



Automatic prediction of rejected edits in Stack Overflow

Saikat Mondal¹ · Gias Uddin² · Chanchal Roy¹

Accepted: 14 September 2022 / Published online: 24 November 2022

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

Abstract

The content quality of shared knowledge in Stack Overflow (SO) is crucial in supporting software developers with their programming problems. Thus, SO allows its users to suggest edits to improve the quality of a post (i.e., question and answer). However, existing research shows that many suggested edits in SO are rejected due to undesired contents/formats or violating edit guidelines. Such a scenario frustrates or demotivates users who would like to conduct good-quality edits. Therefore, our research focuses on assisting SO users by offering them suggestions on how to improve their editing of posts. First, we manually investigate 764 (382 questions + 382 answers) rejected edits by rollbacks and produce a catalog of 19 rejection reasons. Second, we extract 15 texts and user-based features to capture those rejection reasons. Third, we develop four machine learning models using those features. Our best-performing model can predict rejected edits with 69.1% precision, 71.2% recall, 70.1% F1-score, and 69.8% overall accuracy. Fourth, we introduce an online tool named *EditEx* that works with the SO edit system. *EditEx* can assist users while editing posts by suggesting the potential causes of rejections. We recruit 20 participants to assess the effectiveness of *EditEx*. Half of the participants (i.e., treatment group) use *EditEx* and another half (i.e., control group) use the SO standard edit system to edit posts. According to our experiment, *EditEx* can support SO standard edit system to prevent 49% of rejected edits, including the commonly rejected ones. However, it can prevent 12% rejections even in free-form regular edits. The treatment group finds the potential rejection reasons identified by *EditEx* *influential*. Furthermore, the median workload suggesting edits using *EditEx* is *half* compared to the SO edit system.

Keywords Stack overflow · Rejected edits · Classification model · User study · Tool support

Communicated by: Neil Ernst

This article belongs to the Topical Collection: *Registered Reports*

✉ Saikat Mondal
saikat.mondal@usask.ca

Extended author information available on the last page of the article.

1 Introduction

The adoption, growth, and continued success of an online question and answering (Q&A) site such as Stack Overflow (SO) depend on two major factors—(1) participation of users and (2) quality of the shared knowledge (Bagozzi and Dholakia 2006; Lakhani and von Hippel 2003; Parnin et al. 2012). SO thus introduces an edit system to promote quality by allowing its users to communicate on the quality of the posts through editing. In particular, collaborative editing helps to keep posts clear, relevant, and up-to-date. For example, users often edit posts to fix grammar and spelling mistakes, clarify the meaning, and add related resources or hyperlinks. Unfortunately, many suggested edits in SO get rejected because of undesired (i.e., it does not satisfy the post owner) editing or violating edit guidelines (Overflow 2015). However, edits can be rejected in two ways—*rollback* and *expert review*. Rollback reverts a post to a previous version in the edit history (Exchange 2009b) and thereby rejects one/multiple revisions. On the other hand, experts (e.g., users with a reputation score $\geq 2K$) can reject if suggested edits do not improve the quality of the posts. However, manual identification of undesired edits or edits that violate the editing guidelines wastes community time and effort. For example, one user responded to the issue of the identification of undesired edits manually, *“It takes time to read and parse through those questions when I am trying to spend my time more efficiently reading through the actual question and figuring out how to answer it appropriately”* (Exchange 2009a). At least 921 users supported this comment by casting upvotes. It suggests that manually identifying undesired edits wastes users’ valuable time and resources. On the other hand, users who suggest edits and later get rejected become frustrated because many users (especially novices) are unaware of editing guidelines (Mondal et al. 2021a). Unfortunately, the existing editing system of SO does not identify the rejected edits with the potential rejection reasons. Therefore, a study on automatic identification of rejected edits with reasons is warranted to assist SO users.

Realizing the need for an automated tool, some users started writing personal scripts to identify undesired edits programmatically. For example, one user wrote a script to automatically identify greetings (e.g., hello, dear) while reviewing suggested edits (Exchange 2009a). Such a scenario urges a system that identifies the potential rejected edits. However, capturing all rejection reasons using simple rule-based scripts is challenging. Thus, a robust technique (e.g., machine learning classifiers) needs to be introduced to reasonably identify rejected edits and the potential reasons behind those rejections. Wang et al. (2018) investigate the rejected edits in SO. They analyze 369 rejected edits (by rollbacks) of answers and identify 12 reasons (e.g., undesired text formatting). Their study shows the empirical evidence of the complexity and diversity of reasons that can contribute to the rejection of suggested edits. However, we are unaware of any existing edit assistance system that automatically identifies rejected edits with reasons to support the current editing system of SO.

This study focuses on assisting SO users by offering them automated suggestions on how to improve their editing of posts. First, we manually analyzed 764 rejected edits (382 questions + 382 answers). We identified 19 rejection reasons (Table 1), seven of which were not reported by Wang et al. (2018). Second, we extract 15 texts and user-based features to capture those rejection reasons. Third, using those features, we develop four machine learning classifiers (e.g., random forest). According to the experiment, our best-performing

Table 1 Summary of the *nineteen* manually derived reasons behind rollback edits

| Rollback reasons | Description | Question | Answer | Wang et al. | Count (%) |
|---------------------------|---|----------|--------|-------------|-------------|
| Undesired Text Formatting | Users change the format of the texts unnecessarily. Such changes include changing the font, text cases (uppercase/lowercase), emphasizing text by making them bold/italic, adding or removing spaces/newlines, creating bullet/number list, and formatting text term as code element or vice versa. | ✓ | ✗ | ✓ | 157 (20.5%) |
| Undesired Text Add/Remove | Users add texts with less or no impact on the quality/clarification of posts or remove essential texts. | ✓ | ✓ | ✓ | 223 (29.2%) |
| Undesired Text Change | Users make undesired changes of the sentence structures (e.g., simple, complex), tenses (e.g., present, past), voices (e.g., active, passive), rewording, interchanging contractions by root words, acronyms/abbreviations by elaborations and vice versa. | ✓ | ✓ | ✓ | 85 (11.1%) |
| Incorrect Text Change | Users perform rewording with incorrect terms, grammatical and spelling mistakes, incorrect changes in software versions or specifications, changes that alternate the meaning of the sentence. | ✓ | ✓ | ✓ | 86 (11.3%) |
| Undesired Code Formatting | Users make the undesired modification of code indentation (e.g., addition/removal of spaces/newlines), addition/removal of line numbers, split/merge of code segments, changes in text cases (e.g., <i>select</i> to <i>SELECT</i> in a SQL query). | ✓ | ✓ | ✓ | 74 (9.7%) |
| Undesired Code Add/Remove | Unwanted code statements (e.g., alternative solutions) are added, or essential segments are removed. | ✓ | ✓ | ✓ | 89 (11.6%) |
| Undesired Code Change | Users make undesired code changes, such as changing options of a command, changing APIs, refactoring (e.g., variable renaming), and editing comments. | ✓ | ✓ | ✓ | 24 (3.1%) |

Table 1 (continued)

| Rollback reasons | Description | Question | Answer | Wang et al. | Count (%) |
|-----------------------------|--|----------|--------|-------------|------------|
| Incorrect Code Change | Users make incorrect changes to the data type of variables, function return types, function arguments, arithmetic expressions. | ✓ | ✓ | ✓ | 58 (7.6%) |
| Status Update | Users add/remove personal notes to clarify confusion, append messages missed during the submission of their posts, and acknowledge users' responses. | ✓ | ✓ | ✗ | 77 (10.1%) |
| Emotion Add/Remove | Addition/removal of words/sentences/emoticons that represent personal emotion. | ✓ | ✓ | ✓ | 6 (0.8%) |
| Gratitude Add/Remove | Addition/removal of thanksgiving sentences (e.g., thank you, cheers!). | ✓ | ✓ | ✗ | 42 (5.5%) |
| Greetings Add/Remove | Addition/removal of greeting/salutations (e.g., hello, dear). | ✓ | ✗ | ✗ | 4 (0.5%) |
| Undesired Ref. Modification | Users add inactive hyperlinks, inappropriate images, or diagrams with posts. On the contrary, sometimes they remove essential hyperlinks/images or unreasonably modify them. | ✓ | ✓ | ✓ | 85 (11.1%) |
| Signature Add/Remove | Addition/removal of users' names, hyperlinks to personal websites. | ✓ | ✓ | ✗ | 14 (1.8%) |
| Partial Acceptance | Revision is rolled back, but part of the changes are still accepted. Then, the accepted changes are included in later revisions. | ✓ | ✓ | ✓ | 4 (0.5%) |
| Deprecation Note Add/Remove | Addition/removal of deprecation notes inside the body of an answer. | ✗ | ✓ | ✗ | 1 (0.1%) |
| Duplication Note Add/Remove | Addition/removal of duplication notes inside the body of a question. | ✓ | ✗ | ✗ | 11 (1.4%) |
| Community Trust | The reputation score estimates how much the community trusts a user. Users with low reputations often do not follow the guidelines when editing posts. Thus, their edits are rejected more than the highly reputed users. | ✓ | ✓ | ✗ | N/A |
| Other | Other reasons include asking questions inside the answer, adding solutions inside the question, interchanging the position of texts, introducing spam (i.e., users deface text/code to promote a product or service, insert garbage texts), and changing posts completely. | ✓ | ✓ | ✓ | 40 (5.2%) |

model can predict rejected edits with 69.1% precision, 71.2% recall, 70.1% F1-score, and 69.8% accuracy. Fourth, we introduce an online tool named *EditEx* that works with the SO edit system. *EditEx* can assist users while editing posts by identifying the rejected edits and the potential causes of rejections.

Figure 1 shows the overview of *EditEx* workflow. The workflow is described as follows. (1) On the client-side, users need to install Tampermonkey,¹ one of the most popular userscript managers. It offers an effortless way to manage userscripts. In addition, it is available as a browser extension for all the popular browsers, such as Chrome, Firefox, Safari, Microsoft Edge, and Opera. (2) Users require to add two JavaScript scripts for the *EditEx* and *Suggestion* interface with the SO edit system. Tampermonkey enables users to add JavaScript scripts that can be used to modify web pages. (3) *EditEx* enables users to edit the posts. (4) *Suggestion* captures the texts (before & after edits) and user information (e.g., reputation score, name). It then transmits texts and user information to the server where the machine learning classifiers have been deployed. The server-side application extracts the features. It then predicts whether the edit will be rejected or not and identify the potential reasons (if rejected) using the machine learning classifiers. (5) Finally, it shows the decision (rejected/accepted) of the classifier and suggests to users with the potential reasons if rejected.

We recruited 20 participants and divided them into treatment and control groups. The treatment group used *EditEx* and the control group used the SO standard edit system to edit posts. We survey the participants after completing their edits. According to survey results, the treatment group found the potential reasons identified by *EditEx* behind the rejections *influential*. Moreover, 49% rejections (including the commonly rejected ones) were prevented upon following the suggestions offered by *EditEx*. However, *EditEx* is also capable of preventing 12% rejections even in *free-form* regular edits. Free-form edits refer to those edits that are not related to any specific rejection reasons. This tool also significantly reduces the effort of suggesting edits and makes participants more confident. The following major stakeholders in crowd-sourced knowledge-sharing platforms that use collaborative editing features can be benefited from the findings from our study and the tool *EditEx*: (a) forum like SO designers to improve the edit system, (b) forum users to assist their edit behavior, and (c) software engineering researchers to study and improve collaborative editing support in crowd-shared platforms.

Deviations of this study from our registered report (Mondal et al. 2020) are discussed in Section 7.

Structure of the Article The rest of this article is structured as follows. Section 2 represents a catalog of edit rejection reasons. In Section 3, we discuss a model that predicts rejected edits with potential reasons. Section 4 discusses our online tool named *EditEx*, its architecture, and analyzes its effectiveness. Feature ranking, the reasons behind the misclassifications of our model, and the implications of our study are discussed in Section 5. Section 6 focuses on the threats to validity, Section 8 discusses the related work and finally, Section 9 concludes our study.

¹<https://www.tampermonkey.net>

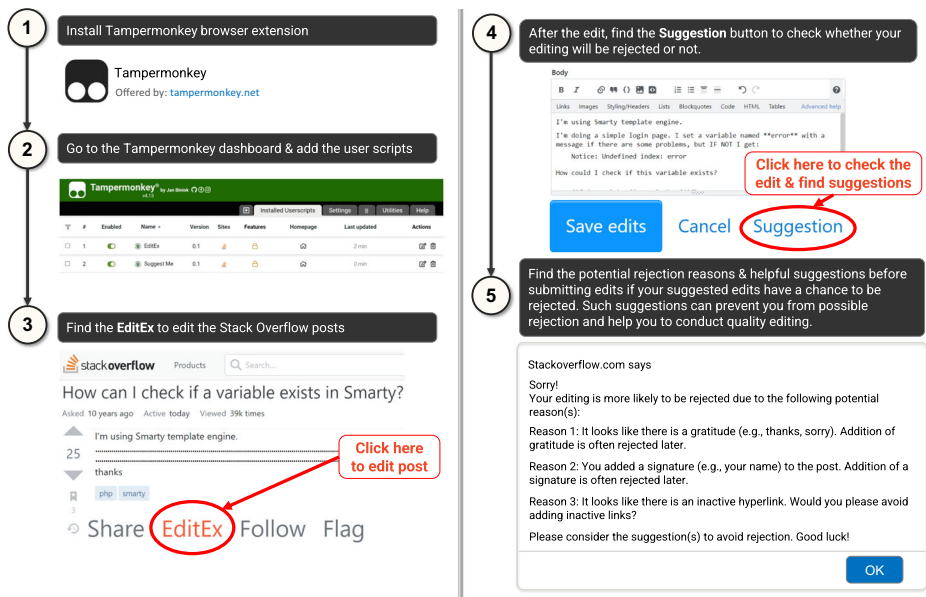


Fig. 1 An overview of the EditEx workflow

2 A Catalog of Edit Rejection Reasons

SO introduces an edit system to improve the quality of the posts. However, edits may not always be satisfactory and thus can be rejected. Wang et al. (2018) conduct an initial investigation on rejected edits of answers by rollback and expose 12 potential reasons behind rejections. We extend their investigation by manually investigating rejected edits by rollbacks of questions and answers. This section discusses our manual investigation process first and later summarizes the identified edit rejection reasons.

2.1 Dataset Preparation

We downloaded the September 2019 (which was the start time of the journey of this study) data dump of SO from the Stack Exchange site (Exchange 2019). The data dump stores the history of all events (e.g., rollback and edit body) of the posts. In this study, we only investigate the revisions where users made edits to the body of posts. Our data dump contains a total of 1,116,473 rejected edits (72,159 questions + 44,314 answers) by rollbacks and 26,604,779 accepted edits (13,624,495 questions + 12,980,284 answers). We manually analyze rejected edits to explore the rejection reasons. However, accepted edits will be used later to train and test our machine learning classifiers in Section 3.2. We randomly sampled a statistically significant sample size from rejected and accepted edits.

To achieve a confidence level of 95% with a confidence interval of 5% (Boslaugh 2012), we randomly sampled—(1) 382 from 72,159 rollback edits of questions, (2) 382 from 44,314 rollback edits of answers, (3) 385 from 13,624,495 accepted edits of questions,

and (4) 385 from 12,980,284 accepted edits of answers. We use the following formula to compute the size of our random sample.

$$\frac{Nz^2p(1-p)}{e^2N + z^2p(1-p)} \quad (1)$$

where N is the population size (e.g., 72,159), z is the Z-score corresponding to a particular confidence level (e.g., 1.96 for a confidence level of 95%), e is the confidence interval (e.g., 5%), and p is population proportion (e.g., 0.5) (Wang et al. 2018).

The data dump only stores the latest contents of the questions and answers after each revision. However, we need both the texts before and after rollback/acceptance to analyze which suggested edits are rejected/accepted. Therefore, we collect the `PostId` and `RevisionGUID` for each of our randomly selected revisions from the data dump. We manually save the web pages using `PostId` from the SO site containing each post's revision history. Next, we find the target revision from the history using the `RevisionGUID`. Note that `RevisionGUID` is a unique ID used to find a particular revision. Each revision contains the text before and after rollback/acceptance. Finally, we extract those target texts.

2.2 Edit Rejection Reasons

Table 1 summarizes the rollback edit reasons. We (two authors of this paper) manually investigate the randomly selected 764 (382 questions + 382 answers) rejected edits by rollbacks. We consider the rollback reasons identified by Wang et al. (2018) as the baseline during our analysis. However, we discuss the rollback edit reasons in multiple interactive sessions. We then analyze 200 rollback edits (100 questions + 100 answers) from our selected dataset and label the reasons. For a given rollback edit, we meticulously analyze the texts before and after rollback to see the edits that cause a rollback. Our in-depth investigation exposes a total of *nineteen* potential reasons. *Twelve* of them were identified by Wang et al., and the remaining *seven* are new reasons. The new reasons are—(1) status update, (2) gratitude add/remove, (3) greetings add/remove, (4) signature add/remove, (5) deprecation note add/remove, (6) duplication note add/remove, and (7) community trust. We then measure the agreement using Cohen's Kappa (Cohen 1968; 1960). The value of κ was 0.98, which means the strength of the agreement is almost perfect. Next, we resolve the remaining few disagreements by discussion. However, the agreement level indicates that any coder can do the rest of the labeling without introducing individual bias. Thus, the first author of this paper analyzes the remaining dataset and manually labels the reasons.

3 A Model to Predict Potential Rejection of Suggested Edits

The identification of undesired editing that causes rejection is essential to promote quality editing. However, manual differentiation of undesired and accepted edits can waste a lot of time and effort of users. Thus, this study aims to assist SO users by offering them automated support while editing a post. Specifically, we attempt to exploit cues in the textual contents of suggested edits to build a classifier that can automatically determine the rejected edits with potential rejection reasons. In this section, we answer the research question as follows.

RQ_1) To what extent do our classifiers predict the rejected edits with the potential reasons?

For RQ_1 , we define the following null hypothesis.

H_1 : The accuracy of our developed classifiers is not better than a random classifier with 50% accuracy.

Figure 2 shows the workflow of how we can detect the potential rejected edits of SO with reasons. Our prediction pipeline includes three main components: *Feature Extractor*, *Rejected Edit Predictor* and *Rejection Reason Classifier*. First, the feature extractor takes the body text of the original post, the post with suggested edits, and the user's information (e.g., reputation) as inputs. Then, it produces a feature vector based on the predictor variables. Second, the rejected edit predictor takes the feature vectors as input and outputs a dichotomous variable, *rejected*. The value of *rejected* is 1 if the predictor determines that the suggested edit will most likely be rejected, and it is 0 otherwise. Third, the rejection reason classifier takes the corresponding feature vector and texts as inputs if the value of *rejected* is 1. Then, it outputs the potential reasons for rejection. Finally, we measure the performances of the rejected edit predictor and rejection reason classifier.

3.1 Feature Extractor

Table 2 summarizes the features employed to predict whether suggested edits will be rejected or not. We extracted *fifteen* texts & user-based features to predict the potential rejected edits. Each feature is connected to one or multiple reasons behind a rejected edit. Note that we discarded *emotion*, which was included in our registered report (Mondal et al. 2020), and added *reputation*. The reasons behind such decisions are discussed in Section 7. However, this section discusses how we extract the features as follows.

Text/Code Formatting Text or code formatting refers to the changes in their presentation styles. For example, consider the text with HTML tags—`<p> I am using C# programming language </p>`. Here, C# is formatted as bold. Someone can reject the bold format of C# by removing the ` ... ` tags. However, in both cases, the content remains unchanged. The extracted content will be “I am using C# programming language”. We thus detect the text/code formatting in the following ways.

Determine Text Formatting

```
textFormat = Tbr != null && Tar != null
&& De(Tbr, Tar) == 0
&& De(Tbrwt, Tarwt) != 0
```




Fig. 2 Workflow to predict rejected edits of SO with reasons

Table 2 Features of our predictor

| Predictor | Features | Covered reasons |
|------------------|-------------------------|---|
| Editing Distance | Text/Code Formatting | Undesired Text/Code Formatting |
| | Text/Code Modification | Undesired Text/Code Add/Remove, Undesired Text/Code Change |
| | Deface Post | Other |
| | Complete Change of Post | Other |
| Status | Status | Status Update |
| Gratitude | Gratitude | Gratitude Add/Remove |
| Greetings | Greetings | Greetings Add/Remove |
| Reference | Reference Modification | Undesired Reference Modification |
| | Inactive Hyperlink | Undesired Reference Modification |
| Signature | Signature | Signature Add/Remove |
| Deprecation | Deprecation Note | Deprecation Note Add/Remove |
| Duplication | Duplication Note | Duplication Note Add/Remove |
| Reputation | Reputation Score | Community Trust |

where T_{brwt} : Texts before rollback with HTML tags (e.g., `<p> I am using C# programming language </p>`), T_{br} : Texts before rollback (e.g., I am using C# programming language), T_{arwt} : Texts after rollback with HTML tags (e.g., `<p> I am using C# programming language </p>`), T_{ar} : Texts after rollback (e.g., I am using C# programming language), and D_e : *Levenshtein* (Yujian and Bo 2007; Wikipedia 2020) editing distance. However, we remove new lines, strip leading & trailing spaces from T_{br} & T_{ar} , and make them lowercase before processing.

Determine Code Formatting

```
codeFormat = Cbr != null && Car != null
            && De(Cbr, Car) == 0
            && De(Cbrwt, Carwt) != 0
```

where C_{brwt} : Code before rollback with HTML tags, C_{br} : Code before rollback, C_{arwt} : Code after rollback with HTML tags, and T_{ar} : Code after rollback. Like text formatting, we remove new lines, strip leading & trailing spaces from C_{br} & C_{ar} , and make them lowercase before processing.

Text/Code Modification Text/code modification refers to text/code addition/removal or change of existing text/code.

Determine Text Modification

```
if (Tbr != null && Tar != null && De(Tbr, Tar) != 0){
    textModification = De(Tbr, Tar)
} else
    textModification = False
```

We normalize *textModification* with respect to the character length of T_{br} since T_{br} varies. We determine the code modification in the same way as text modification.

Deface Post Remove the texts or code entirely.

Deface Text

```
defaceText = ( $T_{br} \neq \text{null} \ \&\& \ T_{ar} == \text{null}$ ) || ( $T_{ar} \neq \text{null} \ \&\& \ T_{br} == \text{null}$ )
```

We determine the deface code in the same way as the deface text. However, deface post will be *True* if *defaceText* or *defaceCode* becomes *True*.

Complete Change of Post Users change texts or code segments entirely.

Determine Complete Change of Text

```
completeChangeText =  $T_{br} \neq \text{null} \ \&\& \ T_{ar} \neq \text{null}$   
    &\& ( $D_e(T_{br}, T_{ar}) == \text{Len}(T_{br})$  ||  $D_e(T_{br}, T_{ar}) == \text{Len}(T_{ar})$ )
```

We determine the complete code change in the same way as the text change. However, complete change of post will be *True* if *completeChangeText* or *completeChangeCode* becomes *True*.

Status To detect the addition/deletion of status, we attempt to find a keyword match to T_{br}/T_{ar} from a keywords list, L_{kw} . Here, $L_{kw} = \{\text{edit}, \text{update}, \text{note}, \text{ps}\}$ (Mondal et al. 2021a).

Determine Status

```
status = False  
for (keyword in  $L_{kw}$ ) {  
    if (( $\text{Match}(\text{keyword}, T_{br}) == \text{True} \ \&\& \ \text{Match}(\text{keyword}, T_{ar}) == \text{False}$ ) ||  
        ( $\text{Match}(\text{keyword}, T_{ar}) == \text{True} \ \&\& \ \text{Match}(\text{keyword}, T_{br}) == \text{False}$ )) {  
        status = True  
        return status  
    }  
}  
return status
```

Gratitude We detect the addition/removal of gratitude similarly to status. However, here the keyword list, $L_{kw} = \{\text{welcome}, \text{thanks}, \text{sorry}, \text{appreciated}, \text{thank}, \text{ty}$ (i.e., thank you), *thx*, *regards*, *tia* (i.e., thanks in advance)\} (Mondal et al. 2021a).

Greeting We detect the addition/removal of greeting similarly to status. However, here the keyword list, $L_{kw} = \{\text{hi}, \text{hello}, \text{hey}, \text{dear}, \text{greetings}, \text{hai}, \text{guys}, \text{hii}, \text{howdy}, \text{hiya}, \text{hay}, \text{heya}, \text{hola}, \text{hihi}, \text{salutations}\}$ (Exchange 2009a).

Reference Modification To detect reference modification, we extract the values of the *href* attribute of <a> tag from T_{brwt} & T_{arwt} . Then insert the hyperlinks into two lists— LST_{br} : list of hyperlinks found from texts before rollback, and LST_{ar} : list of hyperlinks found from texts after rollback.

Determine Reference Modification

```

referenceModification = False
for (reference in LSTbr){
    if (Contain(reference, LSTar) == False){
        referenceModification = True
        return referenceModification
    }
}
for (reference in LSTar){
    if (Contain(reference, LSTbr) == False){
        referenceModification = True
        return referenceModification
    }
}
return status

```

Inactive Hyperlink To detect the inactive (e.g., broken/dead) hyperlink, we check the HTTP response of each of the hyperlinks of LST_{ar} . We decide whether a hyperlink is inactive or not based on the response code.

Signature To detect the addition/removal of the signature, we extract and store the full name, first part, and last part (if any) of two users (who suggested edit & who rolled back) into a list, L_{name} . We then detect the addition/removal of the signature similarly to status.

Deprecation Note We detect the addition/removal of deprecation notes similarly to status. However, here the keyword list, $L_{kw} = \{deprecation, deprecate, oldcode\}$ (Mondal et al. 2021a).

Duplication Note We detect the addition/removal of duplication notes similarly to status. However, here the keyword list, $L_{kw} = \{duplicate, duplication\}$ (Mondal et al. 2021a).

Reputation Score We compute the reputation score of users to estimate—(1) how much the community trusts them and (2) whether they follow the editing guidelines. The official data dump of SO only reports the latest reputation scores of the users, which are not appropriate for our analysis. We thus use the snapshot of user activities (e.g., votes, acceptances, bounties) to compute the reputation score of users during their editing of posts. In particular, we use a standard equation provided by the SO to calculate the reputation score (Exchange 2020).

3.2 Rejected Edit Predictor

In this section, we first describe machine learning classifiers (Section 3.2.1) and their evaluation setup (Section 3.2.2). Then we evaluate the performance of our classifiers in identifying rejected edits (Section 3.2.3). Finally, we construct the baseline models and report their performance in Section 3.2.4.

3.2.1 Machine Learning Models

The relationship between edit categories (rejected/accepted) and their corresponding feature values might be complex. Thus, we choose the following four popular machine learning classification techniques with different learning strategies to identify the potential rejected edits. They are widely used in the relevant studies (Saha et al. 2013; Ponzanelli et al. 2014b; Rahman and Roy 2015a; Beyer et al. 2018).

Decision Trees (DT) is a non-parametric supervised machine learning technique for classification and regression. Non-parametric means it does not make any assumptions about the underlying data distribution. The intuition behind decision trees is that simple decision rules are inferred from the dataset features and continually split the training set until all data points belonging to each class are isolated. In particular, this technique employs different heuristics (e.g., entropy, information gain) to decide which feature to be used for the subsequent split of the training set. The commonly used decision trees are ID3, C4.5, and CART. However, ID3 can only be used when features are categorical. C4.5 and CART are extensions of ID3, which can work with features of both categorical and continuous data. Here, we use CART since our extracted features have both continuous and categorical values.

Random Forest (RF) is a supervised machine learning technique. The ‘forest’ it builds is an ensemble of many decision trees, usually trained with the ‘bagging’ method. The underlying principle behind the ensemble model is that a group of weak learners come together to form a strong learner. Ensemble learners thus improve the performance of single classifiers by inducing several classifiers and combining them to obtain a new classifier that outperforms every one of them (Polikar 2006). RF is scalable to any number of dimensions and usually has acceptable performance. However, it adds additional randomness to the model while growing the trees. For example, instead of searching for the most important feature while splitting nodes, it searches for the best feature among a random subset of features. RF thus prevents the overfitting of datasets by creating random subsets of the features.

K-Nearest Neighbors(KNN) is a non-parametric method employed in classification and regression problems (Goldberger et al. 2005). It does not use the training data points to perform any generalization. Thus, the KNN’s training phase is much faster than other classification algorithms. In KNN, K represents the number of nearest neighbors, which is the core factor in deciding a data point’s label (i.e., class). This technique finds the K closest neighbors of a target point using distance measures (e.g., Euclidean distance). Then, each neighbor votes for their class, and the class with the most votes is taken as the prediction.

eXtreme Gradient Boosting (XGBoost) is a scalable tree boosting technique that predicts a target variable by combining an ensemble of estimates from a set of more simplistic and weaker models (Chen and Guestrin 2016). It is a supervised learning algorithm and can be employed for both classification and regression. XGBoost is an extension to gradient boosted decision trees (GBM) with improved speed and performance. However, it is faster than other algorithms because of its parallel and distributed computing. XGBoost performs well because of its robust handling of various data types, relationships, distributions, and a variety of hyperparameters. In addition, XGBoost has inbuilt cross-validation and a variety of regularizations, which helps reduce overfitting.

3.2.2 Model Evaluation Setup

Dataset Selection We used the randomly sampled dataset (rejected & accepted) from Section 2.1 to train and test our machine learning classifiers. However, we keep the training and testing data separate since we have one time-dependent feature (i.e., reputation score). For example, consider a user suggested edits to multiple posts at different times. The reputation score of users increases over time. Thus, the reputation score of that user will be lower while suggesting prior edits than the later edits. Therefore, we use earlier edits as the training set and later edits as the test set to ensure that past data is not predicted based on future data. In particular, we take 70% samples that were edited relatively earlier to train the machine learning classifiers and use the remaining 30% to test them.

Performance Metrics The selection of the evaluation criteria is vital to guarantee a reliable assessment of the prediction models. In a binary classification problem like rejected edit prediction, a confusion matrix (e.g., Table 3) records the correctly and incorrectly recognized examples of each class. Therefore, we can obtain several metrics from the given confusion matrix to independently evaluate models' performance for both positive and negative classes.

The machine learning community often measures the *classification accuracy* as a simple scalar performance metric for binary classification. Classification accuracy measures the ratio of correctly classified edits into rejected & accepted classes with respect to all classified edits. However, according to He and Garcia (2008), accuracy might lead to an incorrect conclusion as the measure is highly sensitive to changes in data. In such cases, *precision* is a useful metric to capture the effect on a classifier performance of having a larger number of negative examples (Davis and Goadrich 2006). In particular, precision measures the ratio of correctly classified edits into a class (i.e., rejected/accepted) with respect to all edits classified into that class. However, He and Garcia (2008) argued that precision is still sensitive to changes in the data distribution, and it cannot assert how many positive examples are classified incorrectly. Unlike precision, *recall* is not sensitive to data distribution that measures the ratio of correctly classified edits with respect to the actually observed edits as true instances. However, any assessment based solely on recall would be inadequate, as it provides no insight into how many examples are incorrectly classified as positives. Therefore, neither precision nor recall can provide a reliable assessment of classification performance (Calefato et al. 2019). However, these individual scalar metrics can be combined to build more reliable classification performance measures. Specifically, these aggregated performance metrics include the F-measure that represents the harmonic mean of precision and recall. We thus measure precision, recall, F1-score, and overall accuracy to conclude the models' performance better. They can be measured as follows.

$$\begin{aligned} \text{Precision} &= \frac{TP}{TP+FP} & \text{Recall} &= \frac{TP}{TP+FN} \\ \text{F1-Score} &= \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} & \text{Accuracy} &= \frac{TP+TN}{TP+FP+TN+FN} \end{aligned}$$

Table 3 Confusion matrix

| Sample Size (N) | Expected (<i>Rejected</i>) | Expected(<i>Accepted</i>) |
|-------------------------------|------------------------------|-----------------------------|
| Predicted (<i>Rejected</i>) | True Positive (TP) | False Positive (FP) |
| Predicted (<i>Accepted</i>) | False Negative (FN) | True Negative (TN) |

3.2.3 Model Performance Evaluation

We experiment with our models to see how well the classification models perform based on our features. Figure 3 summarizes the performance of our models. Our primary focus is to predict the rejected edits. We see that our models can predict the rejected edits with 62.3%–69.1% precision. The random forest can predict the rejected edits more precisely than the other three models. Its precision is 69.1%. On the other hand, XGBoost shows the highest recall (i.e., 71.8%). However, the highest F1-score is achieved by the random forest model in predicting rejected edits. The precision to predict the accepted edits ranges from 59.1%–70.5%. Like rejected edits, the random forest model achieves the highest precision in predicting accepted edits. On the contrary, the k-nearest neighbors shows the lowest precision.

From Fig. 3, we see that the overall accuracy of the models is more than 60%. The highest accuracy is about 70%. Our experimental result shows that the random forest performs best, whereas the k-nearest neighbors shows the lowest performance. The k-nearest neighbors could suffer from the high dimensional data. However, according to the experimental result, XGBoost slightly outperforms decision trees but underperforms the random forest. We thus further investigate why XGBoost does not outperform random forest. We analyze the predicted class of XGBoost and random forest models against our test dataset. We find




Fig. 3 Performance of our machine learning models

that XGBoost misclassified 31 samples (15 accepted + 16 rejected), which random forests classified correctly.

In our dataset, addition/removal of duplication notes get rejected more than 94% times, and signatures get rejected 100% times. However, a few samples are classified as accepted by XGBoost even after the addition/removal of duplication notes or signatures. On the contrary, random forest classified them correctly as rejected. Besides, XGBoost classified several samples with trivial text changes as rejected, which were incorrect. Our analysis shows that trivial changes have more chance of being accepted than rejected. However, random forest classified the target class correctly in those cases. Hence, the above reasons could explain why XGBoost slightly underperforms random forest. Therefore, we select the random forest model to deploy with our online tool.

Our models could capture unnecessary details or too specific relationships within the training dataset and thus suffer from overfitting. However, overfit models are not very stable since they fail to generalize well to the data. Moreover, overfit models generally perform poorly on unseen (e.g., test) data. In particular, a overfit model has a substantial difference between the accuracy of the training and test dataset. That is why we can never trust an overfit model and put it into deployment. We thus attempt to reduce overfitting before deploying our model. In particular, we tuned the critical parameters of the model to balance the accuracy between training and test datasets (Table 4). Then, we set the values of the parameters to avoid model overfitting.

Figure 4 shows the model accuracies on the training and test dataset in contrast to example critical parameters of the model. For example, we run the random forest model over the depth of the tree from 1 to 20. As shown in Fig. 4a, the training accuracy improves with the depth of the tree. However, the difference between training and test accuracies increases when the depth exceeds five. We thus set the depth value as five while training the random forest model. We set the depth values for the decision trees and XGBoost models similarly. Their depth values are five and three. In particular, we attempt to balance the training and test accuracies when determining depth values.

On the other hand, Fig. 4c shows the training and test accuracies of the k-nearest neighbors in contrast to the number of neighbors. We vary the number of neighbors from 1 to 50, evaluate the model on the train and test datasets for each number of neighbors, and report the accuracy. We see that performance on the test set improves initially and then worsens, and performance on the training set continues to degrade. As shown in Fig. 4c, the training accuracy is dropping to converge with the line for the test set. However, the training and testing accuracies are very close when the number of neighbors is 46. We thus choose the parameter value is 46 (i.e., number of neighbors = 46).

Table 4 Accuracy of the machine learning models

| Model name | Accuracy | |
|---------------------|----------|----------|
| | Testing | Training |
| Random Forest | 69.8% | 71.2% |
| Decision Trees | 66.8% | 66.8% |
| K-Nearest Neighbors | 60.4% | – |
| XGBoost | 67.2% | 67.1% |

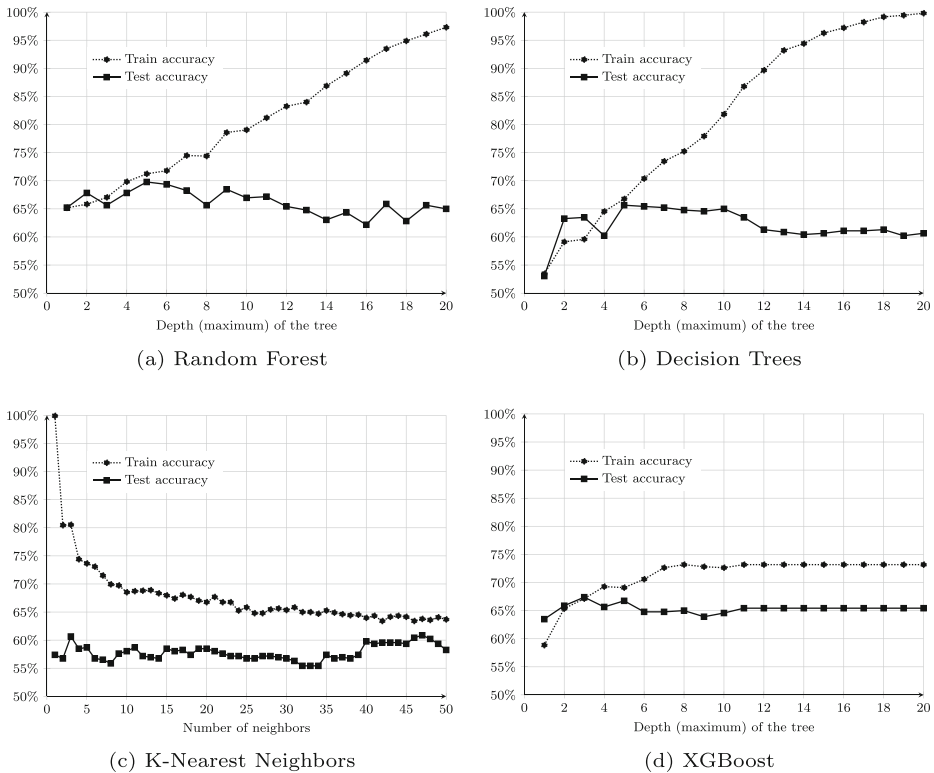


Fig. 4 Parameter tuning to reduce model overfitting

3.2.4 Performance of Baseline Model

To the best of our knowledge, there was no existing machine learning model to identify the rejected edits with reasons at the time when we conducted this study. However, we construct baseline models that reject/accept trivial edits and evaluate their performance. We thus first attempt to categorize the types of edits conducted in the samples of our dataset. We calculate the *Levenshtein* distance between the original post's content (text + code) and the post with suggested edits. We normalize the distance with respect to the character length of the original posts. Then, we classify the edits into four categories based on the edit distance. They are—*trivial* (distance \leq lower quartile), *small* (lower quartile $>$ distance \leq median), *medium* (median $>$ distance \leq upper quartile) and *major* (distance $>$ upper quartile) edit.

Figure 5 shows the percentage and count of rejected and accepted edits for each category. We see that 62.8% of trivial edits get accepted (Fig. 5a) in our dataset. On the contrary, such a statistic is only 37.2% for rejected edits. However, the percentage of rejected edits is higher than the accepted edits for the remaining three categories. We then develop a rule-based classifier that rejects trivial edits. Figure 6a summarizes the performance of the classifier. The precision of identifying rejected edits is below 40%. The classifier also has poor recall (i.e., 19.1%) in identifying rejected edits with 43.7% overall accuracy. Next, we build a classifier that rejects non-trivial (i.e., small, medium & major) edits. That is, it accepts trivial edits. Figure 6b shows its performance. This classifier shows a higher




Fig. 5 Change of post contents by edits

recall in identifying rejected edits than our proposed models. However, the precision, F1-score, and overall accuracy are significantly lower than our best-performing model. Such performances suggest that rejected edits cannot be identified reasonably well only based-on edit categories.

3.3 Rejection Reason Classifier

In the previous section, we evaluate the performance of the machine learning models in predicting the rejected edits. This section analyzes how accurately the rejection reason classifier can identify the potential reasons for those rejections.

Our rejected reason classifier can almost accurately identify several rejection reasons by applying the same approach as we extracted features Section 3.1. For example, our manual investigation identified nine keywords (e.g., thanks, welcome) (Section 3.1) that were utilized to identify gratitude. However, our further analysis finds that the addition or removal of gratitude is rejected 85.5% times (e.g., Fig. 7) in our dataset. We thus used this lightweight keyword-based technique to identify the reason “gratitude add/remove” when the edits are predicted as rejected. However, to avoid multiple computations, we look back at the feature vector to check the feature value of gratitude. Similarly, we identify the following potential edits rejection reasons by analyzing the feature vector.

Undesired text/code formatting, complete change of post, deface post, gratitude add/remove, greetings add/remove, signature add/remove, deprecation note add/remove, duplication note add/remove, undesired reference modification

For a remaining few rejection reasons, such as undesired text/code addition/removal, we primarily apply n-grams and pos-tagging-based techniques. Unfortunately, we did not find satisfactory performance (e.g., accuracy <50%). Our further investigation suggests that both desired and undesired texts contain similar words, phrases, or patterns. Therefore, we cannot distinguish them using words, phrases, or pos-based patterns. However, the character length of added or removed text/code can identify the undesired addition or removal of text/code reasonably well. We thus extract the added/removed text or code from our manually analyzed dataset using the appropriate HTML tags. In particular, we extract the contents of HTML elements with a class attribute with either `diff-add` & `diff-delete` values. We measure the length of added/removed characters of text or code. We then normalize




Fig. 6 Performance of baseline machine learning model

the length with respect to the total length of the text/code of revisions. Next, we separate samples into two classes according to our manual label—(1) undesired vs. desired text addition, (2) undesired vs. desired text removal, (3) undesired vs. desired code addition, and (4) undesired vs. desired code removal. We developed four random forest classifiers to identify the undesired text/code addition/removal. However, we resolve the class imbalance problem using Synthetic Minority Oversampling Technique (SMOTE) (Wang et al. 2006). Table 5 shows the performance of the classifiers. The precision of identifying the undesired text/code addition/removal is about 62.1%–69.2%. The recall is slightly lower for undesired text addition/removal than undesired code addition/removal. However, the overall accuracy is more than 63%, except for undesired text removal.

We then evaluate the overall performance of the rejection reason classifier in identifying the potential rejection reasons. We experiment with the rejection reason classifier using our test dataset. In particular, we label each sample in a file as follows.




Fig. 7 Percentage of rejected and accepted edits

Table 5 Performance of classifiers to identify undesired text/code addition/deletion

| Classifier | Precision | Recall | F1-Score | Accuracy |
|-------------------------|-----------|--------|----------|----------|
| Undesired Text Addition | 69.2% | 48.1% | 56.7% | 63.4% |
| Undesired Text Removal | 62.1% | 50.8% | 55.9% | 58.0% |
| Undesired Code Addition | 65.3% | 64.0% | 64.6% | 65.0% |
| Undesired Code Removal | 69.0% | 74.4% | 71.6% | 68.1% |

(1) **Identified:** the reasons detected by the rejection reason classifier. (2) **Expected:** the actual reasons based on our manual analysis. We then create a confusion matrix to analyze the performance of the rejection reason classifier as follows. (i) *True Positive (TP)* = ‘identified’ reasons = ‘expected’ reasons (ii) *False Positive (FP)* = (‘identified’ reasons \neq ‘expected’ reasons) or (‘identified’ reasons but ‘expected’ no reasons) (iii) *True Negative (TN)* = ‘identified’ no reasons and ‘expected’ no reasons, and (iv) *False Negative (FN)* = ‘identified’ no reasons but ‘expected’ one or more reasons. Using the above matrix, we compute four standard metrics (*Precision*, *Recall*, *F1-score*, and *Accuracy*) to compute the performance of the model to identify the rejected reasons.

Table 6 shows the confusion matrix and performance of rejection reason classifier to identify the potential reasons for edit rejections. Our analysis shows that our rejection reason classifier can identify the potential rejection reasons with 62.3% precision, 67.2% recall, 64.7% F1-score, and 66.7% overall accuracy. We further analyze in which cases our rejection reason classifier fails to identify rejection reasons. We find that our model mainly fails to identify all the reasons when there were multiple reasons to reject suggested edits. In that cases, our model can partly identify the reasons. For example, there are three reasons for rejection. However, our model identifies two of them accurately. Note that we identify the potential rejection reason as community mistrust when—(1) our rejection reason classifier cannot identify any reason and (2) the reputation score of the user who suggested edits is below 2K. We choose a reputation score $< 2K$ since users with a reputation score $< 2K$ have no privilege to edit posts instantly. However, our rejection reason classifier cannot identify a couple of reasons (e.g., incorrect text/code change). We will discuss them in Section 7.

Summary of RQ_1 and hypothesis testing: We developed four machine learning classifiers based on 15 texts & user-based features to predict rejected edits. Our best-performing classifier (i.e., random forest) can predict the rejected edits with 69.1% precision, 71.2% recall, 70.1% F1-score, and 69.8% overall accuracy. We then develop a classifier to identify the potential reasons behind rejections. Our rejected reason classifier can identify the potential rejection reasons with overall 66.7% accuracy. Such findings *reject* our null hypothesis (i.e., H_1), which hypothesized that the accuracy of our developed classifiers is not better than 50%.

4 EditEx: A Recommender for Early Fixes to Suggested Edits

We can assess the actual impact of our developed classifier if it can automatically assist users during their editing of SO posts. Editing is a time-consuming and largely voluntary activity in SO. Therefore, users can be assisted with a tool that can recommend potential

Table 6 Confusion matrix and performance of our model to identify the rejection reasons (TP: True Positive, TN: True Negative, FP: False Positive & FN: False Negative)

| TP | FP | TN | FN | Performance | | | |
|-----|----|-----|----|-------------|--------|----------|----------|
| | | | | Precision | Recall | F1-Score | Accuracy |
| 129 | 78 | 153 | 63 | 62.3% | 67.2% | 64.7% | 66.7% |

fixes to their undesired edits. We thus focus on introducing an online tool, namely *EditEx*, that interacts with our classifier, identifies the potential rejection reasons from the suggested edits, and helps to reduce the likelihood of rejection of suggested edits. We then assess the tool's effectiveness via real-world usage by users and answer the research question as follows.

RQ_2) To what extent can *EditEx* be helpful to users avoid rejection in their edits?

For RQ_2 , we define the following null hypothesis.

H_2 : Our *EditEx* tool is not effective in helping users avoid rejected edits.

4.1 *EditEx* Architecture

Figure 8 shows an overview of the *EditEx* architecture. *EditEx* has two parts: *client* and *server*. On the client-side, users get the *EditEx* interface. *EditEx* interface comprises two buttons: *EditEx* and *Suggestion*. *EditEx* enables users to edit the SO posts. On the other hand, users can check whether the edits will be rejected/accepted and the potential rejection reasons upon rejection by clicking the *Suggestion* button.

When users click the *Suggestion* button, the client-side script captures the necessary data required to extract the features for the machine learning model. In particular, it captures—(1) text before edit, (2) text after edit, (3) reputation, and (4) the name of the user who suggests edits. Then it sends this data to the server-side application. Client-side script is written by JavaScript, and server-side application is developed by Java. However, the server-side application captures all the feature values (Section 3.1) using the data sent by the client. The

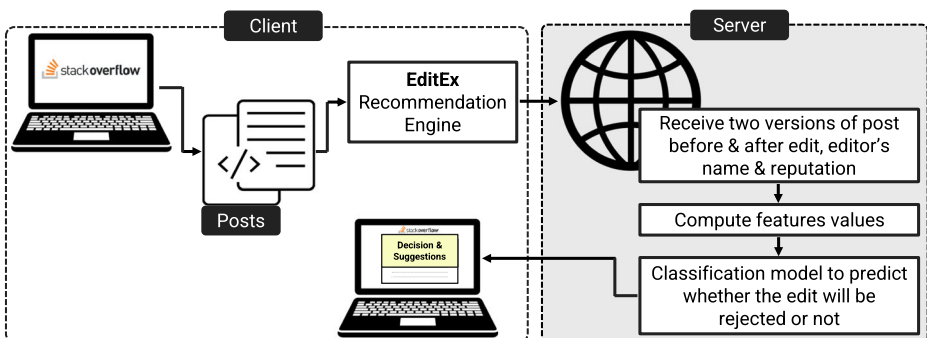


Fig. 8 An overview of the *EditEx* system architecture

feature values form a feature vector. The feature vector is then passed to the classification model to predict whether the edit will be rejected or not.

The potential rejection reasons are identified upon rejection based on the feature vector and texts (before and after edit). Then, the decision of the edit (i.e., rejected/accepted) and the helpful suggestions that notify the rejection reasons (if rejected) are sent to the client-side script. Finally, the client-side script offers users the result and suggestions.

4.2 Effectiveness Evaluation Plan of EditEx

Figure 9 shows the overview of how EditEx is introduced and its effectiveness is evaluated. Users can get the EditEx interface integrated with the SO's existing edit system after installing EditEx. Then, they can suggest edits to posts using EditEx to improve their quality. We measure the effectiveness of EditEx in two ways. First, we recruited 20 participants and divided them into two groups. The control group edits posts using the existing SO edit system of SO, whereas the treatment group uses EditEx to edit posts. When both groups complete their edits, we compare the success rates (e.g., fewer rejection ratios). Second, we survey participants after completing their edits and analyze their feedback.

4.3 Study Design

In this section, we first discuss how we recruit our study participants (Section 4.3.1). We then explain the formation of control vs. treatment groups (Section 4.3.2). Next, we discuss the two phases of our study—task-based evaluation and survey-based feedback collection phase (Section 4.3.3).

4.3.1 Study Participants

We recruit 20 participants who satisfy our constraint (i.e., participants must have editing experience in SO posts) in the following two ways.

- **Snowball Approach:** We use convenience sampling to bootstrap the snowball (Stratton 2021). We first contacted a few software developers who are known to us, easily reachable, and working in software companies worldwide. We discuss our study goals and share the survey with them. We then adopt a snowballing method (Bi et al. 2021) to disseminate the survey to some of their colleagues with similar experiences. We asked

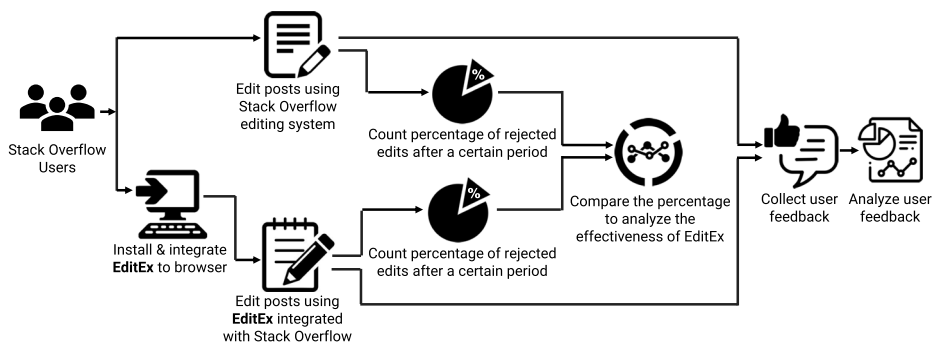


Fig. 9 An overview of the EditEx and its effectiveness evaluation

to share the survey with those who could be interested in editing posts and participating in our survey. In this process, we receive information (e.g., email) from 22 participants who show their interest in our study. However, 12 of them finally confirmed their participation.

- Open Circular: We post a description of this study and our research goals in the specialized Facebook groups to find potential participants. We target the groups where professional software developers discuss their programming problems. We also use LinkedIn as a research tool to reach potential participants because it is one of the largest professional social networks in the world. We get contacts of 20 participants from this open circular who are willing to participate and satisfy our constraints (e.g., must have editing experience). However, some participants did not respond to us when we contacted them. We finally confirmed *eight* participants in this process.

In the end, we were able to recruit a total of 20 developers (12 from snowball + 8 from open circular) who were eligible based on our study constraint. Half of them have 3–5 years, 40% (8 out of 20) have two years (or less), and 10% (2 out of 20) have 6–8 years of software development experience. Note that 35% of participants were from different software industries, and the remaining 65% of them were from academia. They worked as developers, technical leads, grad students, and faculty members worldwide (e.g., Canada, Germany, Bangladesh). In our registered report (Mondal et al. 2020), we proposed to recruit 30 participants using a snowball approach. However, we discussed these discrepancies in the number of participants and recruitment approach in Section 7.

4.3.2 Formation of Control and Treatment Groups

As mentioned above (Section 4.3.1), we recruit 20 participants for our study. We pick 10 of them for the treatment group and 10 for the control group as follows.

- Treatment Group: Each participant in this group was assisted in their editing of SO posts by our developed EditEx tool. However, the participant could also access the standard SO edit system.
- Control Group: Each participant in this group edited SO posts by using the standard SO edit system only.

Table 7 shows the experience and professions of control and treatment groups. We attempt to equate profession and experience between two groups to minimize the subjective

Table 7 Experience and profession of control and treatment groups

| Control group | | | | | Treatment group | | | | |
|--------------------------------|-----|------------|-------------|--------------|--------------------------------|-----|------------|-------------|--------------|
| Development Experience (Years) | | Profession | | | Development Experience (Years) | | Profession | | |
| ≤ 2 | 3-5 | 6-8 | Academician | SW Developer | ≤ 2 | 3-5 | 6-8 | Academician | SW Developer |
| 4 | 5 | 1 | 6 | 4 | 4 | 5 | 1 | 7 | 3 |

biases. For example, control and treatment groups have participants with low and high software development experience. Also, each group contains participants from academia (e.g., faculty members) and software industries.

4.3.3 Execution Plan

First, we conduct a task-based analysis (Section 4.4) by asking each participant (control + treatment group) to edit ten posts. However, we set a task list and asked each participant to suggest edits to posts following the task list. Table 8 lists the tasks. Tasks T1–T6 were set based on the identified reasons that might cause rejections. For example, the treatment and control groups were asked to add gratitude (e.g., thank you) associated with other edits. In this case, EditEx warns the treatment group against adding such gratitude. Therefore, participants of the treatment group proceed with the remaining edits and avoid adding gratitude. On the other hand, the control group did not receive such a warning. Thus, they suggested edits with gratitude. We attempt to see the effectiveness of EditEx in preventing the commonly rejected reasons while suggesting edits from tasks T1–T6. However, the control group might get more rejections for tasks T1–T6. On the contrary, such tasks favor the treatment group. Such a scenario can exaggerate EditEx's effectiveness in reducing edit rejections. Therefore, we also ask the participants to edit posts arbitrarily (e.g., Table 8, T7). That means participants suggest edits to posts that are not related to any rejection reasons to limit the bias of this study.

We circulated the editing guidelines of SO to each participant of the control and treatment groups. We asked participants to follow the guidelines in suggesting edits. After suggesting edits, participants wait until they get the decision (rejected/accepted) on those suggested edits from the edit reviewers. However, users with a reputation score $\geq 2K$ can edit posts instantly. Those edits are neither added to the review queue nor reviewed by experts. However, a rollback can reject their edits. Unfortunately, such a rollback even may take a few months. We thus ask participants with a reputation score $\geq 2K$ to create a new account. Edits suggested by their new account undergo an expert review. This decision also confirms the same privilege level of all the participants. However, our target was to quickly get decisions (rejected/accepted) on the suggested edits and avoid undecided edits.

Table 8 Editing tasks to the participants

| Task ID | Task description |
|---------|---|
| T1 | Add greeting (e.g., hello, dear, good day) at the beginning of a post with other edits. |
| T2 | Add gratitude (e.g., thank you) at the end of a post with other edits. |
| T3 | Add signature (e.g., user name) at the end of a post with other edits. |
| T4 | Format text (e.g., bold/unbold text, make text italic, format text element as code element or vice versa) with other edits. |
| T5 | Format code (e.g., add/remove space, make lowercase letter to uppercase or vice versa) with other edits. |
| T6 | Remove a code segment or a paragraph of text or change them radically. |
| T7 | Other (e.g., fix grammar & spelling, text/code modification) |

Second, we conduct an online survey (Section 4.5) to listen to the participants about their experience in suggesting edits to SO posts with/without our tool *EditEx*. Kitchenham and Pfleeger (2008) suggest considering six main steps for a personal opinion survey. They are—setting survey objectives, designing the survey, developing the survey instrument (i.e., the questionnaire), evaluating the survey instrument, and obtaining and analyzing data. We primarily follow their guidelines to survey participants. However, we also consider the guidance and ethical issues from the established best practices (Groves et al. 2011; Singer and Vinson 2002). For example, we take participants' consent before starting the survey. Besides, we confirm to participants that their provided information must be treated confidentially. Our survey includes different types of questions (e.g., multiple-choice, free-text answers). However, we inform about the estimated time (i.e., approximately 10 minutes) required to complete the survey to the participants. Our survey comprises the parts as follows.

- **Consent and Prerequisite.** This part confirms participants' consent to participate in this survey and agreement to process their data.
- **Participants' Information.** In this part, we collect participants' information such as experience, current profession, organization, country, and editing experience in SO posts.
- **Workload Assessment.** This section assesses the cognitive workload of the control and treatment groups in suggesting edits to SO posts. We leverage the NASA Task Load Index (TLX) (non-weighted) to estimate subjective workload (Cao et al. 2009; Hart and Staveland 1988; Noyes and Bruneau 2007; Sharek 2011). In particular, we assess how much effort participants had to exert mentally and physically to use *EditEx* and the standard edit system of SO. Participants were asked to rate their scores on an interval scale ranging from low (1) to high (10) (Memarian and Mitropoulos 2011) in the following six dimensions—(1) mental demand, (2) physical demand, (3) temporal demand, (4) effort, (5) performance, and (6) frustration.
- **Usefulness Analysis.** In this section, we measure the participants' confidence in suggesting edits using *EditEx* (treatment group) and SO edit system (control group). In addition, we ask the treatment group to rate the usefulness of the suggestions of *EditEx*. In particular, we ask the following two questions to the treatment group participants and employ a 5-point Likert scale (i.e., 1–5) to estimate their consent (Joshi et al. 2015; Vagias 2006).

- (a) How useful did you find the suggestions from *EditEx*? (*5-point Likert scale*)
- (b) How confident were you to follow the *EditEx* suggestions? (*5-point Likert*)

We ask the following question to measure the confidence level of the participants in the control group.

- (c) How confident are you to edit posts using the SO editing system? (*5-point Likert scale*)
- **Suggestions to Improve *EditEx*.** Finally, we seek participants' recommendations to improve the effectiveness of *EditEx*. We ask them the question as follows.
- (a) What are your recommendations to further improve *EditEx*? (*Text*)

We added the *survey form* and its *responses* in anonymized CSV form in our replication package (Mondal et al. 2021b).

4.4 Results from the User Study on Editing of SO Posts

Participants were asked to complete ten edits. However, several participants (especially from the control group) could not edit ten posts due to three main challenges.

- (a) SO's edit queue often remains full, and thus participants could not edit posts according to their schedule.
- (b) Participants could not edit many posts (e.g., more than three) simultaneously. SO restricts them from suggesting further edits before receiving a decision (rejected/accepted) on the pending ones.
- (c) SO does not allow its users to edit for a period when consecutive edits are being rejected.


Figure 10 shows the task completion ratio of control and treatment groups. As mentioned above, we asked each group to edit 100 posts (10 for each). However, participants from the control group were able to edit 83 posts in total (i.e., completion ratio 83%), whereas the treatment group edited 94 posts.

We then attempt to see the rejection ratio of suggested edits. First, we collect information from each participant (treatment & control group) when all of their suggested edits got decisions (e.g., rejected/accepted). In particular, we asked how many edits they suggested and how many of them got rejected. We also collect their editing details to examine whether they follow the given task list (e.g., Table 8) or not. Then we count the total number of suggested and rejected edits of the treatment and control groups. Finally, we calculate the rejection ratio of each group.

As shown in Fig. 11, the rejection ratio of the edits suggested by the treatment group is only 16% (15 out of 94). On the contrary, such a statistic is 65.1% (54 out of 83) for the control group. Overall, the percentage of rejected edits who used EditEx is about 49% lower than those who used the standard editing system of SO. Such a finding gives us a preliminary validation that our tool helps users to prevent their suggested edits from being rejected.

Tasks T1–T6 were set based on our identified rejection reasons. However, EditEx alerts the treatment group participants while suggesting T1–T6. Such alerts might help them to avoid edits that cause rejections. On the contrary, the control group did not receive any alerts from the SO edit system while suggesting T1–T6. Therefore, the overall edit rejection ratio of the control group is much higher than the treatment group. We thus attempt to compare the rejection ratio of task T7 (i.e., free-form editing) between control and treatment groups. According to our analysis, the control group suggests 23 free-form edits. Among them, six were rejected. Since EditEx prevents T1–T6, the treatment group could not suggest T1–T6. Therefore, we consider all of their tasks as T7. However, we found that they suggested

Fig. 10 Task completion of treatment and control groups






Fig. 11 Rejection ratio of treatment and control groups

two edits out of task T7 (i.e., T1–T6). Among the remaining 92, 13 edits were rejected. As shown in Fig. 12, the rejection ratio of the control group for T7 is 26.1%. On the other hand, such ratio for the treatment group is 14.1%. Therefore, EditEx not only assists users in avoiding edits that are usually rejected but also assists them in conducting regular (i.e., free-form) edits.

4.5 Results from Survey of User Study Participants

We received 20 valid survey responses (10 treatment + 10 control). We report the survey responses as follows.

- **Workload assessment** during the completion of the editing tasks,
- **Usefulness ratings** of EditEx suggestions, and
- **Improvement suggestions** by the study participants for EditEx.

4.5.1 Workload Assessment During Edit Task Completion

Figure 13 shows the box plots of the NASA TLX cognitive workload score on a scale of ten. We compute the average workload of each participant. We first sum up the ratings of each of the six dimensions (e.g., mental demand) and then divide it by the number of dimensions. In particular, we use the equation to compute the average workload of each participant as follows.

$$A_{wl} = \frac{1}{D_T} \left[\sum_{i=1}^{D_T} R_i \right] \quad (2)$$

where R_i denotes the rating (1–10) of i^{th} dimension, D_T represents the total dimensions (here, $D_T = 6$).

Figure 13 shows the box plots that represent the average cognitive workload of each participant from the treatment and control groups. We see that the median subjective workload for the treatment group is about half that of the control group. That is, EditEx strongly

Fig. 12 Rejection ratio of T7 between treatment and control groups





Fig. 13 Cognitive workload in editing SO posts using EditEx vs. SO's standard edit system using NASA TLX



supports the SO standard editing system to reduce the users' workload required to suggest edits to posts. We then attempt to find whether the workload difference between treatment and control groups is statistically significant. We use the *Mann-Whitney-Wilcoxon* statistical significance test (McKnight and Najab 2010) and find a statistical significance *p-value* ($p\text{-value} \simeq 0.0 < 0.05$). We also use *Cliff's delta* test (Macbeth et al. 2011) to determine the effect size and find a large *effect size* (Cliff's $d = -0.97$ (large)) with 95% confidence. Given this evidence, EditEx helps the users to edit posts by significantly reducing their workload.

4.5.2 Usefulness of EditEx Suggestions

Table 9 shows the participants' assessment of the effectiveness of EditEx and the SO edit system. We see that participants find the suggestions of EditEx *influential* ($3.41 \leq \text{score} \leq 4.20$) in avoiding the potential rejections. The Likert score (i.e., 4.0) also shows that they were *confident* to follow the suggestions given by EditEx. When we asked the reason behind their confidence level, one participant responded that *EditEx suggests the common reasons behind the unsuccessful attempts of edit. These suggestions help to identify those and fix them*. On the contrary, the participants who use the SO edit system (i.e., control group) were *moderately confident* ($2.61 \leq \text{score} \leq 3.40$) in suggesting their edits. To explain the reasons, one participant stated that *I was not sure whether my edits are good or bad, so my confidence was low. I cannot understand much from the reviewers' comments why my edits were actually rejected. The reasons were too general*. Such findings indicate that EditEx is not only able to provide valuable suggestions but also make users more confident in suggesting edits.

Table 9 Effectiveness analysis of EditEx and SO edit system

| Questions | Target group | Mean value | Interpretation |
|--|--------------|------------|----------------|
| How useful did you find the suggestions from EditEx? | Treatment | 4.1 | Influential |
| How confident were you to follow the EditEx suggestions? | Treatment | 4.0 | Confident |
| How confident are you to edit posts using the SO editing system? | Control | 3.1 | Moderate |

4.5.3 Recommendations for EditEx Improvements

We analyzed the recommendations of all the participants and summarized them into three categories. We see that participants recommended—(1) enhancing existing functionalities, (2) improving Graphical User Interface (GUI), and (3) notification & installation system. We discuss their recommendations below.

Enhance Functionality

- Besides suggesting the rejection reasons, EditEx could also estimate a score based on the quality of the edit.
- EditEx can be enhanced by adding a few features related to natural language processing, such as identifying incorrect spelling, sentence complexity.
- EditEx should detect minor changes (e.g., adding an article) that do not significantly improve the quality of the posts.
- Participants also suggested enhancing the capability of EditEx in such a way that it can identify more potential reasons that might cause rejection.
- Participants recommend paraphrasing the notification sentences (e.g., edits may get rejected due to low reputation) to convey them more positively.

Improve Graphical User Interface

- A few participants recommend improving the GUI of EditEx. For example, the notification system of the potential rejection reasons could be more appealing. In addition, *Suggest Me* button should appear beside the edit window to avoid scrolling.

Add Notification & Improve Installation System

- One of the main barriers to suggesting edit is that the queue remains full most of the time. Participants thus suggest that EditEx should notify them when the queue becomes free to suggest edits to avoid frequent manual checking.
- EditEx uses Tampermonkey to add userscripts for integrating it into the SO edit system. They appreciate it since Tampermonkey is popular and easy to use. However, they suggest deploying EditEx as a standalone browser plug-in in future.

Summary of RQ_2 and hypothesis testing: Results from the user study on the editing of SO posts show that EditEx can support the SO edit system to prevent overall 49% of rejected edits, including the commonly rejected ones. However, it can prevent 12% rejections even in regular free-form edits. On the other hand, results from the survey of user study participants show that users who use EditEx found its suggestions influential while editing posts. Therefore, they were confident to follow the suggestions. As a result, EditEx can reduce the cognitive workload in half compared to the SO edit system. Such findings *reject* the null hypothesis (H_2), which hypothesized that the EditEx tool is ineffective in helping users avoid rejected edits.

5 Discussions

In this section, we first explain the importance of the features used in our machine learning models (Section 5.1). We then discuss the reasons behind the misclassifications of the machine learning models (Section 5.2). Finally, we discuss the implications of our study findings and the developed `EditEx` tool in Section 5.3.

5.1 Ranking of Features in the Machine Learning Models

In Section 3.1, we discussed several features used to develop our machine learning models in predicting whether an edit will be rejected or accepted. However, we do not know which features are more important in machine learning classifiers than others in differentiating rejected and accepted edits. To find the important features, we thus attempt to rank our features using two popular measures as follows.

Information Gain We attempt to determine which features are more robust than others for discriminating between rejected and accepted edits. We thus employ an information gain-based feature ranking technique because it can estimate the discrimination power of each of the given features. In information theory, the information gain of a random variable is the change in information entropy between an initial state and a state that takes some information (Saha et al. 2013). Therefore, the information gained from a particular attribute in classifying rejected edits is as follows.

$$InfoGain(C, a_i) = H(C) - H(C|a_i) \quad (3)$$

where C represents a particular class (i.e., rejected/accepted), a_i denotes the attribute, and H denotes information entropy.

Figure 14a shows the information gain of our selected features. We see that changes in text and reputation score have the highest information gain. It means that they might discriminate the edits more accurately than others. Changes in code and addition/removal of gratitude have the next highest information gain. However, the information gains of the




Fig. 14 Ranking of features

remaining features are minimal. Therefore, those features contribute less to the machine learning models to classify the rejected edits from the accepted ones.

SHapley Additive exPlanations (SHAP) Feature Importance The SHAP value is the average marginal contribution of a feature towards the model's prediction across all possible combinations of features (Molnar 2020). It shows whether a feature value can increase a model's prediction over a random baseline (Lundberg et al. 2020). However, the idea behind SHAP feature importance is that features with larger absolute SHAP values are more important than others. SHAP values can be calculated for any tree-based model. We calculate the SHAP values to rank features from the Random Forest model.

Figure 14b shows the SHAP values of our selected features. We see that the top four features are the same as when using the information gain-based feature ranking technique. According to SHAP values, code change has the second-highest capability of identifying rejected edits following the reputation score. However, the addition/removal of gratitude has more power to discriminate rejected edits from accepted ones over changes in texts. We then attempt to see the actual effect of our selected features on the model's performance. We thus remove the low-ranked features (according to SHAP values) one by one and evaluate the performance of our model. Table 10 shows the experimental results. According to the experiment, the model's performance gradually degrades when we remove features one by one. For example, overall accuracy decreases 2% when we remove deprecation. Interestingly, we see a slightly higher performance when we keep the top six features than the top seven. Then the performance again decreases after removing features. However, the overall accuracy drops from 69.8% to 61.3% when we remove all other features except the top four.

5.2 Analysis of Trained Machine Learning Model & its Misclassifications

Machine learning models can produce accurate/inaccurate predictions. However, their *black box* nature might prevent their easy adoption and enhancement by others. The SHAP

Table 10 Effect of individual feature on predicting rejected edits (f₁: reputation, f₂: code change, f₃: gratitude, f₄: text change, f₅: text format, f₆: deface post, f₇: signature, f₈: status, f₉: duplication, f₁₀: reference modification, f₁₁: greetings, f₁₂: inactive link, f₁₃: code format, f₁₄: complete change, f₁₅: deprecation)

| Features | Rejected | | | Accepted | | | Accuracy |
|---------------------------------|-----------|--------|----------|-----------|--------|----------|----------|
| | Precision | Recall | F1-Score | Precision | Recall | F1-Score | |
| f ₁ –f ₁₅ | 69.1% | 71.2% | 70.1% | 70.5% | 68.4% | 69.5% | 69.8% |
| f ₁ –f ₁₄ | 68.3% | 65.9% | 67.1% | 67.4% | 69.7% | 68.5% | 67.8% |
| f ₁ –f ₁₃ | 65.9% | 71.6% | 68.6% | 69.2% | 63.2% | 66.1% | 67.4% |
| f ₁ –f ₁₂ | 65.5% | 72.9% | 69.0% | 69.8% | 61.9% | 65.6% | 67.4% |
| f ₁ –f ₁₁ | 66.7% | 69.0% | 67.8% | 68.2% | 65.8% | 67.0% | 67.4% |
| f ₁ –f ₁₀ | 62.9% | 74.7% | 68.3% | 69.2% | 56.3% | 62.1% | 65.4% |
| f ₁ –f ₉ | 63.6% | 73.4% | 68.2% | 68.9% | 58.4% | 63.2% | 65.9% |
| f ₁ –f ₈ | 61.5% | 73.8% | 67.1% | 67.6% | 54.1% | 60.1% | 63.9% |
| f ₁ –f ₇ | 62.1% | 71.6% | 66.5% | 66.8% | 56.7% | 61.4% | 64.1% |
| f ₁ –f ₆ | 64.4% | 71.2% | 67.6% | 68.1% | 61.0% | 64.4% | 66.1% |
| f ₁ –f ₅ | 61.3% | 72.1% | 66.3% | 66.5% | 55.0% | 60.2% | 63.5% |
| f ₁ –f ₄ | 59.4% | 70.3% | 64.4% | 64.0% | 52.4% | 57.6% | 61.3% |

(Lundberg and Lee 2017) is a popular model interpretation framework to interpret the classification/misclassification results of the model. In our experiment, we conduct a binary classification where the *rejected edits* was considered as the *positive* class and *accepted edits* as the *negative* class. Thus, our models attempt to predict the rejected edits by default. Therefore, a positive SHAP value indicates an increase in our models' prediction of positive class and vice versa. Figure 15 shows the importance of our selected features using a bee swarm plot from our random forest model. The bee swarm plot visualizes the SHAP value of a feature from each of the training instances on the x-axis. On the y-axis, it sorts all features in descending order according to their sum of SHAP values. The blue color indicates a low feature value, whereas the red indicates a high feature value in our plot. We see that reputation is the most important feature according to our random forest model. That is, community trust, which is estimated by reputation score, is an important predictor of the acceptability of suggested edits. We note that this feature with true response often leads to negative SHAP values, which indicates an increased prediction towards edit acceptance. That is, edits suggested by users with a high reputation score have a higher chance of being accepted and vice versa. We also analyze our dataset for further insights on this. In particular, we conduct a comparative analysis of the reputation scores between users with rejected and accepted edits. We find the median reputation score of users whose suggested edits were rejected is 3428. On the contrary, such a score of users whose suggested edits were accepted is more than double (i.e., 7660). We then attempt to see whether the difference is statistically significant or not. We use the *Mann-Whitney-Wilcoxon* statistical significance test and find a statistical significance *p*-value ($p\text{-value} \simeq 0.0 < 0.05$). We also use *Cliff's delta* test to determine the effect size and find a medium effect size (Cliff's $d = -0.34$ (medium)) with




Fig. 15 Feature importance using bee swarm plot (Random Forest model)

95% confidence. Given this evidence, low-reputed users are less trusted by the community, and thus their edits can be rejected.

The second most important feature is the code change. Code change with true response often leads to positive SHAP values, which indicates an increased prediction towards edit rejections. It refers that suggested edits get rejected when users change the code much. On the contrary, text change with true response leads to negative SHAP values, which means that significant changes in texts are acceptable. Similarly, we see that the true response of text & code format, reference modification, and complete change of posts lead to negative SHAP values, which indicates an increased prediction towards edit acceptance. On the other hand, the true response of gratitude, deface post, signature, status, greetings, and inactive links lead to positive SHAP values, which indicates an increased prediction towards edit rejections. However, according to the SHAP visualization, a few features such as reputation, code & text change, status, and reference modifications confuse our model. For example, reference modification's true response leads to positive and negative SHAP values. Therefore, those features might cause misclassifications of our model.

5.3 Implications of Study Findings

The findings from our study and the tool `EditEx` can guide the following major stakeholders in crowd-sourced knowledge-sharing platforms that use collaborative editing features: (a) **forum designers** to improve the edit system, (b) **forum users** to guide their edit behavior, and (c) **researchers** to study and improve collaborative editing support in crowd-shared platforms. We discuss the implications below.

Forum Designers The quality assurance of shared content is paramount for the usefulness and popularity of a crowd-shared knowledge-sharing platform like SO. While the editing of content allows users to suggest improvements in quality, the lack of proper guidance to the user can lead to unnecessary rejections of the suggested edits. SO can use our developed machine learning model and the `EditEx` tool to offer on-demand and context-aware edit fix recommendations to the users. As we observed in Section 4.5, the SO users that used `EditEx` had significantly less number of rejections compared to the users that did not use `EditEx` tool. The SO edit assessment queue is set up to ensure novice SO users (with less than 2K reputation) can not make bad edits. Many users in SO fall under this novice category, but their inputs to edits are equally as important as the non-novice users. However, as we noted, the SO edit queue can be often congested with many suggested edits that it could take a disproportionate time for an edit to get reviewed. Even after the review, many suggestions can be rejected due to trivial edits like undesired text formatting. A tool like `EditEx` can help SO users by reducing such trivial edits, which then can ultimately help SO and its expert edit reviewers with less rejection, which in turn can reduce the workload on the edit queue.

Forum Users Interactive browser plug-ins like `EditEx` can warn forum users of potential edits that could be rejected. Thus `EditEx` can improve the confidence of SO users during the editing. Indeed, as we reported in Section 4.5, the study participants were more confident while using `EditEx` than using SO editing system (average confidence of 4 while using `EditEx` vs 3.1 while using SO). The SO users had only 16% of their edits rejected while using `EditEx`, while the rejection rate was 65% while using SO. While a rejection ratio of 65% for SO could be biased due to our choice of editing task, the much lower rejection rate while using `EditEx` does indeed highlight that SO users can benefit from a simple tool like

EditEx. EditEx is very easy to install. Given that EditEx is simply a browser plug-in that SO users can easily install, we hope that EditEx will be accepted by the wider SO community.

Indeed, our developed classifiers can be used to detect rejected edit reasons in SO automatically. It can extend current tools and techniques that predominantly use contents from suggested edits to recommend editing suggestions (e.g., see the works of Chen et al. 2017a, b). The tool EditEx, with further modification (as suggested by our study participants in Section 4.5.3), can be influential in the reduction of rejected edit reasons in SO and can improve the overall satisfaction of SO users. In the long term, the tool can promote better content because users will be more motivated. Such high quality contents then can offer better content and recommendation support for tools and techniques that focus on the quality of contents shared in SO (Zhang et al. 2018; Ya et al. 2013; Hudson et al. 2015; Rahman and Roy 2015b; Agichtein et al. 2008; Mondal et al. 2019), the suite of tools and techniques developed to detect and recommend quality posts (Ponzanelli et al. 2014c, d; Ya et al. 2015; Harper et al. 2008; Li et al. 2015b; Calefato et al. 2018).

Researchers The quality of knowledge shared in SO is important because developers worldwide now rely on this shared knowledge. Indeed, knowledge shared in SO can support diverse activities like bug fixing, feature enhancement, API selection, and documentation (Uddin and Khomh 2017a, b, c, 2019; Uddin et al. 2019, 2020a, b; Chakraborty et al. 2021). This sharing of knowledge is important, because official software documentation can often be lacking (Uddin and Robillard 2015; Khan et al. 2021). However, the editing of content is a voluntary activity. SO users can be demotivated to produce quality edits if they become frustrated due to unnecessary/unwanted rejection of their suggested edits. Tools like EditEx or proactive policy assurance by Chen et al. (2017a, 2018) can help SO users with suggestions to improve their edits. The positive survey responses of our tool EditEx show the potential of deploying the tool in SO. Future research can contribute by including more features into EditEx and by conducting more studies to learn how SO users can further benefit from such tools. Such findings can promote quality contents, which then can support all development tasks that rely on SO, as noted above.

6 Threats to Validity

Threats to internal validity relate to experimental errors and biases (Tian et al. 2014). We asked participants to suggest edits using EditEx (treatment group) and SO's standard edit system (control group). The suggested edits were either accepted or rejected by the expert review. However, the accepted edits could be rejected later by a rollback that might affect the edit rejection ratio. We thus further analyzed how many accepted edits were rejected by rollbacks in our dataset. However, such a statistic was less than 1% in our dataset. Thus, it might not affect our results significantly.

Threats to external validity relate to the generalizability of a technique. We agree that there might be rejection reasons that we could not identify. However, we analyze statistically significant samples of rejected edits by rollbacks from questions and answers. We thus believe that our manual investigation exposes all the main rejection reasons. Still, there is scope to analyze more samples to explore additional rejection reasons. However, suggested edits can be rejected by either rollback or expert review. Unfortunately, we could not collect samples of those edits rejected by expert reviews. We could not find any convenient way to collect such samples since their information is not readily available in the SO data dump.

Thus, similar to existing literature, we consider the rejected edits by rollbacks (Wang et al. 2018). However, suggested edits are reviewed in SO by users with at least a 2K reputation score. In our manually analyzed dataset, 90.6% (692 out of 764) of users who rolled back edits have a reputation score $\geq 2K$. The remaining were self-rollback (i.e., rolled back by post owner). Therefore, our intuition is that the main reasons for rejections by expert reviews would be similar to our identified reasons.

Our survey participants range from novice to experienced, mainly software developers and academicians (Table 7). Such diversity in the survey participants offers validity and applicability to the survey findings. Furthermore, we ensure that control and treatment groups have participants with different professions and experience levels to mitigate individual bias.

We set a task list T1–T7 (Table 8) and asked participants (control & treatment group) to suggest edits to posts based on that list. However, tasks T1–T6 were set based on our identified rejection reasons. Thus, EditEx alerts the treatment group while suggesting T1–T6, but the control group does not get such alerts. As a result, the treatment group receives favor from EditEx, which might reduce their rejections. To mitigate this bias, we asked participants to suggest free-form edits (T7). From T7, we attempt to see the effectiveness of EditEx in assisting users in suggesting regular edits besides preventing common rejections. However, EditEx can support the SO edit system to prevent 12% rejections. Such a finding confirms the effectiveness of EditEx in preventing not only common rejections but also rejections from regular edits.

7 Deviations from Registered Report

This section discusses the deviations of this study from our registered report (Mondal et al. 2020) and explains them.

Rollback Reasons & Predictors In our registered report, we wanted to see whether the addition/deletion of emotion influences edit rejection using EmoTxt (Calefato et al. 2017). Our primary analysis found that emotion has almost no effect on edit rejection/acceptance. Furthermore, the overall accuracy of our rejected edit classifier improved only about 1% if we consider emotion as a predictor. However, integrating a complex model to capture emotion and its deployment is costly. Therefore, it can affect the performance of our online tool EditEx. We thus discard emotion from this study. On the other hand, we added *reputation* as a predictor (Table 2). The reputation score estimates how much the community trusts a user (Anderson et al. 2012; Overflow 2022). Therefore, we added *Community Trust* as a rollback reason (which was absent in the registered report) (Table 1) that was estimated by reputation score. Moreover, users with lower reputation scores might violate the editing guidelines more than those with higher reputation scores. Note that violating edit guidelines is one of the causes of edit rejection. We thus consider reputation as a predictor that significantly improves the performance of our rejected edit classifier. Another discrepancy to the registered report is that we included the *Introducing Spam* rollback reason under *Other*.

Manually Investigated Sample Size In the registered report, we manually investigated 777 rollback edits (382 questions + 395 answers). The statistically significant sample size of rollback edits for both question and answer is 382. However, due to a programming problem, we randomly selected 395 samples from rejected answer revisions. We kept 395 samples in the registered report because—(1) we completed our analysis using 395 samples,

and (2) 395 is more than the statistically significant sample size. However, we later randomly selected 382 (among 395) samples (Section 2.1) to equalize it with the statistically significant sample size and analyze them.

Recruitment of Participants We planned to recruit 30 participants (15 for treatment + 15 for the control group) who edited at least 100 posts. After deploying our tool EditEx, we realized that EditEx could be helpful to both expert and novice SO users. Therefore, we relaxed our constraints. We recruit participants who edited any SO post to ensure familiarity with SO editing. However, we struggled to recruit 30 participants due to COVID-19 and the extensive nature of this user study. The study was extensive because each user had to do multiple edits of SO posts. Primarily, we planned to recruit participants using a snowball approach. To recruit more participants, we then extend our approach. Besides snowball, we attempt to recruit participants using an open circular. However, finally, we recruited 20 participants (10 for treatment + 10 for the control group) with different experience levels and diverse professions (Section 4.3.1).

Number of Edits Per User We planned to ask each participant to suggest ten edits. However, several participants (especially the control group) could not edit ten posts for three main challenges, as we discussed in Section 4.4.

EditEx's Functionality of Highlighting Texts We planned to include EditEx's functionality to highlight texts that may cause rejection. However, the current version of EditEx cannot highlight texts. EditEx predicts the edit decisions (accepted/rejected) and alerts users with the potential rejection reasons if rejected. While highlighting texts could be helpful, we found that the EditEx tool with the basic features was usable and effective. Therefore, we leave the highlighting of texts in EditEx as a future extension.

NASA TLX Workload We planned to estimate TLX effort as a task load by combining all the ratings provided by a participant in five dimensions in the TLX metrics. However, most existing studies estimate subjective workload by combining all the ratings in six dimensions (Cao et al. 2009; Hart and Staveland 1988; Noyes and Bruneau 2007; Sharek 2011; Hart 1986). Therefore, we also take ratings from each participant on six dimensions (mental demand, physical demand, temporal demand, effort, performance, and frustration). Finally, we estimate the cognitive workload (Fig. 13). However, we consider a scale of ten with one step size (Memarian and Mitropoulos 2011) to take ratings conveniently from participants for each dimension. It slightly sacrifices granularity in comparison to scale 100 with step size five. However, the results should not be affected much.

Unidentified Reasons Our rejection reason classifier cannot identify a couple of rejection reasons, such as partial acceptance and incorrect text/code change. We can extract added or removed text/code using appropriate HTML tags. However, partly accepted text/code cannot be separated from added text/code. Furthermore, it requires analyzing the future revisions to check partial acceptance, which is impractical. We did not find any patterns that can identify incorrect changes in code/texts. Identification of such reasons demands manual efforts.

8 Related Work

We developed our tool `EditEx` to recommend fixes to suggested edits in SO so that SO users can avoid committing undesired edits that may lead to the rejection of the edits. As such, our research in this paper belongs to a broader area called ‘collaborative editing in social forums’. Major related work can broadly be divided into **Studies** of collaborative editing systems in crowd-sourced forums (see Section 8.1) and **Techniques** to suggest improvements to the editing system (see Section 8.2). In addition, SO data are used extensively in SE research for various tasks (see Section 8.3), all of which could be potentially impacted by having low-quality data due to erroneous/inefficient edits.

8.1 Studies of Collaborative Editing Systems

Editing of content can improve the content quality. As such, it is intuitive that the social Q&A forums offer to edit the post contents. Since social forums can be accessed by many users simultaneously, it is a cost-effective measure for the forums to support collaborative editing by allowing their users to do the editing. Indeed, studies show that collaborating editing in social forums and online collaborative knowledge-sharing portals (e.g., Wikipedia) can positively impact towards the improvement of shared contents (Li et al. 2015a; Kittur and Kraut 2008). The nature of the collaborative editing can be similar across the social forums (e.g., Q&A site) and knowledge portals (e.g., Wikipedia). The research of Li et al. (2015a) looked at the adoption of Wikipedia-style collaborative editing into a Q&A site like SO. They found that users with good edits are rewarded with positive votes by other users. They analyzed five years of historical editing data from SO and found that substantive edits from other users can increase the number of positive votes by 18% for the questions and 119% for answers. This reward can be beneficial for a user who does the edit because the edit may only offer at most 5% improvement over the original post (i.e., the user can be rewarded with mindful but low-cost editing efforts). Indeed, the SO reward system can serve as an added influence to the users to suggest edits. A recent study by Wang et al. (2018) in SO found that users are motivated to edit more when they are closer to getting a badge.

Overall, both Wang et al. (2018) and Li et al. (2015a) conclude that offering incentives as reputation scores is useful to improve post quality within a collaborative editing platform like SO. This finding was also observed in other collaborative editing platforms like webcasts (Munteanu et al. 2008) and Wikipedia (Kittur and Kraut 2008). Munteanu et al. (2008) tested the effectiveness of engaged users in collaborating in a wiki-like webcast platform to edit/correct transcripts that are produced from webcasts through an automated speech recognition system. Collaborative editing can be a cost-effective but useful means to improve the quality of the ASR (Automated Speech Recognition) system in webcasts because ASR systems can have an average error rate of 45%—above the accepted threshold of 25%. The field study carried out by the authors in a real lecture environment found that using students to edit the webcast transcript was useful in reducing the error rate. The editing was supported via a webcast extension that engages users to collaborate in a wiki-like manner. Kittur and Kraut (2008) find that the increase in the number of editors does not guarantee the quality of the articles on Wikipedia.

The quality of the question is important to get an answer: lack of clarity, relatedness, and reproducibility of the problem, as well as the too short question, could dissuade developers from answering the question (Asaduzzaman et al. 2013; Mondal et al. 2019). The reputation and past activity of an asker could also factor into the likelihood of a question getting resolved (Rahman and Roy 2015b). As such factors of good questions are investigated, e.g.,

code to text ratio, etc. (Calefato et al. 2018; Duijn et al. 2015). However, depending on the platforms and user characteristics, these factors can vary (Hudson et al. 2015). As such, it is important to detect content quality automatically (Ponzanelli et al. 2014a, c; Ya et al. 2015). Wang et al. (2018) found that users who make more edits in a short time are likely to get more edits rejected. Thus bad edits can harm the content quality.

Our research on SO rollback edits initially started in 2019 to better understand the edit rejection reasons as reported by Wang et al. (2018). Through our qualitative analysis of SO posts, we also found all the edit rejection reasons reported by Wang et al. (2018). In addition, we found four more edit rejection reasons. We report the edit rejection reasons in Section 2 of this paper. While the above papers, including Wang et al. (2018), focus on analyzing editing mechanisms in collaborative platforms based on empirical studies, our paper focuses on developing techniques to automatically suggest fixes to suggested edits so that the edits will not be rejected upon submission. As such, our developed tool *EditEx* can further contribute to supporting the content quality in social forums by assisting users with guidance on improving the quality of their suggested contents. Thus, our paper offers complementary viewpoints to the above studies by offering tools and techniques that can facilitate improved edit content in a social Q&A site like SO.

8.2 Techniques to Develop to Improve Collaborative Editing Systems

Collaborative editing systems are common in Wikipedia (Li et al. 2015a; Kittur and Kraut 2008), GitHub code editing (Dabbish et al. 2012), webcasts (Munteanu et al. 2008), scientific contents (Lowry et al. 2005; Calvo et al. 2005), and so on. Compared to substantial research on conducting studies on existing collaborative editing systems, we are not aware of much research that focuses on developing tools and techniques to improve the systems. This is perhaps due to the fact that currently available collaborative platforms like Wikipedia seem to work well and are hugely popular. In all these platforms, the focus of collaborative editing is to improve the quality of the shared content based on user engagement (Agichtein et al. 2008).

Chen et al. (2017a) observed that most of the edits in SO are small sentence edits. While developing their SOTorrent database, Baltes et al. (2018) also observed that majority of edits in SO are relatively small. In a follow-up study, Chen et al. (2018) predicted whether a post needs to be edited. Their approach is based on the concept of ‘proactive policy assurance’, which assures that a modification to a suggested edit will satisfy the current ‘reactive policy assurance’ in SO, which accepts/rejects based on the matching of exiting editing policy after an edit is submitted (i.e., reactive). They developed a deep-learning-based policy assurance tool to recommend post owners or other users’ potential mid-level edits to given post content. The deep learning model is a CNN (Convolutional Neural Network). In a large-scale experiment, they find that the tool offers good precision, recall, and F1-score (at least 0.7) while suggesting mid-level edits.

As we noted in Section 8.1, our research of this paper started in 2019 to gain hands-on experience on the edit rejection reasons observed by Wang et al. (2018). Our initial exploration led to an expansion of the edit rejection reasons and to the submission of a registered protocol report in 2020 (Mondal et al. 2020). In the registered protocol report, we outlined our vision of this paper by offering to develop machine learning model to automatically detect the edit rejection reasons and to build our *EditEx* tool that can offer proactive guidance to fix suggested edits. While working on this paper, we observed that some edit rejection reasons could be present both in accepted and rejected edits, resulting in *inconsistencies* in the editing acceptance/rejection process. We reported a catalog of such

inconsistencies in our MSR 2021 paper (Mondal et al. 2021a). In our MSR 2021 paper, we also report several rule-based tools that we developed to detect inconsistencies in SO edits automatically. While developing our *EditEx* tool in this paper, we purposefully did not consider those inconsistencies, given those were not outlined in our 2020 registered protocol report (Mondal et al. 2020). We note that an immediate extension of *EditEx* could investigate whether and how the inclusion of the inconsistencies into the rejection prediction models and the *EditEx* tool could make the overall editing process more effective for the SO users. We leave it as our immediate future work.

8.3 Other SE Research using SO Data

Several studies have been conducted to study developer discussions on different crowd-shared developer platforms, including SO. Seaman et al. studied developer discussion on inspection meetings (Seaman and Basili 1998). Reiner et al. used content analysis to study developer discussions on software processes (Rainer et al. 2003). Gottipati et al. study relevant answers in 3 software forums, Dzone, Tips, and Oracle forums (Gottipati et al. 2011). Several studies has focused on discussions on microblogs, such as, Twitter (Tian et al. 2012; Prasetyo et al. 2012; Wang et al. 2013), and chat communities, such as, HipChat (Alkadhi et al. 2017), IRC messages (Alkadhi et al. 2018; Shihab et al. 2009), and Slack (Chatterjee et al. 2019). Recently, SO Q&A forums have been subject to a number of papers to study various aspects of software development, such as what developers are discussing in general (Barua et al. 2012), or about a particular aspect, e.g., concurrency (Ahmed and Bagherzadeh 2018), big data (Bagherzadeh and Khatchadourian 2019), chatbot development (Abdellatif et al. 2020).

Several studies has been conducted to study developer sentiments on online discussions (e.g., SO data) (Guzman et al. 2014; Murgia et al. 2014; Ortu et al. 2015; Novielli et al. 2014; Uddin and Khomh 2017a, c, 2019; Uddin et al. 2019, 2020a, b; Chakraborty et al. 2021; Lin et al. 2022). Guzman et al. applied sentiment analysis on code comments (Guzman et al. 2014). Islam and Zibran study emotional variations in commit messages (Islam and Zibran 2016). Garcia et al. (2013) study the emotions of developers in the Gentoo community. Guzman and Bruegge studied developer sentiments on mailing lists (Guzman and Bruegge 2013). Novielli et al. conduct sentiment analysis on SO and Github discussions (Novielli et al. 2015). Many of these studies use automated sentiment analysis tools, which are found to provide contradictory results in software engineering research (Jongeling et al. 2017).

All the above research works using SO data could be benefited from improved data quality offered by collaborative editing in SO. As such, our tool *EditEx*, once adopted by the SO users, can help the SO users as well as the SO-based research community with better quality data.

9 Conclusion

SO has become an essential online resource with millions of programming-related problems and solutions. However, the quality of the shared knowledge is vital for the growth and success of SO. To promote quality, SO introduces an edit system so that users can suggest an improvement to posts. Unfortunately, numerous suggested edits are rejected due to either undesired changes of posts or violating edit guidelines. Such a scenario not only hurts the quality of content but also frustrates and demotivates users. We conducted a qualitative analysis of 764 (382 questions + 382 answers) rejected edits by rollbacks and identified 19

rejection reasons. We then extract 15 texts and user-based features to automatically capture those reasons and develop four machine learning models using them. Our best-performing model can predict rejected edits with about 70% accuracy, and the rejection reason classifier can identify the potential rejection reasons with 67% accuracy. We also introduced an online tool named *EditEx* that can be integrated with the SO edit system. It analyzes the edits, predicts whether they will be rejected, and suggests users with the potential rejection reasons. We conduct a survey to assess *EditEx* and SO edit system. According to survey results, the participants find reasons for rejection identified by *EditEx* *influential*. Moreover, *EditEx* can support the SO edit system to prevent 49% rejections, including the commonly rejected reasons. Such a statistic is 12% when users suggest regular free-form edits. Moreover, our tool significantly decreases the subjective workload and increases participants' confidence in suggesting edits.

Acknowledgements This research is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery grants, and by an NSERC Collaborative Research and Training Experience (CREATE) grant, and by two Canada First Research Excellence Fund (CFREF) grants coordinated by the Global Institute for Food Security (GIFS) and the Global Institute for Water Security (GIWS).

Data Availability The datasets analyzed during this study are available in our online appendix (Mondal et al. 2021b).

Declarations

Conflict of Interest The authors have no conflict of interest.

References

- Abdellatif A, Costa D, Badran K, Abdalkareem R, Shihab E (2020) Challenges in chatbot development: a study of stack overflow posts. In: 17th International conference on mining software repositories, October 5–6, 2020, Seoul, Republic of Korea. ACM, New York
- Agichtein E, Castillo C, Donato D, Gionis A, Mishne G (2008) Finding high-quality content in social media. In: Proceedings of the 2008 international conference on web search and data mining, pp 183–194
- Ahmed S, Bagherzadeh M (2018) What do concurrency developers ask about?: a large-scale study using stack overflow. In: Proceedings of the 12th ACM/IEEE international symposium on empirical software engineering and measurement, article no. 30
- Alkadhi R, Lata T, Guzman E, Bruegge B (2017) Rationale in development chat messages: an exploratory study. In: 2017 IEEE/ACM 14th international conference on mining software repositories (MSR), IEEE, pp 436–446
- Alkadhi R, Nonnenmacher M, Guzman E, Bruegge B (2018) How do developers discuss rationale? In: 2018 IEEE 25th international conference on software analysis, evolution and reengineering (SANER). IEEE, pp 357–369
- Anderson A, Huttenlocher D, Kleinberg J, Leskovec J (2012) Discovering value from community activity on focused question answering sites: a case study of stack overflow. In: Proceedings of the 18th ACM SIGKDD international conference on knowledge discovery and data mining, pp 850–858
- Asadzaman M, Mashiyat AS, Roy CK, Schneider KA (2013) Answering questions about unanswered questions of stack overflow. In: Proceedings of the 10th working conference on mining software repositories, pp 87–100
- Bagherzadeh M, Khatchadourian R (2019) Going big: a large-scale study on what big data developers ask. In: Proceedings of the 2019 27th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering, ESEC/FSE 2019. ACM, New York, pp 432–442
- Bagozzi RP, Dholakia UM (2006) Open source software user communities: a study of participation in linux user groups. *J Manag Sci* 52(7):1099–1115
- Baltes S, Dumani L, Treude C, Diehl S (2018) Sotorrent: reconstructing and analyzing the evolution of stack overflow posts. In: Proceedings of the 15th international conference on mining software repositories, pp 319–330

- Barua A, Thomas SW, Hassan AE (2012) What are developers talking about? An analysis of topics and trends in stack overflow. *Empir Softw Eng* 1–31
- Beyer S, Macho C, Di Penta M, Pinzger M (2018) Automatically classifying posts into question categories on stack overflow. In: 2018 IEEE/ACM 26th international conference on program comprehension (ICPC). IEEE, pp 211–21110
- Bi T, Xia X, Lo D, Grundy J, Zimmermann T, Ford D (2021) Accessibility in software practice: a practitioner's perspective. [arXiv:210308778](https://arxiv.org/abs/210308778)
- Boslaugh S (2012) Statistics in a nutshell: a desktop quick reference. O'Reilly Media, Inc
- Calefato F, Lanubile F, Novielli N (2017) Emotxt: a toolkit for emotion recognition from text. In: 2017 Seventh international conference on affective computing and intelligent interaction workshops and demos (ACIIW), pp 79–80
- Calefato F, Lanubile F, Novielli N (2018) How to ask for technical help? Evidence-based guidelines for writing questions on stack overflow. *J Inf Softw Technol* 94:186–207
- Calefato F, Lanubile F, Novielli N (2019) An empirical assessment of best-answer prediction models in technical q&a sites. *ESE* 1–48
- Calvo RA, O'Rourke ST, Jones J, Yacef K, Reimann P (2005) Collaborative writing support tools on the cloud. *IEEE Trans Learn Technol* 41:66–99
- Cao A, Chintamani KK, Pandya AK, Ellis RD (2009) Nasa tlx: software for assessing subjective mental workload. *Behav Res Methods* 41(1):113–117
- Chakraborty P, Shahriyar R, Iqbal A, Uddin G (2021) How do developers discuss and support new programming languages in technical q&a site? An empirical study of go, swift, and rust in stack overflow. *Information and Software Technology (IST)* 19
- Chatterjee P, Damevski K, Pollock L, Augustine V, Kraft NA (2019) Exploratory study of slack q&a chats as a mining source for software engineering tools. In: 2019 IEEE/ACM 16th international conference on mining software repositories (MSR). IEEE, pp 490–501
- Chen T, Guestrin C (2016) Xgboost: a scalable tree boosting system. In: Proceedings of the 22nd ACM sigkdd international conference on knowledge discovery and data mining, pp 785–794
- Chen C, Xing Z, Liu Y (2017a) By the community & for the community: a deep learning approach to assist collaborative editing in q&a sites. In: Proceedings of the ACM on human-computer interaction, Article 32
- Chen C, Xing Z, Liu Y (2017b) By the community & for the community: a deep learning approach to assist collaborative editing in q&a sites. In: Proceedings of the ACM on human-computer interaction, Article No. 32
- Chen C, Chen X, Sun J, Xing Z, Li G (2018) Data-driven proactive policy assurance of post quality in community q&a sites. In: Proceedings of the ACM on human-computer interaction, Article 33
- Cohen J (1960) A coefficient of agreement for nominal scales. *Educ Psychol Meas* 20(1):37–46
- Cohen J (1968) Weighted kappa: nominal scale agreement provision for scaled disagreement or partial credit. *Psychol Bull* 70(4):213
- Dabbish L, Stuart C, Tsay J, Herbsleb J (2012) Social coding in github: transparency and collaboration in an open software repository. In: Proceedings of the ACM conference on computer supported cooperative work, pp 37–46
- Davis J, Goadrich M (2006) The relationship between precision-recall and roc curves. In: Proceedings of ICML, pp 233–240
- Duijn M, Kucera A, Bacchelli A (2015) Quality questions need quality code: classifying code fragments on stack overflow. In: Proceedings of the IEEE/ACM 12th working conference on mining software repositories, pp 410–413
- Exchange S (2009a) Should 'hi', 'thanks', taglines, and salutations be removed from posts? <https://meta.stackexchange.com/questions/2950/> online; Last accessed February 2020
- Exchange S (2009b) What is a 'rollback'? <https://meta.stackexchange.com/questions/17038/what-is-a-rollback> online; Last accessed February 2020
- Exchange S (2019) StackExchange API. <http://data.stackexchange.com/stackoverflow>
- Exchange S (2020) (Accessed on: December 2021) How does reputation work? <https://meta.stackexchange.com/questions/7237/how-does-reputation-work>
- Garcia D, Zanetti MS, Schweitzer F (2013) The role of emotions in contributors activity: a case study on the gentoo community. In: 2013 International conference on cloud and green computing. IEEE, pp 410–417
- Goldberger J, Hinton GE, Roweis ST, Salakhutdinov RR (2005) Neighbourhood components analysis. In: Advances in neural information processing systems, pp 513–520
- Gottipati S, Lo D, Jiang J (2011) Finding relevant answers in software forums. In: 2011 26th IEEE/ACM international conference on automated software engineering (ASE 2011). IEEE, pp 323–332
- Groves RM, Fowler JFJ, Couper MP, Lepkowski JM, Singer E, Tourangeau R (2011) Survey methodology

- Guzman E, Bruegge B (2013) Towards emotional awareness in software development teams. In: Proceedings of the 2013 9th joint meeting on foundations of software engineering, pp 671–674
- Guzman E, Azócar D, Li Y (2014) Sentiment analysis of commit comments in github: an empirical study. In: Proceedings of the 11th working conference on mining software repositories, pp 352–355
- Harper FM, Raban D, Rafaeli S, Konstan JA (2008) Predictors of answer quality in online q&a sites. In: Proceedings of the SIGCHI conference on human factors in computing systems, pp 865–874
- Hart SG (1986) Nasa task load index (tlx)
- Hart SG, Staveland LE (1988) Development of nasa-tlx (task load index): results of empirical and theoretical research. In: Advances in psychology, vol 52. Elsevier, pp 139–183
- He H, Garcia EA (2008) Learning from imbalanced data. *IEEE Trans Knowl Data Eng* (9):1263–1284
- Hudson N, Chilana PK, Guo X, Day J, Liu E (2015) Understanding triggers for clarification requests in community-based software help forums. In: Proceedings of the IEEE symposium on visual languages and human-centric computing, pp 189–193
- Islam MR, Zibran MF (2016) Towards understanding and exploiting developers' emotional variations in software engineering. In: 2016 IEEE 14th international conference on software engineering research, management and applications (SERA). IEEE, pp 185–192
- Jongeling R, Sarkar P, Datta S, Serebrenik A (2017) On negative results when using sentiment analysis tools for software engineering research. *Empir Softw Eng* 22(5):2543–2584
- Joshi A, Kale S, Chandel S, Pal DK (2015) Likert scale: explored and explained. *CJAST*
- Khan JY, Khondaker MTI, Uddin G, Iqbal A (2021) Automatic detection of five api documentation smells: practitioners' perspectives. In: IEEE International conference on software analysis, evolution and reengineering (SANER), p 12
- Kitchenham BA, Pfleeger SL (2008) Personal opinion surveys. In: Guide to advanced empirical software engineering
- Kittur A, Kraut RE (2008) Harnessing the wisdom of crowds in wikipedia: quality through coordination. In: Proceedings of the ACM conference on Computer supported cooperative work, pp 37–46
- Lakhani KR, von Hippel E (2003) How open source software works: free user-to-user assistance. *J Res Policy* 32(6):923–943
- Li G, Zhu H, Lu T, Ding X, Gu N (2015a) Is it good to be like wikipedia?: exploring the trade-offs of introducing collaborative editing model to q&a sites. In: Proceedings of the 18th ACM conference on computer supported cooperative work & social computing, pp 1080–1091
- Li L, He D, Jeng W, Goodwin S, Zhang C (2015b) Answer quality characteristics and prediction on an academic q&a site: a case study on researchgate. In: Proceedings of the 24th international conference on World Wide Web, pp 1453–1458
- Lin B, Cassee N, Serebrenik A, Bavota G, Novielli N, Lanza M (2022) Opinion mining for software development: a systematic literature review. *ACM Trans Softw Eng Methodol* 31(3). <https://doi.org/10.1145/3490388>
- Lowry PB, Curtis AM, Lowry MR (2005) A taxonomy of collaborative writing to improve empirical research, writing practice, and tool development. *J Bus Commun* 41:66–99
- Lundberg SM, Lee SI (2017) A unified approach to interpreting model predictions. In: Proceedings of the 31st international conference on neural information processing systems
- Lundberg SM, Erion G, Chen H, DeGrave A, Prutkin JM, Nair B, Katz R, Himmelfarb J, Bansal N, Lee SI (2020) From local explanations to global understanding with explainable ai for trees. *Nat Mach Intell* 2(1):56–67
- Macbeth G, Razumiejczyk E, Ledesma RD (2011) Cliff's delta calculator: a non-parametric effect size program for two groups of observations. *Univ Psychol* 10(2):545–555
- McKnight PE, Najab J (2010) Mann-whitney u test. The Corsini Encyclopedia of Psychology 1–1
- Memarian B, Mitropoulos P (2011) Work factors affecting task demands of masonry work. In: Proceedings of the 47th annual international conference of associated schools of construction
- Molnar C (2020) Interpretable machine learning. Lulu.com
- Mondal S, Rahman MM, Roy CK (2019) Can issues reported at stack overflow questions be reproduced?: an exploratory study. In: Proceedings of the 16th international conference on mining software repositories, pp 479–489
- Mondal S, Uddin G, Roy CK (2020) Automatic identification of rollback edit with reasons in stack overflow q&a site. In: 36th IEEE international conference on software maintenance and evolution (ICSME)—registered protocol report, pp 856–856
- Mondal S, Uddin G, Roy CK (2021a) Rollback edit inconsistencies in developer forum. In: 2021 IEEE/ACM 18th international conference on mining software repositories (MSR), pp 380–391
- Mondal S, Uddin G, Roy CK, Schneider K (2021b) Editex: prediction of rejected edits with reasons in so q&a site. <https://bit.ly/3Hp1xJa>

- Munteanu C, Baecker R, Penn G (2008) Collaborative editing for improved usefulness and usability of transcript-enhanced webcasts. In: Proceedings of the SIGCHI conference on human factors in computing systems, pp 373–382
- Murgia A, Tourani P, Adams B, Ortu M (2014) Do developers feel emotions? An exploratory analysis of emotions in software artifacts. In: Proceedings of the 11th working conference on mining software repositories, pp 262–271
- Novielli N, Calefato F, Lanubile F (2014) Towards discovering the role of emotions in stack overflow. In: Proceedings of the 6th international workshop on social software engineering, pp 33–36
- Novielli N, Calefato F, Lanubile F (2015) The challenges of sentiment detection in the social programmer ecosystem. In: Proceedings of the 7th international workshop on social software engineering, pp 33–40
- Noyes JM, Bruneau DP (2007) A self-analysis of the nasa-tlx workload measure. *Ergonomics* 50(4):514–519
- 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. In: 2015 IEEE/ACM 12th working conference on mining software repositories. IEEE, pp 303–313
- Overflow S (2015) How do i make a good edit? <https://meta.stackoverflow.com/questions/303219/how-do-i-make-a-good-edit>, online; Last accessed February 2020
- Overflow S (2022) What is reputation? How do I earn (and lose) it? <https://stackoverflow.com/help/whats-reputation>
- Parnin C, Treude C, Grammel L, Storey MA (2012) Crowd documentation: exploring the coverage and the dynamics of api discussions on stack overflow. Tech. rep. Georgia Tech.
- Polikar R (2006) Ensemble based systems in decision making. *IEEE Circ Syst Mag* 6(3):21–45
- Ponzanelli L, Mocci A, Bacchelli A, Lanza M (2014a) Improving low quality stack overflow post detection. In: Proceedings of the 30th international conference on software maintenance and evolution, pp 541–544
- Ponzanelli L, Mocci A, Bacchelli A, Lanza M (2014b) Understanding and classifying the quality of technical forum questions. In: 2014 14th International conference on quality software. IEEE, pp 343–352
- Ponzanelli L, Mocci A, Bacchelli A, Lanza M (2014c) Understanding and classifying the quality of technical forum questions. In: Proceedings of the 14th international conference on quality software, pp 343–352
- Ponzanelli L, Mocci A, Bacchelli A, Lanza M, Fullerton D (2014d) Improving low quality stack overflow post detection. In: Proceedings of the IEEE international conference on software maintenance and evolution, pp 541–544
- Prasetyo PK, Lo D, Achananuparp P, Tian Y, Lim EP (2012) Automatic classification of software related microblogs. In: 2012 28th IEEE international conference on software maintenance (ICSM). IEEE, pp 596–599
- Rahman MM, Roy CK (2015a) An insight into the unresolved questions at stack overflow. In: Proceedings of the 12th working conference on mining software repositories. IEEE Press, pp 426–429
- Rahman MM, Roy CK (2015b) An insight into the unresolved questions at stack overflow. In: Proceedings of the 12th working conference on mining software repositories, pp 426–429
- Rainer A, Hall T, Baddoo N (2003) Persuading developers to “buy into” software process improvement: a local opinion and empirical evidence. In: 2003 International symposium on empirical software engineering, 2003. ISESE 2003. Proceedings. IEEE, pp 326–335
- Saha RK, Saha AK, Perry DE (2013) Toward understanding the causes of unanswered questions in software information sites: a case study of stack overflow. In: Proceedings of FSE, pp 663–666
- Seaman CB, Basili VR (1998) Communication and organization: an empirical study of discussion in inspection meetings. *IEEE Trans Softw Eng* 24(7):559–572
- Sharek D (2011) A useable, online nasa-tlx tool. In: Proceedings of the human factors and ergonomics society annual meeting, vol 55. SAGE Publications Sage, Los Angeles, pp 1375–1379
- Shihab E, Jiang ZM, Hassan AE (2009) Studying the use of developer irc meetings in open source projects. In: 2009 IEEE International conference on software maintenance. IEEE, pp 147–156
- Singer J, Vinson NG (2002) Ethical issues in empirical studies of software engineering. *TSE*
- Stratton SJ (2021) Population research: convenience sampling strategies. *Prehosp Disaster Med* 36(4):373–374
- Tian Y, Achananuparp P, Lubis IN, Lo D, Lim EP (2012) What does software engineering community microblog about? In: 2012 9th IEEE working conference on mining software repositories (MSR). IEEE, pp 247–250
- Tian Y, Lo D, Lawall J (2014) Automated construction of a software-specific word similarity database. In: Proceedings of CSMR-WCRE, pp 44–53
- Uddin G, Khomh F (2017a) Automatic summarization of API reviews. In: Proceedings of 32nd IEEE/ACM international conference on automated software engineering, p 12
- Uddin G, Khomh F (2017b) Mining api aspects in api reviews. Tech. rep. <https://swat.polymtl.ca/data/opinionvalue-technical-report.pdf>

- Uddin G, Khomh F (2017c) Opiner: a search and summarization engine for API reviews. In: Proceedings of 32nd IEEE/ACM international conference on automated software engineering, p 6
- Uddin G, Khomh F (2019) Automatic opinion mining from API reviews from stack overflow. *IEEE Trans Softw Eng* 35
- Uddin G, Robillard MP (2015) How api documentation fails. *IEEE Softw* 32(4):76–83
- Uddin G, Baysal O, Gueroj L, Khomh F (2019) Understanding how and why developers seek and analyze api related opinions. *IEEE Trans Softw Eng* 40
- Uddin G, Khomh F, Roy CK (2020a) Automatic api usage scenario documentation from technical q&a sites. *ACM Trans Softw Eng Methodol* 43
- Uddin G, Khomh F, Roy CK (2020b) Automatic mining of api usage scenarios from stack overflow. *Inf Softw Technol (IST)* 16
- Vagias WM (2006) Likert-type scale response anchors. Clemson International Institute for Tourism & Research Development, Department of Parks, Recreation and Tourism Management Clemson University
- Wang J, Xu M, Wang H, Zhang J (2006) Classification of imbalanced data by using the smote algorithm and locally linear embedding. In: 2006 8th International conference on signal processing, vol 3. IEEE
- Wang X, Kuzmickaja I, Stol KJ, Abrahamsson P, Fitzgerald B (2013) Microblogging in open source software development: the case of drupal and twitter. *IEEE Softw* 31(4):72–80
- Wang S, Chen THP, Hassan AE (2018) How do users revise answers on technical Q&A websites? A case study on stack overflow. *IEEE Trans Softw Eng* 19
- Wikipedia (2020) Levenshtein distance. https://en.wikipedia.org/wiki/Levenshtein_distance, online; Last accessed February 2020
- Ya Y, Tong H, Xie T, Akoglu L, Xu F, Lu J (2013) Want a good answer? Ask a good question first! Tech. rep. arXiv:1311.6876
- Ya Y, Tong H, Xie T, Akoglu L, Xu F, Lu J (2015) Detecting high-quality posts in community question answering sites. *J Inf Sci* 302(1):70–82
- Yujian L, Bo L (2007) A normalized levenshtein distance metric. *IEEE Trans Pattern Anal Mach Intell* 29(6):1091–1095
- Zhang T, Upadhyaya G, Reinhardt A, Rajan H, Kim M (2018) Are code examples on an online q&a forum reliable?: a study of api misuse on stack overflow. In: Proceedings of the 40th international conference on software engineering, pp 886–896

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

Affiliations

Saikat Mondal¹  · Gias Uddin² · Chanchal Roy¹

Gias Uddin
gias.uddin@ucalgary.ca

Chanchal Roy
chanchal.roy@usask.ca

¹ Software Research Lab, Department of Computer Science, University of Saskatchewan, Saskatoon, Canada

² Data Intensive Software Analytics (DISA) Lab, Department of Electrical and Software Engineering, University of Calgary, Calgary, Canada