliefs and actions is noteworthy, as the value-action gap is a well-documented bias (Barr, 2006). Secondly, Table 7 also shows that the perception of Urgency (OR: 1.95) is more likely to be influenced by negative sentiment than Effort (OR: 1.40) and Importance (OR: 1.50). Finally, 20% of respondents who had no opinion on whether they interpreted negativity as a signal of importance are likelier to assign a higher score for effort when negativity is expressed.

Findings – Negativity

Up to 57% of developers estimate the priority of self-admitted technical debt as higher when negativity is used to describe it. These developers are between 1.4 and 2.0 times more likely to increase, instead of decrease, their prioritization scores for self-admitted technical debt expressing negativity.

3.3 Perceptions and Self-Reported Behavior

Figure 6 shows the responses to the questions asking participants about how often they write or come across SATD that expresses negative sentiment. These results indicate that almost half of the developers encounter SATD expressing

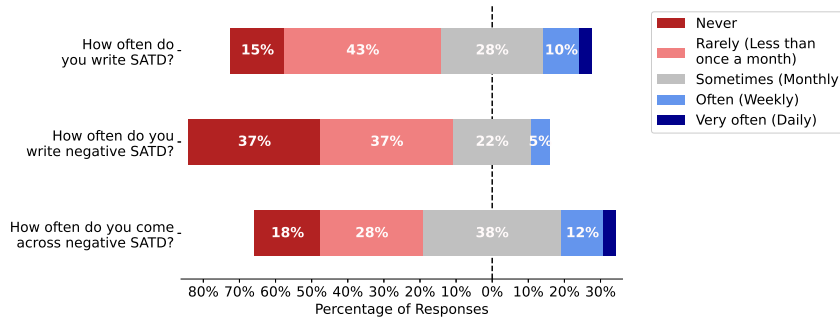


Fig. 6: Responses to survey questions Q4, Q5, Q6.

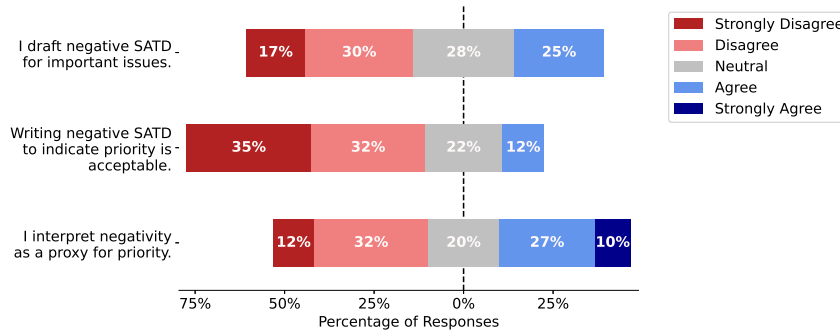


Fig. 7: Responses to survey questions Q7, Q8, Q9.

negative sentiment monthly or even more often. However, most developers, almost 75%, never, or rarely, write any SATD expressing negative sentiment. The distributions observed in this study are similar to the study of Cassee et al. (2022). The most important difference between this study and the findings of Cassee *et al.* is that respondents to the experiment tend to come across SATD expressing negativity less frequently. This could be explained by the fact that we used convenience sampling through our industrial contacts for this experiment, and because industrial developers annotate SATD differently (Zampetti et al., 2021).

Figure 7 outlines the results to the questions on respondents perceptions. Compared to Cassee et al. (2022) the proportion of respondents indicating that they strongly disagree with the statement that they interpret negativity as a proxy for priority is smaller. For the other two statements the results for this experiment are similar to those of Cassee *et al.*

From the combination of Figure 6 and Figure 7 we generally confirm the findings of Cassee et al. (2022). A large majority of developers disagree that writing SATDs that express negative sentiment to signal priority is acceptable. However, even if most developers believe this, they still express negativity in

SATDs to indicate priority or state that they interpret negativity as a proxy for priority.

Findings – Perceptions

Between 25% to 50% of developers draft and/or encounter self-admitted technical debt expressing negativity. Meanwhile, 67% of developers state that using negativity as a proxy for priority is unacceptable.

4 Discussion

Our study shows that developers use source-code comments to communicate not only a description of technical debt but also priority. From previous work, we know that source code comments have always been used for various purposes, including the description of technical-debt (Pascarella and Bacchelli, 2017). Furthermore, descriptions of technical debt in source-code comments cover a wide range of technical issues (Cassee et al., 2022; Maldonado et al., 2017). Based on the results of our experiment, we add to the existing body of knowledge on the role of source-code comments in addressing technical debt. Specifically, we find that negativity in these descriptions results in the same technical issue being prioritized differently. Therefore, we conclude that developers use source code comments not just to describe technical issues but also use negativity to communicate priority. Negativity can be seen as a hidden message or additional dimension, as a SATD comment’s primary purpose is to describe the technical debt. This study and the results of the previous study (Cassee et al., 2022) show that developers purposefully express negativity to emphasize priority and interpret negativity as a proxy for priority.

Does this mean developers should use negativity to describe technical-debt they think is important? For individual developers, using negative expressions to annotate the technical debt they are working on might be an effective strategy. It could result in the issue they care about being prioritized and, therefore, fixed before other technical debt is addressed. However, from a team perspective, this is a potentially problematic strategy. Software projects are often characterized by only having limited capacity available to work on technical debt (Zampetti et al., 2021), and developers using negativity to describe technical debt lay claim to more of this capacity. Secondly, whether developers are influenced by negativity is conditional to their perceptions, so not everyone in a team will interpret the priority similarly. Additionally, two-thirds of developers believe that using negativity as a proxy for priority is unacceptable. Finally, technical debt demoralizes developers (Besker et al., 2020), and expressions of negativity could be particularly demoralizing; therefore, we would not recommend developers use negativity as a proxy for priority as there are alternative prioritization methods. From Cassee et al. (2022), we know devel-

opers prefer the use of other channels to communicate the priority of SATD, such as issue trackers or work-management systems.

Does this mean developers should stop expressing negativity? Emotions and sentiment are a part of everyday life (Kleef, 2009). It has long been found that developers express emotions in many different software engineering activities (Calefato et al., 2018; Lin et al., 2022; Mäntylä et al., 2016), and technical debt appears to be able to trigger negative emotions in software engineers (Olsson et al., 2021). This makes it unrealistic to expect developers to stop expressing their emotions or opinions completely. Instead, this study shows how negativity can affect the management of technical debt, where prioritization is one of the many different management activities of technical debt (Li et al., 2015). However, we believe it is important to emphasize that any negative emotions expressed are not directed towards developers, as they sometimes are (Gachechiladze et al., 2017). Especially as negativity directed towards other developers could be considered toxic, with far-reaching consequences (Miller et al., 2022; Raman et al., 2020).

Generalizability & Future work: We expect the findings of this study to apply to technical debt in general and not just to self-admitted technical debt described in source-code comments. Because of the high level of control in the experiments and the experienced participants, we expect the observed effect of negativity on prioritization to translate to technical debt described on other platforms. Notably, this includes places like issue-trackers, as these are places where technical debt is often described (Cassee et al., 2022; Xavier et al., 2022; Zampetti et al., 2021).

Through the experiment, we conclude that negativity affects prioritization, and we describe how developers use negativity in source-code comments to communicate priority. However, a consequence of the high control of the experiments is the reduced realism. Therefore, there is still an opportunity for future work to understand how negativity in technical debt influences prioritization in the “field” (Storey et al., 2020). Research methods or strategies like ethnographic observations can describe the prioritization of technical debt in software projects, as opposed to the contrived setting of the experiments, like the work of Aranda and Venolia (2009) on bugs. Or, research methods such as mining software repositories can be used to study whether SATD in which negativity is expressed is removed quicker than neutral SATD.

Secondly, in this study, we only focus on the effect of negativity, and we consider the effect of positivity to be out of scope. This is done purposefully, as previous work has found that positivity rarely occurs in SATD (Cassee et al., 2022). Therefore, understanding its effect on prioritization is not as relevant. However, prior literature has described other software engineering activities where developers express positivity in their communication (Asri et al., 2019; Calefato et al., 2018; Cheruvelil and da Silva, 2019). We believe that these activities (change requests, asking questions on SO) are better suited to study the effect of positivity.

5 Threats to Validity

Notwithstanding our effort to ensure our results were valid, factors outside of our control might still have influenced the results. To discuss these threats to validity and our mitigations, we use the framework of Wohlin et al. (2012), as the research method of this study is an experiment.

Internal Validity: We carefully designed the experiments to reduce the effect of confounders as much as possible, for instance, by randomizing the order of vignettes, using existing instances of SATD, and by making sure the instances were understandable. However, to create the vignettes we manipulated descriptions to create either a negative or a neutral counterpart. We took several precautions to reduce the bias of the manipulation of the obtained results. These include ensuring a balance by creating two negative and two neutral manipulated descriptions, generating several manipulated descriptions, and voting for the best one. Furthermore, we added manipulation as a parameter to the model to adjust for any remaining bias.

Secondly, we did not state the exact purpose of the experiments to prevent the respondents from guessing the experiment's hypothesis and responding in line with their *attitude* towards the hypothesis. Similarly, the vignettes did not indicate whether participants viewed the neutral or the negative instance. However, there is still a chance that some respondents might have recognized the purpose of the experiments after scoring two or more vignettes and were, therefore, influenced by their attitudes.

Finally, there is a risk that the wording of the perception-related questions might have biased participants. However, re-using the exact wording from the survey administered by Cassee et al. (2022) allows us to compare our results to those of Cassee *et al.*

External Validity: The respondents sampled for the experiment were diverse with respect to age and working experience, with both industry developers and open-source developers partaking in the experiment. Therefore, we expect our results to generalize to both open and closed-source developers, especially because the respondent's demographics are similar to those of the respondents to the Stack Overflow survey. However, the most important threat to external validity we identify is related to the choice of vignettes. In the experiments, we only showed participants four instances, or vignettes, of technical debt, and those instances all belonged one type of SATD, *Poor implementation choices*. Therefore, the results of this experiment might not generalize over other types or categories of SATD, such as functional issues, or SATD described in issue tracking systems. However, we tried to minimize the threat by using four different vignettes instead of one, and therefore, trying to balance participant fatigue or disengagement with generalizability.

Construct Validity: Most importantly, we recognize that *Importance*, *Urgency*, and *Effort* as Likert scale questions are not operationalizations of priority used

in practice by developers to score technical debt. However, we opted for these three constructs because they are easy to explain to participants and because there is no universally accepted construct to measure the priority of technical debt (Bellotti et al., 2004; Gavidia-Calderon et al., 2021; Middleton et al., 2018). While the usage of these constructs might influence results, they allow us to measure perceived priority consistently.

6 Related Work

The below section discusses the related work on SATD, its management, and sentiment analysis in software engineering.

6.1 Self-Admitted Technical Debt

Many different aspects of technical debt have been studied previously (Tom et al., 2013), such as, for instance, its management (Li et al., 2015), or the financial aspects (Ampatzoglou et al., 2015). This work focuses on the self-admission of Technical Debt (SATD), often present in software systems as source-code comments (Maldonado and Shihab, 2015; Potdar and Shihab, 2014). As for technical debt, many different aspects of SATD have been studied. Such as the annotation practices of industry developers (Zampetti et al., 2021), the places where SATD has been described (Kashiwa et al., 2022; Xavier et al., 2020, 2022), and the curation of datasets of SATD instances (Guo et al., 2021; Sridharan et al., 2023).

Because of the negative connotations of technical debt, research has focused on the classification of SATD in source code. Liu et al. (2018) found that a SATD detector based on text-mining techniques can automatically identify SATD. Meanwhile, Ren et al. (2019) have used deep-learning models to classify whether source code snippets contain SATD, finding that these models, trained on large datasets, can be competitive. However, what automatic detection technique has the best predictive performance is not entirely known, as Guo et al. (2021) found that simple rule-based detection techniques can still outperform existing deep-learning tools. Cassee et al. (2022) has found that different types of issues are described in SATD, including *functional issues*, *poor implementation choices*, and *documentation issues*. Furthermore, from surveying developers, Cassee *et al.* find that a group of developers state they use negativity to prioritize SATD.

Another often studied aspect of SATD is its removal; both Palomba et al. (2017) and Peruma et al. (2022) studied the link between refactorings and SATD. Palomba et al. (2017) found that code refactorings occur in locations where SATD is present, while Peruma et al. (2022) finds that code refactoring often coincides with removing SATD. By studying a dataset of SATD instances Maldonado et al. (2017) found that 75% of SATD instances are removed in subsequent source code revisions. Furthermore, Zampetti et al. (2018) found

that the removal of SATD is acknowledged in commit messages. Additionally, Zampetti *et al.* find that in 20% to 50% of cases, SATD is removed when entire methods or classes are removed. By surveying developers Tan et al. (2021) found that self-removal of SATD is often a conscious decision. Moreover, they found that more experienced developers are more often concerned when it comes to the removal of SATD. However, none of the studies on the removal of SATD focus on whether the presence of negativity influences whether SATD is removed.

6.2 Sentiment in Software Engineering

Because this study focuses on the effect of sentiment on the prioritization of SATD, we describe literature related to the role of sentiment in software engineering. In general, sentiment analysis has been used to study a wide variety of software engineering activities (Lin et al., 2022), such as live meeting analysis (Herrmann and Klunder, 2021) or to design and build recommender systems (Lin et al., 2018).

Different studies have tried to understand whether expressions of negative sentiment are correlated with suboptimal development practices, such as unresolved issues, design smells, or bugs. Valdez et al. (2020) found that unresolved issues in Jira tend to express more negative sentiment than closed issues. Similarly, based on a study of issue reopenings Cheruvilil and da Silva (2019) finds that issues that have been reopened once or more than once tend to have more comments expressing negative emotions. For commit messages Huq et al. (2020) reports that commits related to bug fixing express more negative sentiment. While for code reviews Asri et al. (2019) finds that code reviews in which negative sentiment is expressed tend to take longer to complete. Olsson et al. (2021) find that some design smells cause developers to feel negative emotions.

Wrt. positive sentiment, or the expression of positive emotions, Ortu et al. (2015) has found that more positive sentiment in the description appears to correlate with a shorter issue resolution time. Similarly, Sanei et al. (2021) has studied the sentiment in issue comments, finding that more positive comments correlate with faster resolution.

After studying questions on Stack Overflow Calefato et al. (2018) reports that successful questions on SO tend to be neutral, *i.e.*, express no positive or negative sentiment. Implying that neutral or more factual questions are more likely to receive a quicker response.

The existing studies focus on the impact of sentiment on software engineering through correlational analysis. However, the studies don't focus on causation (*i.e.*, "did the expression of sentiment polarity cause the observed effect?"). Through the controlled experiment reported in this study, we further understand how negativity influences the perception of priority and explain how expressions of negativity impact software development.

7 Conclusion

In this paper, we report on a controlled experiment to study whether developers interpret negativity in self-admitted technical debt as a proxy for priority. We exposed participants to instances of self-admitted technical debt with or without negativity and asked them to estimate the priority of each instance. By analyzing the prioritization scores of the experiments using Bayesian statistics, we describe how negativity influences the perception of priority. Furthermore, to better describe the role negativity plays in the management of self-admitted technical debt, we asked respondents a series of questions about their perceptions towards the usage of negativity in SATD.

Based on the participation of 59 experienced industrial software engineers, we find that one-third to half of the developers are more likely to increase their estimation of priority if negativity is used to describe the SATD. Specifically, they're between 1.5 and 2 times as likely to increase, as opposed to decrease, their prioritization score if negativity is present. Secondly, we confirm previous findings and conclude that two-thirds of developers believe writing SATD in which negativity is expressed or interpreting negativity as a proxy for priority is unacceptable. Our findings show that developers who believe the usage of negativity as a proxy for priority is unacceptable still admit to writing negative SATD and use negativity to estimate priority.

Based on the experiment, we learn how developers use negativity in SATD to communicate priority, in particular, we find that developers use negativity as an additional dimension, in addition to descriptions of technical issues. Furthermore, our results show why developers' expression of negativity is unavoidable and help explain the purpose of negative expressions. However, our results also show why using negativity to describe technical debt might not be advisable because of both developers' perceptions, and not all developers are influenced by it.

Conflict of interest

The authors declare that they have no conflict of interest.

Data Availability Statement

The data used for this study and the notebook used to analyze the data are publicly available in a figshare repository.⁷

⁷ <https://doi.org/10.6084/m9.figshare.26863435>

References

- Aguinis H, Bradley KJ (2014) Best practice recommendations for designing and implementing experimental vignette methodology studies. *Organizational Research Methods* 17(4):351–371, DOI 10.1177/1094428114547952
- Ampatzoglou A, Ampatzoglou A, Chatzigeorgiou A, Avgeriou P (2015) The financial aspect of managing technical debt: A systematic literature review. *Information and Software Technology* 64:52–73, DOI 10.1016/j.infsof.2015.04.001
- Aranda J, Venolia G (2009) The secret life of bugs: Going past the errors and omissions in software repositories. *IEEE*, pp 298–308, DOI 10.1109/ICSE.2009.5070530
- Asri IE, Kerzazi N, Uddin G, Khomh F, Idrissi MAJ (2019) An empirical study of sentiments in code reviews. *Information and Software Technology* 114:37–54, DOI 10.1016/j.infsof.2019.06.005
- Atzmüller C, Steiner PM (2010) Experimental vignette studies in survey research. *Methodology* 6(3):128–138, DOI 10.1027/1614-2241/a000014
- Barr S (2006) Environmental action in the home: Investigating the ‘value-action’ gap. *Geography* 91:43–54, DOI 10.1080/00167487.2006.12094149
- Bellotti V, Dalal B, Good N, Flynn P, Bobrow DG, Ducheneaut N (2004) What a to-do: studies of task management towards the design of a personal task list manager. *ACM*, pp 735–742, DOI 10.1145/985692.985785
- Besker T, Ghanbari H, Martini A, Bosch J (2020) The influence of technical debt on software developer morale. *Journal of Systems and Software* 167:110586, DOI 10.1016/j.jss.2020.110586
- Calefato F, Lanubile F, Novielli N (2018) How to ask for technical help? evidence-based guidelines for writing questions on stack overflow. *Information and Software Technology* 94:186–207, DOI 10.1016/j.infsof.2017.10.009
- Cassee N, Zampetti F, Novielli N, Serebrenik A, Penta MD (2022) Self-admitted technical debt and comments’ polarity: an empirical study. *Empirical Software Engineering* 27:139, DOI 10.1007/s10664-022-10183-w
- Cheruvellil J, da Silva BC (2019) Developers’ sentiment and issue reopening. *IEEE*, pp 29–33, DOI 10.1109/SEmotion.2019.00013
- Cunningham W (1992) The wycash portfolio management system. *SIGPLAN OOPS Mess* 4(2):29–30, DOI 10.1145/157710.157715
- Danilova A, Naiakshina A, Horstmann S, Smith M (2021) Do you really code? designing and evaluating screening questions for online surveys with programmers. *IEEE*, pp 537–548, DOI 10.1109/ICSE43902.2021.00057
- Elwert F (2013) Graphical causal models. DOI 10.1007/978-94-007-6094-3_13
- Frattini J, Fucci D, Torkar R, Montgomery L, Unterkalmsteiner M, Fischbach J, Mendez D (2024) Applying bayesian data analysis for causal inference about requirements quality: A replicated experiment. 2401.01154
- Furia CA, Torkar R, Feldt R (2022) Applying bayesian analysis guidelines to empirical software engineering data: The case of programming languages and code quality. *ACM Transactions on Software Engineering and Methodology* 31:1–38, DOI 10.1145/3490953

- Gachechiladze D, Lanubile F, Novielli N, Serebrenik A (2017) Anger and its direction in collaborative software development. *IEEE*, pp 11–14, DOI 10.1109/ICSE-NIER.2017.18
- Gavidia-Calderon C, Sarro F, Harman M, Barr ET (2021) The assessor’s dilemma: Improving bug repair via empirical game theory. *IEEE Transactions on Software Engineering* 47(10):2143–2161, DOI 10.1109/TSE.2019.2944608
- Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB (2013) *Bayesian Data Analysis*. Chapman and Hall/CRC, DOI 10.1201/b16018
- Gelman A, Vehtari A, Simpson D, Margossian CC, Carpenter B, Yao Y, Kennedy L, Gabry J, Bürkner PC, Modrák M (2020) Bayesian workflow. URL <http://arxiv.org/abs/2011.01808>
- Ghorbani A, Cassee N, Robinson D, Alami A, Ernst NA, Serebrenik A, Wasowski A (2023) Autonomy is an acquired taste: Exploring developer preferences for github bots. *IEEE*, pp 1405–1417, DOI 10.1109/ICSE48619.2023.00123
- Guo Z, Liu S, Liu J, Li Y, Chen L, Lu H, Zhou Y (2021) How far have we progressed in identifying self-admitted technical debts? a comprehensive empirical study. *ACM Transactions on Software Engineering and Methodology* 30:1–56, DOI 10.1145/3447247
- Herrmann M, Klunder J (2021) From textual to verbal communication: Towards applying sentiment analysis to a software project meeting. *IEEE*, pp 371–376, DOI 10.1109/REW53955.2021.00065
- Hughes JL, Camden AA, Yangchen T (2016) Rethinking and updating demographic questions: Guidance to improve descriptions of research samples. *Psi Chi Journal of Psychological Research* 21:138–151, DOI 10.24839/2164-8204.JN21.3.138
- Huq SF, Sadiq AZ, Sakib K (2020) Is developer sentiment related to software bugs: An exploratory study on github commits. *IEEE*, pp 527–531, DOI 10.1109/SANER48275.2020.9054801
- Johnson MK, Hashtroudi S, Lindsay DS (1993) Source monitoring. *Psychological Bulletin* 114:3–28, DOI 10.1037/0033-2909.114.1.3
- Juristo N, Moreno AM (2001) Experiments with undesired variations. DOI 10.1007/978-1-4757-3304-4_9
- Kashiwa Y, Nishikawa R, Kamei Y, Kondo M, Shihab E, Sato R, Ubayashi N (2022) An empirical study on self-admitted technical debt in modern code review. *Information and Software Technology* 146:106855, DOI 10.1016/j.infsof.2022.106855
- Kleef GAV (2009) How emotions regulate social life. *Current Directions in Psychological Science* 18:184–188, DOI 10.1111/j.1467-8721.2009.01633.x
- Kruschke JK, Liddell TM (2018) The bayesian new statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a bayesian perspective. *Psychonomic Bulletin & Review* 25:178–206, DOI 10.3758/s13423-016-1221-4
- Lenarduzzi V, Besker T, Taibi D, Martini A, Fontana FA (2021) A systematic literature review on technical debt prioritization: Strategies, pro-

- cesses, factors, and tools. *Journal of Systems and Software* 171:110827, DOI 10.1016/j.jss.2020.110827
- Li Z, Avgeriou P, Liang P (2015) A systematic mapping study on technical debt and its management. *Journal of Systems and Software* 101:193–220, DOI 10.1016/j.jss.2014.12.027
- Lim E, Taksande N, Seaman C (2012) A balancing act: What software practitioners have to say about technical debt. *IEEE Software* 29:22–27, DOI 10.1109/MS.2012.130
- Lin B, Zampetti F, Bavota G, Penta MD, Lanza M, Oliveto R (2018) Sentiment analysis for software engineering. *ACM*, pp 94–104, DOI 10.1145/3180155.3180195
- Lin B, Cassee N, Serebrenik A, Bavota G, Novielli N, Lanza M (2022) Opinion mining for software development: A systematic literature review. *ACM Transactions on Software Engineering and Methodology* 31:1–41, DOI 10.1145/3490388
- Liu Z, Huang Q, Xia X, Shihab E, Lo D, Li S (2018) Satd detector: a text-mining-based self-admitted technical debt detection tool. *ACM*, pp 9–12, DOI 10.1145/3183440.3183478
- Maldonado EDS, Shihab E (2015) Detecting and quantifying different types of self-admitted technical debt. 2015 IEEE 7th International Workshop on Managing Technical Debt, MTD 2015 - Proceedings pp 9–15, DOI 10.1109/MTD.2015.7332619
- Maldonado EDS, Abdalkareem R, Shihab E, Serebrenik A (2017) An empirical study on the removal of self-admitted technical debt. *IEEE*, pp 238–248, DOI 10.1109/ICSME.2017.8
- Mäntylä M, Adams B, Destefanis G, Graziotin D, Ortu M (2016) Mining valence, arousal, and dominance - possibilities for detecting burnout and productivity? *Proceedings - 13th Working Conference on Mining Software Repositories, MSR 2016* pp 247–258, DOI 10.1145/2901739.2901752
- McElreath R (2018) *Statistical Rethinking*. Chapman and Hall/CRC, DOI 10.1201/9781315372495
- Middleton S, Charnock A, Forster S, Blakey J (2018) Factors affecting - individual task prioritisation in a workplace setting. *Future Healthc J* 5(2):138–142
- Miller C, Cohen S, Klug D, Vasilescu B, KaUstner C (2022) "did you miss my comment or what?": understanding toxicity in open source discussions. *ACM*, pp 710–722, DOI 10.1145/3510003.3510111
- Molokken K, Jorgensen M (2003) A review of software surveys on software effort estimation. *IEEE Comput. Soc*, pp 223–230, DOI 10.1109/ISESE.2003.1237981
- Novielli N, Serebrenik A (2023) Emotion analysis in software ecosystems. DOI 10.1007/978-3-031-36060-2_5
- Okon-Singer H, Hendler T, Pessoa L, Shackman AJ (2015) The neurobiology of emotion-cognition interactions: Fundamental questions and strategies for future research. *Frontiers in Human Neuroscience* 9, DOI 10.3389/fnhum.2015.00058

- Olsson J, Risfelt E, Besker T, Martini A, Torkar R (2021) Measuring affective states from technical debt. *Empirical Software Engineering* 26:105, DOI 10.1007/s10664-021-09998-w
- Ortu M, Adams B, Destefanis G, Tourani P, Marchesi M, Tonelli R (2015) Are bullies more productive? empirical study of affectiveness vs. issue fixing time. *IEEE*, vol 2015-Augus, pp 303–313, DOI 10.1109/MSR.2015.35
- Palomba F, Zaidman A, Oliveto R, Lucia AD (2017) An exploratory study on the relationship between changes and refactoring. *IEEE*, pp 176–185, DOI 10.1109/ICPC.2017.38
- Pascarella L, Bacchelli A (2017) Classifying code comments in java open-source software systems. *IEEE*, pp 227–237, DOI 10.1109/MSR.2017.63
- Peruma A, AlOmar EA, Newman CD, Mkaouer MW, Ouni A (2022) Refactoring debt: myth or reality? an exploratory study on the relationship between technical debt and refactoring. *ACM*, pp 127–131, DOI 10.1145/3524842.3528527
- Potdar A, Shihab E (2014) An exploratory study on self-admitted technical debt. *IEEE*, pp 91–100, DOI 10.1109/ICSME.2014.31
- Raman N, Cao M, Tsvetkov Y, Kästner C, Vasilescu B (2020) Stress and burnout in open source. *ACM*, pp 57–60, DOI 10.1145/3377816.3381732
- Raykov PP, Varga D, Bird CM (2023) False memories for ending of events. *Journal of Experimental Psychology: General* 152:3459–3475, DOI 10.1037/xge0001462
- Ren X, Xing Z, Xia X, Lo D, Wang X, Grundy J (2019) Neural network-based detection of self-admitted technical debt. *ACM Transactions on Software Engineering and Methodology* 28:1–45, DOI 10.1145/3324916
- Robillard MP, Arya DM, Ernst NA, Guo JL, Lamothe M, Nassif M, Novielli N, Serebrenik A, Steinmacher I, Stol KJ (2024) Communicating study design trade-offs in software engineering. *ACM Transactions on Software Engineering and Methodology* DOI 10.1145/3649598
- Sanei A, Cheng J, Adams B (2021) The impacts of sentiments and tones in community-generated issue discussions. *IEEE*, pp 1–10, DOI 10.1109/CHASE52884.2021.00009
- Siegmund J, Kästner C, Liebig J, Apel S, Hanenberg S (2014) Measuring and modeling programming experience. *Empir Softw Eng* 19(5):1299–1334, DOI 10.1007/S10664-013-9286-4
- Singh N, Singh P (2018) How do code refactoring activities impact software developers’ sentiments? - an empirical investigation into github commits. *Proceedings - Asia-Pacific Software Engineering Conference, APSEC 2017-Decem*:648–653, DOI 10.1109/APSEC.2017.79
- Sridharan M, Rantala L, Mäntylä M (2023) Pentacet data-23 million contextual code comments and 250,000 satd comments. DOI 10.1109/MSR59073.2023.00063
- Steele CM, Aronson J (1995) Stereotype threat and the intellectual test performance of african americans. *Journal of Personality and Social Psychology* 69:797–811, DOI 10.1037/0022-3514.69.5.797

- Stefan AM, Gronau QF, Schönbrodt FD, Wagenmakers EJ (2019) A tutorial on bayes factor design analysis using an informed prior. *Behavior Research Methods* 51:1042–1058, DOI 10.3758/s13428-018-01189-8
- Stol KJ, Fitzgerald B (2018) The abc of software engineering research. *ACM Transactions on Software Engineering and Methodology* 27, DOI 10.1145/3241743
- Storey MA, Ernst NA, Williams C, Kalliamvakou E (2020) The who, what, how of software engineering research: a socio-technical framework. *Empirical Software Engineering* 25:4097–4129, DOI 10.1007/s10664-020-09858-z
- Swillus M, Zaidman A (2023) Sentiment overflow in the testing stack: Analyzing software testing posts on stack overflow. *Journal of Systems and Software* 205:111804, DOI 10.1016/j.jss.2023.111804
- Tan J, Feitosa D, Avgeriou P (2021) Do practitioners intentionally self-fix technical debt and why? *IEEE*, pp 251–262, DOI 10.1109/ICSME52107.2021.00029
- Tom E, Aurum A, Vidgen R (2013) An exploration of technical debt. *Journal of Systems and Software* 86:1498–1516, DOI 10.1016/j.jss.2012.12.052
- Torkar R, Furia CA, Feldt R, de Oliveira Neto FG, Gren L, Lenberg P, Ernst NA (2022) A method to assess and argue for practical significance in software engineering. *IEEE Transactions on Software Engineering* 48:2053–2065, DOI 10.1109/TSE.2020.3048991
- Valdez A, Oktaba H, Gomez H, Vizcaino A (2020) Sentiment analysis in jira software repositories. *IEEE*, pp 254–259, DOI 10.1109/CONISOFT50191.2020.00043
- Wehaibi S, Shihab E, Guerrouj L (2016) Examining the impact of self-admitted technical debt on software quality. *IEEE*, pp 179–188, DOI 10.1109/SANER.2016.72
- Wohlin C, Runeson P, Höst M, Ohlsson MC, Regnell B, Wesslén A (2012) Planning. DOI 10.1007/978-3-642-29044-2_8
- Xavier L, Ferreira F, Brito R, Valente MT (2020) Beyond the code: Mining self-admitted technical debt in issue tracker systems. *ACM*, pp 137–146, DOI 10.1145/3379597.3387459
- Xavier L, Montandon JE, Ferreira F, Brito R, Valente MT (2022) On the documentation of self-admitted technical debt in issues. *Empirical Software Engineering* 27:163, DOI 10.1007/s10664-022-10203-9
- Yeung DY, Wong CK, Lok DP (2011) Emotion regulation mediates age differences in emotions. *Aging & Mental Health* 15:414–418, DOI 10.1080/13607863.2010.536136
- Zampetti F, Serebrenik A, Penta MD (2018) Was self-admitted technical debt removal a real removal? *ACM*, pp 526–536, DOI 10.1145/3196398.3196423
- Zampetti F, Serebrenik A, Penta MD (2020) Automatically learning patterns for self-admitted technical debt removal. *IEEE*, pp 355–366, DOI 10.1109/SANER48275.2020.9054868
- Zampetti F, Fucci G, Serebrenik A, Penta MD (2021) Self-admitted technical debt practices: a comparison between industry and open-source. *Empirical Software Engineering* 26:131, DOI 10.1007/s10664-021-10031-3