

Human Values Violations in Stack Overflow: An Exploratory Study

Sara Krishtul¹, Mojtaba Shahin², Humphrey O. Obie¹
Hourieh Khalajzadeh¹, Fan Gai¹, Ali Rezaei Nasab³, John Grundy¹

¹Faculty of IT, Monash University, Melbourne, Australia

²School of Computing Technologies, RMIT University, Melbourne, Australia

³School of Electrical and Computer Engineering, Shiraz University, Shiraz, Iran

{skri0002, fgai0001}@student.monash.edu, mojtaba.shahin@rmit.edu.au, {humphrey.obie, hourieh.khalajzadeh, john.grundy}@monash.edu, alirezaei@hafez.shirazu.ac.ir

ABSTRACT

A growing number of software-intensive systems are being accused of *violating or ignoring human values* (e.g., privacy, inclusion, and social responsibility), and this poses great difficulties to individuals and society. Such violations often occur due to the solutions employed and decisions made by developers of such systems that are misaligned with user values. Stack Overflow is the most popular Q&A website among developers to share their issues, solutions (e.g., code snippets), and decisions during software development. We conducted an exploratory study to investigate the occurrence of human values violations in Stack Overflow posts. As comments under posts are often used to point out the possible issues and weaknesses of the posts, we analyzed 2,000 Stack Overflow comments and their corresponding posts (1,980 unique questions or answers) to identify the types of human values violations and the reactions of Stack Overflow users to such violations. Our study finds that 315 out of 2,000 comments contain concerns indicating their associated posts (313 unique posts) violate human values. Leveraging Schwartz's theory of basic human values as the most widely used values model, we show that *hedonism* and *benevolence* are the most violated value categories. We also find the reaction of Stack Overflow commenters to perceived human values violations is very quick, yet the majority of posts (76.35%) accused of human values violation do not get downvoted at all. Finally, we find that the original posters rarely react to the concerns of potential human values violations by editing their posts. At the same time, they usually are receptive when responding to these comments in follow-up comments of their own.

CCS CONCEPTS

• **Software and its engineering** → **Collaboration in software development**; • **Social and professional topics** → *Socio-technical systems*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EASE '22, June 13–15, 2022, Göteborg, Sweden

© 2022 Association for Computing Machinery.

ACM ISBN AAAAAAAAAA...\$AAAA

<https://doi.org/A>

KEYWORDS

Human Values, Violations, User, Stack Overflow

ACM Reference Format:

Sara Krishtul¹, Mojtaba Shahin², Humphrey O. Obie¹, Hourieh Khalajzadeh¹, Fan Gai¹, Ali Rezaei Nasab³, John Grundy¹. 2022. Human Values Violations in Stack Overflow: An Exploratory Study. In *EASE '22: International Conference on Evaluation and Assessment in Software Engineering, June 13–15, 2022, Göteborg, Sweden*. ACM, New York, NY, USA, 11 pages. <https://doi.org/A>

1 INTRODUCTION

“Every line of code represents an ethical and moral decision; every bit of data collected, analyzed, and visualized has moral implications.”
Grady Booch [10]

Over the last few years, the media has reported many examples of software-intensive systems that intentionally or unintentionally ignored or violated ethical and human values [1, 17]. These systems have sometimes posed irreversible damages and challenges to end-users (e.g., loss of life), society (e.g., ignoring or being biased against a particular gender), and the software industry (e.g., damaging the software creator's reputation). For example, Amazon's "Prime same-day delivery service" was developed to provide all US citizens an equal and fair shopping experience [22]. However, it has been found that it prevents black neighborhoods from receiving such a shopping experience. In another recent example, the Facebook AI-based feature recommendation system mistakenly labeled a video of Black men as 'Primates' [1, 5].

Human values such as privacy, inclusion, power, fairness, and pleasure are defined as something that is deemed important for an individual, a group of people, or a society [37]. Many definitions and models have been proposed for human values in social science [14]. However, the most well-known and widely used one is Schwartz's theory of basic human values [37, 38]. Schwartz's theory of basic values includes ten universal values (*self-direction, stimulation, hedonism, achievement, power, security, conformity, tradition, benevolence, universalism*), which were identified through a survey of participants across 80 countries. While ethics are referred to as moral expectations that all individuals in a society agree upon, each person's values may differ from those of another [16, 46]. Hence, human values are *“enduring beliefs that a specific mode of conduct or end state of existence is personally or socially preferable to an opposite or converse mode of conduct or end state of existence”* [36].

Question

How to leave a message for a github.com user

Asked 8 years, 11 months ago · Active 4 days ago · Viewed 199k times

Answer

315 This method was working as of June 2021

1. Copy and paste the next line into your browser (feel free to bookmark it): <https://api.github.com/users/xxxxxx/events/public>
2. Find the GitHub username for which you want the email. Replace the xxxxxx in the URL with the person's GitHub username. Hit Enter.
3. Press Ctrl+F and search for "email".

As suggested by [qbolec](#), the above steps can be done by using this snippet:

```
<input id=username type="text" placeholder="github username or repo lin
<button onclick="fetch('https://api.github.com/users/${username.value}.r
<div id=results"></div>
```

[Run code snippet](#) [Expand snippet](#)

Comments

10 In my case it does not contain an Email – [TheKitMurkit](#) Jul 11 '17 at 14:18

8 This method does require that the user has some public activity – [ovinoophile](#) Jul 26 '17 at 15:55

7 Whoa, this answer is just a treasure for spammers... "How to find any github's user email adress" – [Alex from Jibiti](#) Sep 6 '17 at 10:24

8 The email address in this api is coming from the user's local git setup, not what they've set in Github. Github just recently added an option, under Settings/Emails, to block commits that accidentally include an email address the user has marked as 'private'. And, yes, the api-revealed content is only applicable to public repos; private repo data won't show. – [leanne](#) Dec 1 '17 at 22:22

2 @pbjh somebody should tell github, probably, this is a privacy-leak. – [Gewure](#) Jun 11 '19 at 0:26

[Show 12 more comments](#)

Figure 1: A question (ID: 12686545) with one of its answers and some of the comments under the answer. One comment criticizes a proposed answer as spammers can use it, e.g., for unethical purposes.

Human values have extensively been studied in the human-computer interaction field since the late 1980s [46]. However, the software engineering community has recently attempted to develop new software engineering practices and techniques or adapt the existing ones (e.g., the *value-based requirements engineering* method [42] and the *fairness-aware programming* technique [6]) to operationalize human values in software (e.g., [15, 20, 34]). Operationalizing human values in software is defined as “*the process of identifying human values and translating them to accessible and concrete concepts so that they can be implemented, validated, verified, and measured in software*” [39]. The ultimate goal is to develop software systems that better reflect and respect human values. Other lines of research aimed to understand which types of human values are discussed by developers in GitHub’s issue tracking systems [30] and analyzed app reviews to determine which human values are ignored or violated by apps [31, 40].

Inspired by [10, 19, 46], arguing that the codes developed or decisions made by developers may have a moral and ethical implication or a values implication, the objective of this study is to understand the human values-sensitive implications of developers’ solutions. Given Stack Overflow is the most active online platform for developers to share their programming issues and programming solutions (e.g., a code snippet or a design pattern as a solution), we focus on Stack Overflow. Each Stack Overflow post includes a question and a list of answers posted to the question [50]. Stack

Overflow provides mechanisms for users to comment on both questions and answers. This enables them to have further discussions (e.g., discuss the weaknesses of a solution) on the posted questions and answers.

In this study, we analyzed 2000 comments and their associated questions or answers (i.e., 1980 unique questions or answers). We first studied which types of human values in Schwartz’s theory of basic values are violated by Stack Overflow posts by manually analyzing comments. Then, we investigated the reactions of Stack Overflow users to human value violations. Key contributions of the work include:

- To our knowledge, the first detailed study of human values violations in Stack Overflow posts
- A large number of Stack Overflow comments are manually analyzed against the Schwartz human values framework
- Identification of 315 comments using the framework that raise concerns that their 313 unique associated posts violate human values
- A set of recommendations for practitioners and researchers to address human value violations in SO posts

We first introduce our research motivation and questions in Section 2. Section 3 provides the background and summarizes the related studies. In Section 4, we describe our data collection process, followed by reporting our findings in Section 5. We reflect on our findings in Section 6. Possible threats of our study are reported in Section 7. We conclude our paper in Section 8.

2 MOTIVATION AND RESEARCH QUESTIONS

Posts (questions and answers) in Stack Overflow may come with a number of issues. For example, a proposed solution in an answer might be obsolete [51] or even incorrect [50]. Stack Overflow provides different mechanisms such as downvoting [3] and commenting [2] to enable Stack Overflow users to indicate the possible issues associated with posts. Stack Overflow recommends that question and answer comments be used to provide constructive criticism, to request clarification for a question or an answer from the poster, or to add relevant information [2]. A question owner can also comment under their question, and comments under an answer can be added by the answer owner and the owner of its associated question. Any user with at least 50 reputation points can post comments under any question and answer [2, 50].

Previous research has shown that more than 50% of both hidden and displayed comments are informative and can enhance their associated answers [50]. Zhang et al. [50] classified comments in Stack Overflow in seven categories, in which comments falling in *advantage*, *improvement*, *weakness*, *inquiry*, and *addition* categories are considered informative. Hence, we argue that comments (in particular *improvement*, *weakness*, and *inquiry* comments) can be the ideal place to point out problems and weaknesses of a post from a human values perspective. The reason behind this is that these types of comments try to challenge a question or answer.

Figure 1 shows a question (**QUESTION ID: 12686545**) in Stack Overflow with one of its highly voted answers and some of the comments posted under the answer. As shown in Figure 1, a user posted a comment to criticize the proposed answer because spammers can use it, for example, to send thousands of unsolicited emails

to GitHub’s users resulting in upsetting or annoying them. Hence, we argue that this answer has a values implication (e.g., violating the value of *hedonism* - value item of *pleasure*). We developed the following research questions that we wanted to answer in this study:

RQ1. *Which types of human values are perceived to be violated in Stack Overflow posts?*

Motivation. As discussed, comments under a post can be used by Stack Overflow users to point out problems and weaknesses of a post from different perspectives, such as a human values perspective. This research question aims to identify types of values that are violated in Stack Overflow posts, and the frequency with which each value type is violated, by analyzing the claims made in Stack Overflow comments. This question can provide deep insights for Stack Overflow askers and answerers on potential human values-implications of their posts and the Stack Overflow community to develop mechanisms to recognize such posts.

RQ2. *How quick are commenters to raise concerns over human values violations?*

Motivation. A post (question or answer) might get several comments. Comments under a post are sorted and displayed by their creation time. As discussed earlier, the nature of comments varies from praising a post to criticizing and pointing out possible issues. This research question aims to understand how quickly comments citing human values violations are added to a post and determine their position among all comments under a post.

RQ3. *Are posts accused of violating human values downvoted by Stack Overflow users?*

Motivation. Apart from the commenting mechanism, the downvoting mechanism in Stack Overflow also enables the community to indicate the posts (questions and answers) with deficiencies, e.g., incorrect answers [3]. While voting down questions is free, voting down answers comes with some costs [3]. A limited number of downvotes can be made by a user (voter) per day. Voting down an answer diminishes two reputations from the post owner and one reputation from the voter. Hence, voting down an answer should be done carefully. In this research question, we want to know whether Stack Overflow users vote down posts that violate human values despite the associated costs.

RQ4. *How does the original poster react to concerns of potential human values violations?*

Motivation. Once human values violations occur in a software system, its creators are expected to react to such violations adequately (e.g., fixing values violations immediately). In this research question, we aim to understand the reactions of Stack Overflow posters (askers and answerers) once issues of potential human values violations are raised concerning their posts. Such reactions can range from completely denying the violation to modifying the post to mitigate the violation.

3 BACKGROUND AND RELATED WORK

3.1 Human Values in Software Engineering

Human values are the guiding principles for what people consider important in life [14]. Although these principles are often unarticulated using formal terminologies, they undergird people’s decisions, technologists and non-technical people alike. Hence, the influence

of human values can be detected in people’s preferences, from the choice of end-user applications [32], to the technical design decisions of developers in software engineering projects [48].

The study of human values in software engineering (SE) is often based on Schwartz’s theory of basic human values [33]. This theory organizes values into 10 broad values and is established on surveys conducted in multiple countries covering a wide range of ages, genders, occupations, cultural backgrounds, and geography [37]. Table 1 shows the 10 value categories. The 10 value categories comprise 58 value items, e.g., the value category of *benevolence* is comprised of the value items of *responsible, helpful, forgiving, honest, loyal, mature love, a spiritual life, meaning in life, and true friendship* (c.f. [37]). Additionally, the theory has been widely accepted and adopted in several areas, including the social sciences, computer science, and software engineering [33].

Recent research on human values in SE has underscored the need for software companies to directly cater to issues of human values in their software development processes [46], as the resulting software artefacts have a direct and indirect impact on end-users and society at large. Hussain et al. [21] argue that the maturity levels of companies in addressing human values may very well depend on their awareness and overall organisational culture. They propose that incorporating human values should be done through the evolution of already established software practices, i.e., adapting existing processes to include human values considerations, e.g., the inclusion of values in the development of personas, rather than through a revolution of the field of SE. Winter et al. [47] proposed the values Q-sort method for measuring human values in SE. Applying the values Q-sort method to 12 software engineers shows 3 “software engineer” values prototypes. In a similar study, Shams et al. [41] applied the portrait values questionnaire (PVQ) to 193 Bangladeshi female farmers to elicit their values. Their study reported *conformity* and *security* as the most important value categories while *power, hedonism, and stimulation* were the least important for Bangladeshi female farmers.

Other studies have applied indirect approaches by using app reviews as a proxy for eliciting values requirements. Shams et al. [40] analysed 1,522 reviews from 29 Bangladeshi agricultural apps to understand both the desired and missing values that should be addressed in the development of such apps. Furthermore, Obie et al. [31] introduced a dictionary-based natural language processing technique for detecting the violation of human values in app reviews. The result of their study showed that 26.5% of the analyzed 22,119 app reviews contained perceived violations of human values by the end-users. In addition, *benevolence* and *self-direction* were the most violated categories while *conformity* and *tradition* were the least violated categories.

As important steps towards addressing the violation of the values of mobile apps users in society, Obie et al. [31] proposed the mining of values requirements from rich data sources, the alignment of values between stakeholders in SE projects, and the adoption of critical technical practice in mobile SE. A recent study further contends that careful consideration of domain context in the design and application of values instruments should be made during values requirements gathering, as the hierarchy of end-users values may vary depending on the end-users domain context [32].

Table 1: Value categories and descriptions [37]

Value Category	Description
Self-direction	Independent thought and action - choosing, creating, exploring
Stimulation	Excitement, novelty, and challenge in life
Hedonism	Pleasure or sensuous gratification for oneself
Achievement	Personal success through demonstrating competence according to social standards
Power	Social status and prestige, control or dominance over people and resources
Security	Safety, harmony, and stability of society, of relationships, and of self
Conformity	Restraint of actions, inclinations, and impulses likely to upset or harm others and violate social expectations or norms
Tradition	Respect, commitment, and acceptance of the customs and ideas that one's culture or religion provides
Benevolence	Preserving and enhancing the welfare of those with whom one is in frequent personal contact
Universalism	Understanding, appreciation, tolerance, and protection for the welfare of all people and for nature

Furthermore, Mougouei [28] proposed a framework for accounting for human values at the level of source code. This framework established a relationship between human values and Android APIs and includes the following aspects: annotating APIs with the relevant human values, inspecting source code to detect potential sources of values violations, and recommending fixes to mitigate the violations. Building on the work of Mougouei [28], Li et al. [26] proffered 6 algorithms for detecting potential violation of values in 6 Android APIs. Their analysis applying these algorithms to 10,000 Android apps shows a correlation between the violation of human values and the presence of viruses in these apps.

As these studies have shown, the reflection of, support, and violation of human values may impact individual end-users and society as a whole. The research area of human values in SE is still in its early stages, and more work needs to be done. However, we present this work to further the discussion of human values in SE from the perspective of software developers as captured in their Stack Overflow posts.

3.2 Mining of Stack Overflow Posts

Question and Answer (Q & A) websites such as Stack Overflow are a rich source of information and provide insights into understanding developers' behaviours, interactions, and viewpoints on specific topics amongst others [43]. Several studies have mined Stack Overflow posts and comments to shed light on key areas. For example, Novielli et al. [29] focused on the social aspect of Stack Overflow and showed that the emotional lexicons in technical questions have an impact on the probability of obtaining satisfying responses to questions. Similarly, another study introduced multi-label classifiers for classifying the emotions encapsulated in Stack Overflow posts [12]. Wang et al. [45] analysed 100,000 questions from Stack Overflow to understand developer interactions on the platform.

Key findings from their study show that developers are keen on contributing to the community and not just getting their questions answered; developers extend a helping hand to others whether or not their gesture is reciprocated. Similarly, another study found that being prompt and being the first to respond to questions helps quickly build a reputation on Stack Overflow [11].

Other studies on Stack Overflow have focused on specific technical topics, e.g., Bangash et al. [9] analysed posts to understand what developers know about machine learning, while Bajaj et al. [8] reported on the challenges and misconceptions among web developers by mining posts related to HTML, CSS, and Javascript. Kavalier et al. [24] examined the relations between the Android marketplace usage of APIs to the occurrence of these APIs in Stack Overflow questions.

Some studies have relied on the rich dataset from Stack Overflow to develop tools for supporting software development. For instance, Ponzanelli et al. [35] introduced PROMPTER, an Eclipse plugin. Given a context in the IDE, PROMPTER automatically retrieves and analyses relevant discussions from Stack Overflow and then notifies the developer about available help. To support automatic source code documentation, Vassalo et al. [44] proposed CODES, a tool that extracts candidate method documentation from Stack Overflow discussions and creates Javadoc descriptions from it.

The studies discussed above have been vital in understanding the various themes discussed on Stack Overflow and have also shown Stack Overflow as a rich data source for understanding these varied themes in discussions related to developers and the software development practice. We build on this body of knowledge by investigating developers' discussion from the important lens of human values and how it affects society.

4 DATA COLLECTION

To understand human values violations in Stack Overflow and the possible reactions of its users to such violations, we first needed to identify a sufficient number of posts that have a values implication (e.g., violating a human value) (See Figure 1). In the first data collection step, we executed a random SQL SELECT query on the Stack Overflow publicly available dataset¹ with Google BigQuery. This dataset, hosted through the Google Cloud Public Dataset Program², is updated weekly by Stack Overflow and contains information about posts, comments, and voting, among other kinds of site activity. The SQL SELECT statement returned a random sample of 10000 comments and their corresponding posts, of which the first four authors (the analysts) individually manually analyzed the first 300. We imported the ID and content of these 300 comments and their corresponding posts into a spreadsheet. This spreadsheet was then shared between the analysts. The spreadsheet included 10 columns to enable the analysts to indicate which of the ten human value categories they judged were violated by the corresponding posts. The analysts then read each comment and its associated post (question or answer) to identify which of the ten value categories in Schwartz's theory (if any) were violated in the parent post according to the comment.

¹<https://tinyurl.com/4c74uz5n>

²<https://tinyurl.com/y2x9rzch>

Once the analysts finished this labeling process, they held several meetings and used a negotiated agreement method [13, 27] to resolve any disagreements and conflicts. Using the negotiated agreement method, all analysts collaboratively agreed on the label (coding) of an item under review. This approach is particularly useful for addressing reliability issues of codes when there are multiple categories as opposed to a binary category where a Cohen’s Kappa measure would suffice. At the end of this step, we found a very low prevalence of comments (8 of 300 comments) that raised concerns about at least one of Schwartz’s value categories.

In the second step, we designed a query to identify more posts that may potentially contain human values violations. To do so, we developed a list of regex keywords and phrases that are likely to be associated with human values violations concerns, such as “moral”, “ethical”, “human-cent”, and “society”. This list was designed based on our observations from the 300 comments and their corresponding posts labeled in the first step and consulting a dictionary of human values-related keywords and phrases developed in [31] for identifying human values violations in app user reviews. This list was adjusted over several query runs, reducing to 21 regex keywords, and was used to design a SQL SELECT query to identify comments that contained one of the keywords in the list. The SQL SELECT query returned 10144 comments. Note that keywords have also been used in previous studies (e.g., [31, 51]) to fine-tune data collection and minimize false positives. The list of 21 keywords and phrases is available in our replication package [25].

In the final step, we randomly selected 2000 comments from 10144 comments, which is well above a significance level of 99% and a significance interval of 3%. We use these 2000 comments and their associated posts (1980 unique posts) to answer our research questions.

5 FINDINGS

5.1 RQ1. Which types of human values are perceived to be violated in Stack Overflow posts?

Approach. To identify human values violations in Stack Overflow, we qualitatively analyzed 2,000 comments and their associated posts (1,980 unique questions or answers) collected in the Data Collection section (See Section 4). We used Schwartz’s theory of basic values [38], specifically Schwartz’s ten value categories (*self-direction, stimulation, hedonism, achievement, power, security, conformity, tradition, benevolence, universalism*), as a reference point to identify human values violations in Stack Overflow posts. This decision was made because Schwartz’s theory of basic values is the most widely used and cited values model in social science and software engineering [31, 33, 37]. The 10 value categories, in turn, comprise 58 value items. In this study, we focus on the 10 value categories. Note that the treatment of the 58 individual value items is beyond the scope of this work. However, where appropriate, we refer to the relevant value items associated with the value categories to increase the clarity of the results.

We first created a spreadsheet and shared it with all authors. The spreadsheet included 15 columns. The first four columns recorded the Comment ID and the content of the 2000 comments and their

associated post ID (question ID or answer ID) and the link to the posts. The next ten columns were the ten value categories. We also added a column called “remark” to allow the analysts to point out what they thought was important about a given comment/post.

The data analysis process was conducted in two steps. In the first step, the first author (the first analyst) followed an iterative process to label the first 1,000 comments and their associated posts. The first analyst selected approximately 100 comments and their associated posts and labeled them in each iteration. The reason behind investigating the associated posts was to thoroughly understand the context, meaning, and rationale behind comments. The first analyst was asked to indicate whether a comment discussing its associated post violated at least one human value. If so, they had to specify which of the ten value categories the given comment violated and put “1” in the corresponding columns in the spreadsheet. Comments could be labeled as violating more than one category of human value. After each iteration, three other authors (the validators) cross-checked the comments labeled in that iteration. In total, 400 comments (out of the 1,000 comments labeled by the first analyst) and their associated posts were cross-checked by the first validator, 400 by the second validator, and the rest by the third validator. This distribution was based on the availability of the validators. A negotiated agreement method [13, 27] was used to resolve conflicts and disagreements between the first analyst and the validators. The validators had extensive experience in human values and software engineering.

In the next step, the second analyst (the fifth author) labeled the rest of the comments (1,000 comments) and their corresponding posts. The second analyst conducted an iterative labeling process similar to the first analyst for this purpose. Then, the three validators in the previous step and the first analyst (acting as the fourth validator in this step) cross-checked the comments labeled by the second analyst in each iteration. In total, the first and second validators cross-checked 300 comments each, the third validator checked 100 comments, and the fourth validator cross-checked the rest. Similar to the previous step, the second analyst held several meetings with the validators to resolve disagreements and conflicts using the negotiated agreement method [13, 27].

Results. Our analysis of 2,000 Stack Overflow comments and their associated posts (1980 unique questions and answers) indicates that 315 comments (15.75%) complained their corresponding posts (313 unique posts) violated human values (See our replication package [25]). Out of 10 Schwartz theory’s value categories, we only found violations related to *self-direction, hedonism, security, conformity, tradition, benevolence, and universalism* in the 315 comments. Our analysis did not find any violations regarding *power, achievement, and stimulation*. The vast majority of the comments (270 out 315) include concerns that explained their corresponding post violated the value of *hedonism*. An example of a comment highlighting the violation of the value of *hedonism* (value item of *pleasure*) is:

👍 “Interesting approach and a possible solution, but **not very user-friendly**. The user won’t see the proper day values when he is moving the slider (unless I write another JS function to take care of that).” **COMMENT ID: 35155704**

The value of *benevolence* is the second most frequently violated value (reported in 41 comments). For example, a Stack Overflow user raised the issue of the violation of *benevolence* (value item of

responsibility) and criticized the irresponsibility of another user because their proposed solution does not care about the sensitivity of information but only about the cost of the solution.

“What I meant was that when it comes to sensitive information, your attitude of “SSL is too expensive” is **unethical and irresponsible**. When you have sensitive information in your hands, you must do everything you can to secure it. You say that you don’t see its worth in “common business cases”, but “common business cases” often involve sensitive information (addresses, phone numbers, email messages, trade secrets, etc). A business has much to lose in a breach of data integrity.” **COMMENT ID: 4919033**

We found only 10 comments in which Stack Overflow users complained that the proposed solutions violate *security*.

“Still, **passwords are private to the person that fills it in during registration. Not encrypting them is not very ethical, but I guess that’s another subject to discuss.**” **COMMENT ID: 3334429**

Eight comments were found complaining about violating *conformity*. In the following example, a user criticized the poster as their approach to terminate an app (proposed in a question, **QUESTION ID: 3318806**) violates Apple’s User Interface standards and guidelines.

“I think your app may get rejected if you terminate it within the app (unless due to unrecoverable error/fault handling). **Apple doesn’t like you to mess with their user experience, and pressing the Home button to exit/suspend an app is a big part of that user experience.**” **COMMENT ID: 3442304**

For *self-direction* and *universalism*, we found five and four examples of human values violation, respectively, and two for *tradition*. For example, the comment below is mapped to violation of the value of *self-direction* as its associated post violates the *freedom* and *independence* of end-users (i.e., *freedom* and *independence* are two value items in the value category of *self-direction*).

“Not to mention it strikes me as **ethically dicey at best to grab a user’s location without their permission.**” **COMMENT ID: 16973544**

5.2 RQ2. How quick are commenters to raise concerns over human values violations?

Approach. To answer this RQ, we extracted the date of the 315 comments about values violations and the date of their corresponding posts and calculated the time difference. To further understand how quickly comments about values violations are raised, we compared the position of a post’s comment that voices out about values violation with that of all of its associated comments.

Results. Figure 2 depicts the time differences in hours and days. It is shown that almost 55% of comments (173 out of 315) about values violations were received less than one hour after the corresponding posts were made. Our analysis shows that only 59 comments were raised 24 hours after the date of their corresponding posts. As shown in Figure 3, comments that raised concerns about human values mostly were the first or second comments (202 out of 315 comments, 64.12%) of their corresponding questions or answers. Only 25 comments about values violation appeared after the sixth comment.

5.3 RQ3. Are posts accused of violating human values downvoted by Stack Overflow users?

Approach. We developed an SQL Query to count the downvotes that were cast on the 313 unique posts associated with the 315 human values violation comments following the creation of those comments.

Results. We found that Stack Overflow users did not vote down most posts accused of violating human values violation. In fact, out of 313 posts with human values violation comments, only 74 (23.64%) posts were downvoted after comments complaining about human values violations were raised. Out of these 74 downvoted posts, most were downvoted once (46 posts), followed by twice (10 posts) and three times (10 posts). The rest were downvoted four times (7 posts) and five times (1 post).

For example, a user asked the question (**QUESTION ID: 6854611**): “How to return value when AJAX request is succeeded.” The accepted answer suggested setting parameter `async` to `false` to address the problem. However, a Stack Overflow user criticized this solution because it negatively impacts user experience. The accepted answer received three downvotes after the comment was made.

“`Async: false` will **LOCK UP THE BROWSER** for the duration of the ajax call. **This is often a horrible user experience.** The good solution requires refactoring the code to use a callback or a function call from the success/error functions to continue execution of the process and pass on the result of the ajax call.” **COMMENT ID: 8151749**

In another example, a user sought a solution to automatically turn on the user phone’s GPS as the given android app launches (**QUESTION ID: 17723347**). Another user indicated that such solutions violate *user privacy*. Such solutions can also violate the *self-direction* value (specifically, the value item of *independence*) as it denies the *independence* of the user to choose and set their security and privacy settings. This question got downvoted once.

“You can’t and you shouldn’t - it’s **ethically questionable to override something the user has set that has security/privacy considerations.**” **COMMENT ID: 25833312**

5.4 RQ4. How does the original poster react to concerns of potential human values violations?

5.4.1 RQ4.1 Does the original poster respond to concerns about potential human values violations in further comments? If so, how?

Approach. To understand if and how the original poster responds to a comment raising concerns about human values violation in a post, we conducted a qualitative study on all comments made by the original posters after the date of the 315 human value violations comments. We found that the original posters added 288 comments after the 315 human value violations comments. The first author applied the open coding technique [18] on all the 288 posters’ comments to decide which one(s) were added to respond to the reported human value violations and categorized the original posters’ responses (e.g., denying the post violates a human value). To reflect the distribution of original posters’ reaction on Stack Overflow, we limited tagging to at most one response comment one per human

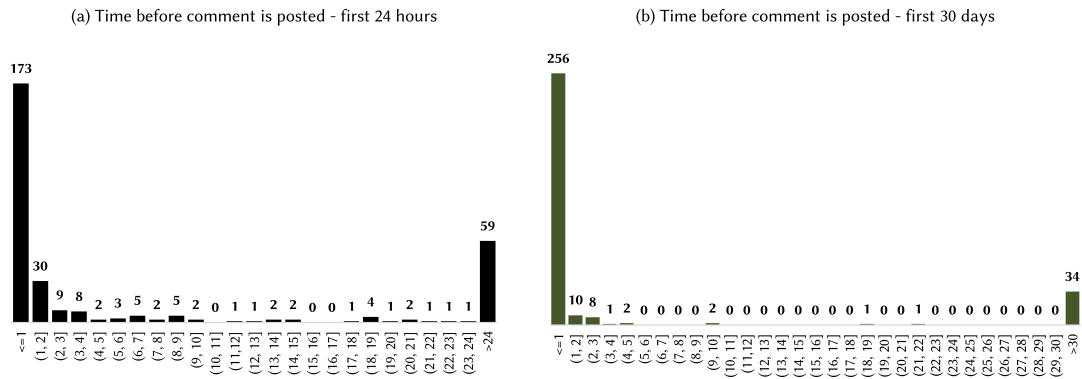


Figure 2: The time taken for concerns about violated human values to be raised (N=315)




Figure 3: Positions of values violations comments (N=315)

values violation comment; this way, each original poster’s response to their associated human values concern would be counted once, if, e.g., the poster created multiple comments after the concern was raised. In the next step, the outcomes of the open coding process (i.e., identified codes and categories) were shared with the second author for review. Then, the first and second authors held several Zoom meetings to discuss disagreements and inconsistencies and reached a consensus on the final list of codes and categories.

Results. Our analysis shows that out of the 288 comments added by original posters after a comment citing a human values violation, 140 did not address the violation concern. The remaining 148 comments addressed the human values violation concern directly, but the exact nature of the responses varied widely. The qualitative analysis of these 148 comments using open coding led to their classification as one of six response types.

The first four response types - namely “Acknowledging their violation”, “Adding redemptive detail”, “Proposing an alternative”, and “Asking for clarification” - were all responses, in which the original poster was receptive to commenter’s concern about human values violations in their post. Receptive responses totaled 105 of the 148 relevant responses.

Acknowledging their violation (n=44): In these responses, the original poster accepted the commenter’s concern wholly as described. This was often accompanied by an attempt to mitigate the violation, either by adapting their original proposed solution or by abandoning their approach altogether.

“@David, I think you are right. The user experience of editing a Rich TextBox within a DataGridView is rather bad, so I followed your hint on providing a separate edit mask. However, I’m also stuck here. See: <http://stackoverflow.com/questions/10224556/how-to-edit-a-dataset-in-a-new-form>.” COMMENT ID: 13134970 replied to COMMENT ID: 13112316

Adding redemptive detail (n=40): In this type of response, the original poster added further information about the context and/or requirements of their solution to explain why their post does not violate the human value claimed by the commenter.

“I need to do it. It is a special app that is meant to change desktop background, get system settings, disable icons, auto-hide task bar etc. It is used in one of the computer stores to display computer sale sticker directly on the screen.” COMMENT ID: 9302848 replied to COMMENT ID: 9302833

Proposing an alternative (n=13): In these responses, the original poster responded with an alternative solution in an attempt to mitigate the violation while still achieving their objective.

“In that case then you can queue users selection and updated it one by one, or if it is possible to stop the current operation, stop it and start the new fresh thread.” COMMENT ID: 28900232 replied to COMMENT ID: 28887851

Asking for clarification (n=8): In these responses, the original poster asked for further information about the commenter’s concern. These responses reflected a willingness to engage with the concern raised and to further inquire as to whether their solution is violating human values.

“@JPReddy: Sorry, I didn’t understand your problem, the changing must affect only the ComboBox control, so it must not affect any other cell in this column or other column, so can you explain more.” COMMENT ID: 4324670 replied to COMMENT ID: 4324160

The remaining two response types - “Denying” and “Conceding and pressing on with the issue” - were dismissive, rather than receptive, to the human values concerns raised earlier in the comment thread.

Denying (n=28): In these responses, the original poster denied that their post violated human values at all.

“@Stephen: No. ‘Expensive’ isn’t ‘unethical’. It’s just free market economy. The school didn’t have to give the job to that guy. It CHOSE to do so. They could always look for alternatives and choose the cheaper offer.” COMMENT ID: 1749489 replied to COMMENT ID: 1748956

Conceding and pressing on with the issue (n=15): In these responses, the original poster wholly accepted that there was a human values violation in their post but asserted that they were nonetheless going to persist with their approach.

“Preventing me from viewing a site simply because I’m on desktop because they want to feed me more ads and JS spam isn’t particularly ethical either. Oh well..” **COMMENT ID: 59368548** replied to **COMMENT ID: 59304750**

5.4.2 RQ4.2 Does the original poster modify their original post in light of concerns about human values violations? If so, how?

Approach. Some original posters may go beyond a comment response to a values violation claim, and may modify their original post itself in light of the claim. We first collected posts associated with the 315 comments which cited human values violations. Then we checked how many of them were modified after those comments were added. We found that 103 posts had at least one modification made on or after the time that the human values violation comment was created. Next, the first author manually checked all activities carried out on the 103 posts after the time of the 315 human value violations comments to understand which ones were related to the reported values violations.

Results. Out of 103 posts modified by its poster after the reported human values violations, only 14 posts were edited by the original poster in response to the claim about human values violations. The rest (89 posts) were modified for other reasons. These 14 posts can be grouped into two categories. Nine posts were edited to mitigate human value violations. For example, an original poster changed the source code in their post (**Post ID: 15135545**) after receiving criticism from a user who indicated the proposed solution “will cause the GUI to completely freeze, which is not a good user experience” **COMMENT ID: 21307680**.

In the other category of edit, the original poster edited their post to reinforce that no human values violation existed - this was true of five posts. For example, although an original poster edited their post (**Post ID: 16660109**), the solution in the post was not modified, and the original poster only added more information to clarify why there was no user experience issue with their solution.

6 DISCUSSION AND IMPLICATIONS

6.1 Prevalence of Hedonism Violations

The distribution of the types of values violations in our dataset was largely dominated by *hedonism* violations. The Schwartz model defines the *hedonism* value category as “*pleasure or sensuous gratification for oneself*”, and associated with it are the value items *pleasure, enjoying life, and self-indulgent* [38]. As such, any comment in the dataset which voiced concern about a post’s negative impact on users’ pleasure and experience would be tagged as a *hedonism* violation. Moreover, the value of a programming solution is largely determined by how the user feels when interacting with it, and user experience requirements are typically included within the technical requirements for software projects. Thus, when commenters voice concerns about a violation of *hedonism*, they may be addressing the technical requirements of the software rather than incidental, unintended consequences on human values. As a result, any objections on technical user experience grounds to a solution would be considered a complaint about a violation of the *hedonism* value category.

Implications: This finding raises the need to re-think the paradigms used for analyzing human values discourse in software contexts.

When the concern for human values falls directly in line with software technical requirements –as in the case of *hedonism*– it is important to account for the distinction of this dynamic from other cases where human values come at the expense of satisfying and delivering technical requirements. We claim that values conforming to technical goals is precisely the objective of values discourse: the more the software community voices their opinions about the importance of human values, the more those values become a fixed feature of software quality. Regardless, the dynamics between human values and technical requirements should be directly addressed in further research in these fields.

6.2 Reactions to Posts Violating Values

We investigated the reactions to posts that violated human values from the perspectives of Stack Overflow commenters, users, and the original posters. We found that commenters are quick in terms of raising concerns about human values violations. However, our study shows that Stack Overflow users did not downvote most posts (76.35%) accused of violating values. This may not be surprising as Stack Overflow users are encouraged to use upvotes and downvotes to report on the quality of the information in posts. A downvote on a question means “*this question does not show any research effort; it is unclear or not useful*”, and, on an answer, “*this answer is not useful*” [3]. As such, downvotes are not intended to reflect perceived violations of human values per se. A question post that is well-crafted and thought-out, while containing blatant human values violations, could avoid any downvotes, and so too with an answer post. Conversely, any downvotes on a post containing a human values violation may have nothing to do with the values violations in question but rather with the quality and usefulness of the post.

Nevertheless, it is useful to observe Stack Overflow users’ actual behavior in casting downvotes. Our results indicate that downvotes are, to some extent, cast in the wake of a human values violation being voiced in the comments. Indeed, it is possible that users are employing downvotes as a way to cast disapproval in light of perceived values violations, despite official site guidelines; how well users comply with site guidelines in their site activity is a major topic of discussion on the Meta Stack Exchange site [4].

Implications: Further research is needed to measure the significance of our findings on downvotes in human values contexts against general patterns of downvoting on Stack Overflow posts. This would shed light both on users’ reactions to perceived values violations on Stack Overflow and their regard for site voting guidelines in general.

6.3 Reactions to the Accusation of Human Values Violation in Posts

We observed while the original posters usually acknowledge commenters who criticize their posts, they tended to downplay the severity of the issue, either by minimising the impact of the violation, or by justifying their decision in spite of the severity of the issue. While being receptive is a good characteristic, we emphasize that the original posters need to mitigate values violations in their posts actively. This can potentially avoid the possible risks that such posts may have on the end-user and society. There is also the need for (Stack Overflow) developers to consider human

values and the potential violation of these values in the technical solutions that they proffer in these platforms. We also recommend the consideration of other currently less investigated human values such as *achievement*, *tradition*, and *conformity* (as categorised by Schwartz [37]) beyond the well-researched values of *privacy* and *security*.

Implications: Investigating the nature of values discourse on Stack Overflow could find relationships between the language and format used to voice concerns about value violations and the types of reactions they evoke from authors of the software in question. This would allow researchers to understand what makes users more or less receptive to criticism when it comes to the implications of their work on human values.

6.4 Towards an Automated Tool

Given our study is an exploratory study, we chose a manual approach to identify comments containing concerns about human values violations and categorize the response comments from the original posters to such values violations comments. This limited our sample size, as manual methods become more time-consuming as the amount of data scales up.

Implications: Automated approaches using machine learning and natural language processing methods should be developed to detect comments raising concerns about human values violations. In this study, we mainly used the contents of comments under posts to recognize posts violating values. AI-based techniques could leverage other features (e.g., votes, response comments by the original poster, the reputation score of the commenter and poster) to detect such posts and the types of values violated by these posts. Such automated methods will inform Stack Overflow users of a possible values implications of a post and let them decide if they want to use solutions proposed in the post in their software systems [7, 30].

7 THREATS TO VALIDITY

External Validity refers to what extent our findings can be generalized to other contexts [49]. This study collected and analyzed a random sample of only 2000 comments and their associated posts (questions or answers) from a dataset of comments and their associated posts described in Section 4. So, we acknowledge that our findings may not be generalized to all posts and comments in Stack Overflow and other question and answer websites such as Reddit³ and Gitter⁴. Further research needs to be conducted to explore how posts in other question and answer websites violate human values and how their users react to such violations.

Internal Validity is defined as threats that may have impacted our findings [49]. The random selection of the 2000 comments from a dataset of 10144 comments and the qualitative analysis processes conducted for RQ1 and RQ4 may have threatened our findings. First, our decision to build a dataset of 10144 comments from millions of comments in Stack Overflow was motivated to reduce the number of false-positive comments as much as possible. Furthermore, it was not possible for us to manually analyze all 10144 comments. Hence, we analyzed a random sample of 2000 comments and their corresponding posts. Therefore, we may have

missed some important types of value violations because of our dataset in Section 4 and the random selection of the 2000 comments from the dataset. Furthermore, some of the phrases used to build the dataset are particularly hedonism-related terms, which may have allowed the *hedonism* category to be over-represented in our dataset, resulting in a skew towards hedonism comments in RQ1.

The qualitative analysis processes to answer RQ1 and RQ4 might be subjective and error-prone. In RQ1, we employed two approaches to reduce these issues. First, the assigned analysts were asked to analyze the data iteratively (in each iteration, only 100 comments and their corresponding posts were analyzed by the analyst). Second, once each set of 100 comments and their associated posts were analyzed, the validators cross-checked these comments labeled by each analyst. Furthermore, several meetings were organized between the analysts and validators to resolve disagreements and conflicts using the negotiated agreement method [13, 27]. In RQ4, the second author checked all categories and their corresponding codes, and any disagreements and conflicts were resolved through meetings. In both RQ1 and RQ4, there were comments and their associated posts that made it difficult for us to precisely identify the type of human values violations (RQ1) or the type of the poster's responses (RQ4). In such cases, we labeled the comment as a non-human values comment (RQ1) or as an irrelevant response (RQ4) to avoid possible risks and mistakes. Hence, we can be reasonably confident that our findings are credible with minimum mislabelled comments.

Construct Validity. In RQ1, our decision to use ten value categories in the Schwartz theory may have introduced two threats. First, other values models such as Rokeach's Value Survey [36] and List of Values [23] with different types and numbers of values could be used instead of the Schwartz theory. While none of them have been developed for software engineering, the Schwartz theory is widely used in software engineering (e.g., [30, 33]). Furthermore, the definition of human values might have been vague for the analysts and validators. So, they might have struggled to map Stack Overflow comments to human values. To mitigate this threat, apart from reading seminal papers [37, 38] on the Schwartz theory, we consulted the previous research that leveraged (the definition of) human values in the software engineering context, e.g., app reviews [31, 40], GitHub issue discussions [30], and source codes [28]. Finally, the quantitative measures used in RQ2, RQ3, and R4 to determine the characteristics of comments citing values violations or reactions to posts being accused of values violations would not capture all aspects of these comments and posts. Future research is encouraged to further characterize these comments and posts.

8 CONCLUSION AND FUTURE WORK

In this work, we conducted an exploratory study investigating the potential violation of human values in Stack Overflow. Adopting the widely accepted Schwartz model of basic human values, we analyzed 2,000 Stack Overflow comments and their associated posts (1980 unique questions or answers) to identify posts and comments containing perceived human values violations, the categories of the values violated, and the reactions of Stack Overflow users to concerns related to these violations. Our results show that 315 (out of 2,000) comments raised issues concerning the violation of

³<https://www.reddit.com/>

⁴<https://gitter.im/>

7 out of the 10 value categories in the Schwartz model. We find that Stack Overflow commenters react quickly to issues of values violations; 203 (out of 315) comments raising the concerns of values violations were made less than 2 hours after the corresponding posts. Also, most posts (76.35%) accused of human values violation did not get downvoted at all. Furthermore, only 148 of the original posters responded to the concerns of values violations made by other commenters in follow-up comments of their own.

In the future, we plan to build upon our exploratory study by diving deeper into specific value categories and their associated value items to understand the different factors that cause their violations. In addition, due to the limitations of a manual approach to categorizing values and their violations, we plan to build machine learning models to automate this process.

ACKNOWLEDGMENTS

Support for this work from ARC Laureate Program FL190100035 and Discovery Project DP200100020 is gratefully acknowledged.

REFERENCES

- [1] [n. d.]. *AI Incident Database*. <https://incidentdatabase.ai/>
- [2] [n. d.]. Comment Everywhere. <https://stackoverflow.com/help/privileges/comment>
- [3] [n. d.]. Vote Down. <https://stackoverflow.com/help/privileges/vote-down>
- [4] [n. d.]. When is it justifiable to downvote a question? <https://tinyurl.com/54eep4rz>
- [5] 2021. Facebook Apologizes After A.I. Puts ‘Primates’ Label on Video of Black Men. <https://tinyurl.com/ytynjntn>
- [6] Aws Albarghouthi and Samuel Vinitzky. 2019. Fairness-aware programming. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 211–219.
- [7] Eman Abdullah AlOmar, Wajdi Aljedaani, Murtaza Tamjeed, Mohamed Wiem Mkaouer, and Yasmine N El-Glaly. 2021. Finding the Needle in a Haystack: On the Automatic Identification of Accessibility User Reviews. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [8] Kartik Bajaj, Karthik Pattabiraman, and Ali Mesbah. 2014. Mining Questions Asked by Web Developers. In *Proceedings of the 11th Working Conference on Mining Software Repositories (Hyderabad, India) (MSR 2014)*. ACM, 112–121.
- [9] Abdul Ali Bangash, Hareem Sahar, Shaiful Chowdhury, Alexander William Wong, Abram Hindle, and Karim Ali. 2019. What do Developers Know About Machine Learning: A Study of ML Discussions on StackOverflow. In *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. 260–264.
- [10] Grady Booch. 2014. The human and ethical aspects of big data. *IEEE software* 31, 1 (2014), 20–22.
- [11] Amianguhu Bosu, Christopher S. Corley, Dustin Heaton, Debarshi Chatterji, Jeffrey C. Carver, and Nicholas A. Kraft. 2013. Building reputation in StackOverflow: An empirical investigation. In *2013 10th Working Conference on Mining Software Repositories (MSR)*. 89–92.
- [12] Luis Adrián Cabrera-Diego, Nik Bessis, and Ioannis Korkontzelos. 2020. Classifying emotions in Stack Overflow and JIRA using a multi-label approach. *Knowledge-Based Systems* 195 (2020), 105633.
- [13] John L Campbell, Charles Quincy, Jordan Osserman, and Ove K Pedersen. 2013. Coding in-depth semistructured interviews: Problems of unitization and inter-coder reliability and agreement. *Sociological Methods & Research* 42, 3 (2013), 294–320.
- [14] An-Shou Cheng and Kenneth R Fleischmann. 2010. Developing a meta-inventory of human values. *Proceedings of the American Society for Information Science and Technology* 47, 1 (2010), 1–10.
- [15] Maria Angela Ferrario, Will Simm, Stephen Forshaw, Adrian Gradinar, Marcia Tavares Smith, and Ian Smith. 2016. Values-first SE: research principles in practice. In *2016 IEEE/ACM 38th International Conference on Software Engineering Companion (ICSE-C)*. IEEE, 553–562.
- [16] James Fieser. 2016. Ethics. *The Internet Encyclopedia of Philosophy*. ISSN 2161–0002.
- [17] Sainyam Galhotra, Yuriy Brun, and Alexandra Meliou. 2017. Fairness testing: testing software for discrimination. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*. 498–510.
- [18] Barney G Glaser, Anselm L Strauss, and Elizabeth Strutzel. 1968. The discovery of grounded theory; strategies for qualitative research. *Nursing research* 17, 4 (1968), 364.
- [19] DW Gotterbarn, Bo Brinkman, Catherine Flick, Michael S Kirkpatrick, Keith Miller, Kate Vazansky, and Marty J Wolf. 2018. *Acm code of ethics and professional conduct*. (2018).
- [20] Maaïke Harbers, Christian Detweiler, and Mark A Neerincx. 2015. Embedding stakeholder values in the requirements engineering process. In *International Working Conference on Requirements Engineering: Foundation for Software Quality*. Springer, 318–332.
- [21] Waqar Hussain, Harsha Perera, Jon Whittle, Arif Nurwidyantoro, Rashina Hoda, Rifat Ara Shams, and Gillian Oliver. 2020. Human Values in Software Engineering: Contrasting Case Studies of Practice. *IEEE Transactions on Software Engineering* (2020), 1–15.
- [22] David Ingold and Spencer Soper. 2016. Amazon Doesn’t Consider the Race of Its Customers. Should It? <https://tinyurl.com/tryfrxwu>
- [23] Lynn R Kahle and Patricia Kennedy. 1988. Using the list of values (LOV) to understand consumers. *Journal of Services Marketing* (1988).
- [24] David Kavalier, Daryl Posnett, Clint Gibler, Hao Chen, Premkumar Devanbu, and Vladimir Filkov. 2013. Using and Asking: APIs Used in the Android Market and Asked about in StackOverflow. In *Social Informatics*. Adam Jatowt, Ee-Peng Lim, Ying Ding, Asako Miura, Taro Tezuka, Gaël Dias, Katsumi Tanaka, Andrew Flanagan, and Bing Tian Dai (Eds.). Springer International Publishing, Cham, 405–418.
- [25] Sara Krishtul, Mojtaba Shahin, Humphrey O. Obie, Hourieh Khalajzadeh, Fan Gai, Ali Rezaei Nasab, and John Grundy. [n. d.]. The replication package of the paper. <https://doi.org/10.5281/zenodo.5919444> October, 2021.
- [26] Conghui Li, Humphrey O. Obie, and Hourieh Khalajzadeh. 2021. A First Step Towards Detecting Values-violating Defects in Android APIs. arXiv:2109.14359 [cs.SE]
- [27] Elizabeth R Morrissey. 1974. Sources of error in the coding of questionnaire data. *Sociological Methods & Research* 3, 2 (1974), 209–232.
- [28] Davoud Mougouei. 2020. Engineering Human Values in Software through Value Programming. *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops* (2020), 133–136.
- [29] Nicole Novielli, Fabio Calefato, and Filippo Lanubile. 2014. Towards Discovering the Role of Emotions in Stack Overflow. In *Proceedings of the 6th International Workshop on Social Software Engineering (Hong Kong, China) (SSE 2014)*. 33–36.
- [30] Arif Nurwidyantoro, Mojtaba Shahin, Michel RV Chaudron, Waqar Hussain, Rifat Shams, Harsha Perera, Gillian Oliver, and Jon Whittle. 2021. Human values in software development artefacts: A case study on issue discussions in three Android applications. *Information and Software Technology* (2021), 106731.
- [31] Humphrey O. Obie, Waqar Hussain, Xin Xia, John Grundy, Li Li, Burak Turhan, Jon Whittle, and Mojtaba Shahin. 2021. A First Look at Human Values-Violation in App Reviews. In *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*. IEEE, 29–38.
- [32] Humphrey O. Obie, Mojtaba Shahin, John Grundy, Burak Turhan, Li Li, Waqar Hussain, and Jon Whittle. 2021. Does Domain Change the Opinion of Individuals on Human Values? A Preliminary Investigation on eHealth Apps End-users. arXiv:2110.01832 [cs.SE]
- [33] Harsha Perera, Waqar Hussain, Jon Whittle, Arif Nurwidyantoro, Davoud Mougouei, Rifat Ara Shams, and Gillian Oliver. 2020. A study on the prevalence of human values in software engineering publications, 2015–2018. In *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*. IEEE, 409–420.
- [34] Harsha Perera, Gunter Mussbacher, Waqar Hussain, Rifat Ara Shams, Arif Nurwidyantoro, and Jon Whittle. 2020. Continual human value analysis in software development: A goal model based approach. In *2020 IEEE 28th International Requirements Engineering Conference (RE)*. IEEE, 192–203.
- [35] Luca Ponzanelli, Gabriele Bavota, Massimiliano Di Penta, Rocco Oliveto, and Michele Lanza. 2014. Mining StackOverflow to Turn the IDE into a Self-Confident Programming Prompter. In *Proceedings of the 11th Working Conference on Mining Software Repositories (Hyderabad, India) (MSR 2014)*. ACM, 102–111.
- [36] Milton Rokeach. 1973. *The nature of human values*. Free press.
- [37] Shalom H Schwartz. 1992. Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. In *Advances in experimental social psychology*. Vol. 25. Elsevier, 1–65.
- [38] Shalom H Schwartz. 2012. An overview of the Schwartz theory of basic values. *Online readings in Psychology and Culture* 2, 1 (2012), 2307–0919.
- [39] Mojtaba Shahin, Waqar Hussain, Arif Nurwidyantoro, Harsha Perera, Rifat Shams, John Grundy, and Jon Whittle. 2021. Operationalizing Human Values in Software Engineering: A Survey. arXiv preprint arXiv:2108.05624 (2021).
- [40] Rifat Ara Shams, Waqar Hussain, Gillian Oliver, Arif Nurwidyantoro, Harsha Perera, and Jon Whittle. 2020. Society-oriented applications development: Investigating users’ values from bangladeshi agriculture mobile applications. In *2020 IEEE/ACM 42nd International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*. IEEE, 53–62.
- [41] Rifat Ara Shams, Mojtaba Shahin, Gillian Oliver, Waqar Hussain, Harsha Perera, Arif Nurwidyantoro, and Jon Whittle. 2021. Measuring Bangladeshi Female Farmers’ Values for Agriculture Mobile Applications Development. In *54th Hawaii International Conference on System Sciences, HICSS’21*. 1–10.

- [42] Sarah Thew and Alistair Sutcliffe. 2018. Value-based requirements engineering: method and experience. *Requirements engineering* 23, 4 (2018), 443–464.
- [43] Christoph Treude, Ohad Barzilay, and Margaret-Anne Storey. 2011. How Do Programmers Ask and Answer Questions on the Web? (NIER Track). In *Proceedings of the 33rd International Conference on Software Engineering* (Waikiki, Honolulu, HI, USA) (*ICSE '11*). ACM, 804–807.
- [44] Carmine Vassallo, Sebastiano Panichella, Massimiliano Di Penta, and Gerardo Canfora. 2014. CODES: Mining Source Code Descriptions from Developers Discussions. In *Proceedings of the 22nd International Conference on Program Comprehension* (Hyderabad, India) (*ICPC 2014*). ACM, 106–109.
- [45] Shaowei Wang, David Lo, and Lingxiao Jiang. 2013. An Empirical Study on Developer Interactions in StackOverflow. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing* (Coimbra, Portugal) (*SAC '13*). ACM, 1019–1024.
- [46] Jon Whittle, Maria Angela Ferrario, Will Simm, and Waqar Hussain. 2021. A Case for Human Values in Software Engineering. *IEEE Software* 38, 1 (2021), 106–113.
- [47] Emily Winter, Steve Forshaw, and Maria Angela Ferrario. 2018. Measuring Human Values in Software Engineering. In *2018 ACM/IEEE 12th International Symposium on Empirical Software Engineering and Measurement*. 1–4.
- [48] Emily Winter, Stephen Forshaw, Lucy Hunt, and Maria Angela Ferrario. 2019. Advancing the Study of Human Values in Software Engineering. In *12th International Workshop on Cooperative and Human Aspects of Software Engineering* (Montreal, Quebec, Canada) (*CHASE '19*). 19–26.
- [49] Claes Wohlin, Per Runeson, Martin Höst, Magnus C Ohlsson, Björn Regnell, and Anders Wesslén. 2012. *Experimentation in software engineering*. Springer Science & Business Media.
- [50] Haoxiang Zhang, Shaowei Wang, Tse-Hsun Chen, and Ahmed E Hassan. 2019. Reading answers on stack overflow: Not enough! *IEEE Transactions on Software Engineering* (2019).
- [51] Haoxiang Zhang, Shaowei Wang, Tse-Hsun Peter Chen, Ying Zou, and Ahmed E Hassan. 2019. An empirical study of obsolete answers on Stack Overflow. *IEEE Transactions on Software Engineering* (2019).