Using StackOverflow content to assist in code review

Balwinder Sodhi and Shipra Sharma

Dept. of Computer Science and Engineering, Indian Institute of Technology Ropar

Nangal Rd Ropar PB 140001 India

Email: {sodhi, shipra.sharma}@iitrpr.ac.in

Abstract

An important goal for programmers is to minimize cost of identifying and correcting defects in source code. Code review is commonly used for identifying programming defects. However, manual code review has some shortcomings: a) it is time consuming, b) outcomes are subjective and depend on the skills of reviewers. An automated approach for assisting in code reviews is thus highly desirable.

We present a tool for *assisting in* code review and results from our experiments evaluating the tool in different scenarios. The tool leveraged content available from professional programmer support forums (e.g. StackOverflow.com) to determine potential defectiveness of a given piece of source code. The defectiveness is expressed on the scale of {`Likely defective`, `Neutral`, `Unlikely to be defective`}.

Basic idea employed in the tool is to: *a)* Identify a set P of discussion posts on StackOverflow such that each $p \in P$ contains source code fragment(s) which sufficiently resemble the input code C being reviewed. *b)* Determine the likelihood of C being defective by considering all $p \in P$. A novel aspect of our approach is to use *document fingerprinting* for comparing two pieces of source code. Our choice of *document fingerprinting* technique is inspired by source code plagiarism detection tools where it has proven to be very successful. In the experiments that we performed to verify effectiveness of our approach source code samples from more than 300 GitHub open source repositories were taken as input. A precision of more than 90% in identifying correct/relevant results has been achieved.

Index Terms

Code Review, StackOverflow, Software Development, Crowd Knowledge, Automated Software Engineering

I. Introduction

We present a novel tool that assists in carrying out effective code reviews. Identifying and fixing buggy code consumes significant time and resources in a software development project. Code reviews by peers [1] and experienced programmers is an effective method [2], [3] for identifying potentially buggy code. However, the process of code review is slow and quality of results depends on skills and experience of the reviewers involved. Moreover, a code review carried out by an individual expert is always subjective and hence open for questioning. An automated tool which can improve the quality of code reviews is thus highly desirable. Especially if such a tool can reduce or eliminate the subjectivity associated with individual experts' code reviews, it will be a considerable gain. This is exactly what our tool helps with by leveraging "crowd expertise".

How do programmers acquire "expertise" to become "expert reviewers"?
Mostly one learns from other's experience which may be available in variety of forms such as in a text book, a best practices guide or available in the on-line question and answer forums. It is observed that often times a programmer, when faced with a problem or a bug in some source code, turns to searching an on-line professional programmer forum such as StackOverflow® for assistance. In fact, most software vendors now use StackOverflow[1] as a programmer support channel for their software. As such the information available on professional programmer forums has become a valuable source of experiential knowledge – or "crowd expertise" – about several aspects of software design and development.

Challenges in using "crowd expertise" for assisting in software development tasks
Search engines and information retrieval (IR) technologies have made it easy to locate relevant information on professional programmer support forums, however, in order to zero-in on a suitable solution for the problem at hand a programmer has to manually sift through the content presented by IR tools. One of the reasons why a programmer needs manual sifting through the search results is that normal IR tools may not always take into account the semantic context in which a programmer is operating. As such the ability to derive benefit from the content available on professional programmer support forums is limited by the domain expertise of a programmer and his/her fluency in relevant technical vocabulary.

To address the issues concerning consumption of raw "crowd expertise" researchers have leveraged [4], [5] technologies such as Natural Language Processing (NLP) and Knowledge Discovery (KD). Particularly, the application of NLP techniques for deriving sentiment (on a selected scale) has been a popular direction. However, even the existing text analysis techniques such as CoreNLP [6], [7], Vader [8] etc. may not give accurate results when used for determining defectiveness "sentiment" about a piece

---

[1]Interest in StackOverflow.com has been steadily increasing over the last decade: https://trends.google.com/trends/explore?date=all&q=%2Fm%2F05mw61p

TABLE I: Scale for defectiveness score, α

Value of α	Meaning
-1	Likely to be defective
0	Neutral
1	Not likely to be defective

of source code by analysing the narrative associated with the source code. This is because such text analysis techniques, with their commonly used models, seem to perform poorly when applied to domain specific narrative. For example consider the example post content shown in the Sidebar-1.

It is highly unlikely that a comment such as the narrative in above example about a source code will be considered positive. However, if you run the above narrative through a sentiment analysis tools such as [6], [8] they report the narrative sentiment as "positive", which is misleading if defectiveness of the code that accompany the narrative is judged from the sentiment of that narrative.

In this paper we present a system which bridges such gaps. To understand general working approach of our system consider the following (very simplified) scenario involving professional programmer support forum such as StackOverflow. Suppose that a programmer has encountered a subtle bug in his/her program. He/she then seeks assistance on StackOverflow by describing the issue in a post, p, there. In the post he/she has attached the relevant source code, c, from his/her program. Suppose this question has also been up-voted sufficient number of times by viewers, thus confirming the validity of issue described in the question. Now, if another piece of source code, c', which is being reviewed, sufficiently resembles c, then by association one can infer that c' is *highly likely* to encounter the same issue as reported in the post p.

> ### Sidebar-1 (Example post)
>
> *Consider the following code. This function always goes into infinite loop at sixth line even for the normal inputs.*
>
> ```
> 1.   check = True
> 2.   count = 1
> 3.
> 4.   def flagged_sum(flag, count):
> 5.      sum = 0
> 6.      while check or flag:
> 7.          sum = sum + count
> 8.          print "OK\n"
> 9.      return sum
> 10.
> 11. flagged_sum(True, 2)
> ```

The proposed tool works by identifying discussion posts on StackOverflow such that each post, p, contains source code, c_p, which sufficiently resembles the input code, c_{input}, being reviewed. By analyzing the content (text narrative, meta-data and attached source code c_p) of the post p we determine whether c_p represents a defective code or not. Because c_p sufficiently resembles c_{input}, we can infer the defectiveness of c_{input} itself. Table-I depicts the scale used for specifying/measuring code's defectiveness scores.

The proposed system thus *assists in* code reviews and identification of potentially defective source code. It improves the confidence in a code review by identifying similar scenarios on professional programmer support forums such as StackOverflow.

Paper is organized as follows: We discuss the related work in Section-I-A. Design of the proposed tool is presented in Section-II and its implementation details in Section-III. In Section-IV we provide in-depth discussion on empirical evaluation of our tool and present key observations in Section-V. We have also discussed in Section-V-B some threats to the validity of our approach.

A. Related work

The use of *crowd expertise* for addressing software engineering issues is not entirely new. There have been recent works such as [9], [10] which aim to assist programmers by leveraging content available on professional programmer support forums. Both of these are recommender systems which analyze the source code being typed in an IDE, and generate suitable queries on-the-fly to retrieve relevant discussion threads from StackOverflow. [9] proposes a ranking model to select the most relevant discussion posts from StackOverflow based on the current context (i.e. the source code being typed) in the IDE. [10] is a more primitive version of [9]. They both use Apache Solr[2] document indexing engine to index the StackOverflow posts that are retrieved from the dump of StackOverflow.

Another related approach has been proposed by [11] which suggests API invocation sequences to programmers who are looking to obtain a target object starting with a different object. They depend on code search engines to retrieve relevant code samples from the Web, and then statically analyze such code snippets to identify the method invocation sequences. Yet

---

[2]http://lucene.apache.org/solr/

another novel approach that assists a programmer by suggesting him/her the code snippets is presented in [12]. It suggests to a programmer how an API should be used in a given programming context.

Similarly, [13] proposes a *"Context-Aware IDE-Based Meta Search Engine for Recommendation about Programming Errors and Exceptions"*. They implemented an IDE (Eclipse) plug-in that makes use of the APIs provided by popular web search engines (viz. Google, Yahoo, Bing) to search the entire Web using keywords from a programming context in an IDE. They determine relevant results by taking into consideration the programming context as well as popularity of the candidate search results.

There have also been works such as [14] which infers possible links/relatedness among artifacts of a software development project by analyzing the projects artifact repositories.

Main goal in case of most the existing works such as the ones highlighted above is to suggest code samples or relevant Q&A discussions which can assist the programmer in completing a coding task in an IDE. Emphasis is on assisting a programmer in *writing the code* by leveraging code snippets (or discussions about code) from different sources. Secondly, most of the existing works leverage standard information retrieval (IR) techniques to fetch relevant code/content from various sources such as the Web. For instance, [9], [10] make use of *term frequency–inverse document frequency* based text mining techniques, whereas [11], [13] make use of existing search engines/APIs for retrieving relevant content.

Specifically, in the area of code review there have been sentiment analysis of reviews [9], study on parameters that affect code review [10], and semi-automation in the review process [15]. Here, although [9] differs from our work as it just reviews the comments of three OSS and characterizes them on the basis of their sentiment and [15] proposes a tool based on static analysis of code to automate the code review process. [16] gives a detailed description of various (semi-) automated tools of code reviews avialble for use but concerning specifically to security aspect of the software. These tools are based on static analysis and their performance depends on code quality.

Main aim of the work presented in this paper, however, is to assist in reviewing a given piece of source code. The proposed tool supports the reviewer by identifying relevant content on professional programmer support forums. Secondly, instead of using regular IR approach to identify matching code present in Web content, we leverage a proven and robust *document fingerprinting* technique called Winnowing [17] to identify the matching code. Winnowing has been used in one of the most successful code plagiarism detection tools MOSS[3].

Algorithm 1 Steps in our approach

Require: $PostsDB$:= Database containing posts data (i.e. code, meta-data, text, fingerprints etc.).
Ensure: α := Overall defectiveness score for an input source code unit/file.
    P := Ranked set of relevant posts supporting the obtained probability value.
1: C := Set of self-contained code blocks (strings) in a source code unit.
2: Δl := Code size error tolerance to use when searching for fingerprint matches in $PostsDB$. It is % of input code string length.
3: δ := Fingerprint match threshold to use when searching in $PostsDB$. It is % of fingerprint length of the input code.
4: $\alpha := 0$ /* See Table-I */
5: **for** Every code block string c in C **do**
6:    F_c := Fingerprint calculated on c
       /* F_c is a set of hash values. */
7:    l_c := Length of code string c.
8:    **for** Every record r in $PostsDB$ **do**
9:      **if** $(|l_{r.code} - l_c| \leq \Delta l)$ AND $(|r.fingerprint \cap F_c| \geq \delta)$ **then**
10:        $P := P \cup \{r\}$
11:        $AggregateDefectivenessScore(\alpha, r.defScore)$
12:      **end if**
13:    **end for**
14: **end for**
15: Sort P to select top k matches having highest degree of code match.

II. PROPOSED APPROACH

Association between software development and crowdsourced knowledge has been studied and confirmed by Bogdan et. al. [18] where they studied data from GitHub (a popular repository of open source software) and StackOverflow. The type of

---

questions that are asked and get answered or remain unanswered on StackOverflow has been explored by [19], [20]. One can identify primarily three types of posts on StackOverflow:

Type-1 Questions posted by programmers soliciting help/solution for a programming problem that they are facing with code/API etc.

Type-2 Questions posted by programmers soliciting recommendations/suggestions about a design decisions choice. For example, whether one should use MongoDB or MySQL as data store in a specific application development scenario.

Type-3 Responses posted by programmers to other's questions of the above types.

We also find that there are more than 1.5 times[4] as many posts of Type-3 as there are of Type-1 and 2 combined. Further, we find[5] that more than 73% of Type-1 posts contains source code snippet(s). Each of such posts typically describes some problem involving the source code snippet included in the post.

In view of the above it can be argued that: i) A piece of code accompanying a StackOverflow question is quite likely to be involved in a defect [19], [20], and ii) The code accompanying accepted or high scoring answers to such a question is quite likely to be free from the problem described in associated question post.

Our approach relies on the above results. Main idea of our approach can be stated as follows:

- Professional programmer support forums such as StackOverflow contain numerous posts each containing the description of a software program related problem. Each such post, $p_{question}$, typically contains relevant source code, $c_{defective}$, in full or in part.
- Associated with each post $p_{question}$ there are zero or more reply posts, p_{reply}, which potentially solve the problem described in $p_{question}$. The post p_{reply} also may contain source code, $c_{correct}$, which is a correct version of $c_{defective}$ or somehow solves the problem(s) described in $p_{question}$.
- If a piece of source code, $c_{reviewed}$, is to be reviewed and it resembles sufficiently with the code present in a post $p_{question}$ or p_{reply} then by association we can infer about the quality of $c_{reviewed}$. Depending on whether $c_{reviewed}$ matches significantly with $c_{defective}$ or with $c_{correct}$, we may deduce whether $c_{reviewed}$ is likely to be defective or not.

Sidebar-2 (Winnowing Algorithm for Document Fingerprinting)

Basic goal of the Winnowing algorithm for document fingerprinting is to find substring matches between a given set of document such that they satisfy two properties:

1) All substring matches that are at least as long as some threshold, t, are guaranteed to be detected.
2) Any matches shorter than the noise threshold, k, are not detected.

The constants t and $k \leq t$ are chosen by the user. High-level steps in the algorithm are as follows:

1) Take the document text T to be fingerprinted. E.g. `A do run run run, a do run run.`
2) Transform the T to T' by removing irrelevant features. E.g. `adorunrunrunadorunrun`
3) Generate a sequence of $k-$grams from T'. E.g. A sequence of 5-grams: `adoru dorun orunr runru unrun nrunr runru unrun nruna runad unado nador adoru dorun orunr runru unrun`
4) Generate hashes for each of the k-grams. E.g. `77 74 42 17 98 50 17 98 8 88 67 39 77 74 42 17 98`
5) Identify windows of hashes of length m. E.g. Windows of length 4 for the above example text are: `(77, 74, 42, 17) (74, 42, 17, 98) (42, 17, 98, 50) (17, 98, 50, 17) (98, 50, 17, 98) (50, 17, 98, 8) (17, 98, 8, 88) (98, 8, 88, 67) ( 8, 88, 67, 39) (88, 67, 39, 77) (67, 39, 77, 74) (39, 77, 74, 42) (77, 74, 42, 17) (74, 42, 17, 98)`
6) In each window select the minimum hash value. If there is more than one hash with the minimum value, select the rightmost occurrence. Now save all selected hashes as the fingerprints of the document. E.g. `{17 17 8 39 17}`

A. Design

Our approach is aimed at reviewing source code files, one at a time, that a programmer creates using an IDE or a text editor. Code blocks from source code to be reviewed are used as inputs to a *document fingerprinting* [17] driven search algorithm as described in Algorithm-1 which does the following:

- Identifies posts that contain source code resembling *sufficiently* with the input.
- Determine the defect probability for the code attached with each such post. We use the steps as described in Algorithm-2 for this task.

---

[4]On StackOverflow, there are more than 14m questions and 22m answers as of July 2017. See https://data.stackexchange.com/.
[5]Our query for finding this number is available here: https://data.stackexchange.com/stackoverflow/query/557623.

- Use the above probability scores to assess the defect probability for input code.

A key idea employed in our approach is to calculate "fingerprints" for the code which is attached with posts. Later we search for relevant posts using these fingerprints. Fingerprinting of code is done using Winnowing [17] which is a local algorithm for "document" fingerprinting. Overall process is described in Algorithm-1, and the logical structure of the overall system is shown in Fig. 1a.

B. Deriving defectiveness score for source code attached with a post

Each post on StackOverflow contains useful meta-data (discussed further in Section-III-B1) in addition to text narrative and attached source code. Interesting meta-data attributes of the post are:

1) `Score`: It is an integer value that reflect how many viewers have endorsed or disapprove the post. Each endorsement increments the score while a disapproval decrements it. Higher value of score for a post indicates the genuineness of its contents.

2) `ViewCount`: It is an integer value which reflects how many times this post has been viewed by users. ViewCount of a post can be taken to represent prevalence of the issues being discussed in the post.

To derive the defectiveness score α for a post we consider: a) sentiment of the post's narrative text as computed with NLP tools, b) the `Score` and `ViewCount` values for the post. Algorithm-2 shows the steps involved in determining α for the code present in a given post. Function `CalculateNarrativeSentiment()` makes use of CoreNLP [6] to determine the sentiment score $s_{narrative}$ for text of the post. Function `CalcDefectivenessScore()` applies $s_{narrative}$ to α in order to arrive at final value of the defectiveness score of the post.

Choice of $s_{threshold}$ is made by considering statistical spread of `Score` values for Q&A posts available on StackOverflow. We observed[6] that the average value of `Score` for questions is 1 and that for replies is 2. The standard deviation for both being about 20 indicates that the `Score` values are fairly spread out. Therefore, to be on the conservative side, we chose the value of $s_{threshold}$ to be 1.

Algorithm 2 Calculating defectiveness score α for the code present in a post.

Require: M := Set of meta-data attributes for a post.
$\quad T$:= Text narrative present in the post.
$\quad s_{threshold}$:= Threshold for minimum score for a post.
Ensure: α_p := Defectiveness score for code present in a post, p.
1: $\alpha_p \leftarrow 0$ /* 0 means neutral. See Table-I */
2: **if** $M.PostType = QUESTION$ AND $M.Score > s_{threshold}$ **then**
3: $\quad \alpha_p \leftarrow -1$
4: **else if** $M.PostType = ANSWER$ **then**
5: $\quad$ **if** $M.Score > s_{threshold}$ **then**
6: $\quad\quad \alpha_p \leftarrow +1$
7: $\quad$ **else**
8: $\quad\quad \alpha_p \leftarrow -1$
9: $\quad$ **end if**
10: **end if**
11: $s_{narrative} \leftarrow CalculateNarrativeSentiment(T)$
12: $\alpha_p \leftarrow CalcDefectivenessScore(\alpha_p, s_{narrative})$
13: **return** α_p

---

III. IMPLEMENTATION

Here we describe the salient features of our implementation of the approach presented in previous section (i.e. Section-II). The necessary software program for implementation of our proposed design has been developed using Java programming language. Java was chosen primarily because of: **a)** availability of expertise in Java with us. **b)** Ease of integrating with variety of database engines and other middleware. **c)** Reasonably good performance for the effort devoted to coding.

A. Details of processing flow

The proposed tool relies on availability of certain data artifacts for its correct functioning. These data artifacts are derived from the StackOverflow data dump. Once the data artifacts are setup, the tool can serve code review requests. Thus, broadly, there are two tasks in the processing flow:

---

[6]Our query is available here: https://data.stackexchange.com/stackoverflow/query/759376
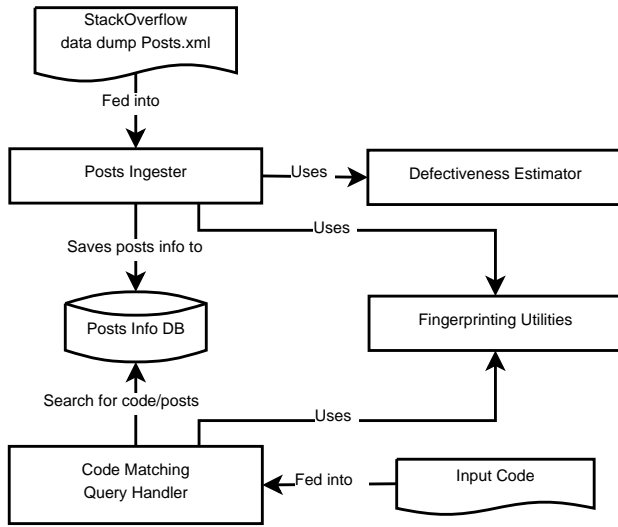
(a) Logical structure of the system.
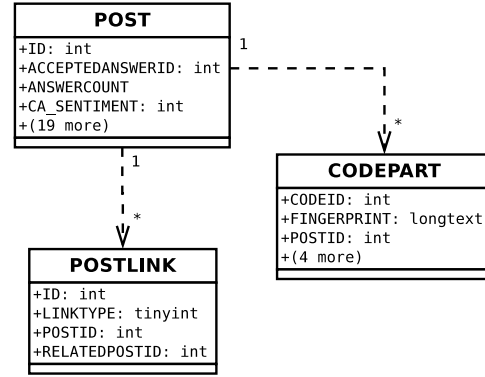
(b) Schema of PostsDB.

Fig. 1: Structure of the complete system and Schema of our database

1) Data preparation: Main task here is to extract the relevant information from posts available in downloaded data dump of StackOverflow. The extracted information includes post's meta-data, textual description of the problem as present in a post and source code snippets if present in the post's body. In addition to meta-data we also need to calculate fingerprints and defectiveness score for the source code snippets extracted from a post. We also extract the information about links (i.e. related and duplicate posts) for a given post. Links/relatedness data is available in StackOverflow data dump itself. Finally all this information is persisted in a database.

2) Code review requests handling: Following are the main steps:

i) Take the source code input by a user. A user may supply a single file or a directory containing multiple files. Each file is processed independently.

ii) For each input source file perform the steps as described in Algorithm-1.

B. Main constituent elements

Overall architecture of the system is shown in Fig. 1a. Following elements are the main constituents of our implementation:

1) StackOverflow Data Dump: Content contributed by professional programmers on Stackoverflow.com is made available to general public under Creative Commons License. We downloaded a snapshot of posts from `https://archive.org/download/stack` (file is $\approx 9.2GB$ in size at the time of download). This archive contains an XML file named Posts.xml whose structure looks like:

```
<posts>
<row ... />
<row ... />
..
</posts>
```

Detailed description of each XML element and its attributes is available here: `https://ia800500.us.archive.org/22/items/` The most relevant attributes of the `row` element for us are: `Id`, `PostTypeId`, `Score`, `ViewCount`, `Body`, `Title` and `Tags`. Values of `Score` and `ViewCount` are useful for deriving the sentiment value of the post. Code fragments to be fingerprinted are present inside the value of `Body` attribute. Value of `Tags` is useful for filtering the query results when searching for posts with specific tags.

2) Posts info database (PostsDB): Information extracted from StackOverflow posts is persisted in an RDBMS for querying by the tool. Schema of such a database is shown in Fig. 1b. It consists of three tables:

i) `POST`: This table stores the meta data found in a post.

ii) `CODEPART`: This table stores the calculated attributes such as the code and text parts extracted from a post, code size, fingerprint for the code etc.

iii) `POSTLINK`: Information about any links among different posts is held in this table. For instance, the information whether a post is duplicate of another post is kept here.

TABLE II: Code size statistics for our posts sample.

ITEM	VALUE
Number of posts considered, N	8339939.00
Average code size (char), c_{avg}	2216.01
Max. code size (char), c_{min}	41802.00
Min. code size (char), c_{max}	1001.00
Std. dev. of code size (char), c_{sd}	2110.38

3) Posts Ingester: Purpose of this module is to extract relevant information from posts found in StackOverflow data dump and persist in the `PostsDB`. Relevant information includes entities such as: *a)* embedded source code fragments in the posts, *b)* fingerprints of the embedded source code and defectiveness score for the post and *c)* meta data such as `ViewCount`, `Score` etc. available with each post.

The information which this module extracts and persists in `PostsDB` will be searched by the `Code Matching Query Handler` (see Sec. III-B4) module when looking for matches with input source code being reviewed.

Determining the input code size threshold: StackOverflow posts also contain very short pieces of text that are marked as code. Examples are function or variable names etc. Such small pieces of code are not useful for our tool. In order for the `Code Matching Query Handler` to produce meaningful results only those posts which have source code above certain minimum size are useful. Therefore the `Post Ingester` persists data of only those posts that contain source code whose size is above certain minimum value.

The `Code Matching Query Handler` module uses a self-contained code block (chosen as a function definition in most programming languages) as the input for its code search operation. Therefore, the size of such a self-contained code block is used as a threshold when ingesting posts data into `PostsDB`. For determining value of this threshold we considered two aspects concerning how programmers typically write source code:

i) The coding best practice recommendations[7] across different programming languages, and

ii) Empirically observed size of code fragments present in the StackOverflow posts.

A common best practice recommendation for writing functions/methods in most programming languages states that length of a function body be such that a programmer can view the entire function within a single page/screen, and that each line of code does not spill beyond 80 columns/characters on screen. Considering the common display resolutions, size of such a function works out to be about 40 lines (each 80 characters long) of code. Assuming that on average 20% of a line is whitespace, the total non-whitespace characters length for a function works out to be $40 \times 80 \times (1 - 0.20) = 2560$ characters.

Table II shows distribution of code fragment sizes that were found in the `PostsDB`. Observed average size (2216) of a code fragment is in close vicinity of the value (2560) obtained based on best practices parameters. However, as shown in Table II, the standard deviation of code size is very close to the mean code size. This implies that the code sizes have large variation and are not centered around the mean. Therefore, the input code size for running evaluation experiments can be safely chosen to be within any range that gives adequate number of samples. Input code size in our experiments lies in 1 – 10kB range.

4) Code Matching Query Handler: Main purpose of this module is to identify relevant posts based on the input source code and the source code fragments found in the posts. Algorithm-1 outlines the processing steps performed by this module.

It can take input in the form of a directory containing multiple source code files or as a single source code file. By default this module breaks an input source file into a collection of self-contained code blocks. Such a code block here typically is a function/method definition. It uses one code block at a time to search for matches in the `PostsDB`. Overall result for a given input source file is determined by combining the results for each self-contained code block.

One problem that this module has to handle is that the code which we have extracted from StackOverflow posts is not always in the form of self-contained code blocks such as a complete function/method definition. In many cases the code that we find in a post consists of one or more snippets taken out of function definition(s). Luckily, the document fingerprinting method (see Winnowing Algorithm side box in Section-II-A) that we use for comparing the source code does not impose any structural/formatting requirements on the pieces of code being compared. Therefore our tool automatically takes care of the problem.

5) Fingerprint Utilities: In order to search the relevant posts, source code matching is performed using *document fingerprints*. Main purpose of this module is to provide implementation of document fingerprinting functions using Winnowing [17] algorithm. An illustration of the Winnowing algorithm is presented in the side box in Section-II-A.

6) Defectiveness Estimator: This module implements our approach for calculating the defectiveness score, α, for a piece of source code which is attached with a StackOverflow post. In order to estimate the value of α for a StackOverflow post p we make use of p's meta-data and the narrative text description from `PostsDB`. The narrative text description of p is used for deriving natural language based sentiment about the post via CoreNLP [6]. Table-I shows the scale used for α.

Section-II-B outlines the approach for estimating value of α, and Algorithm-2 describes the calculation steps.

---

[7]A popular guide from Google Inc. is: https://google.github.io/styleguide/

IV. EVALUATING EFFECTIVENESS OF THE TOOL

The proposed tool can be considered effective if it forms its code review assessment using sufficiently relevant posts, and limits the use of irrelevant posts in its assessment. Once the correct posts which match the input source code have been identified by the tool, estimating defectiveness score of the posts and input code becomes relatively easier (Algorithm-2). Therefore, one of the goals of our evaluation experiments has been to evaluate the performance of our tool in identifying correct posts. In order to determine the efficacy of our tool we used multiple approaches.

First we assessed tool's efficacy based on the quality of code matching (Section-IV-B2). Secondly, we evaluate the proposed tool by using crowd-sourced pairs of $\langle original, duplicate \rangle$ posts (Section-IV-B3). Finally, the tool was also evaluated (Section-IV-B4) by a group of experienced programmers on real projects.

Algorithm 3 Setting up verification data.

Require: PostsDB := Database containing posts data (i.e. code, meta-data, text, fingerprints etc.).
$\quad$ $I :=$ The proposed tool's programmatic interface.
$\quad$ $\psi_{min} :=$ Minimum size of the input code.
$\quad$ $\psi_{max} :=$ Maximum size of the input code.
$\quad$ $\delta_{min} :=$ Fingerprint matching threshold.
Ensure: $T_v :=$ Database table containing verification results.
1: $P_{input}^{unique} =$ Select N $\langle c, p \rangle$ tuples from PostsDB
$\qquad\qquad | \; \psi_{min} \leq \psi \leq \psi_{max}$ and p is unique.
2: $P_{input}^{dup} =$ Select N $\langle c', p', r \rangle$ tuples from PostsDB
$\qquad\qquad | \; \psi_{min} \leq \psi \leq \psi_{max}$ and p' has a duplicate r.
$\quad$ /* Here, p, p', r are posts, and c, c' are code parts belonging to p, p' respectively. ψ is length of code c or c'. */
3: **for** Every code part $c \in P_{input}^{unique}$ **do**
4: $\quad$ $M = I.findMatchingPosts(c, \delta_{min})$
5: $\quad$ **for** Each match $m \in M$ **do**
6: $\quad\quad$ Save match info from $\langle c, p, \delta_{actual}, \psi, m \rangle$ into T_v
7: $\quad$ **end for**
8: **end for**
9: **for** Every code part $c' \in P_{input}^{dup}$ **do**
10: $\quad$ $M' = I.findMatchingPosts(c', \delta_{min})$
11: $\quad$ **for** Each match $m' \in M'$ **do**
12: $\quad\quad$ Save match info from $\langle c', p', r, \delta_{actual}, \psi, m' \rangle$ into T_v
13: $\quad$ **end for**
14: **end for**

A. Key questions considered in empirical evaluation

Primarily, following variables affect the overall performance of our tool:

i) Number and *quality* of posts available in StackOverflow data dump. The *quality* of a post would be considered good if it contains sufficient amount of properly written source code, and also has reliable meta-data such as score, classification tags etc.
ii) Ability of the code matching algorithm to identify correct matches.
iii) Size of the input source code that is being reviewed.
iv) Correct setting of various parameters of the tool and its algorithms.

Contents of the posts available in StackOverflow data dump is fixed (at a given time). Therefore, when testing for the effectiveness of our tool we mainly experiment with the last three variables, i.e., parameters of the code matching algorithm, the input source code and the settings of various parameters of the tool. As such, following are some of the key questions around which we designed our empirical evaluation experiments:

i) What level of precision can be expected for the tool's output? Is there any correlation of precision with other metrics such as level of fingerprint match etc.?
ii) How does the choice of fingerprint matching threshold, δ, affect tool's ability to detect relevant posts?
iii) Does the size of input code size, ψ, affect quality of tool's output? In other words, does the input code size, ψ, affect the number of relevant posts and the defectiveness score reported by the tool?
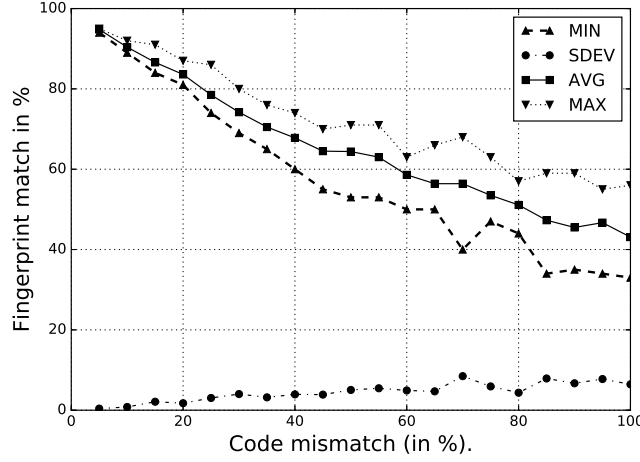
Fig. 2: Fingerprint match vs. Degree of code mismatch.

iv) What is the minimum value of δ which is required to correctly detect a post that has a code part matching, up to a given %age, with the input code?

Two parameters play an important role when using the proposed tool: a) The fingerprint match threshold, δ, and b) The size, ψ, of input source code. Answers to the above questions would help in suitably tuning the tool's parameters for getting the best results.

For carrying out empirical evaluation experiments we developed the necessary programs using Java programming language. Algorithm-3 shows the steps performed by this program.

B. Empirical evaluation

Here we discuss the details of important experimental measurements that we performed to evaluate: **a)** how effective is the proposed tool in identifying StackOverflow posts which are relevant for reviewing input source code, and **b)** effectiveness of the tool when used on real projects by the programmers.

1) Correlation between degree of code mismatch and δ_{actual}: An important question we wanted to answer is that for a given amount of textual difference, ψ_{diff}, between two samples C_1 and C_2 of source code, what is the actual observed amount of fingerprint match δ_{actual}. An answer to this question would offer a guideline for selecting the fingerprint match threshold for different code matching scenarios.

For determining the answer we measured the effect of ψ_{diff} on δ_{actual}. The textual difference ψ_{diff} is measured in terms of aggregate length of the text regions that do not match in the two source code samples.

The code samples C_1 and C_2 are strings synthesized from an alphabet Θ_{code}. For introducing mismatching regions into one of the samples we used short strings synthesized from another alphabet Θ_{noise}. The locations of mismatching regions in a code sample were chosen randomly.

Fig. 2 depicts the observed correlation between δ_{actual} and the degree of code mismatch ψ_{diff}. An important observation is that:

> On average, it is possible to have $\approx 40\%$ match of fingerprints for two completely different code samples. For detecting two code samples which are at least 50% similar, their fingerprints match should be $\geq 60\%$.

2) Code matching performance: An important part of our tool is the algorithm that we employ for handling contextual search of source code parts in `PostsDB`. Effectiveness of the proposed tool depends to a great extent on accuracy of the method used for source code matching. This is because the relevance of search results is proportional to the degree to which an input source code matches with code parts[8] found in `PostsDB`. For performing source code matching we have employed a well known technique called Winnowing [17] which is a local algorithm for "document" fingerprinting. The code matching effectiveness and limitations that such algorithm has will also be applicable to our tool.

---

[8]A post has one or more pieces of code embedded in it. Each such embedded piece of code is called a "part". A "part" may or may not be a complete program.

(a) Effect of fingerprint match, δ_{actual}.   (b) Effect of input code size, ψ.   (c) Precision, ξ.
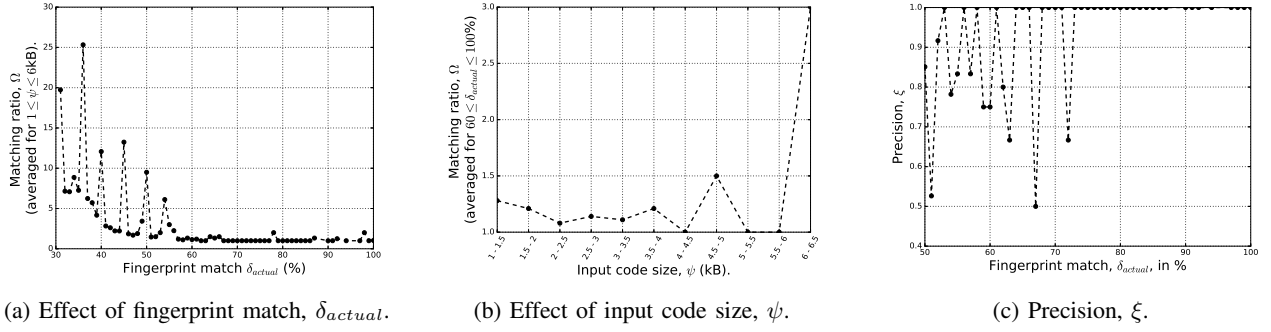
Fig. 3: Precision and matching ratio performance (when input taken from unique posts).

In order to evaluate and compare the performance of proposed tool for various scenarios in our experiments we use the following metrics:

i) We define a metric, Ω, that we call matching ratio, as in Equation-1.
ii) We use the standard metric of *precision*, as defined in Equation-2.

$$\Omega = \frac{\sum_{i=1}^{N} |M_i|}{N} \tag{1}$$

$$\xi = \frac{tp}{tp + fp} \tag{2}$$

Here, $|M_i|$ is the number of matched posts for i^{th} input, and N is the total number of inputs. M (or M') and N are as described in Algorithm-3. Ω gives an average measure of the matches found per input. tp is the number of true positive results, and fp is the number of false positives observed in an experiment. A string token matching algorithm was used to identify *false positive* code matches (we also manually verified several randomly selected false positives).

To ensure the correct implementation of Winnowing based code matching algorithm we have evaluated its performance in two scenarios:

i) When code input to the tool was substantially similar to the code present in a post.
ii) When code input to the tool matched only partially with the code present in a post.

In our experiments we observed that for both of the above mentioned scenarios the expected matches were always identified by the tool. That is, *false negatives* were zero. Section-V gives a detailed discussion of the results.

3) Evaluation using input code taken from StackOverflow posts: Rationale for this approach is based on a pattern that we observed in posts on StackOverflow. We found that there are sufficient number[9] of posts on StackOverflow which are flagged by the moderators (at StackOverflow) as duplicate of some other post. A rigorous crowd-sourced review approach[10] is employed by StackOverflow before labeling a post as duplicate. Availability of such programmer verified $\langle original, duplicate \rangle$ pairs of posts in StackOverflow data offers us very good quality test data set for evaluation of our tool. Correctness of our tool is verified if the tool reports in its results both the original post and its duplicate(s), for an input which included the code from only the original post.

Following are the steps that we performed in this scenario:

i) Select a random set P_{test} of $\langle p^{orig}, p^{dup} \rangle$ tuples of posts from `PostsDB`.
ii) For each tuple $\langle p_i^{orig}, p_i^{dup} \rangle \in P_{test}$ extract the code parts C_i from p_i^{orig}.
iii) Run the tool on C_i to get a set M of relevant posts.
iv) If the set M contains both p_i^{orig} and p_i^{dup}, then we deduce that the tool worked correctly to identify the desired posts.

Results are depicted in Fig. 3 and 4 and further discussed in Section-V.

---

[9] We composed a query (available here: https://data.stackexchange.com/stackoverflow/query/620299) to estimate the number of posts on StackOverflow which have a tag `Java` and have been marked to have a possible duplicate. The query reports about 398409 posts to be duplicate as of Nov 2017.

[10] Please refer to this document: http://meta.stackexchange.com/questions/10841/how-should-duplicate-questions-be-handled and the database schema here: https://data.stackexchange.com/stackoverflow/query/new

4) Evaluation by programmers on real projects: The tool was also evaluated by a group of programmers who were actively involved in developing code for real projects. Evaluation was performed on source code of open source projects available on GitHub. Table-III shows the relevant information about those projects. Following are the steps in our evaluation process:

i) Select a set, G, of open source projects from GitHub repositories. *In our current evaluation set up we considered mostly Java source code.*

ii) For each project $g \in G$, randomly select a set S_g of source files.

iii) Run the tool for each source file $s \in S_g$ to produce a set M of relevant posts.

iv) A team of experienced programmers checks the relevance of posts in M.

We developed a program to fetch source code files from open source GitHub repositories. The program made use of GitHub search API to search and fetch the source files from repositories. We downloaded only those source files whose size was between 1 to 10kB. Other than source code file size we did not impose any specific constraints on size of teams or the size of code bases for respective projects.

Observations from this evaluation are discussed in Section-V-A2.

TABLE III: Evaluation of the tool on GitHib OSS projects.

No. of repositories	Average repository size in MB	Average No. of contributors	Main language
303	4.34	2.47	Java
21	3.24	1.81	JavaScript
6	6.10	0.67	HTML
5	18.19	7.60	C++
4	2.14	1.00	CSS
3	16.30	1.33	PHP
3	89.84	1.00	C

5) Defectiveness score validation: As described in Section-II-B, determining defectiveness score for the input code is relatively simple once the correct matching posts have been identified. To assess relevance of defectiveness scores as determined by our tool we have compared the defectiveness score α^p of a post p determined by our tool with the score $\alpha^p_{reference}$ calculated for the same post using different NLP based tools such as CoreNLP [6] and VADER [8]. We adopted the following approach for carrying out this comparison:

For each matched post p we do:

i) Calculate defectiveness score α^p as described in Algorithm-1 and 2.

ii) Calculate[11] defectiveness score $\alpha^p_{reference}$ for the <u>narrative text</u> found in post p.

iii) Compare the values $\alpha^p_{reference}$ with α^p to find the overall degree of concurrence. We calculate the fraction Γ as:

$$\Gamma = \frac{\text{No. of posts for which } \alpha^p \text{ matches } \alpha^p_{reference}}{\text{Total number of posts.}} \tag{3}$$

V. Key observations from evaluation experiments

A major goal of our tool is to determine defectiveness score for code parts present in StackOverflow posts so that we can estimate the defectiveness score for a similar input source code. Also, identifying StackOverflow posts which contain source code snippets that sufficiently resemble the input code remains an important step in our approach. As such we designed our experiments to evaluate our tool for code matching performance as well as correctness of reported defectiveness scores for the input code. Evaluation has been carried out for a wide range of inputs under different parameter settings of the tool as discussed next.

TABLE IV: Key statistics for evaluation using input from StackOverflow posts.

ITEM	VALUE (for $\delta_{actual} \geq 50\%$)	
	For unique inputs	For inputs with a duplicate
Total number of input files	500	500
Input files having a match	500	500
	⟨Average, Minimum, Maximum, Std. Deviation⟩	
Fingerprint match, δ_{actual} (in %)	⟨83.71, 50, 100, 21.62⟩	⟨64.26, 50, 100, 17.34⟩
Input code size, ψ (kB)	⟨2.0, 1.0, 6.2, 1.0⟩	⟨2.0, 1.0, 6.4, 0.9⟩
Matching ratio, Ω	⟨1.72, 1.00, 141.00, 6.49⟩	⟨7.61, 1.00, 758.00, 43.73⟩
Precision, ξ (in %)	⟨96.78, 66.67, 100, 8.48⟩	⟨97.77, 50.55, 100, 7.98⟩

---

[11]Calculated using the implementation of CoreNLP from: https://stanfordnlp.github.io/CoreNLP/ and VADER from https://github.com/cjhutto/vaderSentiment

(a) Effect of fingerprint match, δ_{actual}. (b) Effect of input code size, ψ. (c) Precision, ξ.
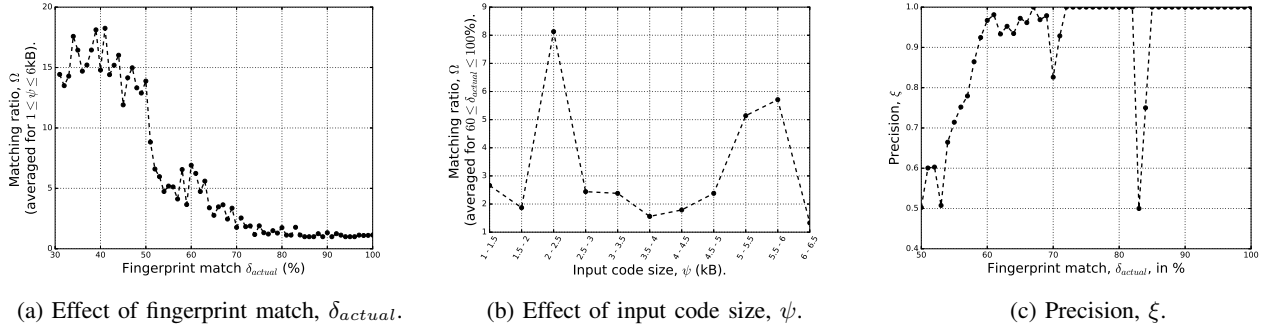
Fig. 4: Precision and matching ratio performance (when input taken from posts having a duplicate).

A. Observations from empirical evaluations

1) Effect of δ and ψ on matching ratio Ω: We experimented with two scenarios: **i)** where input code was chosen from existing posts that did not have a duplicate post, and **ii)** where input code was chosen from posts that had duplicate posts. Important statistics seen from our experiments are summarized in Table-IV.

Fig. 3 depicts observations for scenario **i** where we see following salient points:

- The cut-off point for fingerprint match δ_{actual} after which the matching ratio Ω remains stable at its expected value of 1 is $\approx 57\%$. The average value of precision is about 96% for $\delta_{actual} \geq 60\%$. In other words, number of false positives becomes very small for $\delta_{actual} > 60\%$.
- Ω and number of false positives rise sharply as the value of δ_{actual} falls below 60%.
- Overall average precision, ξ, of the results remain at $\approx 93\%$ (trend is shown in Fig. 3c).

Fig. 4 shows observations for scenario **ii**, and the salient points are:

- Ω settles at about 2 (expected value) when fingerprint match $\delta_{actual} \geq 72\%$. The average precision also remain about 97% for $\delta_{actual} \geq 70\%$.
- Overall average precision, ξ, in this case is $\approx 91\%$ (trend is shown in Fig. 4c).

It is difficult to deduce any clear correlation between the size, ψ, of input code and the matching ratio, Ω, as seen from Fig. 3b and 4b.

Possible reasons for observed trend: A near independence of Ω from size ψ of input code, is likely due to the large standard deviation observed in the sizes of source code present in StackOverflow posts. As we see in Table-II, the value of standard deviation is as large as the mean. Such a large value of standard deviation implies that there is significant dispersion in code sizes. Therefore, the variation of input code size seems to have little effect on the matching ratio.

2) Evaluation on real projects: Effectiveness of the proposed tool was also tested by using input source code from open source project repositories at GitHub. We used 370 source code files (size between 1 - 10kB each) from 345 different repositories hosted at GitHub. Key statistics from our testing are summarized in Table-V.

TABLE V: Key statistics for GitHub code evaluation.

ITEM	VALUE	
	For $\delta_{actual} \geq 30\%$	For $\delta_{actual} \geq 50\%$
Total number of input files	370	370
Input files having a match	50	2
	$\langle$**Average, Minimum, Maximum, Std. Deviation**$\rangle$	
Fingerprint match, δ_{actual} (in %)	$\langle 34.2, 31.0, 53.0, 3.6 \rangle$	$\langle 51.4, 50.0, 53.0, 0.97 \rangle$
Input code size, ψ (kB)	$\langle 7.3, 1.1, 9.9, 2.3 \rangle$	$\langle 9.6, 9.2, 9.9, 0.35 \rangle$
Matching ratio, Ω	$\langle 20.6, 1.0, 208.0, 38.2 \rangle$	$\langle 5.5, 5.0, 6.0, 0.5 \rangle$
Precision, ξ (in %)	$\langle 97.34, 89.66, 100, 3.68 \rangle$	$\langle 100, 100, 100, 0.0 \rangle$

Fig. 5 depicts the variation in matching ratio, Ω, w.r.t. observed fingerprint match, δ_{actual}, and size, ψ, of the input source code. The observed general trends are similar to the ones seen in previous scenarios, though the absolute values are different. Notable observations are:

- The largest value observed for fingerprint match, δ_{actual}, is about 53%.
- Less than 1% of the input files find a relevant matching post having $\delta_{actual} \geq 50\%$.
- For $\delta_{actual} \geq 50\%$ the value of $\Omega \leq 3$.
- Average precision, ξ, of the results is about 97%.

(a) Effect of fingerprint match, δ_{actual} ($\geq$ 30% only.)

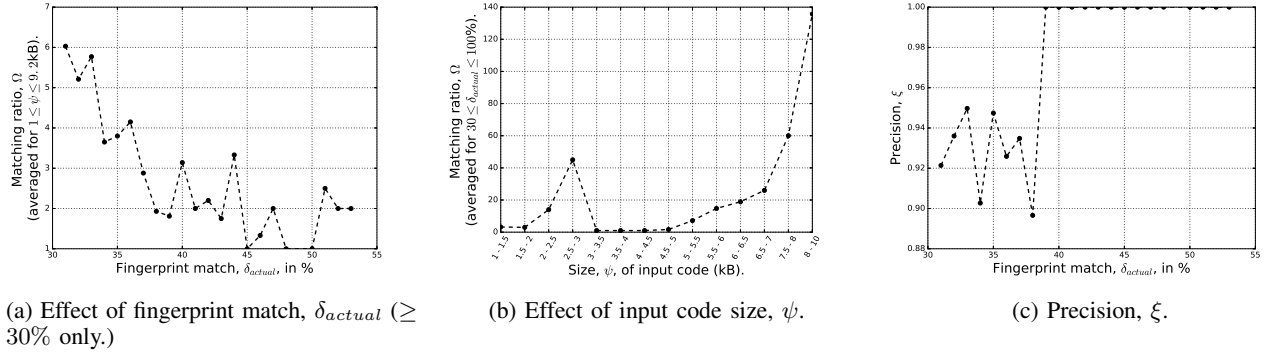(b) Effect of input code size, ψ.

(c) Precision, ξ.

Fig. 5: Evaluation with input code taken from GitHub repositories.

Interestingly, our observation that very few input files find a sufficiently relevant/matching post seems to validate the results reported by Bogdan et. al. in [18] where one of their observation is:

"Active GitHub committers ask fewer questions on StackOverflow than others."

3) Correctness of defectiveness score: A post's `SCORE` value (available in the post's mete-data) indicates its acceptance by the professional programmers. Likelihood of a post being representative of a real and genuine programming issue/scenario is proportional to the `SCORE` value achieved by the post. This is the main reason we have used `SCORE` value of a post for calculating its defectiveness score in Algorithm-2. Table-VI shows important statistics about the `SCORE` values observed for different types of posts available in StackOverflow data dump that we used.

TABLE VI: Statistics about calculated defectiveness scores.

Item	Item value for different post types		
	Questions	Replies	Accepted replies
Average `SCORE` value	4.23	4.08	7.11
Minimum `SCORE` value	-147	-58	-54
Maximum `SCORE` value	9432	11055	11055
Std. Deviation of `SCORE`	28.49	26.8	41.08
Γ in % from Eq.-(3)	42	29	32

We observe that $\Gamma \approx 42\%$ (Equation-3) for the posts of type *question*. That is, the defectiveness scores α and $\alpha_{reference}$ matched for about 42% of such posts. This number was 29% and 32% for all posts of type *reply* and *accepted reply* respectively. These numbers were observed when we set the $\alpha_{threshold}$ value in Algorithm-2 equal to average value of `SCORE` observed in the StackOverflow data.

Possible reasons for observed trend: We analysed the low values for Γ by manually examining a subset of recommendations produced by our tool. We observed that the text narrative available in StackOverflow posts makes use of vocabulary which is specific to the domain of software development. A polar word/phrase in such a narrative and vocabulary may or may not be considered polar from normal English language perspective. For example consider the following sentence:

```
``The following piece of code takes a huge amount of memory and CPU, and takes very
long to produce results.''
```

A narrative similar to the above sentence about a piece of code is highly likely to be considered negative. However, most existing NLP based tools would label this sentence as either a "neutral" or "positive" comment about the code it refers to. For instance, as of December 2017, the CoreNLP [6] Sentiment Analysis tool (available live at http://corenlp.run) labels this sentence as positive. Similarly, Vader [8] too labels it as positive.

We manually checked a random subset of instances where the defectiveness scores α and $\alpha_{reference}$ did not match. In all such cases the defectiveness score assigned by our algorithm was found to be correct.

B. Threats to validity

The basic premise on which our tool works is that:

- the code present in a question post or a poorly rated reply post is highly likely to be of poor quality, and
- the code present in an accepted answer post or a highly rated reply post is highly likely to be of good quality.

Often there are more than one correct ways to write code for certain scenarios. Sometimes a solution coded in a particular style gets higher rating/acceptance [21], [22] mainly because it matches the personal/organizational coding style of the programmer

who posted the original question. In such cases even though the code present in low rated reply posts may be defect free, there are chances that such code may get a poor defectiveness score in relative terms.

Another possible shortcoming may be in the manner by which the `Code Matching Query Handler` processes input code. Irrespective of whether the tool uses an input code file as a single code block or breaks it into smaller basic blocks, there remains a possibility of leaving out relevant matches available in `PostsDB`. The reason why this may happen is as follows: Code present in a post may be anything from a complete and compilable unit (e.g. a Java class definition) pasted as-is, to a small fragment of it (e.g. a while loop from a function definition). Thus, if the tool uses a full compilation unit as-is to search for matches in `PostsDB`, then those posts which contain only a small fragment from such code may not show up as a match. Similar situation occurs for the reverse scenario, i.e., when input is a small fragment but a post contains a complete program.

Also, in our current implementation we do not pre-process the source code to neutralize the effect of coding style variations such as different naming style for variables etc. Incorporating such pre-processing can enhance the effectiveness of our tool.

Lastly, limitations and parameter settings of the code fingerprinting algorithm [17] that we use can also affect the quality of results of our tool. For example, when implementing this fingerprinting algorithm the choice of a hash function that the algorithm uses affects the quality of code matching.

VI. Conclusion

Lowering the cost of creating and maintaining good quality software is undoubtedly an important goal in software development. Researchers have shown that improving code quality can significantly lower the overall cost of software development and maintenance. In this paper we have proposed a tool that helps a programmer in writing better code by pointing out potentially buggy portions of the code.

The discussions available on professional programmer support forums such as StackOverflow are used by our tool for assessing defectiveness of input code. Defectiveness of the input source code is estimated by inferring the "sentiment" of discussion posts that contain any code that is similar to the input code. A key idea in our approach is the use of a document fingerprinting technique for efficient and accurate comparison of source code when searching for similar code in discussion posts. The document fingerprinting techniques have been very successfully used in source code plagiarism detection tools. They make a good choice for us because the code matching scenario that our tool faces is similar to the source code plagiarism detection.

Efficacy of the proposed tool has been verified by checking the correctness of results along three important dimensions: a) by measuring code matching accuracy, b) by verifying results for known $\langle input, output \rangle$ pairs, and c) by verifying the results for code taken from real projects such as from GitHub.

Our experiments have shown that the document fingerprinting based search approach performs well in identifying relevant posts that contain code similar to an input code. We have shown that our tool performs better than existing NLP based techniques for determining defectiveness of a given piece of source code.

References

[1] I. Sommerville, *Software engineering*. Pearson, 2010.

[2] S. McConnell, *Code complete: a practical handbook of software construction*. Microsoft Press. (Redmond, WA), 1993.

[3] D. Huizinga and A. Kolawa, *Automated defect prevention: best practices in software management*. John Wiley & Sons, 2007.

[4] E. Cambria and B. White, "Jumping nlp curves: A review of natural language processing research," *IEEE Computational intelligence magazine*, vol. 9, no. 2, pp. 48–57, 2014.

[5] R. McEntire, D. Szalkowski, J. Butler, M. S. Kuo, M. Chang, M. Chang, D. Freeman, S. McQuay, J. Patel, M. McGlashen *et al.*, "Application of an automated natural language processing (nlp) workflow to enable federated search of external biomedical content in drug discovery and development," *Drug discovery today*, vol. 21, no. 5, pp. 826–835, 2016.

[6] C. D. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. J. Bethard, and D. McClosky, "The Stanford CoreNLP natural language processing toolkit," in *Association for Computational Linguistics (ACL) System Demonstrations*, 2014, pp. 55–60. [Online]. Available: http://www.aclweb.org/anthology/P/P14/P14-5010

[7] R. Socher, A. Perelygin, J. Y. Wu, J. Chuang, C. D. Manning, A. Y. Ng, C. Potts *et al.*, "Recursive deep models for semantic compositionality over a sentiment treebank," in *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*, vol. 1631, 2013, p. 1642.

[8] C. J. Hutto and E. Gilbert, "Vader: A parsimonious rule-based model for sentiment analysis of social media text," in *Eighth international AAAI conference on weblogs and social media*, 2014.

[9] L. Ponzanelli, G. Bavota, M. Di Penta, R. Oliveto, and M. Lanza, "Mining stackoverflow to turn the ide into a self-confident programming prompter," in *Proceedings of the 11th Working Conference on Mining Software Repositories*, ser. MSR 2014. New York, NY, USA: ACM, 2014, pp. 102–111.

[10] L. Ponzanelli, A. Bacchelli, and M. Lanza, "Seahawk: Stack overflow in the ide," in *35th International Conference on Software Engineering (ICSE)*, May 2013, pp. 1295–1298.

[11] S. Thummalapenta and T. Xie, "Parseweb: A programmer assistant for reusing open source code on the web," in *Proceedings of the Twenty-second IEEE/ACM International Conference on Automated Software Engineering*, ser. ASE '07. New York, NY, USA: ACM, 2007, pp. 204–213.

[12] A. Mishne, S. Shoham, and E. Yahav, "Typestate-based semantic code search over partial programs," in *ACM SIGPLAN Notices*, vol. 47, no. 10. ACM, 2012, pp. 997–1016.

[13] M. M. Rahman, S. Yeasmin, and C. K. Roy, "Towards a context-aware ide-based meta search engine for recommendation about programming errors and exceptions," in *2014 Software Evolution Week - IEEE Conference on Software Maintenance, Reengineering, and Reverse Engineering (CSMR-WCRE)*, Feb 2014, pp. 194–203.

[14] D. Čubranić and G. C. Murphy, "Hipikat: Recommending pertinent software development artifacts," in *Proceedings of the 25th International Conference on Software Engineering*, ser. ICSE '03, 2003, pp. 408–418.

[15] T. Di Noia, M. Mongiello, and E. Di Sciascio, "Ontology-driven pattern selection and matching in software design," in *Software Architecture*. Springer, 2014, pp. 82–89.

[16] G. McGraw, "Automated code review tools for security," *Computer*, vol. 41, no. 12, 2008.

[17] S. Schleimer, D. S. Wilkerson, and A. Aiken, "Winnowing: local algorithms for document fingerprinting," in *Proceedings of the 2003 ACM SIGMOD international conference on Management of data*. ACM, 2003, pp. 76–85.

[18] B. Vasilescu, V. Filkov, and A. Serebrenik, "Stackoverflow and github: Associations between software development and crowdsourced knowledge," in *2013 International Conference on Social Computing*, Sept 2013, pp. 188–195.

[19] C. Treude, O. Barzilay, and M. A. Storey, "How do programmers ask and answer questions on the web?: Nier track," in *33rd International Conference on Software Engineering (ICSE)*, May 2011, pp. 804–807.

[20] S. Wang, D. Lo, and L. Jiang, "An empirical study on developer interactions in stackoverflow," in *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, ser. SAC '13. New York, NY, USA: ACM, 2013, pp. 1019–1024.

[21] D. Movshovitz-Attias, Y. Movshovitz-Attias, P. Steenkiste, and C. Faloutsos, "Analysis of the reputation system and user contributions on a question answering website: Stackoverflow," in *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, ser. ASONAM '13. New York, NY, USA: ACM, 2013, pp. 886–893.

[22] V. Honsel, S. Herbold, and J. Grabowski, "Intuition vs. truth: Evaluation of common myths about stackoverflow posts," in *2015 IEEE/ACM 12th Working Conference on Mining Software Repositories*, May 2015, pp. 438–441.