



HAL
open science

Fine-Grained Static Detection of Obfuscation Transforms Using Ensemble-Learning and Semantic Reasoning

Ramtine Tofighi-Shirazi, Irina Mariuca Asavoae, Philippe Elbaz-Vincent

► **To cite this version:**

Ramtine Tofighi-Shirazi, Irina Mariuca Asavoae, Philippe Elbaz-Vincent. Fine-Grained Static Detection of Obfuscation Transforms Using Ensemble-Learning and Semantic Reasoning. Software Security, Protection, and Reverse Engineering Workshop (SSPREW9), Dec 2019, San Juan, United States. hal-02355528

HAL Id: hal-02355528

<https://hal.science/hal-02355528v1>

Submitted on 12 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fine-Grained Static Detection of Obfuscation Transforms Using Ensemble-Learning and Semantic Reasoning

Ramtine Tofighi-Shirazi
Univ. Grenoble Alpes, CNRS, Institut
Fourier
Trusted Labs, Thales Group
Meudon, France
ramtine.tofighishirazi@thalesgroup.com

Irina Măriuca Asăvoae
Trusted Labs, Thales Group
Meudon, France
irina-
mariuca.asavoae@thalesgroup.com

Philippe Elbaz-Vincent
Univ. Grenoble Alpes, CNRS, Institut
Fourier
F-38000 Grenoble, France
philippe.elbaz-vincent@univ-
grenoble-alpes.fr

ABSTRACT

The ability to efficiently detect the software protections used is at a prime to facilitate the selection and application of adequate deobfuscation techniques. We present a novel approach that combines semantic reasoning techniques with ensemble learning classification for the purpose of providing a static detection framework for obfuscation transformations. By contrast to existing work, we provide a methodology that can detect multiple layers of obfuscation, without depending on knowledge of the underlying functionality of the training-set used. We also extend our work to detect constructions of obfuscation transformations, thus providing a fine-grained methodology. To that end, we provide several studies for the best practices of the use of machine learning techniques for a scalable and efficient model. According to our experimental results and evaluations on obfuscators such as Tigris and OLLVM, our models have up to 91% accuracy on state-of-the-art obfuscation transformations. Our overall accuracies for their constructions are up to 100%.

KEYWORDS

machine learning, ensemble learning, deobfuscation, obfuscation, reverse engineering, symbolic execution

1 INTRODUCTION

Code obfuscation is a widely used software protection technique to mitigate the risks of reverse-engineering. It aims at protecting intellectual property by hiding the logic and data of a code. The use of code obfuscation transformations depends on the sensitivity of the application. Its applications are mainly digital right management, software licensing code or white-box cryptography, among others. Malicious codes also use extensively code obfuscation to hide their intent, evade detection and hinder analyses.

In order to properly evaluate obfuscation transformations, or to efficiently analyze malwares, many deobfuscation techniques have emerged. Their goal is to remove the protection layers applied on the code. The deobfuscation process can be seen as different strategies such as reverting, simplifying, or gathering information about the obfuscated code. In this paper we mainly focus on information gathering, particularly the static detection of obfuscation transformations. We also study an extension to the transformations constructions, namely the different methods employed for a specific obfuscation transformation to be achieved (*e.g.* dispatch-methods for control-flow flattening or code virtualization). This approach is previously known as *metadata recovery attacks* [53].

State-of-the-art deobfuscation techniques are often specific to obfuscation transformations. For example, the work of Udupa *et al.* [64] targets control-flow transformations, whereas others [6, 42, 47, 62] aim at removing opaque predicates. Generic deobfuscation techniques, however, make no assumption about the applied protections [54, 71]. These techniques are based on dynamic symbolic execution and may lack in code coverage and scalability.

Though obfuscation transformations are semantic-preserving, they may introduce side effects to the code [14]. Each transformation has its own construction methodology, thus specific patterns. Recent works try to tackle the detection of software protections using machine learning or deep learning techniques. Ugarte-Pedrero *et al.* [65] propose a semi-supervised learning approach in order to classify packed and unpacked binaries. Sun *et al.* [59], and more recently Biondi *et al.* [7], aim at detecting and identifying packers using machine learning techniques. Tofighi-Shirazi *et al.* [61] propose a deobfuscation methodology for invariant opaque predicates based on machine learning techniques.

From the variety of obfuscation techniques, as well as deobfuscation methodologies, the ability to efficiently detect the software protections used is at a prime. To that end, the recent work of Salem *et al.* [53] focuses on the detection of obfuscation transformations. Their goal is to facilitate the selection and application of adequate deobfuscation techniques. To the best of our knowledge, their work is the first to tackle code obfuscation detection using machine learning. However, their methodology is also prone to some limitations as explained next.

Current limitations. Existing detection technique for code obfuscation [53] based on machine learning techniques comes with the following limitations:

- (1) *Code dependency:* machine learning and syntax-reasoning used for the detection of obfuscation transformations can lead to code dependency. Namely, the trained model becomes dependent to the analyzed code used in the training set, thus lowering its accuracy.
- (2) *Multi-class problem:* the methodology used relies on multi-class problems for classification. Namely, they consider that one binary cannot be obfuscated with more than one obfuscation transformation. However, transformations can be combined, thus the necessity to be able to detect the several applied layers.
- (3) *Granularity:* the detection technique has a high-level of granularity. They may detect an obfuscation transformation, but they do not focus on their constructions. The latter is of

importance in order to decide which analysis to apply on obfuscated code. Many transformations constructions are made to prevent existing deobfuscation techniques.




Figure 1: Control-flow graph of a quick-sort function obfuscated using several Tigriss transformations.

Motivation. When applying obfuscation transformations for software protections, stealth is sometimes not desired. Many applications aim for dissuasion in order to prevent reverse-engineering. In any case, the goal of our methodology is to provide a static and automated framework to help reverse-engineers. By detecting obfuscation transformations, and more specifically their constructions, an analyst will gain an important amount of time. The selection of the deobfuscation process to apply requires such knowledge beforehand. A motivating example is illustrated in Figure 1. It represents the obfuscated control-flow graph of a quick-sort function. Based on the previously introduced problems, our goal is to answer the following questions:

- *Complexity:* can we detect all applied layers of obfuscation transformation?
- *Granularity:* can we detect the constructions of applied obfuscation transformations?
- *Efficiency:* can we create accurate and generic enough models for unknown data?

As previously discussed in [53], metadata recovery attacks are usually manual tasks, therefore a potential bottleneck in the reverse engineering process. Our methodology, which could be plugged-in a disassembler framework, provides all applied transformation and construction and allows reverse-engineers to setup automated deobfuscation strategies. As an example, several opaque predicates constructions prevent SMT-solver based deobfuscation techniques [70]. Other recent works prevent the application of dynamic symbolic execution techniques [5, 44]. Thus, knowing which transformations and constructions analysts are facing may prevent using unadapted techniques for the deobfuscation process.

Contributions. In order to face the above limitations and answer our motivating questions, we bring the following contributions:

- (1) A novel methodology that combines semantic reasoning with ensemble learning techniques applied for a multi-label and multi-output ensemble model. We believe that semantic reasoning will prevent our model from code dependency limitations, and provides us with the ability to detect several combined layers of obfuscation transformations.

- (2) An extension of our methodology for a fine-grained detection. Based on our main approach, a second classification model is used for the detection of the transformations constructions, based on a multi-class classification model (*i.e.* one unique label per instances).
- (3) Several studies and experiments that justify the constructions of our methodology. We compare different machine learning approaches and techniques in order to build efficient and scalable models. We also evaluate our methodology against state-of-the-art obfuscators such as Tigriss [12] and Obfuscator-LLVM [29] (*i.e.* OLLVM).

Our paper is organized as follows. Section 2 presents the background information about code obfuscation and targeted transformations. We also introduce related work, as well as notions of supervised machine learning. Section 3 describes our methodology which combines semantic reasoning with ensemble learning. Section 4 contains our studies and experiments towards an efficient implementation of our methodology. Section 5 illustrates our evaluations on state-of-the-art and publicly available obfuscators. Section 6 briefly discuss the application of our methodology to setup deobfuscation strategies. Then, we discuss our design limitations in Section 7, as well as our perspectives in Section 8. Finally, Section 9 presents our conclusions.

Limitations. While our results illustrate the interest of the methodology, evaluating the exact gain of the different components of the approach and experimental comparison to related contributions are left as future work.

2 BACKGROUND

We briefly present code obfuscation and some of the employed transformations. Then we introduce several notions related to supervised machine learning and *metadata recovery attacks* introduced in [53].

2.1 Code obfuscation

Collberg *et al.* [14] define code obfuscation as follows:

Let $P \xrightarrow{T} P'$ be a transformation T of a source program P into a target program P' . We call $P \xrightarrow{T} P'$ an *obfuscating transformation* if P and P' have the same observable behavior, P' is harder to analyze than P , and P' is no more than polynomially slower than P . Consequently, the following conditions must be fulfilled for an obfuscating transformation : if P fails to terminate, or terminates with an error condition, then P' may or may not terminate; otherwise, P' must terminate and produce the same output as P .

2.2 Obfuscation transformations

An obfuscation transformation T can be classified into different categories such as data obfuscation, static code obfuscation, and dynamic code obfuscation. Early techniques are given by Collberg *et al.* [13, 14]. A classification of all these obfuscations, as well as known deobfuscation methods has been provided by S. Schrittwieser *et al.* [56]. The following paragraphs present a non-exhaustive list of obfuscation transformations.

2.2.1 Encodings. Static data within binaries, such as strings or constant values, contain useful information for an analyst. Encoding,

as an obfuscation transformation T , converts data to a different representation. To this end, special encoding functions are employed to mitigate the need of storing the static data in clear text within the binary. During execution, the inverse function is used to decode the obfuscated data. To prevent pattern-matching attacks, the obfuscated representation must be parameterized in order to have a *family* of representations. In other words, each representation renders different-looking obfuscated variables. However, they are all based on the same obfuscating algorithm.

2.2.2 Instructions substitutions. Each program behavior can be implemented in multiple ways [68]. In other words, instructions or sequences of instructions can be replaced with syntactically different, yet semantically equivalent code. As an example, complex instruction substitution include the replacement of call instructions with a combination of push and ret instructions [35]. De Sutter *et al* [60] replaced infrequently used opcodes with blocks of more frequently used instructions in their work. This transformation reduced the total number of different opcodes used in the code and normalizes their frequency.

2.2.3 Opaque predicates. An opaque predicate [15] represents an obfuscated predicate with its outcome known at obfuscation time, but difficult to determine for a deobfuscator. Opaque predicates are used to make static reverse-engineering more complex. They introduce an analysis problem which is difficult to solve without running the program. There are two types of *invariant* opaque predicates and the *two-ways* opaque predicates. Collberg *et al.* defined these predicates by, respectively, P^T , P^F and $P^?$ opaque predicates. Several works use two-ways opaque predicates constructs, either referred to as range-dividers [4], or as correlated opaque predicates [42, 69]. Moreover, regardless of their output, *e.g.* their type, there exists many different kinds of construction that produce the opaque predicates.

2.2.4 Control-flow flattening. This obfuscation transformation aims at obscuring links between basic-blocks by flattening the control-flow. Wang *et al.* [67] describe as *chenxification* this transformation, which puts the basic-blocks of a program into a large switch-statement. A dispatcher decides then where to jump next. Control-flow flattening using a central dispatcher is also described by Chow *et al.* [11]. A similar concept by Lynn and Debray [38] uses what is called *branch functions*, which directs the control-flow to the actual target based on a call table. Further control-flow obfuscation constructions are described in [10, 17, 36, 46, 55].

2.2.5 Code virtualization. Code virtualization describes the concept of converting a program functionality into byte-code for a custom virtual machine interpreter that is bundled with the program [23, 31]. This obfuscation transformation can also be combined with *polymorphism* by implementing custom virtual machine interpreters and payloads for each instance of the program [2]. Other work [66] proposes the combination of fine-granular encryption and code virtualization to hide the virtual machine code from analysis. Collberg *et al.* [14] describe a variant of this concept under the term *table interpretation*. A similar concept by Monden *et al.* [43] uses a finite state machine-based interpreter to dynamically map between instructions and their semantics. Thus, code virtualization proposes many constructions, as for previous transformations.

2.2.6 Dynamic code modification. In this technique, similar functions are obfuscated by providing a general template in memory that is patched right before its execution [14]. Static analysis techniques fail to analyze the program, as its functionality is available at runtime only. Other concepts of dynamic code modification [30, 40] implement the idea of correcting intentionally erroneous code at runtime, right before execution.

Our goal in this paper is to evaluate our methodology against the previously presented obfuscation transformations. Beforehand, the next section will recall some notions about supervised machine learning techniques for classification.

2.3 Supervised machine learning

Supervised machine learning [26, 33] provides a dedicated methodology to produce general hypotheses from external supplied instances via a given algorithm. From these hypotheses, predictions about future instances are possible. The aim of a supervised machine learning is to build a *classification model* which will be used to assign *labels* to unknown instances. In other words, let X be an input (*i.e.* instance) and Y the output (*i.e.* predicted label). A supervised machine learning algorithm will be used to learn the mapping function f such that $Y = f(X)$. The goal is to approximate f such that for any new instance X we can predict its label Y . In our case the inputs are represented by n -dimensional vectors of numerical features for which the extraction is described in the following paragraph. The traditional *single-label* classification associates an instance X with a unique label Y' from a previously *known* finite set of labels L . This approach is then considered a *binary* classification problem if $|L| = 2$, or a *multi-class* classification problem if $|L| > 2$. Other approaches exist, such as the *multi-label* classification. In this case, an instance X is associated with a **set** of labels $S_{Y'} \subset L$. Moreover, if the model is based on a mapping function f that can return a set of multiple labels, we have a *multi-output* classification model. In our work, we use all these classification problems as described in Section 3.

2.3.1 Feature extraction. Instances of a machine learning model are usually derived from what is called *raw data*, *i.e.* the data samples we want to classify or predict. These data samples cannot be directly given to a classification model and need to be processed beforehand. This processing step is called *feature extraction* [25] and consists in combining the raw data variables into numerical features. It allows to effectively reduce the amount of data that must be processed, while accurately describing the original dataset of raw data. In our case, raw data are text documents (*e.g.* disassembly code, symbolic execution state, etc.). Therefore, one practical use of feature extraction consists in extracting the *words* (*i.e.* the features) and classify them by frequency of use (*i.e.* weights). Different approaches exist for understanding what a word is and to compute its weight. *In this paper we use the bag of words approach* [41], which identifies terms with words using term frequency, in order to extract the features for our model. It is an efficient and simple approach which fits adequately our semantic reasoning approach.

2.3.2 Classification algorithms. The choice of which specific learning algorithm to use is a critical step. Many classification algorithms

exist [28], each of them having different mapping functions. Classification is a common application of machine learning. As such, there are many metrics that can be used to measure and evaluate our models. In order to compute these metrics, *k-Fold Cross-Validation* [32] is a frequently used technique. The definition of *k-fold cross-validation* consists in reserving a particular set of samples on which the model does not train. The limited set of samples allows to estimate how the model is expected to perform on data not used during the training phase. The parameter *k* refers to the number of groups that a given dataset of samples is split into, in order to calculate the mean of our models *accuracy* as well as the *F1-score* based on the value of *k*. While the accuracy of the model represents the ratio of correctly predicted labels to the total of labels, F1-score takes both false positives and negatives into account. In our experimentations and evaluations, the accuracies and F1-scores are calculated using *k-fold cross-validation*, with $k = 10$ for a better generalization of our model to unknown instances.

Another application of cross-validation, introduced in [53], consists in a functionality-based folding. In other words, the learning set and training set are divided based on the functionality of the samples from which the raw data are generated. The goal of such evaluation methodology is to measure if the model is dependent to the underlying code functionality, independently of the obfuscation transformation applied. The next paragraph introduce furthermore the work of Salem *et al.* [53], known as *metadata recovery attack*.

2.4 Related work: Metadata recovery attack

Salem *et al.* [53] introduce the use of machine learning techniques to evaluate the stealth of obfuscation transformations throughout their detection (known as *metadata recovery attack*). Their primary hypothesis is that machine learning techniques are capable of implementing these attacks by classifying obfuscated programs according to the transformations applied. Their experiments are based on two learning algorithms, namely Naive Bayes [22] and Decision trees [52]. Their raw data are based on static disassembly or dynamic instruction traces, either stripped or not. Thus, we refer to such raw data generation as *syntax-reasoning*. The evaluation of their models is made with two classification techniques. The first one is a traditional *k-fold cross validation*, with $k = 10$. The second one is more fine-tuned since it discriminates the training and test dataset on program functionality. In other words, the test dataset is excluded of any raw data that have been used in the training dataset, based on the functionality they implement. Such process is also repeated 10 times, to calculate the average accuracy for each fold. Their results are promising, showing up to 100% of accuracy for obfuscation transformations detection with decision trees, on dynamic traces. However, these results are obtained with the conventional cross-validation, whereas the second classification mode provides lower results (up to 61% of accuracy) with decision trees. This indicates that their model is dependent of the functionality implemented in their raw data. Moreover, their work is not implemented yet to cover several layers of obfuscation transformations, as it can be the case in most obfuscated programs. In our work, we also used both cross-validation approach to compare our results

with their work. This gives an brief idea about the advantages of semantic reasoning over syntax-based approaches.

Our goal in this paper is to combine semantic reasoning and more advanced machine learning classification techniques in order to improve the accuracy. We want to have a static analysis tool, based on symbolic execution, in order to have a model that does not depend on the functionality of the program. The models are used to detect several layers of obfuscation transformations, thus having a multi-label and multi-output classification problem. Then, we extend our detection not only to the obfuscation transformations but also to their constructions. To this end, in the next section, we present our approach and methodology.

3 METHODOLOGY

In this section we present our methodology composed of several steps, as illustrated in Figure 2. **I.** In order to create our models, we generate obfuscated as well as clean samples. This generation is done using publicly available obfuscators, specifically Tigress and OLLVM. **II.** We employ then semantic reasoning via symbolic execution¹ to extract our raw data, from the generated samples. This step is presented in Section 3.1. **III.** We create two different datasets for two different kinds of classifications. Using labeled raw data, we build our datasets for the detection of obfuscation transforms, including several combinations. Another dataset is made for the detection of specific constructions related to the transformations. These steps are introduced in Section 3.4. **IV.** The previous datasets are used to train our models. In order to select the most relevant approach and learning algorithms, several studies and experiments are provided in Section 4. **V.** The final step consists in their evaluation and their application on unknown instances, as presented in Section 5.

3.1 Semantic reasoning

Static symbolic execution is a binary analysis technique that captures the semantics (*i.e.* logic) of a program. An interpreter is used to trace the program, while assuming symbolic values for inputs rather than obtaining concrete values as a normal execution would. A symbolic state \mathbb{S} is built and consists in a set of symbolic expressions S for each variables (*i.e.* registers, memory, flags, etc.). Several techniques exist for symbolic execution [3]. In order to avoid path explosions in static symbolic execution, we use an intra-procedural and bloc-centric approach, as summarized next.

3.1.1 Bloc-centric intra-procedural symbolic execution. We use semantic reasoning for the generation of our *raw data*. The symbolic representation helps to efficiently detect obfuscation transformations and constructions. Raw data refers to the representation of data samples, containing noisy features, which need to be processed in order to extract the informative characteristics to train the models. For the detection of obfuscation transformations, we choose to work on disassembled functions of binary code. On these functions, we apply static symbolic execution to retrieve their semantic representation. In our work we use disassembled functions to collect the symbolic expressions from the code, as illustrated in Algorithm

¹In our work we consider semantic retrieval only. We are not interested in generating inputs for program exploration.




Figure 2: Design steps for fine-grained static detection of obfuscation transformations and constructions.

1. First, the semantic reasoning part of our methodology is given a disassembled function F as input. For the learning phase of our methodology, F needs to be labeled. In other words, we need to know which transformations are applied in order to properly train our model. However, in order to use our methodology as a static and automated detection framework, F does not require to be labeled once the models are trained. Based on F , we iterate over each basic

Algorithm 1 semantic reasoning for raw-data generation

```

1: procedure SEMANTIC REASONING( $F$ : a disassembled function)
2:   Initialize a dictionary  $L$ 
3:   for each basic block  $B$  in  $F$  do
4:      $I_B \leftarrow \text{getInstructions}(B)$ 
5:      $IR_B \leftarrow \text{getIntermediateLanguage}(I_B)$ 
6:      $S_B \leftarrow \text{symbolicExecution}(IR_B)$ 
7:      $NS_B \leftarrow \text{normalizeSemantics}(S_B)$ 
8:      $L[B] \leftarrow NS_B$ 
9:   end for
10:   $\text{textField} = \text{generateRawData}(L, F)$ 
11:  return textField
12: end procedure
    
```

block B . We then collect the instructions of B , denoted by I_B , with the function $\text{getInstructions}()$. I_B is translated into an intermediate language, denoted by IR_B , using $\text{getIntermediateLanguage}()$. Finally, IR_B is being used for the bloc-centric symbolic execution function $\text{symbolicExecution}()$. The latter will return the symbolic state S_B , in order words, expressions of each modified variables in a static single assignment form, based on the intermediate representation IR_B previously used. The generated semantics S_B is then normalized using $\text{normalizeSemantics}()$ function. Finally, the normalized semantics NS_B is added to the dictionary L containing all normalized semantics for each processed basic block B . The content of L will be used to generate our raw data as text file. Our normalization step has the crucial role of making the model scale to unknown data. Next, Section 3.2 describes this step, along with the content of our raw data.

3.2 Semantic-based raw data

Intermediate representations often use concrete values within their generated expressions. This causes raw data to depend on addresses that are specific to some binaries and prevents our models to scale on unknown data. Some intermediate representations also use identifiers in order to express modified registers or memory areas.

```

1 ExprMem(ExprOp('+', ExprId('RSP_init', size=64), ExprInt(0xffffffffffffd0,
2 size=64), size=64) = ExprId('RDI_init', size=64)
3 ExprId('af', size=1) = ExprSlice(ExprOp('+', ExprOp('RSP_init',
4 size=64), ExprInt(0xffffffffffffc8, 64)), ExprOp('+', ExprId('
5 RSP_init', size=64), ExprInt(0xfffffffffffff8, 64)), ExprInt(0x30,
6 64), 4, 5)
7 ExprId('RBP', size=64) = ExprOp('+', ExprId('RSP_init', size=64), ExprInt(0
8 xfffffffffffff8, 64))
9 ExprMem(ExprOp('+', ExprId('RSP_init', size=64), ExprInt(0xffffffffffffc8,
10 64)), size=32) = ExprSlice(ExprId('RDX_init', size=64), 0, 32)
11 ExprId('pf', size=1) = ExprOp('parity', ExprOp('&', ExprMem(ExprInt(0x606078,
12 64), size=64), ExprInt(0xff, 64)))
13 ExprId('RAX', size=64) = ExprMem(ExprInt(0x606078, 64), size=64)
14 ExprId('IRDst', size=64) = ExprCond(ExprMem(ExprInt(0x606078, 64), size=64),
15 ExprInt(0x40064b, 64), ExprInt(0x400644, 64))
16 ExprId('zf', size=1) = ExprCond(ExprMem(ExprInt(0x606078, 64), size=64),
17 ExprInt(0x0, 1), ExprInt(0x1, 1))
18 ExprMem(ExprOp('+', ExprId('RSP_init', size=64), ExprInt(0xfffffffffffff8,
19 64)), size=64) = ExprId('RBP_init', size=64)
20 ExprId('of', size=1) = ExprInt(0x0, 1)
21 ExprId('nf', size=1) = ExprSlice(ExprMem(ExprInt(0x606078, 64), size=64), 63,
22 64)
23 ExprId('cf', size=1) = ExprInt(0x0, 1)
24 ExprId('RSP', size=64) = ExprOp('+', ExprId('RSP_init', size=64), ExprInt(0
25 xffffffffffffc8, 64))
26 ExprId('RIP', size=64) = ExprInt(0x400650, 64)
27 ExprId('IRDst', size=64) = ExprInt(0x400650, 64)
    
```

Listing 1: Symbolic state using Miasm2 intermediate language

This notation may further affect the scalability of our trained models. For the purpose of having a model that can scale to unknown data we use a normalization phase. The normalization consists in replacing all identifiers and concrete values by symbols, and non-alphanumerical characters by alphanumerical words. This is a necessary step for a complete features extraction phase that sometimes excludes non-alphanumerical characters when working on text-based raw data. In our methodology, we generate the raw data using the Miasm2 [18] intermediate language. This language is part of the symbolic execution engine that we use for the implementation of our methodology as IDA Pro plug-in. Additionally, the normalized Miasm2 intermediate language has also been successful for the application of machine learning techniques in order to de-obfuscate opaque predicates [61]. Listing 1 illustrates the symbolic

state S of the first basic-block of the function quick-sort, which is illustrated in Figure 1. Note that the complete raw data will contain the symbolic states of each basic-blocks of the quick-sort function. We can see that Miasm2 intermediate language uses several keywords to express the semantics of the basic blocks. For example, ExprId is used for registers and ExprInt for concrete values. The registers and concretes values prevent our model from scaling to unknown data, thus potentially lowering our model accuracy. This underlines the necessity to normalize the intermediate language for an efficient semantic reasoning. Listing 16 illustrates the same basic-block symbolic state, but normalized.

```

1 ExprMem(ExprOp(op+, REG0, v0), size=64) = REG1
2 REG2 = ExprSlice(ExprOp(op+, ExprOp(op+, REG0, v1), ExprOp(op+, REG0, v2), v3), 4, 5)
3 REG3 = ExprOp(op+, REG0, v2)
4 ExprMem(ExprOp(op+, REG0, v1), size=32) = ExprSlice(REG4, 0, 32)
5 REG5 = ExprOp(oparity, ExprOp(op&, ExprMem(v4, size=64), v5))
6 REG6 = ExprMem(v4, size=64)
7 IRDst = ExprCond(ExprMem(v4, size=64), v7, v8)
8 REG8 = ExprCond(ExprMem(v4, size=64), v9, v10)
9 ExprMem(ExprOp(op+, REG0, v2), size=64) = REG9
10 REG10 = v9
11 REG11 = ExprSlice(ExprMem(v4, size=64), 63, 64)
12 REG12 = v9
13 REG13 = ExprOp(op+, REG0, v1)
14 REG14 = v11
15 IRDst = v11

```

Listing 2: Symbolic state using our normalized Miasm2 intermediate language

Additionally, the normalization step also reduces the size of the raw data. This helps enhancing the efficiency of learning and testing phase in terms of execution time. The next sections will present the different machine learning techniques used in our methodology. The purpose is to create automated and efficient models for the detection of obfuscation transformations, as well as their constructions.

3.3 Ensemble learning

In machine learning, ensemble methods [19] use multiple learning algorithms. They are mostly used to obtain better predictive performance than could be obtained from any of the constituent learning algorithms alone [39, 51]. An *ensemble*, in this case, consists of a set of individually trained classifiers whose predictions are combined when processing novel instances. Different families of ensemble learning methods exists, e.g. Bagging [9], Boosting [20, 21] or Stacking [57]. Since every model has its strengths and weaknesses, ensemble models combine individual models to help cope with the weaknesses of each algorithms. In order to select the best possible predictions from our ensemble, we use a *voting* [58] algorithm. Hence, a model is selected to make the final prediction by a simple majority vote for accuracy. Our work aims to study the benefits of ensemble learning approach over individual models. Thus, we base our core methodology on voting classifiers. However, a more in-depth studies of other approaches could provide better insights into the reasons why/if ensemble models get consistently better results for this task.

3.4 Multi-label and multi-class classifications

Multi-label classification methods are increasingly required by modern applications [8, 37]. We use multi-label with multi-output classification, in order to return all the detected obfuscation layers,

specially when combined. We also focus on multi-class classifications which play a key role in our methodology due to the following facts:

- (1) the detection of all the applied obfuscation transformations is a *multi-label classification problem*. For example, if our set of labels are the applied transformations, namely control-flow flattening and code virtualization, then one binary can have both protections. In such case, our methodology needs to return all predicted labels. We then refer to such model as a *multi-output classification*.
- (2) the fine-grained detection of the constructions is a *multi-class classification problem*. For example, if we know that control-flow flattening is applied on a code, then its constructions can only be one unique label (e.g. switch-based, ifnest-based, indirect, call-based, etc.).

Multi-label classification methods differ from binary or multi-class approaches. Tsoumakas *et al.* [63] group multi-label classification methods into two categories: *problem transformation methods* that transform the multi-label classification problem either into one or more single-label classification problems, and *algorithm adaptation methods* that extend specific learning algorithms in order to handle multi-label data directly. In our methodology we use classifier chains [49], where each model is an ensemble of learning algorithm, as presented in Section 3.3. We also study the *binary relevance* methodology [24] in Section 4. These two methodologies are briefly introduced in the following paragraphs.

3.4.1 Problem transformation methods. the *binary relevance* method [72] is a problem transformation technique that transforms any multi-label problem into a binary problem for each label. Hence, it trains several classifiers, one for each class, *i.e.* one per obfuscation transformations. The union of all classes that are predicted is taken as the multi-label output. Binary relevance method is popular due to its easy implementation. However, the main drawback is that it ignores the possible correlations between labels. *Classifier chains* [50] however, as opposed to binary relevance method, take into account the labels correlations. With this methodology we have for n labels also n binary classifiers f_0, f_1, \dots, f_n constructed. The construction is made as a chain where a classifier f_i uses the predictions of all its previous classifiers f_j with $j < i$. The chain order is randomly selected in our design.

3.4.2 Algorithm adaptation methods. Algorithm adaptation extends single label classification to the multi-label context. It is usually done by changing the decision functions. Some learning algorithms support multi-label and multi-output classification (e.g. [73, 74]), whereas others can be extended.

During our experiments, these two classifications approaches, and multi-label problems will be studied in Section 4. Our objective is to provide the best suited algorithms and techniques for an efficient and accurate model.

4 EXPERIMENTS

In this section we present first the dataset used, common with previous related work [53, 61]. Our preliminary studies towards an efficient implementation of a fine-grained detection framework are

also introduced. All our experiments and evaluations are done on a Windows 7 laptop, using 16GB of RAM, and an Intel processor.

4.1 Datasets

Our experiments are made on several C code samples. We use the scikit-learn API [45] for the implementation of the models. The datasets contain various types of code, each of them having different functionalities in order to have models that do not fit to a specific type of program. The used samples are listed below:

- GNU core utilities (*i.e.* core-utils) binaries [48] for normal predicate samples;
- Cryptographic binaries for obfuscated and non-obfuscated predicates [16];
- Samples from [4] containing basic algorithms (*e.g.* factorial, sorting, etc.), non-cryptographic hash functions, small programs generated by Tigress;
- Samples involving the uses of structures and aliases [1, 27].

Our choice is motivated by the samples low ratio of dependencies and their straightforward compilation. This makes their obfuscation possible using tools such as Tigress and OLLVM without errors during compilation. *Furthermore, all datasets used for the studies and evaluations are balanced and contain between 1000 to 5000 samples.* The obfuscation transformations applied are given in Appendix B and A. The next section will present our studies based on these datasets.

4.2 Preliminary studies

Our goal in this section is to provide some answers to the following questions related to our methodology:

- **Study 1:** when only one obfuscation transformation is applied, is a single model more effective than ensemble models for the detection?
- **Study 2:** when several obfuscation transformations are applied, can the model from Study 1 be applied to the multi-label and multi-output classification problems?
- **Study 3:** when several obfuscation transformations are applied, is a multi-label and multi-output model more efficient than one binary model for each transformation, *i.e.* classifier chains?
- **Study 4:** for the fine-grained detection of obfuscation constructions, is a single model more efficient than ensemble models?

Our studies and evaluations present two different types of results based on two different evaluations approaches. One is the traditional k -folds cross-validation with scores in black colored font. The other is made with the functionality-based cross-validation approach in red colored font, used in Salem *et al.* related work [53]. Besides, we use as a traditional single-model random-forest algorithm throughout all our studies. As for the ensemble models, we combined extra-tree and random-forest learning algorithms. These algorithms were selected because they provided the best scores in terms of accuracy. For simplicity, a preliminary evaluation was made between several learning algorithms [34] (*e.g.* decision trees, k -nearest neighbors, support vector machines, neural network, naive Bayes, random forest, etc.). In order to select the best

ensemble models, we combined between 2 to 6 single models, and selected the combination that provided the best scores.

4.2.1 Study 1: In this study we experiment traditional models against ensemble learning for multi-class classification problems. Namely, each sample is assigned with a *unique* label. Thus our model returns only one label per sample. We experiment here if ensemble learning can be more efficient at detecting obfuscation transformation, when only one layer is applied. Therefore, we do not combine obfuscation transformations for this study. Table 1 il-

Obfuscation transformation	Mono-model	Ensemble-learning
Tigress transformations	Random-forest	Extra-tree & Random-forest
EncA	93% / 98%	95% / 100%
EncL	100% / 97%	100% / 100%
EncD	95% / 98%	95% / 100%
AddO	100% / 100%	98% / 100%
Flat	97% / 100%	97% / 100%
Virt	100% / 100%	100% / 100%
Jit	100% / 100%	100% / 100%
clean	91% / 100%	91% / 100%
Overall Accuracy	97% / 99%	100% / 100%

Table 1: Multi-class accuracy and F1-scores per labels for the detection of Tigress obfuscation transformations (1 layer).

lustrates our results where we see that ensemble-learning provides a similar accuracy to random-forest, up to 97%, with traditional cross-validation. The illustrated F1-scores per labels, namely the obfuscation transforms, also points out that most of them are predicted similarly with both approaches. An exception is made for arithmetic encoding, *i.e.* EncA, and opaque predicates, *i.e.* AddO. With the functionality-based cross-validation approach however, the results differs more as observed in red font. Ensemble-learning technique provides 100% accuracy and F1-score for each classes, whereas random-forest achieves slightly lower results, with an average accuracy at 99%. Due to the semantic reasoning of our methodology, the results are better with this approach when having one layer of obfuscation. Yet, these results are not sufficient to select traditional mono-models over ensemble-learning, or the opposite way. Hence, the next study will experiment these two approaches for multi-label and multi-output classification.

4.2.2 Study 2: In the following study, we combine all obfuscation transformations. The goal of our model is to correctly predict all the applied layers of obfuscation transformation. Thus, each sample can have one or more labels. We aim to compare the random-forest algorithm with the ensemble model based on random-forest and extra-trees for multi-label and multi-output classification. Our re-

Obfuscation transformation	Multi-label mono-model	Multi-label ensemble
Tigress transformations	Random-forest	Extra-tree & Random-forest
EncA	95% / 93%	96% / 92%
EncL	90% / 78%	92% / 85%
EncD	95% / 93%	96% / 92%
AddO	96% / 88%	97% / 88%
Flat	98% / 97%	99% / 91%
Virt	99% / 98%	99% / 99%
Jit	100% / 95%	97% / 92%
clean	90% / 90%	91% / 87%
Overall Accuracy	90% / 83%	92% / 82%

Table 2: Multi-label accuracy and F1-scores per labels for the detection of Tigress obfuscation transformations (several layers).

sults in Table 2 illustrate that traditional cross-validation provides

a higher overall accuracy for ensemble learning classifier as opposed to random forest. Our ensemble of models scores 92% as opposed to 90% for random-forest, with F1-scores per labels above 91%. The functionality-based cross-validation provides lower results, with an overall accuracy at 83% and at 82% for respectively random forest and ensemble models. Still, our result indicates that both approaches can efficiently detect several layers of obfuscation transformations. However, we may improve our results using problem transformations methods such as classifier chains.

The next study will experiment this hypothesis.

4.2.3 Study 3: As in the second study, we combine all obfuscation transformations but we use binary classification problem for multi-label and multi-output classification using classifier chains. Our

Obfuscation transformation	Mono-model chain	Ensemble chain
Tigress transformations	Random-forest	Extra-tree & Random-forest
EncA	95% / 92%	97% / 90%
EncL	90% / 80%	93% / 87%
EncD	95% / 92%	97% / 96%
AddO	96% / 92%	97% / 88%
Flat	97% / 97%	99% / 91%
Virt	99% / 98%	99% / 99%
Jit	100% / 90%	100% / 92%
clean	88% / 90%	92% / 90%
Overall Accuracy	90% / 85%	92% / 90%

Table 3: Classifier chain accuracy and F1-scores per labels for the detection of Tigress obfuscation transformations (several layers).

results with standard cross-validation does not differ much from previous Study 2 as illustrated in Table 3. The functionality-based cross-validation provides improved overall accuracies and F1-scores per labels. Ensemble models used in classifier chains provide 90% of overall accuracy, compared to random-forest used in classifier chains that score 85% of overall accuracy. This study led us to select ensemble-learning techniques with classifier chains in our methodology since classifier chains allow us to create an efficient and accurate model for the detection of obfuscation transformations with one or more layers.

4.2.4 Study 4: For this final study, our goal is to evaluate the models for the fine-grained detection of an obfuscation transformation construction. We use in our dataset several Virtualized samples

Code virtualization	Mono-model	Ensemble model
Tigress constructions	Random-forest	Extra-tree & Random-forest
linear-based	100% / 99%	100% / 100%
switch-based	100% / 98%	100% / 100%
if-nest-based	100% / 100%	100% / 100%
Overall Accuracy	100% / 99%	100% / 100%

Table 4: Accuracy and F1-scores per labels for the detection of Virtualized constructions.

with Tigress for our experiment. Tigress allows the user to select different kinds of constructions, such as *switch-based*, *ifnest-based*, *linear-based*, *interpolation-based* for example. This experiment is equivalent to Study 1 in the sense that it is a multi-class classification problem. Namely, each sample has a unique label and the selected model will return one unique label per instance.

Our results in Table 4 show that both random-forest and ensemble models provides the same F1-scores per labels. Their overall

accuracies with standard cross-validation are also with 100% accuracy. With functionality-based cross-validation, ensemble models are slightly more efficient with a 100% accuracy as opposed to 99% for mono-model based on random-forest. This led us to select ensemble models in our methodology also for the classification of constructions, as it allows a fined-grained detection capability.

5 EVALUATIONS

In this section we evaluate our models with respect to the following classification problems:

- (1) *Multi-label and multi-output evaluation:* can our model, based on a classifier chain of ensemble models, efficiently and accurately detect all obfuscation transformations when one or more layers are applied?
- (2) *Multi-class evaluation:* once the obfuscation transformation detected, can our ensemble model efficiently and accurately detect the construction of the latter?

We use both cross-validation evaluation schemes as detailed in Section 2.3.2. Our evaluations are made with publicly available obfuscators, specifically Tigress and OLLVM, in order to combine obfuscation transformations from different tools.

5.1 Transformations detection

Our goal is to evaluate the stealth of obfuscation transformation, either applied as unique layer or combined. We use our multi-label and multi-output model based on ensemble-models and classifier chain to detect all the transformations applied. To measure the efficiency of our model, we used both traditional and functionality-based cross-validation as explained in Section 2.3.2. A list of all combinations of the applied transformations used in our evaluations can be found in Appendices A and B. Additionally, command line options for Tigress and OLLVM are given in A.1 and B.1.

Obfuscation transformation	Classifier Chain
OLLVM	Ensemble model
bcf	98% / 98%
fla	92% / 95%
sub	82% / 80%
clean	94% / 93%
Overall Accuracy	86% / 89%
Cross-validation execution time	11s for 1000 samples

Table 5: Evaluated accuracy and F1-scores per labels for the detection combined OLLVM transformations.

5.1.1 OLLVM. Our first evaluation uses OLLVM. It implements transformations such as opaque predicates (*i.e.* bogus control flow, *bcf*), instruction substitutions (*i.e.* *sub*) and control-flow flattening (*i.e.* *fla*). We built a dataset with several combinations of these transformations (*c.f.* Appendix B) in order to measure the efficiency of our model. Table 5 shows our results. Our model achieves an overall accuracy of 86% with traditional cross-validation and 89% with the functionality-based one. F1-scores for labels *bcf*, *fla*, and *clean* where no transformations are applied, are over 92% and up to 98% for *bcf*. However, the efficiency of our model to detect OLLVM instructions substitutions transformations, labeled as *sub*, achieves a low F1-score at 80%. Further evaluations indicate that *sub* is often considered *clean* by our model. Thus, when combined with other transformations, *sub* transformation is often undetected.

5.1.2 *Tigress*. Our second evaluation is made with the Tigress obfuscator. Tigress can generate state-of-the-art transformations such as dynamic-code generation (i.e. *Jit*), code-virtualization (i.e. *Virt*), control-flow flattening (i.e. *Flat*), opaque predicates (i.e. *AddO*) and several encoding (i.e. *Arithmetics*, *Literals* and *Data*, respectively *EncA*, *EncL* and *EncD*), among others. As illustrated in Table 6, our

Obfuscation transformation	Classifier Chain
Tigress	Ensemble model
EncA	94% / 90%
EncL	90% / 86%
EncD	92% / 91%
AddO	95% / 96%
Flat	96% / 98%
Virt	99% / 100%
Jit	100% / 100%
clean	91% / 89%
Overall Accuracy	90% / 91%
Cross-validation execution time	114s for 4000 samples

Table 6: Evaluated accuracy and F1-scores per labels for the detection combined Tigress transformations.

model accuracy is up to 90% with standard cross-validation. With functionality-based cross-validation, the overall accuracy is at 91%. F1-scores for heavy transformation such as *Virt* and *Jit* are up to 99% and 100%. The lowest F1-score is for i.e. *EncL* which is sometimes considered as a *clean* sample by our model. Regardless, our evaluation underlines the accuracy and efficiency of our methodology against Tigress transformations.

5.1.3 *OLLVM and Tigress*. For this evaluation we combine both OLLVM and Tigress datasets. We aim to see if our model is able to detect common obfuscation transformations. Table 7 shows our

Obfuscation transformation	Classifier Chain
Tigress and OLLVM	Ensemble model
EncA and sub	93% / 90%
EncL	88% / 88%
EncD	90% / 88%
AddO and bcf	95% / 95%
Flat and fla	96% / 99%
Virt	99% / 100%
Jit	100% / 100%
clean	83% / 80%
Overall Accuracy	88% / 86%
Cross-validation execution time	143s for 5000 samples

Table 7: Evaluation accuracy and F1-scores per labels for the detection of both Tigress and OLLVM transformations.

results. F1-scores for heavy transformations such as *Virt*, *Jit* and *Flat* are high, averaging up to 100% for *Jit* as an example. Combined test samples between obfuscators such *EncA-sub*, *AddO-bcf*, and *Flat-fla* have high F1-scores, even when combined with other transformations. These heavy transformations introduce important side-effects to the code, allowing an efficient and accurate detection of our model. The ability to efficiently detect non-obfuscated samples is still low compared to the ability to detect all layers of obfuscation transformations. In that case, our model F1-scores are up to 83% and 80% depending on the cross-validation approach used. Still, our model is averaging an accuracy up to 88% and 86%. These overall accuracies illustrate our model efficiency regarding the detection of obfuscation transformations, even when combined, and between the two different obfuscators.

5.1.4 *OLLVM vs. Tigress*. Our final evaluation aims to compare the accuracies of our model depending on the learning dataset used. First, we use a learning dataset only based on OLLVM transforms. The model will be then evaluated against some similar obfuscation transformations generated by Tigress. Second, we do the opposite, namely train our model on Tigress samples to evaluate it on OLLVM raw data. The results are displayed in Table 8. As we

Training dataset	Testing dataset	Overall accuracy
OLLVM	Tigress (Flat)	100% / 100%
OLLVM	Tigress (Flat, AddO)	68% / 61%
Tigress	OLLVM (fla)	95% / 92%
Tigress	OLLVM (all)	82% / 75%

Table 8: Overall accuracies of our model using either OLLVM or Tigress learning dataset.

can see, our model efficiently detects Tigress *Flat* transformation when training on 1000 samples of all OLLVM transforms, with 100% of accuracy. Results are lower when the training dataset is based on Tigress (4000 samples), against OLLVM *fla* transform, with an overall accuracy up to 95% with a standard cross-validation. Moreover, we can observe that our model cannot efficiently detect Tigress opaque predicates, i.e. *AddO*, when training only on OLLVM transforms. The results, in that case, indicate that our model efficiently detects the *Flat* transformation, but only few *AddO* ones. Finally, when our model is trained on Tigress, the overall accuracy is up to 82% against all OLLVM transforms (c.f. Appendix B). This result indicates that our methodology provides some genericity.

5.2 Constructions detection

In this section we evaluate our model for the detection of specific obfuscation transformations constructions. We use our multi-class model, based on ensemble-models, to provide a fine-grained detection technique. As for previous evaluations, we use traditional and functionality-based cross-validation techniques.

Control-flow flattening	Ensemble model
Tigress and OLLVM	Extra-tree & Random-forest
switch-based	98% / 95%
if-nest-based	98% / 100%
Overall Accuracy	98% / 97%
Cross-validation execution time	12s for 1000 samples

Table 9: Evaluation accuracy and F1-scores per class for the detection of control-flow flattening constructions.

5.2.1 *Control-flow flattening*. As for code virtualization, control-flow flattening can also be constructed in several ways, as introduced in Section 2.2. Facing the same limitations as for code virtualization constructions, we evaluated two constructions namely *switch-based* from the Tigress obfuscator, and *ifnest-based* from OLLVM. The evaluation results are in Table 9. Our model averages high F1-scores and accuracy, the latter being at 98% with standard cross-validation evaluation.

5.2.2 *Opaque predicates*. Many opaque predicates constructions exists, some of them having as purpose preventing the usage of existing deobfuscation techniques based on dynamic-symbolic execution. For the detection of their constructions, we used Tigress,

OLLVM but also novel bi-opaque methods [70]. Our results in Table 10 show that our model is accurately detecting opaque predicates constructions. F1-scores are up to 100% with standard cross-validation. Bi-opaque constructions are however often un-detected when combined with other transformations.

Opaque predicates	Ensemble model
Tigress and OLLVM	Extra-tree & Random-forest
Floats	85% / 89%
Symbolic-memory	87% / 87%
Arithmetic	100% / 100%
Aliasing	100% / 99%
Mixed-boolean and arithmetic	100% / 96%
Overall Accuracy	95% / 93%
Cross-validation execution time	24s for 1000 samples

Table 10: Evaluation accuracy and F1-scores per class for the detection of opaque predicates constructions.

Yet, the overall accuracy of our model is at 95% and 93% depending on the evaluation approach used. This illustrates the efficiency of our methodology towards the detection of obfuscation transformations constructions.

6 LIMITATIONS

One threat to the validity of our results is that we only use datasets of relatively small C programs, except for the core-utils binaries used for non-obfuscated samples. Nevertheless, the samples used in our dataset involve all common programming language constructions and various functionalities (e.g. hash functions, sorts, cryptographic algorithms, etc.). However, our future work will include the evaluation of our methodology on other obfuscators or programs, such as malwares. Our work shows that semantic reasoning combined with advanced machine learning present capabilities for a fine-grained detection of obfuscation transforms.

The capability of detecting *unknown* transformations or constructions represents another limitation of our methodology. If our model did not train on one specific transformation or constructions, it will not predict properly the unknown sample. This can lead to a loss of accuracy when unknown transformations are combined.

Dynamic transformations cause limitations to our model for the *static* detection of obfuscation transforms. Despite from the fact that we are able to accurately detect some of these transformations (i.e. *Jit*, *Virt*), when other obfuscation transformations are applied before them, our model is less efficient. Moreover, other transformations such as packing, or anti-symbolic execution techniques may lower the accuracy of our model. However, as pointed out in the next section, our methodology can scale to dynamically collected traces which allows to thwart some of these limitations.

7 PERSPECTIVES AND FUTURE WORK

First, more in-depth studies of aggregation approaches used in ensemble learning must be done in order to assess if ensemble learning are consistently more efficient for that task compared to mono-models. The hard voting scheme used is a simple approach, but may not achieve the effective benefit of using the ensemble learning approach.

As seen in [61], semantic reasoning and machine learning provides promising results for deobfuscation methodology. The evaluations shown in this paper illustrate that our model does not depend

on the code functionality. A more accurate comparison must be made as future work.

To overcome the dynamic transformations limitations, we can adapt our methodology to dynamically collected instructions traces. With a given instructions trace, we reconstruct each basic-blocks and apply our semantic reasoning approach in order to generate raw data. This step can be done either for the learning or the evaluation phase. Our future work consists in extending the implementation of our current framework and evaluating other combinations of obfuscation transformations based on dynamic traces.

Another issue we need to consider is the application of n layers of the same obfuscation transformations. Presently, our evaluations is done by combining several transformations, but using one time each of them. Our future study should consider the extension of our evaluations to the use of one transformation several times.

We also plan on extending our datasets of C programs with more complex real-world software libraries in the interest of strengthening our experiments.

8 CONCLUSIONS

In this paper we presented the efficiency of semantic reasoning combined with advanced machine learning techniques. This combination is motivated by the construction of a fine-grained detection framework of obfuscation transformations and constructions. By extending our approach to multi-label and multi-output classification, we enhanced metadata recovery attacks to the detection of multiple layers of obfuscation transformations. We proposed a new approach that combines a bloc-centric symbolic execution with machine learning ensemble model and classifier chains. We used our models to evaluate the stealth of both obfuscation transformations and constructions. Our results are promising, with overall accuracies up to 91% for the transformations and 100% for the constructions, showing slight improvements with respect to current mono-models machine learning. The use of static symbolic execution allows us to be dependent on the underlying functionality of the code samples used for the learning phase. Our empirical studies illustrate that our choices conduct towards the implementation of an efficient and accurate evaluation framework against state of the art obfuscators. However, there is still place for improvements with a more in-depth study of learning algorithms used and their parameters. Our work slightly improves metadata-recovery attacks, and paves the way towards the efficient use of advanced machine learning combined with semantic reasoning.

ACKNOWLEDGMENTS

We are much grateful to the referees comments and suggestions from the editorial board. This work is supported by the French National Research Agency in the framework of the Investissements d’Avenir program (ANR-15-IDEX-02).

A TIGRESS TRANSFORMATIONS

We list the combinations of obfuscation transformations used for our datasets, in their application order: AddOpaque (16 or 32 times); AddOpaque, EncodeLiterals; EncodeLiterals; AddOpaque, EncodeArithmetics;

EncodeArithmetics, AddOpaque; EncodeArithmetics; AddOpaque, EncodeData; EncodeData, AddOpaque; EncodeData; AddOpaque, EncodeArithmetics, EncodeLiterals, EncodeData; EncodeData, EncodeArithmetics, EncodeLiterals, AddOpaque; AddOpaque, Flatten; Flatten, AddOpaque; Flatten; Flatten, EncodeData, EncodeArithmetics, EncodeLiterals; Virtualize, AddOpaque; Virtualize, EncodeData, EncodeArithmetics, EncodeLiterals; Virtualize, Flatten; Flatten, AddOpaque, EncodeData, EncodeArithmetics, EncodeLiterals; Virtualize, AddOpaque, EncodeData, EncodeArithmetics, EncodeLiterals; Virtualize, Flatten, AddOpaque, EncodeData, EncodeArithmetics, EncodeLiterals; Jit; Jit, AddOpaque; Jit, AddOpaque, EncodeData, EncodeArithmetics, EncodeLiterals.

A.1 Commands options

```

1 # AddOpaque options
2 tigress --Transform=InitEntropy --Transform=InitOpaque --InitOpaqueStructs=
   list , array , env --Functions=main --Transform=AddOpaque --Functions=${3}
   --AddOpaqueCount=${NM} --AddOpaqueKinds=call , fake , true
3 # Flatten
4 tigress --Transform=Flatten --FlattenDispatch=switch , goto --Functions=${3}
5 # Virtualize
6 tigress --Transform=Virtualize --VirtualizeDispatch=switch , direct , ifnest ,
   linear --Functions=${3}
7 # Jit
8 tigress --include $TIGRESS_HOME/jitter --amd64.c --Transform=Jit --Functions=${
   3} --JitEncoding=hard
9 # EncodeLiterals
10 tigress --Transform=EncodeLiterals --Functions=${3} --EncodeLiteralsKinds=
   integer , string
11 # EncodeArithmetics
12 tigress --Transform=EncodeArithmetic --Functions=${3} --EncodeLiteralsKinds=
   integer
13 # EncodeData
14 tigress --Transform=EncodeData --LocalVariables=${4} --EncodeDataCodecs=poly ,
   xor , add --Functions=${3}

```

Listing 3: Tigress commands for sample generation

B OLLVM TRANSFORMATIONS

We list the combinations of obfuscation transformations used for our datasets, in their application order: bcf; bcf, sub; bcf, sub, fla; bcf, fla, sub; sub; sub, bcf; sub, bcf, fla; fla; fla, bcf; fla, sub, bcf; fla, bcf, sub.

B.1 Commands options

```

1 # Bogus control-flow
2 clang ${1}.c -o ${1} -mllvm -bcf -mllvm -bcf_prob=50
3 clang ${1}.c -o ${1} -mllvm -bcf -mllvm -bcf_prob=100
4 # Control-flow flattening
5 clang ${1}.c -o ${1} -mllvm -fla
6 clang ${1}.c -o ${1} -mllvm -fla -mllvm -split
7 # Instruction substitution
8 clang ${1}.c -o ${1} -mllvm -sub

```

Listing 4: OLLVM commands for sample generation

REFERENCES

- [1] The Algorithms. C. <https://github.com/TheAlgorithms/C/>. [Online; accessed 30-01-2019].
- [2] Bertrand Anckaert, Mariusz H. Jakubowski, and Ramarathnam Venkatesan. Proteus: virtualization for diversified tamper-resistance. In Moti Yung, Kaoru Kurosawa, and Reihaneh Safavi-Naini, editors, *Proceedings of the Sixth ACM Workshop on Digital Rights Management, Alexandria, VA, USA, October 30, 2006*, pages 47–58. ACM, 2006.
- [3] Roberto Baldoni, Emilio Coppa, Daniele Cono D’Elia, Camil Demetrescu, and Irene Finocchi. A survey of symbolic execution techniques. *ACM Comput. Surv.*, 51(3):50:1–50:39, 2018.
- [4] Sebastian Banescu, Christian S. Collberg, Vijay Ganesh, Zack Newsham, and Alexander Pretschner. Code obfuscation against symbolic execution attacks. In Stephen Schwab, William K. Robertson, and Davide Balzarotti, editors, *Proceedings of the 32nd Annual Conference on Computer Security Applications, ACSAC 2016, Los Angeles, CA, USA, December 5-9, 2016*, pages 189–200. ACM, 2016.
- [5] Sebastian Banescu, Christian S. Collberg, and Alexander Pretschner. Predicting the resilience of obfuscated code against symbolic execution attacks via machine learning. In Engin Kirda and Thomas Ristenpart, editors, *26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017*, pages 661–678. USENIX Association, 2017.
- [6] Sébastien Bardin, Robin David, and Jean-Yves Marion. Backward-bounded DSE: targeting infeasibility questions on obfuscated codes. In *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*, pages 633–651, 2017.
- [7] Fabrizio Biondi, Michael A. Enescu, Thomas Given-Wilson, Axel Legay, Lamine Noureddine, and Vivek Verma. Effective, efficient, and robust packing detection and classification. *Computers & Security*, 85:436–451, 2019.
- [8] Matthew R. Boutell, Jiebo Luo, Xipeng Shen, and Christopher M. Brown. Learning multi-label scene classification. *Pattern Recognition*, 37(9):1757–1771, 2004.
- [9] Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- [10] Jan Cappaert and Bart Preneel. A general model for hiding control flow. In Ehab Al-Shaer, Hongxia Jin, and Marc Joye, editors, *Proceedings of the 10th ACM Workshop on Digital Rights Management, Chicago, Illinois, USA, October 4, 2010*, pages 35–42. ACM, 2010.
- [11] Stanley Chow, Yuan Xiang Gu, Harold Johnson, and Vladimir A. Zakharov. An approach to the obfuscation of control-flow of sequential computer programs. In George I. Davida and Yair Frankel, editors, *Information Security, 4th International Conference, ISC 2001, Malaga, Spain, October 1-3, 2001, Proceedings*, volume 2200 of *Lecture Notes in Computer Science*, pages 144–155. Springer, 2001.
- [12] Christian Collberg, Sam Martin, Jonathan Myers, Bill Zimmerman, Petr Krajca, Gabriel Kerneis, Saumya Debray, and Babak Yadegari. The Tigress C Diversifier/Obfuscator. <http://tigress.cs.arizona.edu/index.html>. [Online; accessed 30-01-2019].
- [13] Christian Collberg and Jasvir Nagra. *Surreptitious Software: Obfuscation, Watermarking, and Tamperproofing for Software Protection*. Addison-Wesley Professional, 1st edition, 2009.
- [14] Christian Collberg, Clark Thomborson, and Douglas Low. A taxonomy of obfuscation transformations, 1997.
- [15] Christian S. Collberg, Clark D. Thomborson, and Douglas Low. Manufacturing cheap, resilient, and stealthy opaque constructs. In *POPL ’98, Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Diego, CA, USA, January 19-21, 1998*, pages 184–196, 1998.
- [16] Brad Conte. crypto-algorithms. <https://github.com/B-Con/crypto-algorithms>. [Online; accessed 30-01-2019].
- [17] Bart Coppens, Bjorn De Sutter, and Jonas Maebe. Feedback-driven binary code diversification. *TACO*, 9(4):24:1–24:26, 2013.
- [18] Fabrice Desclaux. Miasm: Framework de reverse engineering. <https://github.com/cea-sec/miasm>, 2012. [Online; accessed 30-01-2019].
- [19] Thomas G. Dietterich. Ensemble methods in machine learning. In *Proceedings of the First International Workshop on Multiple Classifier Systems, MCS ’00*, pages 1–15. London, UK, UK, 2000. Springer-Verlag.
- [20] Yoav Freund. Boosting a weak learning algorithm by majority. In Mark A. Fulk and John Case, editors, *Proceedings of the Third Annual Workshop on Computational Learning Theory, COLT 1990, University of Rochester, Rochester, NY, USA, August 6-8, 1990*, pages 202–216. Morgan Kaufmann, 1990.
- [21] Yoav Freund and Robert E. Schapire. Experiments with a new boosting algorithm. In Lorenza Saitta, editor, *Machine Learning, Proceedings of the Thirteenth International Conference (ICML ’96), Bari, Italy, July 3-6, 1996*, pages 148–156. Morgan Kaufmann, 1996.
- [22] Nir Friedman, Dan Geiger, and Moises Goldszmidt. Bayesian network classifiers. *Mach. Learn.*, 29(2-3):131–163, November 1997.
- [23] Sudeep Ghosh, Jason Hiser, and Jack W. Davidson. A secure and robust approach to software tamper resistance. In Rainer Böhme, Philip W. L. Fong, and Reihaneh Safavi-Naini, editors, *Information Hiding - 12th International Conference, IH 2010, Calgary, AB, Canada, June 28-30, 2010, Revised Selected Papers*, volume 6387 of *Lecture Notes in Computer Science*, pages 33–47. Springer, 2010.
- [24] Shantanu Godbole and Sunita Sarawagi. Discriminative methods for multi-labeled classification. In Honghua Dai, Ramakrishnan Srikant, and Chengqi Zhang, editors, *Advances in Knowledge Discovery and Data Mining, 8th Pacific-Asia Conference, PAKDD 2004, Sydney, Australia, May 26-28, 2004, Proceedings*, volume 3056 of *Lecture Notes in Computer Science*, pages 22–30. Springer, 2004.
- [25] Isabelle Guyon, Steve Gunn, Masoud Nikravesh, and Lotfi A. Zadeh. *Feature Extraction: Foundations and Applications (Studies in Fuzziness and Soft Computing)*. Springer-Verlag, Berlin, Heidelberg, 2006.
- [26] Trevor Hastie, Robert Tibshirani, and Jerome H. Friedman. *The elements of statistical learning: data mining, inference, and prediction, 2nd Edition*. Springer series in statistics. Springer, 2009.
- [27] Simon Howard. c-algorithms. <https://github.com/fraggle/c-algorithms>. [Online; accessed 30-01-2019].
- [28] Mike James. *Classification Algorithms*. Wiley-Interscience, New York, NY, USA, 1985.
- [29] Pascal Junod, Julien Rinaldini, Johan Wehrli, and Julie Michielin. Obfuscator-LLVM – software protection for the masses. In Brecht Wyseur, editor, *Proceedings of the IEEE/ACM 1st International Workshop on Software Protection, SPRO’15, Firenze, Italy, May 19th, 2015*, pages 3–9. IEEE, 2015.

- [30] Yuichiro Kanzaki, Akito Monden, Masahide Nakamura, and Ken-ichi Matsumoto. Exploiting self-modification mechanism for program protection. In *27th International Computer Software and Applications Conference (COMPSAC 2003): Design and Assessment of Trustworthy Software-Based Systems, 3-6 November 2003, Dallas, TX, USA, Proceedings*, page 170. IEEE Computer Society, 2003.
- [31] Samuel T. King, Peter M. Chen, Yi-Min Wang, Chad Verbowski, Helen J. Wang, and Jacob R. Lorch. Subvirt: Implementing malware with virtual machines. In *2006 IEEE Symposium on Security and Privacy (S&P 2006), 21-24 May 2006, Berkeley, California, USA*, pages 314–327. IEEE Computer Society, 2006.
- [32] Ron Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, IJCAI 95, Montréal Québec, Canada, August 20-25 1995, 2 Volumes*, pages 1137–1145, 1995.
- [33] S. B. Kotsiantis. Supervised machine learning: A review of classification techniques. In *Proceedings of the 2007 Conference on Emerging Artificial Intelligence Applications in Computer Engineering: Real World AI Systems with Applications in eHealth, HCL, Information Retrieval and Pervasive Technologies*, pages 3–24. Amsterdam, The Netherlands, The Netherlands, 2007. IOS Press.
- [34] Sotiris B. Kotsiantis. Supervised machine learning: A review of classification techniques. *Informatica (Slovenia)*, 31(3):249–268, 2007.
- [35] Arun Lakhotia, Davidson R. Boccardo, Anshuman Singh, and Aleardo Manacero Jr. Context-sensitive analysis without calling-context. *Higher-Order and Symbolic Computation*, 23(3):275–313, 2010.
- [36] Timea Lazlo and Akos Kiss. Obfuscating c++ programs via control flow flattening.
- [37] Tao Li and Mitsunori Ogihara. Detecting emotion in music. In *ISMIR 2003, 4th International Conference on Music Information Retrieval, Baltimore, Maryland, USA, October 27-30, 2003, Proceedings*, 2003.
- [38] Cullen Linn and Saumya K. Debray. Obfuscation of executable code to improve resistance to static disassembly. In Sushil Jajodia, Vijayalakshmi Atluri, and Trent Jaeger, editors, *Proceedings of the 10th ACM Conference on Computer and Communications Security, CCS 2003, Washington, DC, USA, October 27-30, 2003*, pages 290–299. ACM, 2003.
- [39] Richard Maclin and David W. Opitz. Popular ensemble methods: An empirical study. *CoRR*, abs/1106.0257, 2011.
- [40] Matias Madou, Bertrand Anckaert, Patrick Moseley, Saumya K. Debray, Bjorn De Sutter, and Koen De Bosschere. Software protection through dynamic code mutation. In JooSeok Song, Taekyoung Kwon, and Moti Yung, editors, *Information Security Applications, 6th International Workshop, WISA 2005, Jeju Island, Korea, August 22-24, 2005, Revised Selected Papers*, volume 3786 of *Lecture Notes in Computer Science*, pages 194–206. Springer, 2005.
- [41] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA, 2008.
- [42] Jiang Ming, Dongpeng Xu, Li Wang, and Dinghao Wu. LOOP: logic-oriented opaque predicate detection in obfuscated binary code. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-6, 2015*, pages 757–768, 2015.
- [43] Akito Monden, Antoine Monsifrot, and Clark D. Thomborson. A framework for obfuscated interpretation. In *ACSW Frontiers 2004, 2004 ACSW Workshops - the Australasian Information Security Workshop (AISW2004), Dunedin, New Zealand, January 2004*, pages 7–16, 2004.
- [44] Mathilde Ollivier, Sébastien Bardin, Richard Bonichon, and Jean-Yves Marion. How to kill symbolic deobfuscation for free; or unleashing the potential of path-oriented protections. In *Proceedings of the 35th Annual Computer Security Applications Conference, ACSAC 2019, San Juan, Puerto Rico, USA, December 9-13, 2019*. ACM, 2019.
- [45] 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.
- [46] Igor V. Popov, Saumya K. Debray, and Gregory R. Andrews. Binary obfuscation using signals. In Niels Provos, editor, *Proceedings of the 16th USENIX Security Symposium, Boston, MA, USA, August 6-10, 2007*. USENIX Association, 2007.
- [47] Mila Dalla Preda, Matias Madou, Koen De Bosschere, and Roberto Giacobazzi. Opaque predicates detection by abstract interpretation. In *Algebraic Methodology and Software Technology, 11th International Conference, AMAST 2006, Kuressaare, Estonia, July 5-8, 2006, Proceedings*, pages 81–95, 2006.
- [48] GNU Project. GNU Core Utilities. <https://www.gnu.org/software/coreutils/>, 2002. [Online; accessed 30-01-2019].
- [49] Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. Classifier chains for multi-label classification. *Machine Learning*, 85(3):333–359, 2011.
- [50] Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. Classifier chains for multi-label classification. *Mach. Learn.*, 85(3):333–359, December 2011.
- [51] Lior Rokach. Ensemble-based classifiers. *Artif. Intell. Rev.*, 33(1-2):1–39, 2010.
- [52] Lior Rokach and Oded Maimon. *Data Mining With Decision Trees: Theory and Applications*. World Scientific Publishing Co., Inc., River Edge, NJ, USA, 2nd edition, 2014.
- [53] Aleiaddin Salem and Sebastian Banescu. Metadata recovery from obfuscated programs using machine learning. In *Proceedings of the 6th Workshop on Software Security, Protection, and Reverse Engineering, SSPREW@ACSAC 2016, Los Angeles, California, USA, December 5-6, 2016*, pages 1:1–1:11, 2016.
- [54] Jonathan Salwan, Sébastien Bardin, and Marie-Laure Potet. Symbolic deobfuscation: From virtualized code back to the original. In *Detection of Intrusions and Malware, and Vulnerability Assessment - 15th International Conference, DIMVA 2018, Saclay, France, June 28-29, 2018, Proceedings*, pages 372–392, 2018.
- [55] Sebastian Schrittwieser and Stefan Katzenbeisser. Code obfuscation against static and dynamic reverse engineering. In Tomás Filler, Tomás Pevný, Scott Craver, and Andrew D. Ker, editors, *Information Hiding - 13th International Conference, IH 2011, Prague, Czech Republic, May 18-20, 2011, Revised Selected Papers*, volume 6958 of *Lecture Notes in Computer Science*, pages 270–284. Springer, 2011.
- [56] Sebastian Schrittwieser, Stefan Katzenbeisser, Johannes Kinder, Georg Merz-downnik, and Edgar R. Weippl. Protecting software through obfuscation: Can it keep pace with progress in code analysis? *ACM Comput. Surv.*, 49(1):4:1–4:37, 2016.
- [57] Padhraic Smyth and David H. Wolpert. Linearly combining density estimators via stacking. *Machine Learning*, 36(1-2):59–83, 1999.
- [58] Xiaoyuan Su, Taghi M. Khoshgoftaar, and Russell Greiner. Making an accurate classifier ensemble by voting on classifications from imputed learning sets. *IJIDS*, 1(3):301–322, 2009.
- [59] Li Sun, Steven Versteeg, Serdar Boztas, and Trevor Yann. Pattern recognition techniques for the classification of malware packers. In Ron Steinfeld and Philip Hawkes, editors, *Information Security and Privacy - 15th Australasian Conference, ACISP 2010, Sydney, Australia, July 5-7, 2010, Proceedings*, volume 6168 of *Lecture Notes in Computer Science*, pages 370–390. Springer, 2010.
- [60] Bjorn De Sutter, Bertrand Anckaert, Jens Gregat, Dominique Chanut, and Koen De Bosschere. Instruction set limitation in support of software diversity. In Pil Joong Lee and Jung Hee Cheon, editors, *Information Security and Cryptology - ICISC 2008, 11th International Conference, Seoul, Korea, December 3-5, 2008, Revised Selected Papers*, volume 5461 of *Lecture Notes in Computer Science*, pages 152–165. Springer, 2008.
- [61] Ramtine Tofighi-Shirazi, Irina Măriuca Asăvoae, Philippe Elbaz-Vincent, and Thanh-Ha Le. Defeating Opaque Predicates Statically through Machine Learning and Binary Analysis. In *3rd International Workshop on Software PROtection*, London, United Kingdom, November 2019.
- [62] Ramtine Tofighi-Shirazi, Maria Christofi, Philippe Elbaz-Vincent, and Thanh Ha Le. DoSE: Deobfuscation based on Semantic Equivalence. In *Proceedings of the 8th Software Security, Protection, and Reverse Engineering Workshop, San Juan, PR, USA, December 3-4, 2018*, pages 1:1–1:12, 2018.
- [63] Grigorios Tsoumakas and Ioannis Katakis. Multi-label classification: An overview. *IJDWM*, 3(3):1–13, 2007.
- [64] Sharath K. Udupa, Saumya K. Debray, and Matias Madou. Deobfuscation: Reverse engineering obfuscated code. In *12th Working Conference on Reverse Engineering, WCRE 2005, Pittsburgh, PA, USA, November 7-11, 2005*, pages 45–54, 2005.
- [65] Xabier Ugarte-Pedrero, Igor Santos, Pablo García Bringas, Mikel Gastesi, and José Miguel Esparza. Semi-supervised learning for packed executable detection. In Pierangela Samarati, Sara Foresti, Jiankun Hu, and Giovanni Livraga, editors, *5th International Conference on Network and System Security, NSS 2011, Milan, Italy, September 6-8, 2011*, pages 342–346. IEEE, 2011.
- [66] Zeljko Vrba, Pål Halvorsen, and Carsten Griwodz. Program obfuscation by strong cryptography. In *ARES 2010, Fifth International Conference on Availability, Reliability and Security, 15-18 February 2010, Krakow, Poland*, pages 242–247. IEEE Computer Society, 2010.
- [67] Chenxi Wang, Jonathan Hill, John C. Knight, and Jack W. Davidson. Protection of software-based survivability mechanisms. In *2001 International Conference on Dependable Systems and Networks (DSN 2001) (formerly: FTCS), 1-4 July 2001, Göteborg, Sweden, Proceedings*, pages 193–202, 2001.
- [68] Henry S. Warren. *Hacker's Delight*. Addison-Wesley Professional, 2nd edition, 2012.
- [69] Dongpeng Xu, Jiang Ming, and Dinghao Wu. Generalized dynamic opaque predicates: A new control flow obfuscation method. In *Information Security - 19th International Conference, ISC 2016, Honolulu, HI, USA, September 3-6, 2016, Proceedings*, pages 323–342, 2016.
- [70] Hui Xu, Yangfan Zhou, Yu Kang, Fengzhi Tu, and Michael R. Lyu. Manufacturing resilient bi-opaque predicates against symbolic execution. In *48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2018, Luxembourg City, Luxembourg, June 25-28, 2018*, pages 666–677, 2018.
- [71] Babak Yadegari, Brian Johannesmeyer, Ben Whitley, and Saumya Debray. A generic approach to automatic deobfuscation of executable code. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 674–691. IEEE Computer Society, 2015.
- [72] Min-Ling Zhang, Yu-Kun Li, Xu-Ying Liu, and Xin Geng. Binary relevance for multi-label learning: An overview. *Front. Comput. Sci.*, 12(2):191–202, April 2018.
- [73] Min-Ling Zhang and Zhi-Hua Zhou. Multilabel neural networks with applications to functional genomics and text categorization. *IEEE Trans. on Knowl. and Data Eng.*, 18(10):1338–1351, October 2006.
- [74] Min-Ling Zhang and Zhi-Hua Zhou. Ml-knn: A lazy learning approach to multi-label learning. *Pattern Recogn.*, 40(7):2038–2048, July 2007.