

Adversarial Binaries for Authorship Identification

Xiaozhu Meng, Barton P. Miller, and Somesh Jha
Computer Sciences Department
University of Wisconsin - Madison
Madison, WI, USA, 53706
{xmeng,bart,jha}@cs.wisc.edu

Abstract

Binary code authorship identification determines authors of a binary program. Existing techniques have used supervised machine learning for this task. In this paper, we look this problem from an attacker’s perspective. We aim to modify a test binary, such that it not only causes misprediction but also maintains the functionality of the original input binary. Attacks against binary code are intrinsically more difficult than attacks against domains such as computer vision, where attackers can change each pixel of the input image independently and still maintain a valid image. For binary code, even flipping one bit of a binary may cause the binary to be invalid, to crash at the run-time, or to lose the original functionality.

We investigate two types of attacks: untargeted attacks, causing misprediction to any of the incorrect authors, and targeted attacks, causing misprediction to a specific one among the incorrect authors. We develop two key attack capabilities: feature vector modification, generating an adversarial feature vector that both corresponds to a real binary and causes the required misprediction, and input binary modification, modifying the input binary to match the adversarial feature vector while maintaining the functionality of the input binary.

We evaluated our attack against classifiers trained with a state-of-the-art method for authorship attribution. The classifiers for authorship identification have 91% accuracy on average. Our untargeted attack has a 96% success rate on average, showing that we can effectively suppress authorship signal. Our targeted attack has a 46% success rate on average, showing that it is possible, but significantly more difficult to impersonate a specific programmer’s style. Our attack reveals that existing binary code authorship identification techniques rely on code features that are easy to modify, and thus are vulnerable to attacks.

1 Introduction

The task of binary code authorship attribution is to determine the authors of a binary program, and has significant application to malware forensics, software supply chain risk management, and software plagiarism detection. Recent studies [1, 6, 17, 18, 26] have made significant progress in developing machine learning based techniques to identify authors of binary programs. In this paper, we look at the problem of authorship identification from an attacker’s perspective and attempt to perform authorship evasion, whose goal is to trick machine learning classifiers for authorship identification into making wrong predictions. We show that adversarial machine learning can pose a threat to binary code authorship identification when confronted with a carefully crafted binary code artifact, causing these classifiers to produce misleading results.

Authorship evasion is the application of adversarial machine learning to authorship identification. The field of adversarial machine learning has focused on attacking and defending machine learning systems used in real world applications [4]. A specific threat is called a test time attack, where attackers change a test example to cause misprediction. Researchers have performed successful test time attacks for a wide range of domains, including computer vision [2, 7, 28], audio processing [37], and program analysis tasks [8, 27]. Such test time attacks can have serious security implications. For example, Grosse et al. [8] showed that they can change the manifest file of an Android program to circumvent malware detection; and Simko et al. [27] showed that when given source code from other people, a programmer can change the source code to avoid authorship attribution.

However, currently there are no such attacks against binary code authorship identification. The key challenge for developing such attacks is to modify the binary to not only cause misprediction, but also maintain the structural validity and functionality of the binary. Even flip-

ping one bit of a binary may cause the binary to either be invalid, such as not loadable by the loader, or lose functionality that the attackers care about. Therefore, attacks against binary code are intrinsically more difficult than attacks targeted at domains such as computer vision, where attackers can change each pixel of the input image independently and still maintain a valid image.

In this paper, we present a framework for automatically attacking techniques for binary code authorship identification. The implications of our attack framework are three-fold. First, we show that it is realistic to automatically attack binary in an end-to-end fashion: we take a binary program as input and generate a new, valid binary that has the same functionality as the input binary and causes misprediction. Second, our techniques can be used for adversarial re-training to train more secure classifiers, incorporating the generated adversarial examples into the training set to re-train a classifier with a modified loss function [13, 19]. Third, based on our experiences, we summarize the lessons we learned for designing more secure machine learning systems for binary analysis tasks.

Operationally, there are three different types of attacks: (1) the *confidence-loss attack*, which is to remove any fingerprints and anonymize the program, such that the target classifier rejects to make a prediction; (2) the *untargeted attack*, which is to cause misprediction to any of the incorrect authors; and (3) the *targeted attack*, which is to cause misprediction to a specific incorrect author. To perform a confidence-loss attack, the target classifier must have the capability to reject an input binary based on the lack of confidence in the prediction. However, most existing techniques for binary code authorship identification do not consider the confidence of their prediction [1, 26]. Therefore, we focus on untargeted and targeted attacks.

Stealthiness is an important design goal of our attack. Our attack should not leave obvious footprints that can be easily detected. We aim to improve stealthiness in two dimensions. First, the generated adversarial binary should be similar to the original binary in structure. We prefer small and local modifications over large and global modifications. Second, our attack should be diversified, meaning when running multiple times with the same input, our attack should generate different adversarial binaries. Diversified attacks make hash-based detection strategies ineffective.

We make two main assumptions about the threat model. First, the attackers have perfect knowledge of target authorship identification tool. This assumption allows performing a worst-case evaluation of the security of the target authorship identification tool, common when performing test time attacks [2, 7, 27, 28]. Second, the attackers plan to perform a test time attack, so

they can affect the prediction results only by providing a crafted input binary. Other possible attacks against learning systems such as training set poisoning [3, 15] are not in the scope of this paper.

Authorship identification techniques have a training stage and a testing stage. While we do not directly attack the training stage, three choices made in this stage impact our attacks. First, the design of the binary code features determines the program properties of the binary to modify during attacks. Features are typically defined to describe program properties including machine instructions, program control flow, constant strings, and program meta-data such as function symbols. Second, identification techniques use binary code analysis tools such as Dyninst [23], NDISASM [29] or Radare2 [20] for feature extraction. A key part of our attack is to modify the binary and trick the binary code analysis tools into extracting modified features to cause misprediction. Third, based on the machine learning algorithm used by the identification technique, the attacker may need to use different attack algorithms to determine which features should be modified to cause misprediction. There are existing attack algorithms for a variety of learning models, including Deep Neural Networks (DNNs) [7], Random Forests (RFs) [11], and Support Vector Machines (SVMs) [2, 8].

Our attack ties closely to the testing stage. Figure 1 illustrates the testing stage and the key steps of our attack. The testing stage has two key steps: extracting code features from the input binary to construct a feature vector and applying the pre-trained model on the feature vector to generate the prediction results. Our new attack focuses on developing two interacting attacking abilities: *feature vector modification*, generating an adversarial feature vector that corresponds to a real binary and causes the required misprediction, and *input binary modification*, modifying the input binary to match the adversarial feature vector while *maintaining the functionality of the input binary*. Feature vector modification guides what input binary modification should be performed to cause misprediction, while input binary modification gives feedback to feature vector modification as to which features are difficult to modify, guiding feature vector modification to avoid modifying difficult features.

Our attack framework introduces a large space for generating diversified attacks. Given an input binary and the misprediction target, feature vector modification can generate different adversarial feature vectors to cause the required misprediction. Given an adversarial feature vector, input binary modification can generate different adversarial binaries to match the feature vector.

Our approach to feature vector modification starts with existing attacks for computer vision tasks [7, 22, 37]. As these existing attacks introduced several hyper-

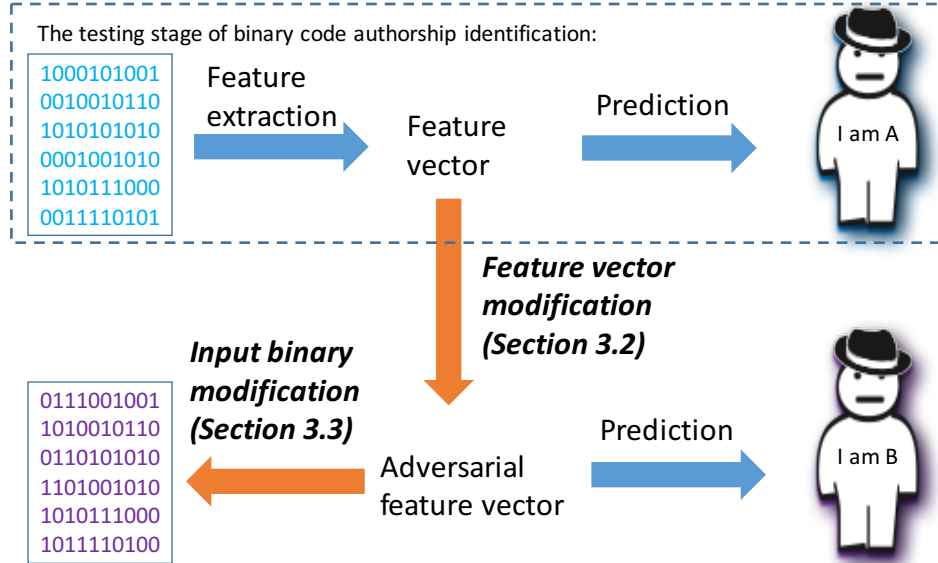


Figure 1: **Overview of our attack.** Our attack includes two key steps. Feature vector modification generates an adversarial feature vector that both corresponds to a real binary and causes the required misprediction. Input binary modification generates a new binary that matches the generated adversarial feature vector. The dashed blue box illustrates the testing stage of binary code authorship identification. A blue arrow represents a step from the testing stage. An orange arrow represents a step of our attack.

parameters, we can generate diversified adversarial feature vectors by using different attacks and running the same attack multiple times with different hyperparameters. We then extend these attacks in two ways to address the structural validity requirement of binary programs:

First, existing attacks modified each feature independently; changing one pixel of an image does not impact other pixels. However, features in our domain can be correlated. Without considering feature correlation, we may generate feature vectors for which there do not exist corresponding valid binaries.

We perform a feature correlation analysis to derive feature correlation from a substitute data set. Note that this data set can be, but does not have to be the training set used for training the target classifier. We can derive useful feature correlation information, as long as this data set is drawn from the same application domain as the training set. We then use the correlation information to ensure that correlated features are modified in a consistent way.

Second, existing attacks did not consider the difficulty of modifying a feature; changing any pixel of an image is equally easy for maintaining the validity of the image. However, for binary analysis, some features are easier to modify than others. For example, local features that describe machine instructions are typically easier to modify than global features that describe program control flow,

because modifying global features can require changing more code, making it more difficult to maintain structural validity.

We categorize binary code features into a small number of feature groups such that features in a group can be modified with the same strategy. We attempt to modify one feature group at a time until causing misprediction. Grouping features also allows generating diversified adversarial feature vectors by modifying different combinations of feature groups.

Our input binary modification removes or injects features according to the results of feature vector modification, with the additional goals of maintaining structural validity and preserving functionality. To remove features, we need to ensure that the program properties that correspond to the removed feature are replaced with semantically equivalent ones. In many cases, we cannot simply remove them because such modification would break the functionality of the binary. On the other hand, the main challenge of injecting features is to ensure that the binary code analysis tools used for feature extraction indeed recognize the injected code, data or meta-data.

We observe that the space of binary modification is large and there could be many different binaries matching the given adversarial feature vector. Therefore, we show the feasibility of our attack by construction. We design injection and removal strategies for each feature group. These modification strategies consist of a se-

quence of binary modification primitives, including inserting, deleting, and replacing code, data, and metadata. Our modification primitives use randomization, thus add another dimension to the diversity of our attack. Binary instrumentation and rewriting tools such as Dyninst support the implementation of these modification primitives.

We evaluate our evasion attacks using five classifiers trained with the techniques presented by Caliskan-Islam et al. [6]. We achieved 96% success rate for untargeted attacks and 46% success rate for targeted attacks. Our results show that we can effectively suppress authorship signal for authorship evasion, but it is significantly more difficult to impersonate the style of another author. Our results also reveal the weakness in current authorship identification techniques. Many features used in current authorship identification techniques are based on program properties that are easy to manipulate. We can automatically modify these features, making such classifiers vulnerable to test time attacks.

2 An Attack Example

We present an example showing how to perform untargeted attacks to a classifier for binary code authorship attribution. The goal of this section is to give an overview of our attack process. In the subsequence sections, we describe the steps in more details.

We first describe the procedures for setting up the target classifier, which is trained with the techniques presented by Caliskan-Islam et al. [6]. We then describe how to generate feature vectors that correspond to real binaries and cause misprediction. Finally, we give examples on how to modify the binary to match the generated feature vectors.

2.1 Binary Code Authorship Attribution

Caliskan-Islam et al. [6] assume that a binary is written by a single author, so, they predict one author for a binary. Their workflow can be summarized in four steps.

1. *Define candidate features*: They used binary code features that describe machine instructions and program control flow. They also included source code features derived from decompiled source code. The source code features include character n-grams and tree n-grams. The tree n-grams are extracted from abstract syntax trees (ASTs) built by parsing the source code. These source code features have been shown to be effective for source code authorship attribution [5].
2. *Extract features*: They used two disassemblers, NDISASM [29] and radare2 [20], to extract binary code features. To derive source code features,

they first used the Hex-Ray decompiler [9], and then used Joern [36] to parse the source code into ASTs. They represent each feature as a string. To derive feature strings, they first split the results of disassembly, decompiling, and source code parsing into tokens and then normalize hex tokens to the generic symbol “hexdecimal” and decimal digit tokens to the generic symbol “number”. They use string matching to count the frequency of a feature string and use the frequencies of feature strings to construct feature vectors.

3. *Select Features*: Typically, hundreds of thousands of features are extracted from a data set. So, feature selection is necessary to avoid overfitting. They selected features that have information gain with respect to the author labels.
4. *Train a classifier*: They compared Random Forests (RFs) with Support Vector Machines (SVMs) and reported that RFs outperformed SVMs.

They used a data set derived from Google Code Jam (GCJ) and evaluated their techniques with binaries compiled by GCC on a 32-bit platform. For binaries compiled with GCC and `-O0`, they achieved 96% accuracy for classifying 100 authors. For binaries compiled with higher optimization levels, they reported slightly lower accuracy.

We obtained the GCJ source files used by Caliskan-Islam et al. [6] and their source code for extracting features. Due to the predominance of 64-bit platforms, we perform attacks on 64-bit platforms. Note that while Caliskan-Islam et al. only evaluated their techniques on 32-bit platforms, their techniques can be directly applied to 64-bit platforms. We compiled the GCJ sources with GCC 5.4.0, using `-O0` optimization on a 64-bit platform, and achieved 90% accuracy for classifying 30 authors.

2.2 Feature Vector Modification

Given the target classifier to attack, the goal of our attack is to modify an input binary to cause the required misprediction. The two key steps for attacking this authorship attribution classifier are generating feature vectors that can correspond to a real binary and cause the required misprediction, and modifying the input binary to match the feature vector. We use examples to illustrate the importance of our feature correlation analysis and feature grouping on generating an adversarial feature vector.

2.2.1 Feature Correlation Analysis

We derive correlations between features to guide feature vector modification to generate feature vectors corresponding to real binaries.

We identify two types of feature correlation for this classifier. First, a feature can contain other features. For example, if feature “push rax; push rbx” is present in a binary, features “push rax” and “push rbx” are also present. So, the frequencies of “push rax” and “push rbx” should be no fewer than the frequency of “push rax; push rbx”. Second, the same properties extracted by different binary analysis tools are treated as different features. For example, the instruction “call fprintf” corresponds to three different features: “call fprintf” extracted by NDISASM, “call fprintf” extracted by radare2, and “fprintf” extracted from decompiled source code. These features should all have the same frequency.

We derive linear correlation between features based on the training set. For each pair of features, we perform linear regression and calculate the correlation coefficient. If the coefficient is larger than a threshold value, such as 0.9, we merge the pair into one feature. While this simple strategy will miss non-linear feature correlation, our experiments showed that capturing linear correlation is sufficient for launching successful attacks against authorship attribution.

2.2.2 Generating Adversarial Feature Vectors

We extend the attacks presented by Carlini and Wagner [7] to generate adversarial feature vectors. Their attacks are designed for DNNs trained for images, and can be readily applied to other gradient based learning algorithms. However, Caliskan-Islam et al. used RFs, which is a non-gradient based learning algorithm. Fortunately, researchers have shown that adversarial examples created for classifiers trained with one type of learning algorithms (such as DNN) are likely to cause misprediction for classifier trained with a different type of learning algorithms (such as RF) [21, 31]. Therefore, we first trained a substitute DNN using the same training data and then applied the adversarial vectors to the RF classifier. The substitute DNN is a simple feed-forward neural network, containing 7 hidden layers with each layer having 50 hidden units. The substitute DNN has 80% accuracy. While the substitute DNN has modestly lower accuracy than the target classifier, as we will show in Section 4, this accuracy gap does not impact the success rate of our attack.

To ensure that the generated adversarial feature vector confuses not only the substitute classifier but also the target classifier, we keep generating new feature vectors until the resulting vectors can mislead the target classifier. Our new attack strategy can generate effective adversarial feature vectors, reducing the accuracy of both the substitute DNN and the RF classifier to 0%.

However, it is difficult to modify the input binary to

completely match the feature vectors generated in this way, as they contain hundreds of modified features.

2.2.3 Categorizing Features

We have observed that while the attacks presented by Carlini and Wagner can make effective changes to the feature vector to cause misprediction, not all changes are necessary for causing misprediction. Therefore, we attempt to modify fewer features to cause misprediction, making it easier to perform binary modification to match the generated feature vector. We categorize features into feature groups, so that features in the same feature group can be modified with the same strategy. And then we modify one feature group at a time until misprediction occurs.

Two important factors for categorizing the features are the program properties that the features describe and the strength of the binary analysis tools. For the first factor, features describing low level code properties such as machine instructions are easier to modify compared to features describing higher level structural properties such as program control flow and data flow. Therefore, we started by attacking instruction features.

For the second factor, recall that Caliskan-Islam et al. used two disassemblers: NDISASM, which disassembles the binary linearly from the first byte of the binary file, and radare2, which understands the layout of the binary, performs binary analysis to identify code bytes, and attempts to disassemble only code bytes. It is easier to modify features extracted by NDISASM, because NDISASM also disassembles non-loadable sections and editing or adding non-loadable sections has no impact on the functionality of the program. On the other hand, instruction features extracted by radare2 typically represent real code. So, we need to ensure that we do not change the functionality when removing a radare2 feature, and ensure that radare2 disassembles the inserted code when injecting a radare2 feature.

After grouping features, we first modify instruction features extracted by NDISASM, reducing the accuracy from 90% to 45%. We then modify instruction features extracted by radare2, further reducing the accuracy from 45% to 7%. Note that only features in the these two feature groups are modified and we can generate new binaries to complete the attack.

2.3 Binary Modification Strategies

Finally, we describe our binary modification strategies for injecting and removing NDISASM and radare2 features, using four typical examples. These examples are extracted from our successful attacks. In each example, we describe the modification primitives that constitute

the modification strategy and explain why our modifications do not change the functionality of the input binary.

2.3.1 Modifying NDISASM Features

We show two examples of modifying NDISASM features. The first example shows the case where we can inject a feature by inserting bytes into the binary. As shown in Figure 2, we need to inject instruction feature “`or [rax], ebp`” into the target binary. Since NDISASM disassembles every bytes in the binary, we can add a new non-loadable section to store the bytes of the corresponding instruction. This simple injecting strategy causes NDISASM to extract this feature and does not change the functionality of the program.

The second example shows the case where we can simply remove a feature, without replacing the removed program property with a semantically equivalent one. As shown in Figure 3, this feature seems to represent an `imul` instruction. However, offset `0x3e09` of the binary is in the `.strtab` section, which stores symbol names for the compile-time symbol table. Therefore, instead of representing an instruction, the feature represents string “`in8.cpp`”. To remove this feature, we can change the string “`in8.cpp`” to any another string. `.strtab` is used at debug-time, and not used at the link-time or run-time (it disappears if the binary is stripped), so changing its content does not impact the functionality of the original program.

In addition, we tried to understand why the string “`in8.cpp`” is a useful feature. We found that the string is extracted from source file name “`1835486_1481492_paladin8.cpp`” and “`paladin8`” is the author’s name. So, this feature turns out to contain three characters of the author’s name. While a string containing three characters of the author’s name is useful for identifying the author, such author name feature is not available in any realistic context. This example teaches us a lesson that machine learning practitioners need to ensure that the feature definition and the extracted features actually match. In this case, instruction features should only be extracted from real code bytes. So, the use of NDISASM is not robust for real world identification because it disassembles all bytes in the binary.

2.3.2 Modifying radare2 Features

We now show two examples of modifying radare2 features. The first example shows the case where we need to insert new code and data. As shown in Figure 4, feature “`number.in`” represents a string. Note that this feature is not present in the target binary, and we need to inject it into the target binary to cause misprediction. We found feature “`number.in`” in another binary, based

on the instruction “`mov $0x400c57,%edi`”. Here, address `0x400c57` points to a string “`number.in`”; radare2 recognizes the string and prints it in the disassembly results.

To inject this feature, we need to (1) insert string “`number.in`” into the target binary, and (2) insert a `mov` instruction that loads the address of the inserted string. However, to trick radare2 to disassemble the inserted instruction, there are two additional steps. First, we create a function symbol pointing to the inserted code. Second, we append a return instruction after the inserted code. Since most binary analysis tools treat function symbols as ground truth for specifying the locations of code bytes, our injection strategies can be also applied to other binary analysis tools.

The second example shows the case where we need to replace existing code with semantically equivalent code to remove a feature. As shown in Figure 5, we need to remove a feature describing an object symbol. The feature is extracted from instruction “`mov 0x20157d(%rip),%rax`”. Here radare2 recognizes that the result of the PC-relative calculation points to an object symbol, so it annotates the instruction with the name of the object symbol in the disassembly results.

To remove this feature, we need to transform the calculation of the symbol address to a semantically equivalent calculation done by one or more instructions, so that radare2 cannot recognize the loading of the symbol address. To do this, we can split the address loading into two instructions: loading the address minus one into the target register and incrementing the target register by one. We cannot just overwrite the symbol name with a different string because this symbol is in the `.dynsym` section and it is used for dynamic linking (Overwriting the name of a dynamic symbol will cause the program to not be loadable).

3 Attack Framework

We describe our attack framework in this section, based on the attack algorithm in Figure 6. The inputs to our algorithm includes an input binary b , a target classifier m , feature groups fg , and a misprediction target label tar . The output of the algorithm is an adversarial binary b' that causes the required misprediction. The main component of our algorithm is an attack-verify loop, where we iterate over feature groups until we generate a new binary that causes misprediction.

Our algorithm relies on two routines from the machine learning application we are attacking: `FeatureExtraction` to extract features and `Prediction` to generate a prediction label from a set of known labels. The meaning of these labels depend on the target application. For example, a label can describe an author for authorship attribution or a compiler for compiler

Feature string	or [rax], ebp
Raw bytes	09 28
Modification	Insert bytes 09 28 into a new non-loadable section

Figure 2: An example of injecting a NDISASM feature. We can insert the bytes into a non-loadable section.

Binary name	1835486.1481492_paladin8
Feature string	imul ebp, [fs:rsi+hexadecimal], dword hexadecimal
Offset in the binary	0x3e09
Raw bytes	64 69 6E 38 2E 63 70 70
Modification	Overwrite bytes to other values

Figure 3: An example of removing a NDISASM feature. This feature seems to represent an instruction, but actually represents a string in the .strtab section.

identification. We now describe the other routines in our algorithm.

3.1 Feature Correlation Analysis

Given a set of features $F = \{f_1, f_2, \dots, f_k\}$ used in the target classifier m , our feature correlation analysis generates a partitioning of the features, $P = \{p_1, p_2, \dots, p_k\}$, where each partition consists of all correlated features. So, $\forall f_x \in p_i$ and $f_y \in p_i$, f_x and f_y are correlated; and $\forall i \neq j$, $f_x \in p_i$, and $f_y \in p_j$, f_x and f_y are not correlated. In addition, feature partitions are disjoint. So, $\forall i \neq j$, $p_i \cap p_j = \emptyset$.

We build a undirected graph to generate the feature partitioning. Let $G = (V, E)$, where each node in the graph represents a feature (so $V = F$), and each edge in the graph represents the correlation between two features. We only capture linear correlation between features, creating an edge between two nodes if the linear correlation coefficient between two features is larger than a pre-specified threshold. In another words, $E = \{(f_i, f_j) : coe_{ij} \geq T\}$, where coe_{ij} is the linear correlation coefficient between f_i and f_j and T is the pre-specified threshold. Finally, each connected component in the graph represents a partition of the correlated features.

An important observation is that we do not have to capture the exact correlation between features to launch successful attacks. For example, suppose we have three features: “ f_1 : push rax; push rbx”, “ f_2 : push rax”, and “ f_3 : push rbx”. The precise correlation is

$$(freq(f_1) \leq freq(f_2)) \wedge (freq(f_1) \leq freq(f_3)) \quad (1)$$

where $freq(f)$ represents the frequency of feature f . Our algorithm will put all three features in the same partition

and derive the following correlation:

$$freq(f_2) = A_1 freq(f_1) + B_1, freq(f_3) = A_2 freq(f_1) + B_2 \quad (2)$$

As we will discuss in the next section, it is straightforward to incorporate correlation (2) into our feature vector modification. In addition, as the linear correlation is derived from a data set drawn from the same domain as the training set for the target classifier, feature vectors satisfying correlation (2) typically also satisfy correlation (1).

3.2 Feature Vector Modification

Given an input feature vector $\mathbf{x} = [x_1, x_2, \dots, x_k]$, where x_i represents the feature value of feature f_i , our feature vector modification outputs a modified feature vector \mathbf{x}' , such that the prediction results for \mathbf{x}' are different from the prediction results of \mathbf{x} (for untargeted attacks) or are the specified results (for targeted attacks).

We use the approach of training a substitute DNN and transferring the adversarial example to the target classifier [21]. We extend the attack presented by Carlini and Wagner [7], denoted as the CW attack, to generate adversarial feature vectors. We first summarize the CW attack and then explain how we extend it to our domain. The CW attack is regarded as a powerful targeted attack. The CW attack can also be used for untargeted attacks, but the *projected gradient descent* (PGD) is regarded as a stronger untargeted attack. As we will show in Section 4, the untargeted version of the CW attack works well for us. Note that our attack framework is not specific to the CW attack, we can also use other existing attacks such as PGD to perform diversified attacks.

3.2.1 CW Attack

The CW attack was designed for a DNN. We describe only the prediction process of the DNN as we are attacking a pre-trained model. Given a feature vec-

Feature string	number.in
Machine instruction	mov \$0x400c57,%edi
Modifications	Insert string “number.in” into a new data section Insert new instructions to load the inserted string

Figure 4: **An example of injecting a radare2 feature.** This feature represents a string. We need to insert the string and insert an instruction to load the address of the string.

Feature string	obj.stdin
Machine instruction	mov 0x20157d(%rip),%rax
Modifications	Load 0x20157d(%rip)-1 into %rax Increment %rax

Figure 5: **An example of removing a radare2 feature.** This feature represents an object symbol. We can split the address loading instruction into two instructions to remove the feature.

tor \mathbf{x} , a pre-trained DNN model can be seen as a function $pred(x)$, which generates a prediction label $y = pred(x) = \text{argmax}(\text{softmax}(Z(\mathbf{x})))$, where:

- $\mathbf{z} = Z(\mathbf{x})$, where $\mathbf{z} = [z_1, z_2, \dots, z_l]$ is a vector of raw (non-normalized) predictions that the DNN generates; l is the total number of labels. \mathbf{z} is also known as the *logits*. The calculation of $Z(\mathbf{x})$ is specified by various hyper-parameters, including the depth and the width of the neural network, the choice of the activation function, and the training parameters for each hidden unit. For a pre-trained model, all these parameters are constant.
- $\mathbf{pr} = \text{softmax}(\mathbf{z})$, where $\mathbf{pr} = [pr_1, pr_2, \dots, pr_l]$ is a probability vector and pr_i is the probability that the input belongs to label i . *softmax* normalizes the raw prediction \mathbf{z} to probability distribution.
- $y = \text{argmax}(\mathbf{pr})$, meaning that the predicted label is the one that has the highest probability.

The CW attack has two variations: one for untargeted attacks, and one for targeted attacks. We first describe the untargeted version. Denote y as the original prediction label for \mathbf{x} . The output of the untargeted CW attack is a new vector \mathbf{x}' such that $pred(\mathbf{x}') \neq y$. \mathbf{x}' is defined as $\mathbf{x} + \boldsymbol{\delta}$, so once we have calculated $\boldsymbol{\delta}$, we know \mathbf{x}' .

Carlini and Wagner formulated an optimization problem to calculate $\boldsymbol{\delta}$, balancing two factors for minimization. First, to cause misprediction, the new logits vector $\mathbf{z}' = Z(\mathbf{x} + \boldsymbol{\delta})$ should satisfy the condition that z'_y is no longer the maximum element in \mathbf{z}' , which in turn means that y is not the predicted label for \mathbf{x}' . Carlini and Wagner defined function $g(\mathbf{x}')$ to measure the difference between z'_y and $\max_{i \neq y}(z'_i)$:

$$g(\mathbf{x}') = \max(z'_y - \max_{i \neq y}(z'_i) + s, 0) \quad (3)$$

Intuitively, the smaller the $g(\mathbf{x}')$, the more likely there will be a misprediction. Here, s is a hyper-parameter to

input : an input binary b ; a pre-trained model m ; feature groups fg ; and a misprediction target tar ($tar = -1$ represents untargeted attacks)

output: an adversarial binary b' that causes misprediction

```

1  $P \leftarrow \text{FeatureCorrelationAnalysis}(m)$ ;
2  $\mathbf{x} \leftarrow \text{FeatureExtraction}(b)$ ;
3  $y \leftarrow \text{Prediction}(m, \mathbf{x})$ ;
  // Keep looping until causing
  // misprediction
4 for  $g$  in  $fg$  do
5    $\mathbf{x}' \leftarrow \text{FeatureVectorModification}(\mathbf{x}, g, P)$ ;
6    $b' \leftarrow \text{InputBinaryModification}(b, \mathbf{x}', g, P)$ ;
7    $y' \leftarrow \text{Prediction}(m, \text{FeatureExtraction}(b'))$ ;
  // Non-targeted attacks succeed
8   if  $tar == -1$  and  $y \neq y'$  then break;
  // Targeted attacks succeed
9   if  $tar \neq -1$  and  $tar == y'$  then break;
```

Figure 6: **The attack algorithm.** The main structure of the algorithm is to iterate over feature groups until we generate a new binary that causes misprediction.

control the separation between z'_y and $\max_{i \neq y}(z'_i)$. s is a positive value, typically ranging from 1 to 1000.

Second, the modification to the original feature vector should be minimized to avoid detection. So, the number of non-zero elements in $\boldsymbol{\delta}$ and the magnitude of individual δ_i should also be part of the optimization function. For this purpose, Carlini and Wagner used the L_q norm, defined as

$$L_q(\boldsymbol{\delta}) = \|\boldsymbol{\delta}\|_q = \left(\sum_{i=1}^n \delta_i^q \right)^{1/q} \quad (4)$$

The attacker chooses a value for q based on the target domain. Common choices for q are 0, 2, and ∞ . L_0 measures only the number of modified features and ignores the magnitude of changes. On the other hand, L_∞ measures only the maximal magnitude of changes and ignores all other changes. L_2 balances the number of changed features and the magnitude of changes.

In general, minimizing (3) and (4) are conflicting; the more changes are made to \mathbf{x} (larger value for L_q), the

more likely the attack can cause misprediction (smaller value for $g(\mathbf{x}')$). Carlini and Wagner introduced a hyper-parameter c to balance these two conflicting optimization targets and defined the final optimization function as

$$L_q(\boldsymbol{\delta}) + c \times g(\mathbf{x} + \boldsymbol{\delta}) \quad (5)$$

When $q = 2$, the optimization function (5) is differentiable. A general purpose optimizer such as Adam [12] can be used to minimize (5) and calculate $\boldsymbol{\delta}$. For $q = 0$, (5) is not differentiable. Carlini and Wagner designed an iterative algorithm to calculate $\boldsymbol{\delta}$. In each iteration, the algorithm uses their L_2 attack to identify some features that do not have much effect on causing misprediction and then fixes those features. The values of the fixed features will not change in later iterations. By iteratively eliminating unimportant features, the algorithm identifies a small (but possibly not minimal) subset of features that can be modified to generate an adversarial example. They designed another iterative algorithm for $q = \infty$.

Finally, after calculating $\boldsymbol{\delta}$, we derive $\mathbf{x}' = \mathbf{x} + \boldsymbol{\delta}$.

For targeted attacks, denote tar as our misprediction target. The only difference between the targeted CW attack and the untargeted version is the definition of g . For targeted attacks, we want to make z'_{tar} the maximal element in the new \mathbf{z}' , so that tar will be the new prediction label. So, Carlini and Wagner defined $g(\mathbf{x}')$ as

$$g(\mathbf{x}') = \max(\max_{i \neq tar}(z'_i) - z'_{tar} + s, 0) \quad (6)$$

3.2.2 Extension to CW Attack

To apply the CW attack to our domain, we need to make two modifications. First, the CW attack may generate feature vectors with non-integer values. However, as discussed in Section 2.1, Caliskan-Islam et al. [6] used feature counts to construct feature vectors. So, \mathbf{x}' should only have integer values. A simple strategy that works well for us is to round values generated by the CW attack to the nearest integer.

Second, we must incorporate the feature correlation information derived in Section 3.1 into the attack. To do this, we normalize each individual feature to a Gaussian with zero mean and unit variance, merge all correlated features into one feature, and let the CW attacks work with only the merged features. Recall that we track linear correlation between features; for two correlated features f_1 and f_2 , $freq(f_1) = A freq(f_2) + B$. After the normalization step, A is normalized to 1 and B is normalized to 0. Therefore, we can merge them into a single feature.

We then need to determine the values of the hyper-parameters used in CW attacks. For c and s , we perform a grid search to find a successful value-pair. For L_q , we use the L_0 norm because we would like to minimize the

number of modified features rather than the magnitude of the modifications.

We found that CW’s L_0 attack often did not generate an adversarial feature vectors with the minimal number of modified features. So, we design a two-step post-processing to further reduce the number of modified features and the magnitude of changes. First, for each modified feature, we undo the modification and set its value to its unmodified value. If we can still cause misprediction, we finalize the undo of the modification. Second, for each modified feature, we enumerate every integer between the unmodified value and the new value. We set the value of this feature to the one that is closest to the unmodified value and causes misprediction. As we will show in Section 4, this simple post-processing strategy can effectively reduce the number of modified features.

3.3 Binary modification strategies

Given a new feature vector \mathbf{x}' that causes misprediction, we describe how to modify the input binary to match \mathbf{x}' , grouping features based on the program properties that the features describe and the binary analysis tool used to extract the feature. We also describe feature injection and removal strategies for feature groups. Our modification strategies consist of binary modification primitives supported by tools such as Dyninst [23]. Finally, we discuss how to determine which modification strategy to use for a specified feature and how to generate diverse adversarial binaries.

3.3.1 Feature injection strategies

Table 1 summarizes our feature injection strategies. The first column lists the program properties we are going to inject, including machine instructions and loading the address of a symbol or data. The second and third columns list the modification primitives needed to inject features that can be extracted by NDISASM and radare2. A cell with “NA” means that the binary analysis tool cannot extract the program property. We discuss the non-NA cells in more details:

- Instructions extracted by NDISASM: The modification primitive `InsertNonCodeBytes(I)` creates a new non-loadable section in the binary to store the bytes representing new instructions. As NDISASM disassembles all bytes in the target binary, `InsertNonCodeBytes(I)` ensures that the features are injected and the functionality is unchanged.
- Machine instructions extracted by radare2: The modification primitive `InsertFunction(I)` creates a new function in which we store the inserted instructions. To ensure that radare2 disassembles the inserted code, `InsertFunction(I)` creates a new

Table 1: **Summary of feature injection strategies.** The first column lists the program properties to inject. The second and third columns list the binary modification primitives used for the injection.

Program Property	NDISASM	radare2
Instructions I	InsertNonCodeBytes(I)	InsertFunction(I)
Loading symbol S	NA	addr = InsertSymbol(S) InsertFunction(loading addr)
Loading data D	NA	addr = InsertData(D) InsertFunction(loading addr)
Calling function F	NA	addr = InsertCall(F) InsertFunction(calling addr)

code section to store the inserted instructions, appends a return instruction at the end, and create a new function symbol to point to the inserted instructions.

- Loading symbol S: The modification primitive `InsertSymbol(S)` inserts the symbol S into the target binary and returns the address pointing to the symbol. It is important to properly fill in all fields of the symbol in the symbol table, including symbol type, symbol visibility, and symbol section index. Binary analysis tools may ignore incomplete symbols, causing the injection to fail. We then use `InsertFunction(I)` to insert code that loads the address of the new symbol.
- Loading data D: The modification primitive `InsertData(D)` inserts the specified data into the target binary. We typically need to create a new data section to hold the injected data. Then, we use `InsertFunction(I)` to insert code that loads the data.
- Calling function F: The modification primitive `InsertCall(F)` inserts the specified function F into the target binary, where F can be a function from an external library. In such case, we also need to add information for dynamic linking into the target binary, including a dynamic function symbol, a relocation entry, and a procedural linkage stub (PLT) for performing the external call. Then, we use `InsertFunction(I)` to insert code that calls F.

3.3.2 Feature removal strategies

Table 2 summarizes our feature removal strategies. The first column lists the program properties we are going to remove or replace. The second and third columns list the binary modification primitives needed for removing a feature group:

- Instructions I from debug-time sections: The modification primitive `Overwrite(I)` overwrites the target instruction bytes to other bytes. This strategy does not change the program’s functionality as debug-time sections are not used at link-time or run-

time.

- Instructions I from code sections: We design two strategies for this feature group. The modification primitive `Swap(I)` checks the operand dependencies and reorders the instructions if there is no dependency. The modification primitive `InsertNop(I)` inserts a nop instruction between the original instructions. Note that to insert a nop instruction, we may need to relocate the original instructions to a different location to create extra space for the nop. Therefore, we prefer `Swap` over `InsertNop` if possible.
- Addressing loading of S: The modification primitive `SplitAddrLoad(S)` splits the address loading instruction into two instructions so that radare2 will not recognize the address loading. We use the following two instructions: loading the address minus one into the target register and incrementing the target register.
- Function call to S: The modification primitive `ConvToIndCall(S)` converts a function call to S to an indirect (pointer-based) function call, so that radare2 will not recognize the call target. `ConvToIndCall(S)` uses `SplitAddrLoad(S)` to load the function call target and then generates an indirect call. Note that we need to save and restore the register used for performing the indirect call if it is live at this point in the code.

3.3.3 Deciding which strategy to apply

We have several criteria to determine which strategy to use for a modified feature. Based on the sign of δ_i , we decide whether we need to inject (see Table 1) or remove (see Table 2) features. Based on the address where the feature was extracted, we determine from which section the feature is extracted, including debug-time sections, code sections, or data sections.

For features extracted from code sections, we determine whether the feature describes a function call, loading a symbol, or loading data. If none of the three cases applies, the feature describes just instructions, and no

Table 2: **Summary of feature removal strategies.** The first column lists the program properties to remove or replace. The second and third columns list the binary modification primitives for feature removal.

Program Property	NDISASM	radare2
Instructions I from debug-time sections	Overwrite(I)	NA
Instructions I from code sections	Swap(I) or InsertNop(I)	
Addressing loading of symbol S	NA	SplitAddrLoad(S)
Addressing loading of data D	NA	SplitAddrLoad(D)
Function call to function symbol S	NA	ConvToIndCall(S)

other program property needs to be modified.

3.3.4 Generating diverse adversarial binaries

Table 1 and Table 2 show one set of feasible modification strategies to inject and remove features, out of a large space for binary modification. Other modification strategies can be designed to achieve the same goals of feature injection and removal. We use two examples to show other possibilities of binary modification. Attackers can add more modification strategies to add diversity to the attack.

Use randomization: Several of our modification strategies can incorporate randomization to generate diverse adversarial binaries. `Overwrite(I)` overwrites the target instruction bytes to other bytes. Here, we can randomly generate the overwritten bytes. Similarly, `SplitAddrLoad(S)` can split the address loading instruction into two instructions with randomization: loading the address a randomly generate integer into the target register and incrementing the target register with generated integer.

Generate semantically equivalent instructions: The natural way to remove instruction features is to replace existing instructions with semantically equivalent instructions. Superoptimizer fits our goal here [14], which takes machine instructions as input, and outputs machine instructions that compute the same functionality as the input. It is expensive to perform superoptimization in a general case. However, as we typically need to replace only short instruction sequences, the search space would be relatively small. Therefore, superoptimization is a promising method for generating semantically equivalent instructions.

4 Evaluations

We evaluate several aspects of our attacks: (1) whether we can effectively perform untargeted attacks to evade authorship identification, (2) whether we can effectively perform targeted attacks to impersonate someone else, (3) which features are modified in our attacks and which binary modification strategies are commonly used, (4)

whether our post-processing steps are effective for reducing the number of modified features, and (5) why some of our attacks failed. Our evaluations show that

- Our untargeted attacks are effective. We achieved 96% success rate in our experiments, showing that we can effectively suppress authorship signals.
- The success rate of our targeted attacks are 46% on average, showing that it is significantly more difficult to impersonate someone else.
- The top modified features describe function calls. This indicates that authorship identification classifiers heavily rely on function calls to identify authors. Therefore, inserting function calls that are associated with other authors is an effectively way to cause misprediction.
- Without our post-processing, there are 80 features to modify on average. With our post-processing, there are only 10 features to modify on average. Therefore, our post-processing procedure can significantly reduce the number of changed features for launching a successful attack.
- For failed untargeted attacks, the lack of strategies for modifying CFG features and decompiled source features is the reason for failure. For failed targeted attacks, about a third of the cases are caused by lack of modification strategies for CFG and decompiled source features; the other two thirds of the cases failed because the targeted CW attack cannot generate a feature vector that both corresponds to a real binary and causes the required misprediction.

4.1 Evaluation Methodology

We evaluated our techniques by attacking classifiers trained with the techniques presented by Caliskan-Islam et al. [6] (described in Section 2.1). Our experiments consist of the following steps:

1. Randomly sample K authors from the Google Code Jam data set of around 1000 authors used by Caliskan-Islam et al. [6]. This data set consists of the source code of single-author programs, each with an author label.
2. Compile all the programs written by the sampled

authors with GCC 5.4.0 and -O0 optimization. Each author had an average of 8 binary programs.

3. Split the binaries into a training set and a testing set, with a size ratio of about 7:1.
4. Train a random forest classifier with the training set.
5. Perform our attack on each binary in the testing set for which the target classifier makes the correct prediction. For each test binary, we perform one untargeted attack, and $K - 1$ targeted attacks. The targeted attacks attempt to cause misprediction for each of the incorrect authors.

We varied K from 5 to 100 to investigate how the number of training authors impact the effectiveness of our attacks. For each value of K , we repeated the experiments five times and report the averaged results. We used Scikit-learn [25] for training random forest classifiers, Tensorflow [30] for training substitute classifiers, and Dyninst [23] for implementing our binary modification strategies.

We use success rate to measure the effectiveness of our attacks, defined as

$$\frac{\# \text{ of successful attacks}}{\# \text{ of total attacks}} \quad (7)$$

An attack is successful if the binary generated by our attack caused the target classifier to make an incorrect prediction. For untargeted attacks, incorrect prediction means any of the incorrect authors. For targeted attacks, incorrect prediction means the specific targeted author.

4.2 Evaluation Results

The first question to answer in our evaluation is how effective is our attack. The results are shown in Table 3. In this table, the second and the third columns are the accuracy of the target classifiers and the substitute classifiers. The fourth and the fifth columns list the the success rate of untargeted and targeted attacks. Our untargeted attack has a *96% success rate on average, showing that we can effectively suppress authorship signal*. However, our targeted attacks did not enjoy the same success as the untargeted ones. Our targeted attack has a 46% success rate on average, showing that it is significantly more difficult to impersonate a specific programmer’s style.

Table 3 also shows how the number of training authors K impacts the effectiveness of our attacks. For untargeted attacks, our success rate increases as K increases. Untargeted attacks only need to cause misprediction against any of the $K - 1$ incorrect authors. The larger the K , the more incorrect authors our attacks can work with, and the higher the success rate. For targeted attacks, our success rate decreases as K increases. Targeted attacks must cause misprediction against a specific

target author. The larger the K , the more non-target authors our target attack must avoid, and the more difficult the targeted attack.

The accuracy gap between the target classifier and the substitute classifier does not obviously impact the success rate of our attack. As shown in Table 3, The accuracy gap ranges from 0% to 20%. The success rates of both untargeted and targeted attacks do not exhibit an obvious correlation with the accuracy gap.

We then investigate what are the commonly used binary modification strategies and what are the commonly modified features. In Table 4, we list the number of times that a modification primitive is used in our untargeted attacks for $K = 30$. The most frequently used primitive is `InsertCall`, indicating that the target classifiers heavily rely on function call features to identify authors. So, inserting function calls that are associated with other authors is an effectively way to cause misprediction. `SplitAddrLoad` ranks second, showing that the target classifiers also rely on features that describe the loading of a symbol to identify authors. `InsertFunction` ranks third, showing that inserting instructions that are typically seen in programs written by other authors is also effective for causing misprediction. `Swap` and `InsertNop` serve the purpose of removing instruction features. These two primitives have an effectiveness similar to `InsertFunction`, indicating that removing distinct instruction sequences associated with an author is effective for causing misprediction. Other strategies including editing debug sections, inserting data, and inserting symbols, all play important roles in our attacks.

Next, we investigate how many features we need to change to cause misprediction. In Table 5, the second column shows the number of changed features generated by the untargeted L_0 CW attack, and the third column shows the number of changed features after our post-processing step. Before post-processing, there are 80 features to modify on average. After the post-processing, there are only 10 features to modify on average. Our results show that our post-processing procedure can significantly reduce the number of changed features for performing a successful attack.

4.3 Analysis of Failed Attacks

Our attack contains two key steps: feature vector modification to generate a vector that both corresponds to a real binary and causes the required misprediction, and input binary modification to generate a new binary that matches the adversarial feature vector. A failure in either of the two steps would lead to a failed attack. Feature vector modification fails when it cannot find such an adversarial feature vector that corresponds to a real binary and causes the required misprediction. Input binary mod-

Table 3: **Evaluation results.** The second and the third columns list the accuracy of the target classifiers and the substitute classifiers. The fourth and the fifth columns show the success rate for untargeted and targeted attacks.

K	Target classifier accuracy	Substitute classifier accuracy	Untargeted attack success rate	Targeted attack success rate
5	100%	100%	88%	88%
15	100%	80%	93%	51%
30	89%	73%	98%	47%
50	86%	69%	100%	31%
100	82%	68%	100%	14%
Average	91%	78%	96%	46%

Table 4: **Number of times that a binary modification primitive is used in untargeted attacks.** The numbers are from the attacks for 30 training authors ($K = 30$). The rows are sorted in a decreasing order.

Modification primitive	Times used
InsertCall	586
SplitAddrLoad	292
InsertFunction	196
Swap & InsertNop	193
ConvToIndCall	115
InsertNonCodeByte	102
Overwrite	85
InsertData	68
InsertSymbol	45

Table 5: **The number of feature changed by our untargeted attacks.** The second column lists the average number of features changed by the L_0 CW attack. The third column shows the average number of features changed after our post-processing.

K	L_0 CW attack	Our post-processing
5	57	9
15	107	11
30	87	11
50	59	9
100	92	11
Average	80	10

ification fails when it does not generate a new binary that causes the required misprediction. We found that feature vector modification accounts for all the failed attacks.

We break down the reasons of why our feature vector modification step would fail to generate an adversarial feature vector. Recall that our feature vector modification is based on the CW attack, which generates a feature vector that causes the required misprediction, without considering whether the generated vector would correspond to a real binary. We adapted the CW attack in three ways to generate vectors that correspond to a real binary. First, as we implemented binary modification strategies for only instruction features, the CFG features and decompiled source code features are not modified during feature vector modification. Second, as the value of an instruction feature represents the number of times that this feature appears in a binary, the feature value is an integer. However, the CW attack does not guarantee to generate integer values. So, we round the results of the CW attacks to the nearest integer values. Third, we capture feature correlation and merge correlated features. We can then divide failed feature vector modification into two categories:

Lack of modification strategies for CFG and decompiled source features: It may not be sufficient to modify only instruction features to evade authorship identification. Failed attacks in this category need binary modification strategies for CFG and decompiled source features.

Insufficient handling of finding vectors corresponding to real binaries: Our techniques for generating feature vectors that correspond to real binaries need further improvement. For example, we currently capture only linear correlation between features.

We found that all the failed untargeted attacks were due to not being able to modify CFG or decompiled source features. For failed targeted attacks, not being able to modify CFG or decompiled source features explained about 34% of the failed cases; not being able to find a feature vector that corresponds to a real binary explained the other failed cases. Our analysis shows that

to improve untargeted attacks, we need to continue to design new modification strategies for CFG and decompiled source features. To improve targeted attacks, we also need to improve the targeted CW attack to find feature vectors that correspond to real binaries.

While our binary modification strategies were able to match all modified features, We found that they sometimes caused side effects and changed features that should not have been changed. Fortunately, such side effects did not impact the prediction results. The number of unintended changes ranged from 0 to 20. Most of the unintended changes were made to NDISASM instruction features. This is because our feature injection strategies often insert new code and data sections, which in turn requires changes to the program header of the binary. As NIDSASM disassembles all the bytes in the binary, the changes in the program header would cause unintended changes to NDISASM instruction features. It is not surprising that such unintended changes did not impact the prediction results as the program header is unlikely to carry authorship signals.

In summary, the our evaluations show that our attack framework is effective for untargeted attacks and we can practically suppress authorship signals. Performing targeted attacks is significantly more difficult than untargeted attacks. Our results also reveal weaknesses in current authorship identification techniques. Many features used in current authorship identification techniques are based on program properties that are easy to fabricate, such as function calls and symbols. We have shown that we can automatically modify these features, making such classifiers vulnerable to test time attacks.

5 Discussion and Related Work

To the best of our knowledge, we are the first project to perform binary code authorship evasion. In this section, we place our work in a broader context and discuss several related research areas.

Stealthy binary rewriting: Stealthiness is an important goal of our attack. We observed that the generated binaries show distinct characteristics such as the presence of additional code sections and springboard jump instructions from original code sections to newly added code sections. These distinct characteristics are introduced by Dyninst and every binary modified by Dyninst exhibits such characteristics. Since Dyninst is also widely used in many benign applications, such as binary hardening techniques [24, 32, 33]. The presence of Dyninst footprints does not necessarily indicate the presence of tamperers.

We are aware of other binary rewriting techniques, such as reassembly [34, 35]. Reassembly disassembles the binary and creates artificial symbols for data and

code references. Binary rewriting is performed by first modifying the assembly code and then re-assembling the code. Reassembly has the advantage that code can be injected or removed in place, thus providing better stealthiness.

We chose to use Dyninst for binary rewriting as it is a mature and widely used tool. We leave the exploration of using reassembly for binary rewriting as future work.

Multi-author binary code identification: We evaluated our attack against single author identification. Recent studies on binary code authorship identification investigated identifying multiple authors in a binary [16, 17, 18]. These multi-author techniques performed predictions at the basic block level, meaning they reported one author for each basic block.

It is significantly more difficult to evade multi-author identification techniques, as it will require fine-grained binary modification strategies to match adversarial feature vectors. Fine-grained modification strategies must inject or remove features within basic blocks. Many of the feature injection and removal strategies presented in Section 3 are not applicable as they will introduce new basic blocks. We believe multi-author identification is intrinsically more complicated than single-author identification and it is the natural next step to perform evasion on multi-author identification.

Adversarial learning on malware detection: While our work is not directly targeted to malware detection, we believe our techniques can contribute to this field in two ways.

First, a common threat model of adversarial learning on malware detection assumes that attackers can only inject features and cannot remove features. Monotonic classification [10] ensures that an adversary will not be able to evade the classifier by adding more features. Our results show that we can effectively remove features, challenging the validity of their threat model.

Second, existing techniques for adversarial learning on malware detection have focused on generating adversarial feature vectors to cause misprediction, but have not focused on generating new binaries that match their feature vectors. We show that it is possible to perform end-to-end attack by generating new binaries.

Evading source code authorship: Simko et al. [27] performed a study of evading source code single-author identification. Their evasion target is a classifier that has 98% accuracy on classifying 250 authors, evaluated on a data set derived from the Google Code Jam [5]. The classifier used lexical features such as variable names and language keywords, layout features such as code indentation, and syntactic features derived the abstract syntax trees parsed from the source code.

28 programmers participated in their study, including undergrad students, former or current software develop-

ers. Each programmer was given code from author X and Y and then was asked to modify source code written by X to look like code written by Y. This manual attack achieved 80% success rate for untargeted attacks and 70% success rate for targeted attacks.

Simko et al. varied the number of training authors K from 5 to 50, and commented that untargeted attacks become easier when K increases, and the ease of targeted attacks does not have an obvious connection to K . For untargeted attacks, our results align with their observation. For targeted attacks, our results do not align with their observation, where the success rate of our targeted attacks decreases as K increases. In Simko et al.'s experiments, they always performed targeted attacks against only 5 of the K authors, regardless of the value of K . In another words, their experiments did not evaluate all possible scenarios of targeted attacks. On the other hand, we attempted to perform targeted attacks against each of the incorrect authors, covering all scenarios of targeted attacks.

Simko et al. then inspected the modified source code and summarized the most common modifications:

1. Copy entire lines of code written by Y into code written by X;
2. Make typographical changes such as brackets, new-lines, space between operators;
3. Modify variable names and the location of variable declarations, typically either from or to a global variable;
4. Add or swap library calls; and
5. Change the source code structure such as adding or removing macros, changing loop types, or breaking up an if-statement

These changes are mostly local, involving a few lines of code. The study participants did not need to understand the structure or the functionality of the code to make such modifications. The modification strategies presented in this study are unlikely to achieve equal success for evading binary code authorship identification, as many of the modifications are irrelevant at the binary code level, such as typographic changes, variables renaming, and modifying macros.

6 Conclusion

We have presented our attack framework for performing authorship evasion. Our attack framework includes components for analyzing feature correlation, generating feature vectors to cause misprediction, and binary modification strategies to match the generated feature vectors. Our evaluations have shown that our attack framework is effective for untargeted attacks, which is to cause misprediction to any of the incorrect authors. Targeted attacks are significantly more difficult to achieve, which is

to cause misprediction to a specific one among the incorrect authors.

Our attack experiences show that it is not secure to rely on features derived from program properties that are easy to modify, such as function calls, symbols, data, and instructions. Authorship identification techniques must consider the trustworthiness of the features.

References

- [1] S. Alrabaee, N. Saleem, S. Preda, L. Wang, and M. Debabi. Oba2: An onion approach to binary code authorship attribution. *Digital Investigation*, 11, Supplement 1:S94 – S103, May 2014.
- [2] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrđić, P. Laskov, G. Giacinto, and F. Roli. Evasion attacks against machine learning at test time. In *2013th European Conference on Machine Learning and Knowledge Discovery in Databases (ECMLPKDD)*, Prague, Czech Republic, Sep. 2013.
- [3] B. Biggio, B. Nelson, and P. Laskov. Poisoning attacks against support vector machines. In *29th International Conference on International Conference on Machine Learning (ICML)*, Edinburgh, Scotland, June 2012.
- [4] B. Biggio and F. Roli. Wild patterns: Ten years after the rise of adversarial machine learning. *arXiv preprint arXiv:1712.03141*, 2017.
- [5] A. Caliskan-Islam, R. Harang, A. Liu, A. Narayanan, C. Voss, F. Yamaguchi, and R. Greenstadt. De-anonymizing programmers via code stylometry. In *24th USENIX Security Symposium (SEC)*, Austin, TX, USA, Aug. 2015.
- [6] A. Caliskan-Islam, F. Yamaguchi, E. Dauber, R. Harang, K. Rieck, R. Greenstadt, and A. Narayanan. When coding style survives compilation: De-anonymizing programmers from executable binaries. In *2018 Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, USA, Feb. 2018.
- [7] N. Carlini and D. Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, San Jose, CA, USA, May 2017.
- [8] K. Grosse, N. Papernot, P. Manoharan, M. Backes, and P. McDaniel. Adversarial examples for malware detection. In *22nd European Conference on Research in Computer Security (ESORICS)*, Oslo, Norway, Sep.
- [9] Hex-Rays. Hex-Rays Decompiler, <https://www.hex-rays.com/products/decompiler/>.
- [10] I. Incer, M. Theodorides, S. Afroz, and D. Wagner. Adversarially robust malware detection using monotonic classification. In *the Fourth ACM International Workshop on Security and Privacy Analytics (IWSPA)*, Tempe, AZ, USA, Mar. 2018.

- [11] A. Kantchelian, J. D. Tygar, and A. D. Joseph. Evasion and hardening of tree ensemble classifiers. In *33rd International Conference on International Conference on Machine Learning (ICML)*, New York, NY, USA, June 2016.
- [12] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [13] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In *6th International Conference on Learning Representations (ICLR)*, Vancouver, Canada, Apr. 2018.
- [14] H. Massalin. Superoptimizer: A look at the smallest program. *SIGPLAN Notices*, 22(10):122–126, Oct 1987.
- [15] S. Mei and X. Zhu. Using machine teaching to identify optimal training-set attacks on machine learners. In *The Twenty-Ninth Conference on Artificial Intelligence (AAAI)*, Austin, Texas, USA, Jan. 2015.
- [16] X. Meng. Fine-grained binary code authorship identification. In *24th ACM SIGSOFT International Symposium on Foundations of Software Engineering, Student Research Competition Track (FSE-SRC)*, Seattle, WA, USA, 2016.
- [17] X. Meng and B. P. Miller. Binary code multi-author identification in multi-toolchain scenarios, 2018. Technical Report, <ftp://ftp.cs.wisc.edu/paradyn/papers/Meng17MultiToolchain.pdf>.
- [18] X. Meng, B. P. Miller, and K.-S. Jun. Identifying multiple authors in a binary program. In *22nd European Conference on Research in Computer Security (ESORICS)*, Oslo, Norway, Sep. 2017.
- [19] T. Miyato, S.-i. Maeda, M. Koyama, K. Nakae, and S. Ishii. Distributional smoothing with virtual adversarial training. *arXiv preprint arXiv:1507.00677*, 2015.
- [20] pancake. Unix-like reverse engineering framework and commandline tools security. <http://www.radare.org/>, visited, May 2018.
- [21] N. Papernot, P. McDaniel, and I. Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*, 2016.
- [22] N. Papernot, P. McDaniel, S. Jha, M. Fredrikson, Z. B. Celik, and A. Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 372–387, Saarbrücken, Germany, Mar. 2016.
- [23] Paradyn Project. Dyninst: Putting the Performance in High Performance Computing, <http://www.dyninst.org>.
- [24] A. Pawlowski, M. Contag, V. van der Veen, C. Ouwehand, T. Holz, H. Bos, E. Athanasopoulos, and C. Giuffrida. Marx: Uncovering class hierarchies in c++ programs. In *24th Annual Symposium on Network and Distributed System Security (NDSS)*, San Diego, California, USA, Feb. 2017.
- [25] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [26] N. Rosenblum, X. Zhu, and B. P. Miller. Who wrote this code? identifying the authors of program binaries. In *16th European Conference on Research in Computer Security (ESORICS)*, Leuven, Belgium, Sep. 2011.
- [27] L. Simko, L. Zettlemoyer, and T. Kohno. Recognizing and imitating programmer style: Adversaries in program authorship attribution. *Proceedings on Privacy Enhancing Technologies*, 2018(1):127–144, 2018.
- [28] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [29] S. Tatham and J. Hall. The netwide disassembler: NDIS-ASM. <http://www.nasm.us/doc/nasmdoca.html>, visited, May 2018.
- [30] TensorFlow. An open-source software library for machine intelligence, 2015. <https://www.tensorflow.org/>.
- [31] F. Tramèr, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel. The space of transferable adversarial examples. *arXiv preprint arXiv:1704.03453*, 2017. <https://arxiv.org/abs/1704.03453>.
- [32] V. v. d. Veen, E. Gktaş, M. Contag, A. Pawoloski, X. Chen, S. Rawat, H. Bos, T. Holz, E. Athanasopoulos, and C. Giuffrida. A tough call: Mitigating advanced code-reuse attacks at the binary level. In *2016 IEEE Symposium on Security and Privacy (SP)*, San Jose, CA, USA, May 2016.
- [33] V. van der Veen, D. Andriess, E. Göktaş, B. Gras, L. Sambuc, A. Slowinska, H. Bos, and C. Giuffrida. Practical context-sensitive cfi. In *22nd ACM SIGSAC Conference on Computer and Communications Security (CCS)*, Denver, Colorado, USA, 2015.
- [34] R. Wang, Y. Shoshitaishvili, A. Bianchi, A. Machiry, J. Grosen, P. Grosen, C. Kruegel, and G. Vigna. Ramblr: Making reassembly great again. In *24th Annual Symposium on Network and Distributed System Security (NDSS)*, San Diego, CA, USA, Feb. 2017.
- [35] S. Wang, P. Wang, and D. Wu. Reassembleable disassembling. In *24th USENIX Security Symposium (USENIX Security 15)*, Washington, D.C., USA, Aug. 2015.
- [36] F. Yamaguchi, N. Golde, D. Arp, and K. Rieck. Modeling and discovering vulnerabilities with code property graphs. In *2014 IEEE Symposium on Security and Privacy (SP)*, San Jose, CA, USA, May 2014.

- [37] G. Zhang, C. Yan, X. Ji, T. Zhang, T. Zhang, and W. Xu. Dolphinattack: Inaudible voice commands. In *2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 103–117, Dallas, TX, USA, 2017.