# Evaluating Complexity, Code Churn, and Developer Activity Metrics as Indicators of Software Vulnerabilities

Yonghee Shin, Andrew Meneely, Laurie Williams, *Member*, *IEEE*, and Jason A. Osborne

**Abstract**—Security inspection and testing require experts in security who think like an attacker. Security experts need to know code locations on which to focus their testing and inspection efforts. Since vulnerabilities are rare occurrences, locating vulnerable code locations can be a challenging task. We investigated whether software metrics obtained from source code and development history are discriminative and predictive of vulnerable code locations. If so, security experts can use this prediction to prioritize security inspection and testing efforts. The metrics we investigated fall into three categories: complexity, code churn, and developer activity metrics. We performed two empirical case studies on large, widely used open-source projects: the Mozilla Firefox web browser and the Red Hat Enterprise Linux kernel. The results indicate that 24 of the 28 metrics collected are discriminative of vulnerabilities for both projects. The models using all three types of metrics together predicted over 80 percent of the known vulnerable files with less than 25 percent false positives for both projects. Compared to a random selection of files for inspection and testing, these models would have reduced the number of files and the number of lines of code to inspect or test by over 71 and 28 percent, respectively, for both projects.

**Index Terms**—Fault prediction, software metrics, software security, vulnerability prediction.

✦

## 1 INTRODUCTION

A single exploited software vulnerability[1] can cause severe damage to an organization. Annual world-wide losses caused from cyber attacks have been reported to be as high as $226 billion [2]. Loss in stock market value in the days after an attack is estimated from $50 to $200 million per organization [2]. The importance of detecting and mitigating software vulnerabilities before software release is paramount.

Experience indicates that the detection and mitigation of vulnerabilities are best done by engineers specifically trained in software security and who "think like an attacker" in their daily work [3]. Therefore, security testers need to have specialized knowledge in and a mindset for what attackers will try. If we could predict which parts of the code are likely to be vulnerable, security experts can focus on these areas of highest risk. One way of predicting

vulnerable modules is to build a statistical model using software metrics that measure the attributes of the software products and development process related to software vulnerabilities. Historically, prediction models trained using software metrics to find faults have been known to be effective [4], [5], [6], [7], [8], [9], [10].

However, prediction models must be trained on what they are intended to look for. Rather than arming the security expert with all the modules likely to contain faults, a security prediction model can point toward the set of modules likely to contain what a security expert is looking for: security vulnerabilities. Establishing predictive power in a security prediction model is challenging because security vulnerabilities and non-security-related faults have similar symptoms. Differentiating a vulnerability from a fault can be nebulous even to a human, much less a statistical model. Additionally, the number of reported security vulnerabilities with which to train a model are few compared to nonsecurity-related faults. Colloquially, security prediction models are "searching for a needle in a haystack."

In this paper, we investigate the applicability of three types of software metrics to build vulnerability prediction models: complexity, code churn, and developer activity (CCD) metrics. Complexity can make code difficult to understand and to test for security. Frequent or large amount of code change can introduce vulnerabilities. Poor developer collaboration can diminish project-wide secure coding practices.

*The goal of this study is to guide security inspection and testing by analyzing if complexity, code churn, and developer activity metrics can be used to 1) discriminate between vulnerable and neutral files and 2) predict vulnerabilities.* For this purpose, we performed empirical case studies on two widely used, large-scale open-source projects: the Mozilla

---

1. An instance of a [fault] in the specification, development, or configuration of software such that its execution can violate an [implicit or explicit] security policy [1].

- *Y. Shin is with the College of Computing and Digital Media, DePaul University, 243 S. Wabash Ave., Chicago, IL 60614.*
  *E-mail: yshin2@ncsu.edu.*
- *A. Meneely, and L. Williams are with the Department of Computer Science, North Carolina State University, 3231 EB II, 890 Oval Drive, Campus Box 8206, Raleigh, NC 27695-8206.*
  *E-mail: apmeneel@ncsu.edu, williams@csc.ncsu.edu.*
- *J.A. Osborne is with the Department of Statistics, North Carolina State University, 5238 SAS Hall, Campus Box 8203, Raleigh, NC 27695-8203.*
  *E-mail: jaosborn@stat.ncsu.edu.*

Firefox[2] web browser and the Linux kernel as distributed in Red Hat Enterprise Linux.[3] We analyzed Mozilla Firefox and Red Hat Enterprise Linux (each of them containing over two million lines of source code), and evaluated the adequacy of using CCD metrics as indicators of security vulnerabilities. We also measured the reduction in code inspection effort using CCD metrics against random file selection.

The rest of the paper is organized as follows: Section 2 defines background terms. Section 3 provides our hypotheses and the methodology for our case studies. In Sections 4 and 5, we provide the results from both case studies. Section 6 provides the summary of results from the case studies and discusses issues in the use of models in practice. Section 7 discusses the threats to validity of our study. Section 8 provides related work. Section 9 contains conclusions and future work.

## 2 BACKGROUND

This section describes the background on discriminative power, predictability, network analysis, and the binary classification evaluation criteria that we will use in this paper.

### 2.1 Discriminative Power and Predictability

This study investigates the *discriminative power* and *predictability* of CCD metrics in the realm of software security. *Discriminative power* [11] is defined as the ability to "discriminate between high-quality software components and low-quality software components." In our study, discriminative power is the ability to discriminate the vulnerable files from neutral files. We classify a file as vulnerable if it has been fixed postrelease for a vulnerability and as a neutral file if no vulnerabilities have been found at the time of our analysis.

*Predictability* [11] is the ability of a metric to identify files that are likely to have vulnerabilities with metric values available prior to software release. Note that discriminative power is measured for a single metric with the vulnerability information available at the present time for the purpose of controlling the metric values by redesign or reimplementation or to investigate the feasibility of building prediction models using the given metrics. On the other hand, predictability can be measured using single or multiple metrics together and the data obtained from past and present are applied to predict the code locations that are likely to have vulnerabilities in the future.

### 2.2 Network Analysis

In this paper, we use several terms from network analysis [12] and provide their meaning with respect to developer and contribution networks that will be discussed in Section 3.1. Network analysis is the study of characterizing and quantifying network structures, represented by graphs [12]. A sequence of nonrepeating, adjacent nodes is a *path*, and a shortest path between two nodes is called a *geodesic path* (note that geodesic paths are not necessarily unique).

*Centrality* metrics are used to quantify the location of one node relative to other nodes in the network. The centrality metrics we use for developer activity are degree, closeness, and betweenness. The *degree* metric is defined as the number of neighbors directly connected to a node. The *closeness* centrality of node $v$ is defined as the average distance from $v$ to any other node in the network that can be reached from $v$. The *betweenness* centrality [12] of node $v$ is defined as the number of geodesic paths that include $v$.

*Cluster* metrics are used to measure the strength of interconnection between groups of nodes. A *cluster* of nodes is a set of nodes such that there are more edges within a set of nodes (intraset edges) than edges between a set and other sets of nodes (interset edges). The cluster metric we use for developer activity is *edge betweenness* [13]. The edge betweenness of edge $e$ is defined as the number of geodesic paths that pass through $e$. Since clusters have many intracluster edges, edges within clusters have a low betweenness; conversely, edges between two clusters have a high betweenness [13].

### 2.3 Binary Classification Evaluation Criteria

In this study, we use a binary classification technique to predict files that are likely to have vulnerabilities. A binary classifier can make two possible errors: *false positives (FPs)* and *false negatives (FNs)*. In this study, an *FP* is the classification of a neutral file as a vulnerable file, and an *FN* is the classification of a vulnerable file as neutral. False positives represent excessive files to inspect or test, and false negatives increase the chances of vulnerabilities escaping to the field without inspection/testing. A correctly classified vulnerable file is a *true positive (TP)*, and a correctly classified neutral file is a *true negative (TN)*.

For evaluating binary classification models, we use probability of detection (PD) and probability of false alarm (PF).

*Probability of detection*, also known as *recall*, is defined as the ratio of correctly predicted vulnerable files to actual vulnerable files:

$$PD = TP/(TP + FN).$$

*Probability of false alarm*, also known as *false positive ratio*, is defined as the ratio of files incorrectly predicted as vulnerable to actual neutral files:

$$PF = FP/(FP + TN).$$

The desired result is to have a high *PF* and a low *PF* to find as many vulnerabilities as possible without wasting inspection or testing effort. Having a high *PD* is especially important in software security considering the potentially high impact of a single exploited vulnerability.

Additionally, we provide *precision (P)*, defined as the ratio of correctly predicted vulnerable files to all detected files:

$$P = TP/(TP + FP).$$

We report *P* with *PD* as they tend to trade off each other. However, security practitioners might want a high *PD* even at the sacrifice of *P* [14] considering the severe impact of one exploited vulnerability. Additionally, *precision* and *accuracy* (another frequently used prediction performance criterion that measures overall correct classification) are known to be poor indicators of performance for highly unbalanced data

TABLE 1
Definitions of Complexity Metrics

| Related Hypothesis | Metric | Definition |
|---|---|---|
| $H_{IntraComplexity\_D}$ | CountLineCode | The number of lines of code in a file |
| | CountDeclFunction | The number of functions defined in a file |
| | CountLineCodeDecl | The number of lines of code in a file devoted to declarations |
| | CountLinePreprocessor | The number of lines of code in a file devoted to preprocessing |
| | SumEssential | The sum of essential complexity, where essential complexity is defined as the number of branches after reducing all the programming primitives such as a  for loop in a function's control flow graph into a node iteratively until the graph cannot be reduced any further. Completely well-structured code has essential complexity 1 [18], [20] |
| | SumCyclomaticStrict | The sum of the strict cyclomatic complexity,  where strict cyclomatic complexity is defined as the number of conditional statements in a function |
| | SumMaxNesting | The sum of the MaxNesting in a file, where MaxNesting is defined as the maximum nesting level of control constructs such as if or while statements in a function |
| | MaxCyclomaticStrict | The maximum of strict cyclomatic complexity in a file |
| | MaxMaxNesting | The maximum of MaxNesting in a file |
| $H_{Coupling\_D}$ | SumFanIn | The sum of FanIn, where  FanIn is defined as  the number of inputs to a function such as parameters and global variables |
| | SumFanOut | The sum of FanOut, where FanOut is defined as the number of assignment to the parameters to call a function or global variables |
| | MaxFanIn | The maximum of FanIn |
| | MaxFanOut | The maximum of FanOut |
| $H_{Comments\ D}$ | CommentDensity | The ratio of lines of comments to lines of code |

set where the number of data instances in one class is much more than the data instances in another class [7], [14]. In our case, less than 1.4 percent of the files are vulnerable in both projects. A naive classification that classifies all files as neutral would have an *accuracy* of 98.6 percent.

## 3 RESEARCH HYPOTHESES AND METHODOLOGY

In this section, we provide the hypotheses for the discriminative power and the predictability of the CCD metrics. Although CCD metrics have been known to be effective for fault prediction [7], [15], [16], we provide the rationale for why these metrics can also work for vulnerabilities. Then, we provide the evaluation criteria for the hypotheses and the methods to test the hypotheses and explain prediction modeling techniques.

### 3.1 Hypotheses for Discriminative Power

#### 3.1.1 Code Complexity

Security experts claim that complexity is the enemy of security [3], [17]. Complexity can lead to subtle vulnerabilities that are difficult to test and diagnose [17], providing more chances for attackers to exploit. Complex code is difficult to understand, maintain, and test [18]. Therefore, complex code would have a higher chance of having faults than simple code. Since attackers exploit the faults in a program, complex code would be more vulnerable than simple code. From this reasoning, we set up the following hypothesis on code complexity:

$H_{IntraComplexity\_D}$: *Vulnerable files have a higher intrafile complexity than neutral files.*

Highly coupled code has a higher chance of having input from external sources that are difficult to trace where the input came from. Moreover, developers can use interfaces to modules implemented by other developers or by a third party without properly understanding the security concerns or assumptions of the modules. Therefore:

$H_{Coupling\_D}$: *Vulnerable files have a higher coupling than neutral files.*

Communication by comments between developers is important, especially in open-source projects for which many developers can contribute on the same code segment without central control. Novice developers who do not understand security concerns and do not follow secure coding practice might comment less often. Furthermore, code developed in a hurry (perhaps directly prior to release) might have fewer comments and be more vulnerable. Therefore:

$H_{Comments\_D}$: *Vulnerable files have a lower comment density than neutral files.*

Table 1 defines the complexity metrics we use in this study. Each individual metric has the same hypotheses defined as the same way as the hypotheses for its group. For example, CountLineCode has the same hypotheses as $H_{IntraComplexity\_D}$.

#### 3.1.2 Code Churn

Code is constantly evolving throughout the development process. Each new change in the system brings a new risk of introducing a vulnerability [15], [19]. Therefore,

$H_{CodeChurn\_D}$: *Vulnerable files have a higher code churn than neutral files.*

Code churn can be counted in terms of the number of check-ins into a version control system and the number of lines that have been added or deleted by code change. Specifically, we break down hypothesis $H_{CodeChurn\_D}$ into the following two hypotheses:

$H_{NumChanges\_D}$: *Vulnerable files have more frequent check-ins than neutral files.*

$H_{ChurnAmount\_D}$: *Vulnerable files have more lines of code that have been changed than neutral files.*

Table 2 defines the code churn metrics that we use in this study.

TABLE 2
Definitions of Code Churn Metrics

| Related Hypothesis | Metric | Definition |
|---|---|---|
| $H_{CodeChurn\_D}$ | NumChanges | The number of check-ins for a file since the creation of a file |
| | LinesChanged | The cumulated number of code lines changed since the creation of a file |
| | LinesNew | The cumulated number of new code lines since the creation of a file |

### 3.1.3 Developer Activity

Software development is performed by development teams working together on a common project. Lack of team cohesion, miscommunications, or misguided effort can all result in security problems [3]. Version control data can be used to construct a developer network and a contribution network based upon "which developer(s) worked on which file," using network analysis as defined in Section 2.

**Developer Network Centrality.** In our developer network, two developers are connected if they have both made a change to at least one file in common during the period of time under study. The result is an undirected, unweighted, and simple graph where each node represents a developer and edges are based on whether or not they have worked on the same file during the same release. Central developers, measured by a high degree, high betweenness, and low closeness, are developers that are well connected to other developers relative to the entire group. Readers may refer to [16] for a more in-depth example of how centrality metrics are derived from developer networks. A central developer would likely have a better understanding of the group's secure coding practices because of his/her connections to the other developers of the team. Therefore:

$H_{DeveloperCentrality\_D}$: *Vulnerable files are more likely to have been worked on by noncentral developers than neutral files.*

The metrics we used for our hypotheses are shown in Table 3. Note that we chose to not study DNMaxDegree, DNMinCloseness, and DNMaxBetweenness because, for example, a high DNMaxDegree means at least *one* central developer worked on the file, which is not as helpful as knowing that high DNMinDegree denotes that *all* developers who worked on a file were central. Also note that a high turnover rate in a project results in many noncentral developers: When new developers are added to the project, they initially lack connections to the other developers.

**Developer Network Cluster.** The metrics of developer centrality give us information about individual developers, but we also consider the relationship between groups of developers. In our developer network, a file that is between two clusters was worked on by two groups of developers, and those two groups did have many other connections in common. Clusters with a common connection but few other connections may not be communicating about improving the security of the code they have in common. Therefore:

$H_{DeveloperCluster\_D}$: *Vulnerable files are more likely to be changed by multiple, separate developer clusters than neutral files.*

Since edges and files have a many-to-many relationship, we use the average and maximum of edge betweenness on the developer network as provided in Table 3.

**Contribution Network.** A *contribution network* [21] is a quantification of the focus made on the relationship between a file and developers instead of relationship between developers as in developer networks. The contribution network uses an undirected, weighted, and bipartite graph with two types of nodes: developers and files. An edge exists where a developer made changes to a file, where the weight is equal to the number of check-ins that developer made to the file. If a file has high centrality, then that file was changed by many developers who made changes to many other files—referred to as an "unfocused contribution" [21]. Files with an unfocused contribution would not get the attention required to prevent the injection of vulnerabilities. Therefore:

$H_{ContributionCentrality\_D}$: *Vulnerable files are more likely to have an unfocused contribution than neutral files.*

Note that developers with high centrality can work with many people, but still work on one small part of the system. Being a central developer means being central in terms of people and not necessarily central in terms of the system. Contribution networks, on the other hand, are structured around the system.

TABLE 3
Meaning of Developer Activity Metrics

| Related Hypothesis | Metric | Problematic When | Meaning |
|---|---|---|---|
| $H_{DeveloperCentrality\_D}$ | DNMinDegree | Low | File was changed by developers who are not central to the network |
| | DNAvgDegree | Low | |
| | DNMaxCloseness | High | |
| | DNAvgCloseness | High | |
| | DNMinBetweenness | Low | |
| | DNAvgBetweenness | Low | |
| $H_{DeveloperCluster\_D}$ | DNMaxEdgeBetweenness | High | File was contributed to by more than one cluster of developers, with few other files being worked on by each cluster |
| | DNAvgEdgeBetweenness | High | |
| $H_{ContributionCentrality\_D}$ | NumDevs | High | File was changed by many developers |
| | CNCloseness | High | File was changed by developers who focused on many other files |

Table 3 provides the meaning of the contribution network metrics.

## 3.2 Evaluation of Hypotheses for Discriminative Power

To evaluate the discriminative power of the metrics, we test the null hypothesis that the means of the metric for neutral files and vulnerable files are equal. Since our data set is skewed and has unequal variance (See Fig. 3 in Section 4), we used the Welch t-test [22]. The Welch t-test is a modified t-test to compare the means of two samples and is known to provide good performance for skewed distributions with unequal variance. To see the association direction (i.e., positively or negatively correlated), we compared the means and medians of the measures of CCD metrics for the vulnerable and neutral files. Our hypotheses for discriminative power are considered to be supported when the results from Welch's tests are statistically significant at the $p < 0.05$ level ($p < 0.0018$ with a Bonferroni correction for 28 hypotheses tests) and when the associations are in the direction prescribed in the hypotheses defined in Section 3.1. In addition to the significance, we present the magnitude of the differences in measures between vulnerable files and neutral files for representative metrics using boxplots.

## 3.3 Evaluation of Hypotheses for Predictability

We hypothesize that a subset of CCD metrics can predict vulnerable files at reasonable prediction performance, specifically 70 percent $PD$ and 25 percent $PF$. Since there are no universally applicable standards for the threshold of high prediction performance, we use these average values found from the fault prediction literature [7], [23], acknowledging that the desirable level of $PD$ and $PF$ depends on varying domains and business goals. We hypothesize that:

$H_{Complexity\_P}$: *A model with a subset of complexity metrics can predict vulnerable files.*

$H_{CodeChurn_P}$: *A model with a subset of code churn metrics can predict vulnerable files.*

$H_{Developer\_P}$: *A model with a subset of developer activity metrics can predict vulnerable files.*

$H_{CCD\_P}$: *A model with a subset of combined CCD metrics can predict vulnerable files.*

Note that we select the subset of the metrics using the information gain ranking method instead of testing every subset of the metrics (see Section 3.5) because prediction performance from a few selected metrics is known to be as good as using a large number of metrics [7].

Additionally, we investigate whether the individual metrics can predict vulnerable files.

$Q_{Individual\_P}$: *Can a model with individual CCD metrics predict vulnerable files?*

## 3.4 Measuring Inspection Reduction

Practitioners must consider how effective the prediction model is in reducing the effort for code inspection even when reasonable $PD$ and $PF$ are achieved. Even though we use $PD$ and $PF$ as the criteria to test our hypotheses on predictability, we further investigate the amount of inspection and reduction in inspection by using the CCD metrics, which is related to cost and effort, to describe the practicality of using CCD metrics. We use the number of files and the lines of code to

inspect as partial and relative estimators of the inspection effort. Such measures have been used as cost estimators in prior studies [24], [25]. In our definition, overall inspection is reduced if the percentage of lines of code or the percentage of files to inspect is smaller than the percentage of faults identified in those files [24], [25]. For example, if we randomly choose files to inspect, we need to inspect 80 percent of the total files to obtain 80 percent $PD$. If a prediction model provides 80 percent $PD$ with less than 80 percent of the total files, the model reduced the cost for inspections compared to a random file selection. The two cost measurements are formally defined below.

The *File Inspection (FI)* ratio is the ratio of files predicted as vulnerable (that is, the number of files to inspect) to the total number of files for the reported $PD$:

$$FI = (TP + FP)/(TP + TN + FP + FN).$$

For example, $PD = 80\%$ and $FI = 20\%$ mean that within the 20 percent of files inspected based on the prediction results, 80 percent of vulnerable files can be found.

The *LOC Inspection (LI)* ratio is the ratio of lines of code to inspect to the total lines of code for the predicted vulnerabilities. First, we define lines of code in the files that were true positives, as $TP_{LOC}$, similarly with $TN_{LOC}$, $FP_{LOC}$, and $FN_{LOC}$. Then, $LI$ is defined below:

$$LI = (TP_{LOC} + FP_{LOC})/(TP_{LOC} + TN_{LOC} + FP_{LOC} + FN_{LOC}).$$

While $FI$ and $LI$ estimate how much effort is involved, we need measures to provide how much effort is reduced. We define two cost-reduction measurements.

The *File Inspection Reduction (FIR)* is the ratio of the reduced number of files to inspect by using the model with CCD metrics compared to a random selection to achieve the same $PD$:

$$FIR = (PD - FI)/PD.$$

The *LOC Inspection Reduction (LIR)* is the ratio of reduced lines of code to inspect by using a prediction model compared to a random selection to achieve the same *Predicted Vulnerability (PV)*:

$$LIR = (PV - LI)/PV.$$

where $PV$ is defined below.

The *Predicted Vulnerability* ratio is the ratio of the number of vulnerabilities in the files predicted as vulnerable to the total number of vulnerabilities. First, we define the number of vulnerabilities in the files that were true positives as $TP_{Vuln}$, similarly with $TN_{Vuln}$, $FP_{Vuln}$, and $FN_{Vuln}$. Then, $PV$ is defined below:

$$PV = TP_{Vuln}/(TP_{Vuln} + FN_{Vuln}).$$

## 3.5 Prediction Models and Experimental Design

We used logistic regression to predict vulnerable files. Logistic regression computes the probability of occurrence of an outcome event from given independent variables by mapping the linear combination of independent variables to the probability of outcome using the log of odds ratio (logit). A file is classified as vulnerable when the outcome

```
next_release_validation (Set R₁, Set R₂, Set R₃, Set R₄) {
        S_train ← R₁ ∪ R₂ ∪ R₃
        S_test ← R4
        performance ← 0
        repeat 10 times {
                S_train_vulnerable ← vulnerable files in S_train
                N ← | S_train_vulnerable |
                S_train_neutral ← N neutral files remaining after random removal from S_train for under-sampling
                S_train ← S_train_vulnerable ∪ S_train_neutral
                V ← select variables using the InfoGain variable selection method from S_train
                Train the model M on S_train and variables V
                performance ← performance + prediction performance of M tested on S_test
        }
        return performance ← performance / 10
}
```

Fig. 1. Pseudocode for next-release validation.

probability is greater than 0.5. We also tried four other classification techniques, including J48 decision tree [26], Random forest [27], Naive Bayes [26], and Bayesian network [26], that have been effective for fault prediction. Among those, J48, Random forest, and Bayesian network provided similar results to logistic regression, while Naive Bayes provided higher $PD$ with higher $FI$ than other techniques. Lessmann et al. [28] also reported that no significant difference in prediction performance was found between the 17 classification techniques they investigated. Therefore, we only present the results from logistic regression in this paper.

To validate a model's predictability, we performed *next-release validation* for Mozilla Firefox, where we had 34 releases and $10 \times 10$ *cross-validation* [26] for the RHEL4 kernel where we only had a single release. For *next-release validation*, data from the most recent three previous releases were used to train against the next release (i.e., train on releases R3 to R1 to test against release R). Using only recent releases was to accommodate for process change, technology change, and developer turnovers. For $10 \times 10$ *cross-validation*, we randomly split the data set into 10 folds and used one fold for testing and the remaining folds for training, rotating each fold as the test fold. Each fold was stratified to properly distribute vulnerable files to both training and the test data set. The entire process was then repeated 10 times to account for possible sampling bias in random splits. Overall, 100 predictions were performed for $10 \times 10$ cross-validation and the results were averaged.

Using many metrics in a model does not always improve the prediction performance since metrics can provide redundant information [7]. We found this issue to be the case in our study. Therefore, we selected only three variables, using two variable selection methods: the information gain ranking method and the correlation-based greedy feature selection method [26]. Both methods provided similar prediction performance. However, the set of chosen variables by the two selection methods were different. We present the results using the information gain ranking method in this study.

We performed the prediction on both the raw and log transformed data. For logistic regression, $PD$ was improved at the sacrifice of precision after log transformation. We present results using log transformation in this paper.

Our data are heavily unbalanced between majority class (neutral files) and minority class (vulnerable files). Prior studies have shown that the performance is improved (or at least not degraded) by "balancing" the data [23], [29]. Balancing the data can be achieved by duplicating the minority class data (oversampling) or removing randomly chosen majority class data (undersampling) until the numbers of data instances in the majority class and the minority class become equal [23]. We used undersampling in this study since undersampling provided better results than using the unbalanced data and reduced the time for evaluation.

Figs. 1 and 2 provide the pseudocode of our experimental design explained above for next-release validation and cross-validation. The whole process of validation is repeated 10 times to account for possible sampling bias due to random removal of data instances. The 10 times of repetition also account for the bias due to random splits in cross-validation. In Figs. 1 and 2, *performance* represents the set of evaluation criteria, cost-reduction measurements, and their relevant measurements defined in Sections 2.3 and 3.4. For next-release validation, the input is three prior releases for training a model and one release for testing. For cross-validation, the input is the whole data set.

We used the Weka 3.7[4] with default options for prediction models and variable selection except for limiting the number of variables to be selected to three for multivariate predictions.

## 4 CASE STUDY 1: MOZILLA FIREFOX

Our first case study is Mozilla Firefox, a widely used open-source web browser. Mozilla Firefox had 34 releases at the time of data collection developed over four years. Each release consists of over 10,000 files and over two million lines of source code.

### 4.1 Data Collection

To measure the number vulnerabilities fixed in a file, we counted the number of bug reports that include the details on vulnerabilities and on how the vulnerabilities have been

4. http://www.cs.waikato.ac.nz/ml/weka.

```
cross_validation (Set R) {
        performance ← 0
        repeat 10 times {
                Randomly split R into 10 stratified bins, B = {b₁, b₂, …, b₁₀}
                for each bin bᵢ {
                        S_train ← B − { bᵢ }
                        S_test ← bᵢ
                        S_train_vulnerable ← vulnerable files in S_train
                        N ← | S_train_vulnerable |
                        S_train_neutral ← N neutral files remaining after random removal from S_train for under-sampling
                        S_train ← S_train_vulnerable ∪ S_train_neutral
                        V ← select variables using the InfoGain selection method from S_train
                        Train the model M on S_train and variables V
                        performance ← performance + prediction performance of M tested on S_test
                }
        }
        return performance ← performance / 100
}
```

Fig. 2. Pseudocode for cross-validation.

### TABLE 4
### Project Statistics for Mozilla Firefox

| No. | Firefox Release | # of Files | LOC | Mean LOC | Files with Vulns. | Total Vulns. | Vulns. per File | % of Files with Vulns. |
|-----|-----------------|-----------|-----|----------|-------------------|--------------|-----------------|------------------------|
| R1 | 1.0 / 1.0.1 / 1.0.2 | 10,320 | 2,060,908 | 200 | 70 | 84 | 0.008 | 0.678 |
| R2 | 1.0.3 / 1.0.4 / 1.0.5 / 1.0.6[*] | 10,321 | 2,063,960 | 200 | 123 | 134 | 0.013 | 1.192 |
| R3 | 1.0.7 / 1.5 / 1.5.0.1 | 10,321 | 2,064,747 | 200 | 93 | 159 | 0.015 | 0.901 |
| R4 | 1.5.0.2 / 1.5.0.3 / 1.5.0.4 | 10,956 | 2,226,540 | 203 | 100 | 138 | 0.013 | 0.913 |
| R5 | 1.5.0.5 / 1.5.0.6 / 1.5.0.7 | 10,961 | 2,230,313 | 204 | 109 | 153 | 0.014 | 0.994 |
| R6 | 1.5.0.8 / 2.0 / 2.0.0.1 | 10,961 | 2,232,890 | 204 | 87 | 124 | 0.011 | 0.794 |
| R7 | 2.0.0.2 / 2.0.0.3 / 2.0.0.4 | 11,060 | 2,294,287 | 207 | 114 | 162 | 0.015 | 1.031 |
| R8 | 2.0.0.5 / 2.0.0.6 / 2.0.0.7 | 11,060 | 2,299,054 | 208 | 55 | 72 | 0.007 | 0.497 |
| R9 | 2.0.0.8 / 2.0.0.9 / 2.0.0.10 | 11,076 | 2,301,398 | 208 | 14 | 15 | 0.001 | 0.126 |
| R10 | 2.0.0.11 / 2.0.0.12 / 2.0.0.13 | 11,077 | 2,301,832 | 208 | 84 | 110 | 0.010 | 0.758 |
| R11 | 2.0.0.14 / 2.0.0.15 / 2.0.0.16 | 11,080 | 2,304,048 | 208 | 27 | 46 | 0.004 | 0.244 |

*. *The vulnerabilities for Firefox 1.0.5 and Firefox 1.0.6 were reported together. Therefore, we considered those two versions as one version.*

fixed for the file. We collected vulnerability information from Mozilla Foundation Security Advisories (MFSAs).[5] Each MFSA includes bug IDs that are linked to the Bugzilla[6] bug tracking system. Mozilla developers also add bug IDs to the log of the CVS version control system[7] when they check in files to the CVS after the vulnerabilities have been fixed. We searched the bug IDs from the CVS log to find the files that have been changed to fix vulnerabilities, similar to the approach found in other studies [30]. The number of MFSAs for Firefox was 197 as of 2 August 2008. The vulnerability fixes for the MFSAs were reported in 560 bug reports. Among them, 468 bug IDs were identified from the CVS log. Although some of the files were fixed for regression, all those files were also fixed for vulnerabilities. Therefore, we counted those files as vulnerable.

To collect complexity metrics, we used Understand C++,[8] a commercial metrics collection tool. We limited our analysis to C/C++ and their header files to obtain complexity metrics, excluding other file types such as script files and make files. We obtained code churn and developer activity metrics from the CVS version control system.

At the time of data collection, Firefox 1.0 and Firefox 2.0.0.16 were the first and the last releases that had vulnerability reports. The gap between Firefox releases ranged from one to two months. Since each release had only a few vulnerabilities (not enough to perform analysis at each Firefox release), we combined the number of vulnerabilities for three consecutive releases, and will refer to those releases as a combined release, denoted by R1 through R11 in this paper. For each release, we used the most recent three *combined releases* (RN3 to RN1) to predict vulnerable files in RN. We collected metrics for 11 combined releases. Table 4 provides the project statistics for the 11 combined releases.

## 4.2 Discriminative Power Test and Univariate Prediction Results

Table 5 shows the results of hypotheses testing for discriminative power and the univariate prediction using logistic regression for individual metrics from release R4,

5. http://www.mozilla.org/security/announce/.
6. http://www.bugzilla.org/.
7. https://developer.mozilla.org/en/Mozilla_Source_Code_Via_CVS.
8. http://www.scitools.com.

TABLE 5
Results of Discriminative Power Test and Univariate Prediction Mozilla Firefox

| Related Hypothesis | Metric | Discriminative Power | | | Predictability | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | Welch-t | Associ-ation | $H_D$* | Mean for R4 | | | Std. for R4 | | |
| | | | | | PD* | PF* | $Q_P$* | PD | PF | $N(Q_P)$* |
| | **Complexity** | | | | | | | | | |
| $H_{IntraComplexity\_D}$ | CountLineCode | √ | + | √ | 82 | 28 | X | 0.4 | 0.9 | 1 |
| | CountDeclFunction | √ | + | √ | 87 | 68 | X | 0.0 | 0.0 | 0 |
| | CountLineCodeDecl | √ | + | √ | 79 | 25 | √ | 0.0 | 0.8 | 47 |
| | CountLinePreprocessor | √ | + | √ | 76 | 26 | X | 0.9 | 0.9 | 0 |
| | SumEssential | √ | + | √ | 86 | 60 | X | 0.3 | 3.0 | 0 |
| | SumCyclomaticStrict | √ | + | √ | 86 | 60 | X | 0.0 | 1.2 | 0 |
| | MaxCyclomaticStrict | √ | + | √ | 87 | 68 | X | 0.0 | 0.0 | 0 |
| | SumMaxNesting | √ | + | √ | 82 | 53 | X | 0.0 | 0.0 | 0 |
| | MaxMaxNesting | √ | + | √ | 82 | 53 | X | 0.0 | 0.0 | 0 |
| $H_{Coupling\_D}$ | SumFanIn | √ | + | √ | 87 | 57 | X | 0.0 | 0.0 | 0 |
| | SumFanOut | √ | + | √ | 86 | 62 | X | 0.0 | 0.0 | 0 |
| | MaxFanIn | √ | + | √ | 87 | 57 | X | 0.0 | 0.0 | 0 |
| | MaxFanOut | √ | + | √ | 86 | 62 | X | 0.0 | 0.0 | 0 |
| $H_{Comments\_D}$ | CommentDensity | X | − | X | 59 | 45 | X | 15.4 | 21.8 | 0 |
| | **Code Churn** | | | | | | | | | |
| $H_{NumChanges\_D}$ | NumChanges | √ | + | √ | 86 | 23 | √ | 0.8 | 1.2 | 80 |
| $H_{ChurnAmount\_D}$ | LinesChanged | √ | + | √ | 85 | 25 | √ | 0.3 | 1.4 | 76 |
| | LinesNew | √ | + | √ | 88 | 58 | X | 0.9 | 9.1 | 0 |
| | **Developer Activity** | | | | | | | | | |
| $H_{DeveloperCentrality\_D}$ | DNMinDegree | √ | − | √ | 74 | 84 | X | 4.2 | 1.0 | 0 |
| | DNAvgDegree | √ | + | X | 98 | 77 | X | 0.0 | 0.9 | 0 |
| | DNMaxCloseness | √ | + | √ | 100 | 95 | X | 0.0 | 0.1 | 0 |
| | DNAvgCloseness | √ | + | √ | 100 | 95 | X | 0.0 | 0.0 | 0 |
| | DNMinBetweenness | √ | − | √ | 53 | 45 | X | 0.5 | 30.7 | 0 |
| | DNAvgBetweenness | √ | − | √ | 99 | 87 | X | 0.3 | 0.8 | 0 |
| $H_{DeveloperCluster\_D}$ | DNMaxEdgeBetweenness | √ | + | √ | 88 | 30 | X | 0.0 | 0.0 | 0 |
| | DNAvgEdgeBetweenness | √ | + | √ | 88 | 30 | X | 0.0 | 0.0 | 0 |
| $H_{ContributionCentrality\_D}$ | NumDevs | √ | + | √ | 84 | 24 | √ | 1.5 | 1.6 | 58 |
| | CNCloseness | √ | + | √ | 82 | 22 | √ | 0.0 | 0.0 | 80 |
| | CNBetweenness | √ | + | √ | 64 | 9 | X | 0.0 | 0.0 | 20 |

*. $H_D$: Hypothesis test for discriminative power, PD: Probability of detection, PF: Probability of false alarm, $Q_P$: Question for predictability, $N(Q_P)$: Number of predictions that provided over 70% PD with less than 25% PF from 80 predictions

the first test data set that uses the most three recent releases to train a prediction model. In Table 5, the plus sign in the Association column indicates that the vulnerable files had a higher measure than neutral files. DNAvgDegree was the only sign that did not completely agree with its hypothesis. We discuss the reason in Section 6. Twenty-seven of the 28 metrics showed a statistically significant difference between vulnerable and neutral files using the Welch t-test after a Bonferroni correction as shown in Table 5. We also compared the mean and median of each measurement for vulnerable and neutral files to find the association direction (positive or negative). Fig. 3 shows the boxplots of comparisons of log scaled measures between vulnerable and neutral files for the three representative metrics from each type of CCD metrics: CountLineCode, NumChanges, and Num-Devs. The medians of the three metrics for the vulnerable files were higher than the ones for the neutral files, as we hypothesized.

In Table 5, the means of PD and PF are averaged over 10 repetitions of predictions for R4 according to the algorithm provided in Fig. 1 and presented with standard deviations. We test whether individual metrics can predict

vulnerable files at over 70 percent PD with less than 25 percent PF as described in Section 3.3. $N(Q_P)$ represents the number of predictions that satisfy our criteria in the 80 predictions (10 repetitions for each of the eight combined releases). NumChanges and CNCloseness satisfied our criteria in all 80 predictions. CountLineCodeDecl, Lines-Changed, and NumDevs satisfied our criteria in over half of the predictions. Most metrics provided very small variations between predictions except for CommentDensity, Linesnew, and DNMinBetweenness. We provide $N(Q_P)$ instead of averaging the prediction results across all releases because averaging the results for different data sets can mislead the interpretation of the results from each data set [31].

Note that a single metric can provide such a high PD and a low PF. We further investigate the predictability of CCD metrics when they are used together in a model in the following section.

## 4.3 Multivariate Prediction Results

Although metrics can have low predictability individually, combining metrics into a model can result in better
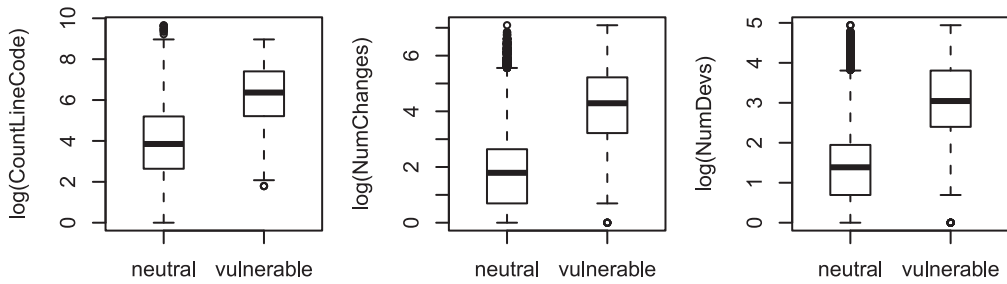
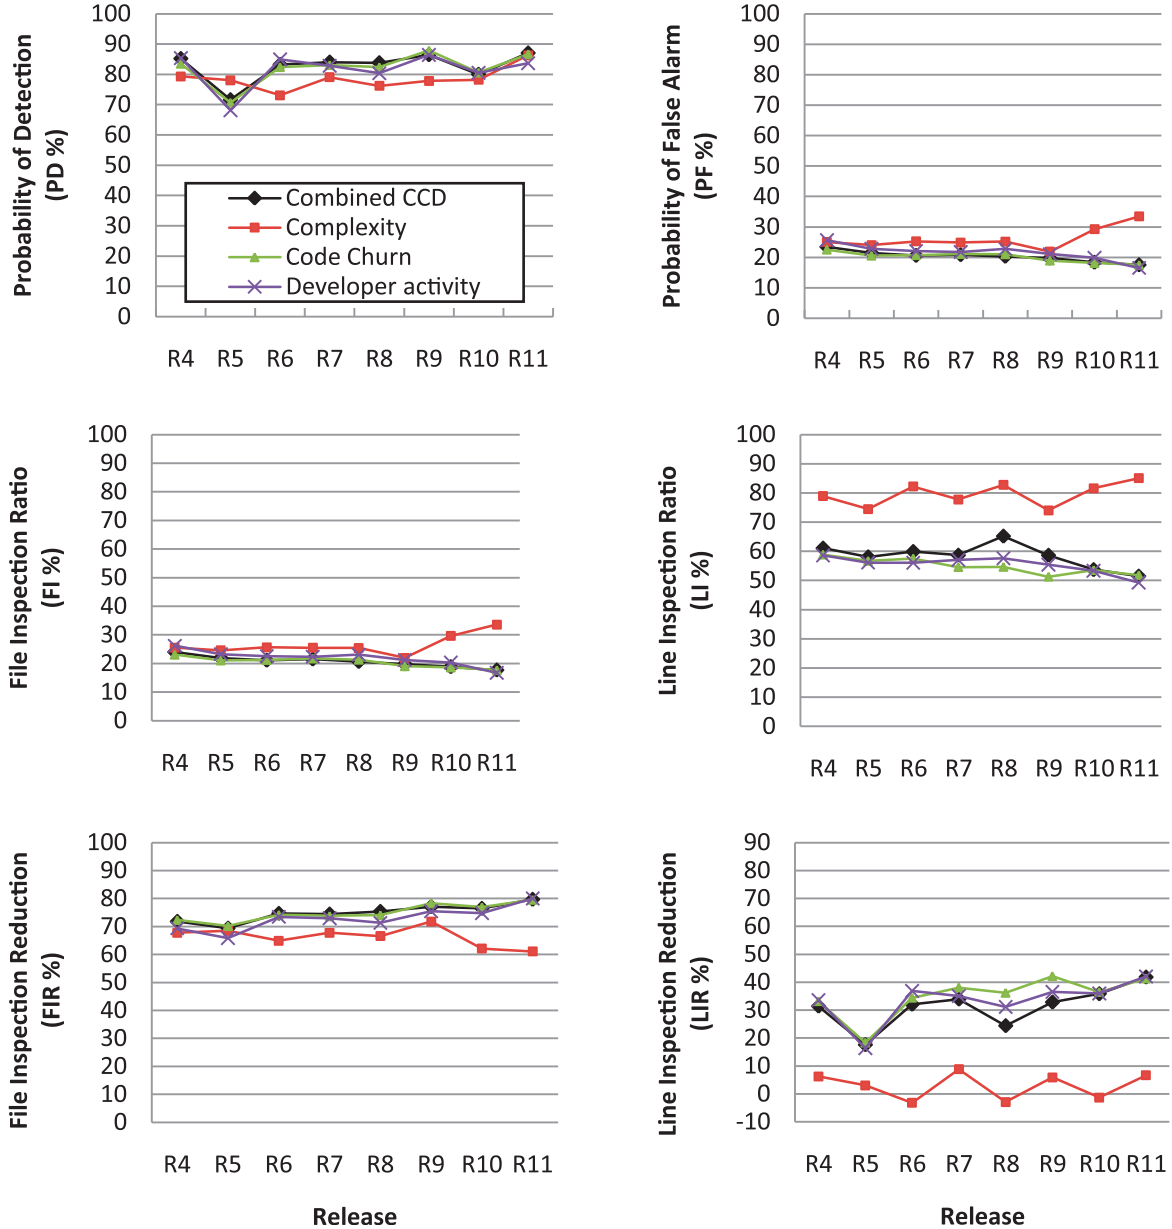Fig. 3. Comparison of measures for vulnerable and neutral files.



Fig. 4. Prediction results for Mozilla Firefox across releases.

predictability [32]. Therefore, we created four types of models using complexity, code churn, developer activity, and combination of the CCD metrics. From each set of the metrics, three variables were selected as we explained in Section 3.5.

Fig. 4 shows *PD, PF, FI, LI, FIR*, and *LIR* across releases where the results from each release were averaged over the results from the 10 repetitions of next-release validation using logistic regression. All the models with code churn, developer activity, and combined CCD metrics provided similar results across releases; *PD* was between 68 and 88 percent and *PF* was between 17 and 26 percent; *FI ratio* was between 17 and 26 percent resulting in 66 to 80 percent reduction in the files to inspect; *Line Inspection (LI) Ratio* was

TABLE 6
Results of Multivariate Prediction for Mozilla Firefox

| | Predictability for R4 | | | | | | Inspection Costs for R4 | | | | |
| | Mean | | | Standard Dev. | | | Costs | | | Cost Reduction | |
| | PD | PF | P | PD | PF | N(H_P) | PV | FI | LI | FIR | LIR |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Complexity | **79** | **25** | 3 | 0.7 | 0.8 | **40** | 84 | 26 | 79 | 68 | 6 |
| Code churn | **84** | **23** | 3 | 1.2 | 1.3 | **76** | 88 | 23 | 59 | 72 | 33 |
| Developer activity | **85** | 26 | 3 | 1.7 | 2.8 | **62** | 88 | 26 | 59 | 69 | 34 |
| Combined CCD | **85** | **24** | 3 | 1.8 | 1.7 | **74** | 89 | 24 | 61 | 72 | 31 |

TABLE 7
Project Statistics for the RHEL4 Kernel

| # of Files | LOC | Mean LOC | Files with Vulns. | Total Vulns. | Vulns. per File | % of Files with Vulns. |
|---|---|---|---|---|---|---|
| 13,568 | 3,068,453 | 226 | 194 | 258 | 0.019 | 1.4% |

between 49 and 65 percent resulting in 16 to 42 percent of reduction in lines of code to inspect. The models using complexity metrics provided 76 to 86 percent *PD* and 22 to 34 percent *PF*. *FI* for the models using complexity metrics was between 22 and 34 percent resulting in 61 to 72 percent reduction in file inspection. *LI* for the models using complexity metrics was much higher than *LI* for the other models (74 to 85 percent), resulting in only 7 percent reduction in code inspection at best. In the worst case, 3 percent more code required to inspect than the percentage of lines of code chosen by random selection. Note that *PF* and *FI* showed almost identical performance for all models. We discuss the similarity between *PF* and *FI* in Section 6.

The three types of models generally provided consistent *PD* except for R5 and R10. The models using complexity metrics provided lower *PD* than other models, but provided consistent results even in R5 and R10. Instead, *PF* from the models using complexity metrics suddenly increased in R10. We conjecture that the sudden decrease in performance for R10 is because the number of vulnerable files in R10 suddenly increases when the model was trained with recent two releases (R8 and R9) with relatively few vulnerabilities, and the model trained from the past trend did not work properly. However, we were not able to find out any particular reason for the sudden decreases in performance for R5. All the prediction results from the three types of models gradually improved as the system matured, but the results from models using complexity metrics stayed constant or degraded.

Table 6 provides detailed results for release R4, including standard deviations. The models with code churn metrics and combined CCD metrics satisfied our prediction criteria in over 90 percent of predictions out of 80 predictions (10 repetitions for each of the eight combined releases). All of the models satisfied our criteria in over 50 percent of the 80 predictions. Standard deviations in *PD* and *PF* were less than 3 in all models for release R4. Interestingly, no multivariate models were noticeably better than some of the best univariate models.

The amount of files to inspect reduced by 68 to 72 percent depending on the models for R4. The amount of lines of code to inspect reduced by 31 to 34 percent with code churn, developer activity, and combined CCD metrics, and only 6 percent with complexity metrics for release R4.

## 5 CASE STUDY 2: RED HAT ENTERPRISE LINUX KERNEL

Our second case study was performed on the Linux kernel as it was distributed in the Red Hat Enterprise Linux 4 (RHEL4) operating system. The RHEL4 kernel consists of 13,568 C files with over three million lines of code. The details of the project, data collection, and prediction results are described in this section.

### 5.1 Data Collection

Gathering security data involved tracing through the development artifacts related to each vulnerability reported in the Linux kernel. We collected our vulnerability data from the Red Hat Bugzilla database and the Red Hat package management system (RPM). Since some vulnerability patches affect only certain releases, we examined each defect report manually to ensure that developers had decided that patch was, in fact, required. Instead of scanning developer commit logs for defect IDs, we used the RPM system to determine the exact patch that was issued to fix each of the 258 known vulnerabilities. Since we are only interested in vulnerabilities that existed at the time of release, we did not include vulnerabilities introduced by postrelease patches (a.k.a. "regressions") in our data set. For the few vulnerabilities that did not have all of the relevant artifacts (e.g., defect reports, patches), we consulted the director of the RHSR team to correct the data and the artifacts. We collected vulnerabilities reported from February 2005 through July 2008.

To obtain code churn and developer activity data, we used the Linux kernel source repository.[9] The RHEL4 operating system is based on kernel version 2.6.9, so we used all of the version control data from kernel version 2.6.0 to 2.6.9, which were approximately 15 months of development and maintenance. Table 7 provides the project statistics.

### 5.2 Discriminative Power Test and Univariate Prediction Results

Table 8 provides the results of hypotheses tests for discriminate power and the results of the univariate prediction for individual metrics. Twenty-seven of the 28 metrics showed a statistically significant difference between vulnerable and neutral files using the Welch t-test after Bonferroni correction

9. http://git.kernel.org/.

TABLE 8
Results of Discriminative Power Test and Univariate Prediction for the RHEL4 Kernel

| Related Hypothesis | Metric | Discriminative Power | | | Predictability | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | Mean | | | Std. | | |
| | | Welch-t | Association | H_D* | PD* | PF* | Q_P* | PD | PF | N(Q_P)* |
| | **Complexity** | | | | | | | | | |
| H$_{IntraComplexity\_D}$ | CountLineCode | √ | + | √ | 87 | 38 | X | 8.7 | 2.8 | 0 |
| | CountDeclFunction | √ | + | √ | 94 | 57 | X | 4.5 | 1.5 | 0 |
| | CountLineCodeDecl | √ | + | √ | 87 | 44 | X | 9.5 | 3.8 | 0 |
| | CountLinePreprocessor | X | + | X | 67 | 41 | X | 9.5 | 1.7 | 0 |
| | SumEssential | √ | + | √ | 93 | 53 | X | 5.6 | 2.9 | 0 |
| | SumCyclomaticStrict | √ | + | √ | 93 | 52 | X | 5.8 | 1.9 | 0 |
| | MaxCyclomaticStrict | √ | + | √ | 94 | 56 | X | 5.5 | 2.1 | 0 |
| | SumMaxNesting | √ | + | √ | 90 | 49 | X | 6.6 | 1.2 | 0 |
| | MaxMaxNesting | √ | + | √ | 90 | 49 | X | 6.6 | 1.2 | 0 |
| H$_{Coupling\_D}$ | SumFanIn | √ | + | √ | 93 | 52 | X | 5.0 | 1.6 | 0 |
| | SumFanOut | √ | + | √ | 92 | 52 | X | 6.4 | 1.7 | 0 |
| | MaxFanIn | √ | + | √ | 93 | 54 | X | 5.0 | 1.3 | 0 |
| | MaxFanOut | √ | + | √ | 93 | 53 | X | 5.2 | 1.2 | 0 |
| H$_{Comments\_D}$ | CommentDensity | √ | − | √ | 88 | 75 | X | 7.9 | 1.9 | 0 |
| | **Code Churn** | | | | | | | | | |
| H$_{NumChanges\_D}$ | NumChanges | √ | + | √ | 83 | 25 | √ | 8.8 | 2.9 | 27 |
| H$_{ChurnAmount\_D}$ | LinesChanged | √ | + | √ | 83 | 39 | X | 9.1 | 2.0 | 0 |
| | LinesNew | √ | + | √ | 90 | 52 | X | 6.0 | 1.3 | 0 |
| | **Developer Activity** | | | | | | | | | |
| H$_{DeveloperCentrality\_D}$ | DNMinDegree | √ | − | √ | 86 | 67 | X | 8.5 | 2.2 | 0 |
| | DNAvgDegree | √ | + | X | 98 | 59 | X | 3.7 | 2.7 | 0 |
| | DNMaxCloseness | √ | + | √ | 98 | 71 | X | 7.3 | 6.9 | 0 |
| | DNAvgCloseness | √ | + | √ | 99 | 74 | X | 1.6 | 1.1 | 0 |
| | DNMinBetweenness | √ | − | √ | 69 | 63 | X | 11.8 | 3.0 | 0 |
| | DNAvgBetweenness | √ | + | X | 97 | 57 | X | 3.7 | 1.5 | 0 |
| H$_{DeveloperCluster\_D}$ | DNMaxEdgeBetweenness | √ | + | √ | 49 | 11 | X | 11.2 | 0.8 | 4 |
| | DNAvgEdgeBetweenness | √ | + | √ | 49 | 11 | X | 11.2 | 0.8 | 4 |
| H$_{ContributionCentrality\_D}$ | NumDevs | √ | + | √ | 80 | 22 | √ | 9.1 | 1.1 | 84 |
| | CNCloseness | √ | + | √ | 99 | 74 | X | 1.6 | 1.1 | 0 |
| | CNBetweenness | √ | + | √ | 92 | 41 | X | 6.2 | 1.3 | 0 |

*. $H_D$: Hypothesis test for discriminative power, PD: Probability of detection, PF: Probability of false alarm, $Q_P$: Question for predictability, $N(Q_P)$: Number of runs in cross-validations that provided over 70% PD with less than 25% PF

except for CommentDensity. The association directions for DNAvgDegree and DNAvgBetweenness did not agree with their hypotheses. We discuss the reason in Section 6.

In vulnerability prediction, only NumDevs provided over 70 percent *PD* with less than 25 percent *PF* in over half of the 100 predictions for $10 \times 10$ cross-validation. RHEL4 had greater standard deviations in *PD* and *PF* than Mozilla Firefox.

## 5.3 Multivariate Prediction Results

Table 9 provides the results for the multivariate predictions by $10 \times 10$ cross-validation using logistic regression. The predictions were performed using the three variables selected by the information gain ranking method. The models using code churn, developer activity, and combined CCD metrics satisfied our prediction criteria in over 50 percent of the 100 runs of cross-validation. However, none of the predictions using complexity metrics satisfied our prediction criteria. When we changed the threshold to 0.4 for binary classification, the models using complexity metrics provided 77 percent *PD* and 29 percent *PF* and 15 of the 100 predictions satisfied our prediction criteria.

The reduction in file inspection compared to a random file selection was between 51 and 71 percent. The reduction

TABLE 9
Results of Multivariate Prediction for the RHEL4 Kernel

| | Predictability | | | | | | Inspection Costs | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | Mean | | | Std. | | | Costs | | | Cost Reduction | |
| | PD | PF | P | PD | PF | N(H_P) | PV | FI | LI | FIR | LIR |
| Complexity | 90 | 43 | 3 | 7.2 | 6.8 | 0 | 92 | 44 | 85 | 51 | 7 |
| Code churn | 82 | 24 | 5 | 9.2 | 1.9 | 59 | 85 | 25 | 58 | 70 | 32 |
| Developer activity | 81 | 23 | 5 | 9.2 | 2.9 | 76 | 84 | 24 | 58 | 70 | 31 |
| Combined CCD | 84 | 24 | 5 | 8.5 | 2.3 | 71 | 87 | 25 | 62 | 71 | 28 |

TABLE 10
Summary of Hypotheses Testing

| | Hypotheses | Firefox | RHEL 4 |
|---|---|---|---|
| $H_{Complexity\_D}$ | Vulnerable files are more complex than neutral files. | Yes for 13 of 14 metrics. | Yes for 13 of 14 metrics. |
| $H_{CodeChurn\_D}$ | Vulnerable files have a higher code churn than neutral files. | Yes for all 3 metrics. | Yes for all 3 metrics. |
| $H_{Developer\_D}$ | Vulnerable files are more likely to have been changed by poor developer activity than neutral files. | Yes for 10 of 11 metrics. | Yes for 9 of 11 metrics. |
| $Q_{Individual\_P}$ | Can a model with individual CCD metric predict vulnerable files?* | 5 of 28 metrics satisfied the prediction criteria in over half of the 80 predictions. | 1 of 28 metrics satisfied the prediction criteria in over half of the 100 predictions. |
| $H_{Complexity\_P}$ | A model with a subset of complexity metrics can predict vulnerable files. * | Supported by 40 of 80 predictions. | Supported by 0 of 100 cross-validations. |
| $H_{CodeChurn\_P}$ | A model with a subset of code churn metrics can predict vulnerable files. * | Supported by 76 of 80 predictions. | Supported by 59 of 100 cross-validations. |
| $H_{Developer\_P}$ | A model with a subset of developer metrics can predict vulnerable files. * | Supported by 62 of 80 predictions. | Supported by 76 of 100 cross-validations. |
| $H_{CCD\_P}$ | A model with a subset of combined CCD metrics can predict vulnerable files. * | Supported by 74 of 80 predictions. | Supported by 71 of 100 cross-validations. |

*. The criteria for the hypotheses tests for vulnerability prediction are 70 % $PD$ and 25% $PF$.

in lines of code inspection was over 28 percent for code churn, developer activity, and combined CCD metrics and only 7 percent for complexity metrics.

## 6 SUMMARY OF TWO CASE STUDIES AND DISCUSSION

Table 10 provides the summary of our hypotheses testing. The hypotheses for discriminative power were supported by at least 24 of the 28 metrics for both projects, except for CommentDensity and DNAvgDegree for Firefox and CountLinePreprocessor, DNAvgDegree, and DNAvgBetweenness for RHEL4. Among these, CountLinePreprocessor and CommentDensity were not discriminative of neutral and vulnerable files. DNAvgDegree and DNAvgBetweenness disagreed with our hypotheses in the direction of association. While DNMinDegree was negatively correlated with vulnerabilities and supported our hypothesis, DNAvgDegree was positively correlated in both projects. This means that files are more likely to be vulnerable if they are changed by developers who work on many other files with other developers on *average*, but when *all* of the developers are central, the file is less likely to be vulnerable. For DNAvgBetweeenness, we were not able to find any clear reason for the hypothesis to not be supported.

Overall, 80 predictions were performed for the eight releases (R4-R11) of Firefox with 10 repetitions to account for sampling bias and 100 predictions for RHEL4 for $10 \times 10$ cross-validation. In the univariate predictions, five of the 28 metrics supported the hypotheses in 50 percent of the total predictions for Firefox and one of the 28 metrics for RHEL4. Considering only the small number of metrics satisfied by the prediction criteria in univariate prediction, relying on a single metric is a dangerous practice. In multivariate predictions, the models using code churn, developer activity, and combined CCD metrics supported the hypotheses in over 50 percent of the total predictions for both projects. Even though the models using complexity metrics supported the hypotheses in over 50 percent of the total predictions for Firefox, none of the predictions

were successful for RHEL4. Considering this result together with the surprisingly low (even negative) inspection cost reduction seen in Fig. 4 of Section 4, although lines of code are an effective cost measurement that depends on situation [25], the effectiveness of complexity metrics as indicators of vulnerabilities is weak for the complexity metrics we collected.

*Precision* from all the models for both projects was strikingly low, with less than five as a result of large numbers of false positives. This result is especially interesting because almost all of the individual metrics had strong discriminative power according to the Welch t-test. This discrepancy can be explained from the boxplots in Fig. 3 of Section 4, where the mean values in the individual metrics show clear difference, but the considerable numbers of neutral files are still in the expected range of vulnerable files, leading to a large number of false positives. Organizations can improve $P$ by raising the threshold for binary classification to reduce false positives. However, $PD$ can become lower in that case as $PD$ and $P$ tend to trade off each other. Whether a model that provides high $PD$ and low $P$ is better than a model that provides high $P$ and low $PD$ is arguable. An organization may prefer to detect many vulnerabilities because it has a large amount of security resources and security expertise. On the other hand, other organizations may prefer a model that provides high $P$ that reduces the waste of effort even with low $PD$.

Preference for high $PD$ and low $P$ requires some caution in terms of cost effectiveness. From the results in Tables 6 and 9, organizations are guided to inspect only less than 26 percent of files to find over around 80 percent of vulnerable files in most models. However, since the projects have a large amount of files, the number of files to inspect is still large. For example, Firefox release R4 has around 11,000 files and the model with combined CCD metrics provided 24 percent $FI$ identifying 2,640 files to inspect. If security inspection requires one person-day per file, over seven full years would be spent for one security engineer to inspect the 2,640 files. If the files to inspect are reduced to 10 percent of the predicted files (264 files) by further manual prioritization, the overall inspection would take essentially a year for one security

engineer only to find a further reduced set of vulnerabilities. In that case, pursuing high $P$ and low $PD$ might be a more cost-effective approach than pursing high $PD$ and low $P$. However, because the inspection time can greatly vary depending on the ability and the number of security engineers involved, organizations should use this illustration and our prediction results only to make an informed decision.

Among the metrics we investigated, history metrics such as code churn and developer activity metrics provided higher prediction performance than the complexity. Therefore, historical development information is a favorable source for metrics and using historical information is recommended whenever possible. However, our result is limited to the metrics that we collected. Other complexity metrics may provide better results.

We also observed sudden changes in prediction performance in a few releases of Firefox where we used next-release validation. This result cannot be observed with cross-validation. Therefore, our study reveals the importance of next-release validation to validate metrics for vulnerability prediction whenever possible.

For both projects, $PF$ and $FI$ are almost equal. We believe the reason for this was that the percentage of files with vulnerabilities is very low (< 1.4 percent) and $P$ is also very low. The low percentage of vulnerable files means that the total number of files ($TP + TN + FP + FN$) is almost the same as the number of nonvulnerable files ($FP + TN$). The low $P$ means that the number of positive predictions ($TP + FP$) is very close to $FP$. Therefore, $FI$ computed by $(TP + FP)/(TP + TN + FP + FN)$ and $PF$ computed by $FP/(FP + TN)$ are almost equal. Knowing this fact provides us a useful hint to guess the number files to inspect when both the percentage of vulnerable files and $P$ are very low.

Interestingly, NumDevs was effective in the vulnerability prediction in our study while another study [33] observed that NumDevs did not improve the prediction performance significantly. The two major differences between the studies are 1) their study was closed source and ours was open source and 2) they were predicting faults and not vulnerabilities.

Further study on the difference in open and closed-source projects and on the difference between fault and vulnerability prediction may further improve our understanding on faults and vulnerabilities on various types of projects and better guide code inspection and testing efforts. In a preliminary study comparing fault prediction and vulnerability prediction for Firefox 2.0, we trained a model for predicting vulnerabilities and a model for predicting faults, both using complexity, code churn, and fault history metrics [34]. The fault prediction provided 18 percent lower $PD$ and 38 percent higher $P$ than vulnerability prediction. The study also showed that the prediction performance was largely affected by the number of reported faulty or vulnerable files. Since only 13 percent of faulty files were reported as vulnerable, further effort is required to characterize the difference between faults and vulnerabilities and to find better way to predict vulnerable code locations.

## 7 THREATS TO VALIDITY

Since our data are based on known vulnerabilities, our analysis does not account for latent (undiscovered) vulnerabilities. Additionally, only fixed vulnerabilities are publicly reported in detail by organizations to avoid the possible attacks from malicious users; unfixed vulnerabilities are usually not publicly available. However, considering the wide use of both projects, we believe that the currently reported vulnerabilities are not too limited to jeopardize our results.

We combined every three releases and predicted vulnerabilities for the next three releases for Mozilla Firefox. Using this study design, the predictions will be performed on every third release. However, considering the short time periods between releases (one or two months), we consider that the code and process history information between the three releases within a combined release is relatively similar and those releases share many similar vulnerabilities. This decision was made to increase the percentage of vulnerabilities in each release because the percentage of vulnerabilities for the subject projects was too low to train the prediction models. In fact, once enough training data are cumulated during a few initial releases, one could predict vulnerabilities in actual releases rather than in combined releases. Alternatively, we could use vulnerability history as a part of metrics together with CCD metrics, as was used in [33], instead of combining releases. We plan to extend our study to accommodate both of the approaches. For Mozilla Firefox, not all of the bug IDs for vulnerability fixes were identified from the CVS log, which could lead to the lower prediction performance.

Actual security inspections and testing are not perfect, so our results are optimistic in predicting exactly how many vulnerable files will be found by security inspection and testing.

As with all empirical studies, our results are limited to the two projects we studied. To generalize our observations from this study to other projects in various languages, sizes, domains, and development processes, further studies should be performed.

## 8 RELATED WORK

This section introduces prior studies on the software vulnerability prediction and usages of CCD metrics in fault prediction.

### 8.1 Vulnerability Prediction

Neuhaus et al. [30] predicted vulnerabilities on the entire Mozilla open-source project (not specific to Firefox) by analyzing the import (header file inclusion) and function call relationship between components. In this study, a component is defined as a C/C++ file and its header file of the same name. They analyzed the pattern of frequently used header files and function calls in vulnerable components and used the occurrence of the patterns as predictors of vulnerabilities. Their model using import and function call metrics provided 45 percent $PD$ and 70 percent precision, and estimated 82 percent of the known vulnerabilities in the top 30 percent components predicted as vulnerable. Our model with CCD metrics provided higher $PD$ (over 85 percent) but lower precision (less than 3 percent) than their work and detected 89 percent of vulnerabilities ($PV$) in 24 percent of files ($FI$) for Mozilla Firefox. We also validated the models

across releases to simulate actual use of a vulnerability prediction in organizations, while their study performed cross-validation.

Gegick et al. [35] modeled vulnerabilities using the regression tree model technique with source lines of code, alert density from a statistic analysis tool, and code churn information. They performed a case study on 25 components in a commercial telecommunications software system with 1.2 million lines of code. Their model identified 100 percent of the vulnerable components with an 8 percent false positive rate at best. However, the model predicted vulnerabilities only at the component level and cannot direct developers to more specific vulnerable code locations.

Shin and Williams [36], [37] investigated whether the code level complexity metrics, such as cyclomatic complexity, can be used as indicators of vulnerabilities at function level. The authors performed a case study on the Mozilla JavaScript Engine written in C/C++. Their results show that the correlations between complexity metrics and vulnerabilities are weak (Spearman $r = 0.30$ at best) but statistically significant. Interestingly, the complexity measures for vulnerable functions were higher than the ones for faulty functions. This observation encourages us to build vulnerability prediction models, even in the presence of faults, using complexity metrics. Results of the vulnerability prediction using logistic regression showed very high accuracy (over 90 percent) and low false positive rates (less than 2 percent), but the false negative rate was very high (over 79 percent). Our study extends these two prior studies by using additional complexity metrics such as fan-in and fan-out, and processes history metrics including code churn and developer activity metrics on two large-size projects.

Walden et al. [38] analyzed the association between the security resource indicator (SRI) and vulnerabilities on 14 open-source PHP web applications. The SRI is measured as a sum of binary values depending on the existence of the four resources in development organizations: a security URL, a security e-mail address, a vulnerability list for their products, and secure development guidelines. SRI is useful to compare security levels between organizations, but does not indicate vulnerable code locations. Additionally, they measured the correlation between three complexity metrics and vulnerabilities. The correlations were very different depending on the projects, which inhibits our ability to generalize the applicability of complexity metrics as an indicator of vulnerabilities. Since the projects were written in PHP and have a different domain than ours, their results cannot be generalized to ours.

## 8.2 Fault Prediction with Complexity, Code Churn, and Developer Metrics

Basili et al. [5] showed the usefulness of object oriented (OO) design metrics to predict fault proneness in a study performed on eight medium-sized information management systems. The logistic regression model with OO design metrics detected 88 percent of faulty classes and correctly predicted 60 percent of classes as faulty. Briand et al. [6] also used OO design metrics to predict defects and their logistic regression model classified fault-prone classes at over 80 percent of precision and found over 90 percent of faulty classes. Nagappan et al. [39] found that sets of complexity

metrics are correlated with postrelease defects using five major Microsoft product components, including Internet Explorer 6. Menzies et al. [7] explored three data mining modeling techniques, OneR, J48, and naive Bayes, using code metrics to predict defects in MDP, a repository for NASA data set. Their model using naive Bayes was able to predict defects with 71 percent $PD$ and 25 percent $PF$.

Nagappan and Ball [15] investigated the usefulness of code churn information on Windows Server 2003 to estimate postrelease failures. The Pearson correlation and the Spearman rank correlation between estimated failures and actual postrelease failures were $r = 0.889$ and $r = 0.929$, respectively, for the best model. Ostrand et al. [9] used code churn information together with other metrics including lines of source code, file age, file type, and prior fault history. They found that 83 percent of faults were in the top 20 percent of files ranked in the order of predicted faults using negative binomial regression. Nagappan et al. [8], [40] also performed empirical case studies on the fault prediction with Windows XP and Windows Server 2003 using code churn metrics and code dependency within and between modules. Both studies used a multiple linear regression model on principal components and the Spearman rank correlations between actual postrelease failures and estimated failures were $r = 0.64$ and $r = 0.68$ in the best cases, respectively.

We use two concepts to measure developer activity: developer networks and contribution networks. The concept of a developer network has come from several sources, including [16], [41]. Gonzales-Barahona et al. [41] were the first to propose the idea of creating developer networks as models of collaboration from source repositories to differentiate and characterize projects. Meneely et al. [16] applied social network analysis to the developer network in a telecommunications product to predict failures in files. They found 58 percent of the failures in 20 percent of the files where a perfect prioritization would have found 61 percent. Pinzger et al. [21] were the first to propose the contribution network as a quantification of the direct and indirect contribution of developers on specific resources of the project. Pinzger et al. found that files that were contributed to by many developers, especially by developers who were making many different contributions themselves, were found to be more failure prone than files developed in relative isolation. Other efforts exist [33], [42], [43] to quantify developer activity in projects, mostly via counting the number of distinct developers who changed a file as we did in our study. The difference between [33] and ours was discussed in Section 6.

## 9 CONCLUSIONS

The goal of this study was to guide security inspection and testing by analyzing if complexity, code churn, and developer activity metrics can indicate vulnerable files. Specifically, we evaluated if CCD metrics can discriminate between vulnerable and neutral files, and predict vulnerabilities. At least 24 of the 28 metrics supported the hypotheses for discriminative power between vulnerable and neutral files for both projects. A few univariate models and the models using development history based metrics such as code churn, developer activity, and combined CCD

metrics predicted vulnerable files with high *PD* and low *PF* for both projects. However, the models with complexity metrics alone provided the weakest prediction performance, indicating that metrics available from development history are stronger indicators of vulnerabilities than code complexity metrics we collected in this study.

Our results indicate that code churn, developer activity, and combined CCD metrics can potentially reduce the vulnerability inspection effort compared to a random selection of files. However, considering the large size of the two projects, the quantity of files and the lines of code to inspect or test based on the prediction results are still large. While a thorough inspection of every potentially vulnerable file is not always feasible, our results show that using CCD metrics to predict files can provide valuable guidance to security inspection and testing efforts by reducing code to inspect or test.

Our contribution in this study is that we provided empirical evidence that CCD metrics are effective in discriminating and predicting vulnerable files and in reducing the number of files and the lines of code for inspection. Our results were statistically significant despite the presence of faults that could weaken the performance of a vulnerability prediction model.

While our results show that predictive modeling can reduce the amount of code to inspect, much work needs to be done in applying models like ours to the security inspection process. Examining the underlying causes behind the correlations found in this paper would assist even further in guiding security inspection and testing efforts.

## ACKNOWLEDGMENTS

## REFERENCES

[1] I.V. Krsul, "Software Vulnerability Analysis," PhD dissertation, Purdue Univ., 1998.

[2] B. Cashell, W.D. Jackson, M. Jickling, and B. Web, "CRS Report for Congress: The Economic Impact of Cyber-Attacks," Congressional Research Service, Apr. 2004.

[3] G. McGraw, *Software Security: Building Security In.* Addison-Wesley, 2006.

[4] N. Fenton, M. Neil, W. Marsh, P. Hearty, D. Marquez, P. Krause, and R. Mishra, "Predicting Software Defects in Varying Development Lifecycles Using Bayesian Nets," *Information and Software Technology,* vol. 49, no. 1, pp. 32-43, 2007.

[5] V.R. Basili, L.C. Briand, and W.L. Melo, "A Validation of Object-Oriented Design Metrics as Quality Indicators," *IEEE Trans. Software Eng.,* vol. 22, no. 10, pp. 751-761, Oct. 1996.

[6] L.C. Briand, J. Wüst, J.W. Daly, and D.V. Porter, "Exploring the Relationships between Design Measures and Software Quality in Object-Oriented Systems," *J. Systems and Software,* vol. 51, no. 3, pp. 245-273, 2000.

[7] T. Menzies, J. Greenwald, and A. Frank, "Data Mining Static Code Attributes to Learn Defect Predictors," *IEEE Trans. Software Eng.,* vol. 33, no. 1, pp. 2-13, Jan. 2007.

[8] N. Nagappan and T. Ball, "Using Software Dependencies and Churn Metrics to Predict Field Failures: An Empirical Case Study," *Proc. First Int'l Symp. Empirical Software Eng. and Measurement,* pp. 364-373, Sept. 2007.

[9] T.J. Ostrand, E.J. Weyuker, and R.M. Bell, "Predicting the Location and Number of Faults in Large Software Systems," *IEEE Trans. Software Eng.,* vol. 31, no. 4, pp. 340-355, Apr. 2005.

[10] T.M. Khoshgoftaar, E.B. Allen, K.S. Kalaichelvan, and N. Goel, "Early Quality Prediction: A Case Study in Telecommunications," *IEEE Software,* vol. 13, no. 1, pp. 65-71, Jan. 1996.

[11] IEEE, "IEEE Standard for a Software Quality Metrics Methodology," IEEE Std 1061-1998 (R2004), IEEE CS, 2005.

[12] U. Brandes and T. Erlebach, *Network Analysis: Methodological Foundations.* Springer, 2005.

[13] M. Girvan and M.E.J. Newman, "Community Structure in Social and Biological Networks," *Proc. Nat'l Academy of Sciences USA,* vol. 99, no. 12, pp. 7821-7826, 2002.

[14] T. Menzies, A. Dekhtyar, J. Distefano, and J. Greenwald, "Problems with Precision: A Response to "Comments on 'Data Mining Static Code Attributes to Learn Defect Predictors'"," *IEEE Trans. Software Eng.,* vol. 33, no. 9, pp. 637-640, Sept. 2007.

[15] N. Nagappan and T. Ball, "Use of Relative Code Churn Measures to Predict System Defect Density," *Proc. 27th Int'l Conf. Software Eng.,* pp. 284-292, May 2005.

[16] A. Meneely, L. Williams, W. Snipes, and J. Osborne, "Predicting Failures with Developer Networks and Social Network Analysis," *Proc. 16th ACM SIGSOFT Int'l Symp. Foundations of Software Eng.,* pp. 13-23, Nov. 2008.

[17] B. Schneier, *Beyond Fear: Thinking Sensibly about Security in an Uncertain World.* Springer-Verlag, 2003.

[18] T.J. McCabe, "A Complexity Measure," *IEEE Trans. Software Eng.,* vol. 2, no. 4, pp. 308-320, Dec. 1976.

[19] T.L. Graves, A.F. Karr, J.S. Marron, and H. Siy, "Predicting Fault Incidence Using Software Change History," *IEEE Trans. Software Eng.,* vol. 26, no. 7, pp. 653-661, July 2000.

[20] A.H. Watson and T.J. McCabe, *Structured Testing: A Testing Methodology Using the Cyclomatic Complexity Metric,* vol. 500, no. 235, Nat'l Inst. of Standards and Technology, Sept. 1996.

[21] M. Pinzger, N. Nagappan, and B. Murphy, "Can Developer-Module Networks Predict Failures?" *Proc. Int'l Symp. Foundations in Software Eng.,* pp. 2-12, Nov. 2008.

[22] M.W. Fagerland and L. Sandvik, "Performance of Five Two-Sample Location Tests for Skewed Distributions with Unequal Variances," *Contemporary Clinical Trials,* vol. 30, pp. 490-496, 2009.

[23] T. Menzies, B. Turhan, A. Bener, G. Gay, B. Cukic, and Y. Jiang, "Implications of Ceiling Effects in Defect Predictors," *Proc. Fourth Int'l Workshop Predictor Models in Software Eng.,* pp. 47-54, May 2008.

[24] E. Arisholm and L.C. Briand, "Predicting Fault-Prone Components in a Java Legacy System," *Proc. ACM/IEEE Int'l Symp. Empirical Software Eng.,* pp. 8-17, Sept. 2006.

[25] T.J. Ostrand, E.J. Weyuker, and R.M. Bell, "Automating Algorithms for the Identification of Fault-Prone Files," *Proc. Int'l Symp. Software Testing and Analysis,* pp. 219-227, July 2007.

[26] I.H. Witten and E. Frank, *Data Mining: Practical Machine Learning Tools and Techniques.* Morgan Kaufmann Publishers, 2005.

[27] L. Breiman, "Random Forests," *Machine Learning,* vol. 45, no. 1, pp. 5-32, 2001.

[28] S. Lessmann, B. Baesens, C. Mues, and S. Pietsch, "Benchmarking Classification Models for Software Defect Prediction: A Proposed Framework and Novel Findings," *IEEE Trans. Software Eng.,* vol. 34, no. 4, pp. 485-496, July/Aug. 2008.

[29] Y. Kamei, A. Monden, S. Matsumoto, T. Kakimoto, and K. Matsumoto, "The Effects of Over and Under Sampling on Fault-Prone Module Detection," *Proc. First Int'l Symp. Empirical Software Eng. and Measurement,* pp. 196-204, Sept. 2007.

[30] S. Neuhaus, T. Zimmermann, and A. Zeller, "Predicting Vulnerable Software Components," *Proc. 14th ACM Conf. Computer and Comm. Security,* pp. 529-540, Oct./Nov. 2007.

[31] J. Demšar, "Statistical Comparisons of Classifiers over Multiple Data Sets," *J. Machine Learning Research,* vol. 7, pp. 1-30, 2006.

[32] I. Guyon and A. Elisseeff, "An Introduction to Variable and Feature Selection," *J. Machine Learning Research,* vol. 3, pp. 1157-1182, 2003.

[33] E.J. Weyuker, T.J. Ostrand, and R.M. Bell, "Do Too Many Cooks Spoil the Broth? Using the Number of Developers to Enhance Defect Prediction Models," *Empirical Software Eng.,* vol. 13, no. 5, pp. 539-559, 2008.

[34] Y. Shin and L. Williams, "Can Fault Prediction Models and Metrics Be Used for Vulnerability Prediction?" Technical Report-2010-6, North Carolina State Univ., Mar. 2010.

[35] M. Gegick, L. Williams, J. Osborne, and M. Vouk, "Prioritizing Software Security Fortification through Code-Level Metrics," *Proc. Fourth ACM Workshop Quality of Protection,* pp. 31-38, Oct. 2008.

[36] Y. Shin and L. Williams, "Is Complexity Really the Enemy of Software Security?" *Proc. Fourth ACM Workshop Quality of Protection,* pp. 47-50, Oct. 2008.

[37] Y. Shin and L. Williams, "An Empirical Model to Predict Security Vulnerabilities Using Code Complexity Metrics," *Proc. Int'l Symp. Empirical Software Eng. and Measurement,* pp. 315-317, 2008.

[38] J. Walden, M. Doyle, G.A. Welch, and M. Whelan, "Security of Open Source Web Applications," *Proc. Int'l Workshop Security Measurements and Metrics,* Oct. 2009.

[39] N. Nagappan, T. Ball, and A. Zeller, "Mining Metrics to Predict Component Failures," *Proc. 28th Int'l Conf. Software Eng.,* pp. 452-461, May 2006.

[40] N. Nagappan, T. Ball, and B. Murphy, "Using Historical In-Process and Product Metrics for Early Estimation of Software Failures," *Proc. 17th Int'l Symp. Software Reliability Eng.,* pp. 62-74, Nov. 2006.

[41] J.M. Gonzales-Barahona, L. Lopez-Fernandez, and G. Robles, "Applying Social Network Analysis to the Information in CVS Repositories," *Proc. Int'l Workshop Mining Software Repositories,* May 2004.

[42] J.P. Hudepohl, W. Jones, and B. Lague, "EMERALD: A Case Study in Enhancing Software Reliability," *Proc. Int'l Symp. Software Reliability Eng.,* pp. 85-91, Nov. 1997.

[43] N. Nagappan, B. Murphy, and V.R. Basili, "The Influence of Organizational Structure on Software Quality: An Empirical Case Study," *Proc. Int'l Conf. Software Eng.,* pp. 521-530, May 2008.

**Yonghee Shin** received the BS degree from Sookmyung Women's University in Korea and the MS degree from Texas A&M University. She received the PhD degree in computer science from North Carolina State University (NCSU) under the supervision of Dr. Laurie Williams. She is currently working as a postdoctoral researcher at DePaul University. Her research interests are in software engineering focusing on software metrics, software reliability and security, empirical software engineering, requirements traceability, and software testing. She worked for Daewoo telecommunications and Samsung SDS in Korea for eight years before returning to academia.

**Andrew Meneely** received the dual BA degree in computer science and mathematics from Calvin College in 2006 and the MS degree from North Carolina State University (NCSU) in 2008. He is currently working toward the PhD degree in the Computer Science Department at NCSU under the supervision of Dr. Laurie Williams. His research interests include empirical software engineering, software security, and developer collaboration.

**Laurie Williams** received the BS degree in industrial engineering from Lehigh University, the MBA degree from the Duke University Fuqua School of Business, and the PhD degree in computer science from the University of Utah. She is an associate professor in the Computer Science Department at North Carolina State University (NCSU). Her research focuses on agile software development practices and processes, software reliability, software testing and analysis, software security, open-source software development, and broadening participation and increasing retention in computer science. She is the director of the North Carolina State University Laboratory for Collaborative System Development and the Center for Open Software Engineering, and an area representative for the Secure Open Systems Initiative. She is the technical codirector of the Center for Advanced Computing and Communication (CACC). She worked for IBM Corporation for nine years in Raleigh, North Carolina, and Research Triangle Park, North Carolina, before returning to academia. She is a member of the IEEE.

**Jason A. Osborne** received the BS degree in mathematics from the University of California Santa Barbara and the PhD degree in statistics from Northwestern University. He is an associate professor in the Department of Statistics at North Carolina State University. Most of his effort goes toward providing internal statistics consulting to the university and his research is driven by problems that arise from that work. Examples include estimation of linear Boolean models for particle flow, estimation of population size in software reliability, and heterogeneity of variance models for gene expression experiments.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.