Using Machine Learning to Generate Test Oracles: A Systematic Literature Review

Downloaded from: <https://research.chalmers.se>, 2023-04-26 10:53 UTC

Citation for the original published paper (version of record):

Fontes, A., Gay, G. (2021). Using Machine Learning to Generate Test Oracles: A Systematic Literature Review. TORACLE 2021 - Proceedings of the 1st International Workshop on Test Oracles, co-located with ESEC/FSE 2021: 1-10. <http://dx.doi.org/10.1145/3472675.3473974>

N.B. When citing this work, cite the original published paper.

Using Machine Learning to Generate Test Oracles: A Systematic Literature Review

Afonso Fontes

afonso.fontes@chalmers.se

Chalmers and the University of Gothenburg
Gothenburg, Sweden

Gregory Gay

greg@greggay.com

Chalmers and the University of Gothenburg
Gothenburg, Sweden

ABSTRACT

Machine learning may enable the automated generation of test oracles. We have characterized emerging research in this area through a systematic literature review examining oracle types, researcher goals, the ML techniques applied, how the generation process was assessed, and the open research challenges in this emerging field.

Based on a sample of 22 relevant studies, we observed that ML algorithms generated test verdict, metamorphic relation, and—most commonly—expected output oracles. Almost all studies employ a supervised or semi-supervised approach, trained on labeled system executions or code metadata—including neural networks, support vector machines, adaptive boosting, and decision trees. Oracles are evaluated using the mutation score, correct classifications, accuracy, and ROC. Work-to-date show great promise, but there are significant open challenges regarding the requirements imposed on training data, the complexity of modeled functions, the ML algorithms employed—and how they are applied—the benchmarks used by researchers, and replicability of the studies. We hope that our findings will serve as a roadmap and inspiration for researchers in this field.

CCS CONCEPTS

• **Software and its engineering** → **Software verification and validation**; • **Computing methodologies** → **Machine learning**.

KEYWORDS

Test Oracle, Automated Test Generation, Automated Test Oracle Generation, Machine Learning

ACM Reference Format:

Afonso Fontes and Gregory Gay. 2021. Using Machine Learning to Generate Test Oracles: A Systematic Literature Review. In *The 1st International Workshop on Test Oracles*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Software testing is invaluable in ensuring the reliability of the software that powers our society [25]. It is also notoriously difficult and expensive, with severe consequences for productivity, the environment, and human life if not conducted properly [11]. New tools and

methodologies are needed to control that cost without reducing the quality of the testing process. Automation has a critical role to play in this effort by controlling testing costs and focusing developer attention on important tasks [2, 27].

Consider test creation, an effort-intensive task that requires the selection of sequences of program input and *oracles* that judge the correctness of the resulting execution [4]. Automated test oracle creation is a topic of particular interest—and has earned the title “the test oracle problem” [4]. In current practice, oracles are often test-specific and require dedicated human effort to create. Advances have been made, but *the test oracle problem remains unsolved*. If oracle creation could be even partially automated, developers’ effort and cost savings could be immense.

Advances in the field of machine learning (ML) have shown that algorithms can match or surpass human performance across many problem domains [7]. Machine learning has been used to advance the state-of-the-art in virtually every field. Automated test generation is no exception. We are interested in understanding and characterizing emerging research around the use of ML to generate or to support the creation of test oracles. Specifically, we are interested in understanding the types of oracles generated, the researchers’ goals using ML, which specific ML techniques were applied, how such techniques were trained and validated, and how the success of the generation process was assessed. We also seek to identify limitations that must be overcome and open research challenges in this emerging field.

To that end, we have performed a systematic literature review. Following a search of relevant databases and a rigorous filtering process, we have gathered a sample of 22 relevant studies. We have examined each study, gathering the data needed to answer our research questions. The findings of this study include:

- ML has been used to generate test verdict (18%), metamorphic relation (27%), and expected output (55%) oracles.
- ML algorithms train predictive models that serve either as a stand-in for an existing test oracle—predicting a test verdict—or as a way to learn information about a function—either the expected output or metamorphic relations—that can be used as part of issuing a verdict.
- Almost all studies (96%) employ a supervised or semi-supervised ML approach, trained on labeled system execution logs or source code metadata and validated based on the accuracy of the trained model.
- 59% of the approaches employed a neural network (NN)—including Backpropagation NNs, Multilayer Perceptrons, RBF NNs, probabilistic NNs, and Deep NNs. 23% of approaches adopted support vector machines. 5% adopted decision trees,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

TORACLE 2021, 23 - 27 August, 2021, Athens, Greece

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>

```

@Test
public void testPrintMessage() {
    String str = "Test_Message";
    TransformCase tCase = new TransformCase(str);
    String upperCaseStr = str.toUpperCase();
    assertEquals(upperCaseStr, tCase.getText());
}

```

Figure 1: Example of a unit test. The `assertEquals` statement is an oracle, comparing the expected and actual output.

and another 5% adopted adaptive boosting. The remaining 5% did not specify a technique.

- Results were most often evaluated using the mutation score (55%), followed by number of correct classifications (18%), classification accuracy (18%), and ROC (5%). One study did not perform evaluation.
- The sampled studies show great promise, but there are still significant limitations and open challenges:
 - Oracle generation is limited by the required quantity, quality, and content of training data. Assembling training data may require significant human effort. Models should be retrained over time.
 - Applied techniques may be insufficient for modeling complex functions with many possible outputs. Varying degrees of output abstraction should be explored. Deep learning and ensemble techniques, as well as hyperparameter tuning, should be explored.
 - Research is limited by overuse of toy examples, the lack of common benchmarks, and the inavailability of code and data. A benchmark should be created for evaluating oracle research, and researchers should be encouraged to provide replication packages and open code.

Our study is the first to summarize this emerging research field. We hope that our findings will serve as a roadmap and inspiration for researchers interested in automated oracle generation.

2 BACKGROUND AND RELATED WORK

Testing and Test Oracles: Before complex software can be trusted, it is important to verify that the code is functioning as intended. Verification is often performed through the process of *testing*—the application of *input* to the system, and analysis of the resulting *output*, to identify visible failures or other unexpected behaviors [25].

During testing, a *test suite* containing one or more *test cases* is applied to the SUT. A test case consists of a *test sequence (or procedure)*—a series of interactions with the SUT—with *test input* applied to some component of the SUT. Input can range from a method call, to an API call, to an action taken within a graphical interface, depending on the granularity of the testing effort. Then, the test case will validate the output of the called components against a set of encoded expectations—the *test oracle*—to determine whether the test passes or fails [25]. An oracle can be a predefined specification—encoded in a form usable by the test case—the output of another program, a past version of the SUT, or a model, or even manual inspection performed by humans. Most commonly, the oracle is formulated as a series of assertions on the values of output and stateful attributes [4].

An example unit test is shown in Figure 1. The test passes a string to the constructor of the `TransformCase` class, then calls

its `getText()` method to transform the string to upper-case. An assertion is used as an oracle to check whether the output is an upper-case version of the provided string.

Machine Learning: ML approaches construct models from observed data—and the structure of that data—to make decisions [3]. Instead of being explicitly programmed with a set of instructions like in traditional software, ML algorithms “learn” from observations using statistical analyses, facilitating the automation of decision-making processes. The learning process begins with the search for patterns in a given dataset and, depending on the algorithm employed, may improve through new interactions over time.

ML approaches largely fall into three categories: supervised, unsupervised, and reinforcement learning [3]. In supervised learning, algorithms use previously labeled “training” data to infer a model that makes predictions about newly encountered data. In contrast to supervised methods, unsupervised algorithms do not make use of previously labeled data. Instead, approaches identify patterns based on the similarities and differences between data items. Rather than labeling items, unsupervised approaches are often used to cluster data and detect anomalies. Reinforcement learning algorithms select actions given their estimation of their ability to achieve some in-built goal, using feedback on the effect of the actions taken to improve their estimation of how to maximize achievement of this goal [28]. Such algorithms are often the basis of automated processes, such as game bots or autonomous driving.

Recent research often focuses on “deep learning” (DL). DL approaches—often supervised—can make complex and highly accurate inferences from massive datasets that would be impossible in traditional ML approaches. This is because DL has an architecture inspired by organic neural networks that attempts to mimic how the human brain works [13] using nonlinear processing layers where one layer’s output serves as the successive layer’s input. Deep learning requires a computationally intense training process and larger quantities of data than traditional supervised ML, but can learn highly accurate models, extract features and relationships from data automatically, and potentially apply models across applications.

Related Work: To date, we are aware of no other systematic literature reviews dedicated to the use of ML to generate test oracles. However, there are secondary studies that cover overlapping topics. Most relevant is the survey on test oracles by Barr et al. [4]. Their survey thoroughly summarizes research on test oracles up to 2014. They divide test oracles into four broad types, including those specified by human testers, those derived automatically from development artifacts, those that reflect implicit properties of all programs, and those that rely on a human-in-the-loop to judge test results. Approaches based on ML belong to the “derived” category, as they learn automatically from project artifacts to replace or augment human-written oracles. They discuss early approaches to using ML to derive oracles.

Durelli et al. performed a systematic mapping study on the application of ML to software testing [8]. Their scope is broader, but they do note that ML has been applied to support test oracle construction. They find that supervised learning is the most-used family of ML techniques overall software testing topics and that Artificial Neural Networks are the most used algorithm.

Our study differs from the above through its focus specifically on the use of ML in oracle generation. This focus allows detailed analysis of this research area that is absent from broader surveys and mapping studies. Our study is also able to reflect more recent research than that covered in older studies.

3 METHODOLOGY

Our concern in this work is to understand how researchers have used machine learning (ML) to perform, or otherwise enhance, automated test oracle generation. We have investigated contributions to the literature related to this topic and seek to understand their methodology, results, and insights. To achieve this task, it is necessary to carry out a secondary study—specifically a Systematic Literature Review (SLR) [19]. This section describes how we conducted our SLR.

We are interested in assessing the *effect* of integrating ML into the oracle generation process, understanding the *adoption* of these techniques—how and why they are being integrated, and which specific techniques are being applied, and identifying the potential *impact* and *risks* of this integration. Table 1 lists the research questions we are interested in answering, briefly defines why those questions are important, and lists the properties extracted from primary studies to answer them (defined in Section 3.3).

The first three questions allow us to understand how ML techniques have enhanced oracle generation, why they were applied, and which specific oracle types were targeted. **RQ2** is motivational, covering the authors’ primary objectives. In contrast, **RQ3** expressly is a technical question, examining the specific roles of the included ML techniques, as well as its training and validation processes.

RQ4 examines which ML techniques were used to perform the generation task, as well as *why* that specific method was adopted, if the authors provide such information. **RQ5** focuses on how the oracle generation approach is evaluated. Finally, **RQ6** aims to cover the limitations of the proposed approaches, open issues, and insights that we have uncovered in this area. To answer these questions, we have done the following:

- (1) Formed a list of studies (Section 3.1).
- (2) Filtered this list for relevance (Section 3.2).
- (3) Extracted data from each study, guided by a set of properties of interest (Section 3.3).
- (4) Identified trends in the extracted data (Section 4).

3.1 Initial Study Selection

To locate studies for consideration, a search was conducted using four databases: IEEE Xplore, ACM Digital Library, Science Direct, and Scopus. We created a search string to narrow the results by combining terms of interest regarding automated test generation and machine learning. Note that our search was purposefully broad, intended to capture studies using ML to enhance both input and oracle generation. This approach allowed us to capture a wide range of studies, including those that a narrow search would miss. We then filtered the pool for relevancy. Each database uses a different search engine, and the search options and search formulation slightly vary between them. In general, the search string used was:

(“test case generation” OR “test generation” OR “test oracle” OR
 “test input”) AND (“machine learning” OR “reinforcement learning”
 OR “deep learning” OR “neural network”)

Figure 2: Steps taken to determine the final list of studies.

These keywords are not guaranteed to capture all existing relevant articles. However, they are designed to capture a sufficiently wide sample to answer our research questions. Specifically, we combine terms related to test case generation—including specific test components—and terms related to machine learning—including common technologies.

Our focus is specifically on the use of ML in oracle generation, not on any form of automated oracle generation. To obtain a representative sample, we have selected ML-related terms that we expect will capture a wide range of studies. These terms may omit some oracle generation techniques that could be in-scope, but allow us to obtain a representative sample while controlling the number of studies that require manual inspection.

Before exporting the results, we applied an initial filter to the results using the advanced search option in each database, which consists of the following selection criteria: (a) published studies in conferences and journals (excluding grey literature such as pre-prints, technical reports, abstracts, editorials, and book chapters); (b) studies published before November 2020 (when we conducted the search); (c) studies written in the English language. After exporting all results, a total of 1936 studies were identified. This is shown as the first step in Figure 2.

To evaluate the search string’s effectiveness, we conducted a three-step verification process. First, we randomly sampled ten entries from the 73 studies that remained following the manual filtering. Then we looked in each article for ten citations that might also be in scope, resulting in a list of 100 citations. We checked whether the search string also retrieved the citations in the list, and all 100 were retrieved by the string (pre-filtering). Although this is a small sample, it indicates the robustness of the search string.

After the search, the next step was to identify whether secondary studies already existed on this topic. If so, the need for this SLR would be reduced. We found no previous secondary studies focusing specifically on ML-based oracle generation. However, we identified a small number of related studies. These are discussed in Section 2.

Table 1: List of research questions, along with motivation for answering the question.

ID	Research Question	Objective
RQ1	Which oracle types have been generated using ML?	Highlights test oracle types (e.g., information used to issue verdicts) targeted for ML-enhanced oracle generation.
RQ2	What is the goal of using machine learning as part of oracle generation?	To understand the reasons for applying ML techniques to perform or enhance oracle generation (e.g., potential benefits, expected result).
RQ3	How was machine learning integrated into the process of oracle generation?	Identifies how the ML technique was applied as part of the process of oracle generation, and specify its training and validation steps.
RQ4	Which ML techniques were used to perform or enhance oracle generation?	Identify specific ML techniques used in the process, including type, learning method, and selection mechanisms.
RQ5	How is the oracle generation process evaluated?	Describe the evaluation of the oracle generation process, highlighting any artifacts (programs or datasets) they relied on.
RQ6	What are limitations and open challenges in ML-based oracle generation?	Highlights the limitations of oracle generation, such as data dependency, accuracy, or training time, and challenges that must be overcome to apply oracle generation in the field.

Table 2: List of properties used to answer the research questions. For each property, we include a name, the research questions the property is associated with, and a short description.

ID	Property Name	RQ	Description
P1	Test Oracle Type	RQ1, RQ2	The specific type of oracle focused on by the approach. It helps to categorize the studies, enabling comparison between contributions.
P2	Proposed Research	RQ2	A short description of the approach proposed or research performed.
P3	Hypotheses and Results	RQ1, RQ3	Highlights the differences between expectations and conclusions of the proposed approach.
P4	ML Integration	RQ3	Covers how ML techniques have been integrated into the oracle generation process. It is essential to understand what aspects of generation are handled or supported by ML.
P5	ML Technique Applied	RQ4	Name, type, and description of the ML technique used in the study.
P6	Reasons for Using the Technique	RQ4	The reasons stated by the authors for choosing this particular ML technique.
P7	ML Training Process	RQ4	How the approach was trained, including the specific data sets or artifacts used to perform this training. Helps us understand how contributions could be replicated or extended.
P8	External Tools or Libraries Used	RQ4	External tools or libraries used to implement the ML technique.
P9	ML Objective and Validation Process	RQ4, RQ5	The objective of the ML technique (i.e., validation metric), and how it is validated, including data, artifacts, and metrics used (if any).
P10	Oracle Creation Evaluation Process	RQ5	Covers how the ML-enhanced oracle generation process, as a whole, is evaluated (i.e., how successful are the generated oracles at detecting faults or meeting some other testing goal?). Allows understanding of the effects of ML on improving the testing process.
P11	Potential Research Threats	RQ6	Notes on the threats to validity that could impact each study.
P12	Strengths and Limitations	RQ6	Used to understand the general strengths and limitations of enhancing oracle creation with ML.
P13	Future Work	RQ6	Any future extensions proposed by the authors, with a particular focus on those that could overcome identified limitations.

Figure 3: Growth in the number of publications in ML-based oracle generation from 2002-2020.

3.2 Selection Filtering

The initial search resulted in 1,936 studies. It is unlikely that all would be relevant. Therefore, we applied a series of filtering steps to obtain a focused list. Figure 2 presents the filtering process and the number of entries after applying each filter. The tag in the center of box 1 represents the 1,936 studies exported from the search and added to the list. The tags in the other boxes represent the number of entries removed in that particular step. The numbers between boxes show the total number of articles that resulted after applying the previous step. Finally, the highlighted box at the end shows the final number of studies used to answer our research questions.

To ensure relevancy, we used a set of keywords to filter the list. We first searched the title and abstract of each study for the keyword “test”. This step removed 834 articles. We then searched the resulting

list for either “learning” or “neural”—representing the application of machine learning. Every article from IEEE Xplore and Scopus passed these filters. However, the number of articles from the ACM Digital Library and Science Direct was significantly reduced. We merged the filtered lists for both keywords. Some studies contained both keywords in the title or abstract. To remove these, as well as any studies that were returned by multiple databases, we removed all duplicate entries, which resulted in 626 remaining studies. We then removed 22 secondary studies, leaving 604 studies.

We examined the remaining studies manually, removing all not in scope following an inspection of the title and abstract. We removed any studies not related to software test generation or that do not apply ML during the test generation process (i.e., the ML element is related to a particular activity such as test suite reduction). This determination was made by first reading the abstract of the paper, then the introduction, then the full paper, until a clear determination could be made of the relevancy of the study. Both authors independently inspected studies during this step to prevent the accidental removal of relevant studies. In cases of disagreement, the authors discussed and came to a conclusion. This left 73 studies. Finally, we divided these studies into those related to input or oracle generation. This step resulted in a final total of 22 studies related to oracle generation for consideration.

Figure 3 shows the rate of growth in this emerging research area. The first study in our sample was published in 2002 and the most recent in 2020. Interest in this topic is growing with the emergence of new and more powerful ML approaches, with over half of the studies having been published since 2016.

3.3 Data Extraction

To answer the questions listed in Table 1, we have examined each study. We have focused on a set of key properties, identified in Table 2. Each property listed in the table is briefly defined and is associated with the research questions that it will help answer. In many cases, several properties are collectively used to answer a RQ. For example, the answer to RQ2, which aims to cover the goals of using ML as part of the automated test generation process, can be extracted from property P2 in many cases. However, P1 is related because it provides context to the research and the particular type of test oracle may dictate how ML is applied. Each property is important in capturing the essential details of the study and how it contributes to answering our RQs.

In reported experiments, the proposed approach either exceeded or failed to meet the initial hypotheses. This is covered by the third property, P3, which could lead to or be part of the answer for RQ1 and RQ3. The fourth property targets RQ3 and notes how the adopted ML technique is integrated into the testing process. To understand how ML techniques can enhance automated test generation, it is important to understand which techniques are applied as well as the motivation behind adopting a specific technique. These aspects are covered by P5 and P6, which are used to answer RQ4. We also note whether the project analyzed is new or the continuation of prior research as part of collecting data for these properties.

The following three properties focus on understanding the application of ML in the study, including a partial assessment of the potential to replicate the research, by covering core characteristics of the ML technique—the training process (P7), external tools used to implement the technique (P8), and the validation process (P9). P7 focuses on the datasets or other information sources used to train the learning technique. Our primary focus with P8 is to cover how external tools, environments, or ML libraries—such as TensorFlow or Keras—are used to train, build, or execute the ML technique. The combination of properties P7, P8, and P9 will answer RQ4, which examines how the ML technique is trained, validated, and assessed as part of its integration. RQ5 examines how the entire oracle generation process is evaluated. P10 is primarily used to answer this research question. However, P9 may also help answer this question.

Research question RQ6 covers open challenges. Properties P11–P13 contribute to answering this question, including limitations and threats to validity—either disclosed by the authors or inferred from our analysis—and future work.

Data extraction was performed primarily by the first author of this study. However, to ensure the accuracy of the extraction process, the second author performed an independent extraction for a randomly-chosen sample of the studies. We compared our findings, and found that we had near-total agreement.

4 RESULTS AND DISCUSSION

We divide the examination of the results as follows: the types of oracles generated using ML and why ML was applied (RQ1-2, Section 4.1), how ML was applied in the examined studies (RQ3-5, Section 4.2), and the limitations and open challenges in this emerging research field (RQ6, Section 4.3).

Figure 4: The types of oracles generated, and the number of studies where this type of oracle is generated.

4.1 Test Oracle Types and Motivation

Before examining which ML techniques have been integrated into oracle generation, or how they have been integrated, it is first crucial to understand *why* they have been integrated. A test oracle is a broad, high-level concept—simply *some* means to judge the correctness of the system given test input [4]. Therefore, our first two research questions are intended to give an overview of the specific types of oracle that have been the focus of the collected studies (RQ1) and to provide motivation for why ML was applied as part of creating these oracles. Figure 4 shows our results. Broadly, three types of oracles have been generated in the examined studies:

- **Test Verdicts:** The approach directly predicts the final test verdict, given provided input. For example, this type of oracle might directly issue a “pass” or “fail” verdict for the test case.
- **Expected Output:** The approach predicts specific system behavior that should result from applying the provided input [10]. The predicted behavior can vary in its level of abstraction, from a concrete output to a broad *class* of output—generally leaning more towards the abstract, given the challenges of making specific predictions for complex systems.
- **Metamorphic Relations:** A metamorphic relation is a necessary property of a function of the software, relating input to the output produced [14]. For example, a metamorphic relation for a *sine* function is $\sin(x) = \sin(\pi - x)$. Such relations allow us to infer expected results for different input values to a function, and violations of such properties identify potential faults. Approaches in this category attempt to learn metamorphic relations for new systems from provided data.

Of the 22 collected studies, a majority—12 approaches—produce expected output oracles. Six produce metamorphic relations, and four produce direct test verdicts.

The goal of ML is to automate or support a decision process. Given an observation, a ML technique can make a prediction. That prediction can either be the final decision to be made, or it can relate to a piece of information needed to make that decision. Test oracles follow a similar conceptual model. Test oracles consist of two core components—the oracle *information*, or a set of facts used to issue the verdict on the test case, and the oracle *procedure*, the actions taken to issue a verdict based on the embedded information and observations of system behavior [29]. Motivationally, we can see that ML offers a natural means to replace either the oracle information—which typically requires human effort to specify—or the oracle as a whole. Test verdict oracles perform the entire

Table 3: Data on the sampled studies, including the type of ML approach, specific ML technique, training data used, the targeted goal of the ML approach, how the approach was evaluated, and the type of application used in the evaluation.

Ref	Year	Oracle Type	ML Approach	Technique	Training Data	ML Objective	Evaluation Metric	Evaluated On
[5]	2018	Test Verdicts	Supervised	Adaptive Boosting	System Executions	Regression	Mutation Score	Shopping Cart
[12]	2018	Test Verdicts	Supervised	Backpropagation NN	System Executions	Regression	Mutation Score	Embedded Software
[21]	2016	Test Verdicts	Supervised	Multilayer Perceptron	System Executions	Regression	Accuracy	User Creation
[33]	2010	Test Verdicts	Supervised	Backpropagation NN	System Executions	Regression	Mutation Score	Student Registration
[1]	2004	Expected Output	Supervised	Backpropagation NN	System Executions	Regression	Correct Classifications	Triangle Classification
[6]	2016	Expected Output	Supervised	SVM	System Executions	Label Propagation	Mutation Score	Image Processing
[16]	2008	Expected Output	Supervised	Backpropagation NN	System Executions	Regression	Correct Classifications	Triangle Classification
[24]	2019	Expected Output	Supervised	Deep NN	System Executions	Regression	Mutation Score	Mathematical Functions
[31]	2011	Expected Output	Supervised	RBF NN	System Executions	Regression	Correct Classifications	Triangle Classification
[32]	2011	Expected Output	Supervised	Multilayer Perceptron	System Executions	Regression	Mutation Score	Insurance Application
[34]	2012	Expected Output	Supervised	Multilayer Perceptron	System Executions	Regression	Mutation Score	Insurance Application
[35]	2016	Expected Output	Supervised	Backpropagation NN + Cascade	System Executions	Regression	Accuracy	Credit Analysis
[37]	2002	Expected Output	Supervised	Not Specified	System Executions	Regression	Mutation Score	Credit Analysis
[38]	2014	Expected Output	Supervised	Backpropagation NN	System Executions	Regression	Mutation Score	Triangle Classification
[39]	2006	Expected Output	Supervised	Multilayer Perceptron	System Executions	Regression	Mutation Score	Mathematical Functions
[41]	2019	Expected Output	Supervised	Probabilistic NN	System Executions	Regression	Correct Classifications	Prime, Triangle Class
[14]	2018	Metamorphic	Supervised	SVM	Code Features	Label Propagation	Accuracy	Various Functions
[15]	2020	Metamorphic	Reinforcement	Not Specified	System Executions	Discovered Relations	Not Evaluated	Ocean Modeling
[17]	2013	Metamorphic	Supervised	SVM, Decision Trees	Code Features	Regression	Mutation Score	Various Functions
[18]	2016	Metamorphic	Supervised	SVM	Code Features	Regression	Mutation Score	Various Functions
[26]	2019	Metamorphic	Supervised	SVM	Code Features	Label Propagation	ROC	Matrix Calculation
[40]	2017	Metamorphic	Supervised	RBF NN	Code Features	Multi-label Regression	Accuracy	Various Functions

decision process, directly issuing a verdict. The other two oracle types, expected outputs and metamorphic relations, replace human specification of oracle information with a model that predicts that information instead. The procedure can then act on that prediction rather than relying on human-specified facts.

RQ1 (Oracle Types): Machine Learning algorithms have been used to generate test verdict (18%), metamorphic relation (27%), and expected output (55%) oracles.

RQ2 (Goal of ML): ML algorithms train models that serve either as a stand-in for a test oracle or to learn information about a function (e.g., expected output or metamorphic relations) that can be used as part of issuing a verdict.

4.2 Application of Machine Learning

Table 3 summarizes relevant data gathered from the 22 studies where ML was used to generate test oracles. Immediately, we can see that almost all approaches adopted a supervised approach, where a model is trained and used to make predictions about new input. Unsupervised and reinforcement learning (RL) have been used as part of input generation. These approaches may also be applicable as part of oracle generation—e.g., an oracle modeled as a RL agent could make predictions and get feedback on their accuracy, or an unsupervised clustering approach could be used as part of an oracle that detects anomalies. One study did propose the use of RL-like techniques as part of metamorphic relation generation. However, the focus has been firmly on supervised learning.

The sampled studies train oracles using a set of previously-captured and labeled system executions or metadata about source code features. The model is then used to predict the correctness of new behaviors or to predict the type of behavior that will result from applying the input. We will discuss each oracle type in turn.

Test Verdicts: All studies within this category applied a ML technique to associate patterns in the training data with the resulting test verdict (i.e., they trained a model for the purpose of regression). This approach enables the oracles generated to assert whether a test passes or fails without running the SUT.

Makondo et al. [21] utilize a Multilayer Perception (MLP) Neural Network (NN)—a basic NN, often constructed with a single hidden layer. Shahamiri et al. [33] and Gholami et al. [12] utilized Feed-forward Backpropagation (BP) NNs to create their test oracles. In essence, a BP NN “learns” by reducing error rates by tuning the weights in each neuron after computing the error, making the model more stable. Braga et al. [5] use a classifier based on adaptive boosting.

Braga et al. [5] gather usage data from a shopping website by inserting several specific capture components into the site. The data then goes through a preprocessing step and then is finally used for training the ML. Shahamiri et al. [33] focus on a student registration-verifier application that checks whether a students’ records satisfy the minimum requirements for enrollment. Gholami et al. [12] focus on embedded systems in their evaluation. Makondo et al. [21] examined a user creation function. Braga et al. [5], Gholami et al. [12] and Shahamiri et al. [33] evaluate their approaches using the mutation score. They insert synthetic faults, and measure how many of these faults that the generated oracle can detect. Makondo et al. [21] evaluate using the accuracy of the classification model.

Expected Output: More than half of the studies generate expected output oracles. The approaches train on system executions, and then predict the output given a new input. Often, the level of detail of the output generated is constrained or abstracted to a small set of representative values, rather than attempting to predict highly specific output. For example, rather than yielding a specific integer for integer output, the approach might constrain the output to a limited set of representative values (classifications) and predict one of those values. Otherwise, evaluation is limited to code that issues output from an enumerated set of values. A common application is the “triangle classification problem,” also known as TRITYP [1, 31, 31, 38, 41]. The program receives three numbers representing the lengths of a triangle’s sides and outputs a classification of the type of triangle as scalene, isosceles, equilateral, or not a triangle. This is a problem that can prove challenging given its branching behavior. However, it still has a limited set of output possibilities. This makes it a reasonable starting point for oracle generation.

Zhang et al. [41] also model a function that judges whether an integer is prime or not. This is an even more straightforward application—a two-class classification problem. Shahamiri et al. [32, 34] generate oracles for a car insurance application, while Singhal et al. [35] and Vanmali et al. [37] generate oracles for a credit analysis at a bank. Ding et al. [6] generate oracles for an image processing function that classifies a type of cell from image sections. All of these applications produce output from an enumerated set of values, easing the difficulty of generating an oracle.

Ye et al. [39] and Monsefi et al. [24] generate oracles for functions with integer output. Some of the cases they examine have a limited range of produced outputs (e.g., a function that predicts the length of a route). Still, the remaining functions offer some indication that deep learning can model more complex functions or predict more detailed expected output.

Ding et al. [6] used a support vector machine (SVM) to perform label propagation. Label propagation is a semi-supervised learning technique, where a mixture of labeled and unlabeled training data is used to train the model, and the algorithm attempts to propagate labels from the labelled data to similar, unlabeled data. This can reduce the quantity of training data needed.

The other approaches follow a more traditional supervised, regression-based learning process, and generally make use of different NNs. Four of the examined studies adopt a Backpropagation NN [1, 16, 35, 38]. Three other studies employ the Multilayer Perceptron technique [32, 34, 39]. Sangwan et al. uses a Radial Basis Function (RBF) NN [31]. RBF is a specific activation function applied to the inputs of the network. Monsefi et al. [24] adopt a Deep NN, which has more input and output layers than a regular NN, with a fuzzy encoder + decoder. Finally, Zhang et al. adopt a probabilistic NN [41].

In terms of evaluation, five of the studies are focused on the accuracy of the oracle in a set of cases where the ground truth is known—measuring the percentage of correct classifications [1, 16, 31, 41] or the accuracy of the model [35]. The remaining seven used the mutation score as the evaluation metric [6, 24, 32, 34, 37–39].

Metamorphic Relations: Six approaches generate metamorphic relations—necessary properties of a function that explain how a particular input links to its corresponding output [14]. Such relations allow us to infer expected results for different input values to a function, and violations of such properties identify potential faults.

Several of the examined studies build on the initial ideas of Kanewala et al. [17], where they proposed an approach that (a) converts the source code of functions into control-flow graphs, (b) selects source code elements as features for a data set, (c) train a model that can predict whether a feature exhibits a particular metamorphic relation (selected from a pre-compiled list of relations). This requires a set of training data, where features are labeled with a binary classification based on whether or not they exhibit that particular relation. A SVM and Decision Trees are used to train the predictive model. Kanewala et al. extended this work by adding a graph kernel to the process [18]. Hardin et al. adapted this approach to work with a semi-supervised label propagation algorithm [14]. Finally, Zhang et al. [40] experimented with the use of a RBF NN. They extended the approach to a multi-label classification that can handle multiple metamorphic relations at once instead of predicting

one at a time. All four of these studies are evaluated on a variety of functions, from mathematical functions, to data structures, to sorting operations. They were evaluated either using the mutation score or accuracy measurements.

Nair et al. [26] extended this work by demonstrating how data augmentation can be used to enlarge the training dataset by using mutants as the source of the additional training data. They compared the enlarged dataset to the original dataset on a set of 45 matrix calculation functions in terms of the Receiver Operating Characteristic, or the ratio of true positive to false positive classifications.

Hiremath et al. [15] propose an approach for using an ML algorithm to predict metamorphic relations for an ocean modeling application. The approach would post a set of relations, evaluate whether they hold, and attempt to minimize a cost function based on the validity of the set of proposed relations. They do not specify an approach, but this maps to common applications of Reinforcement Learning. They do not evaluate their approach, but plan to develop and evaluate it in future work.

We can answer RQ3-5 as follows:

RQ3 (Integration of ML): Almost all studies (96%)

employed a supervised or semi-supervised approach, trained on labeled system execution logs or source code metadata and validated using the accuracy of the trained model.

RQ4 (ML Techniques): 59% of the approaches employed a NN—including Backpropagation NNs (27%), Multilayer Perceptrons (18%), RBF NN (9%), probabilistic NN (5%), and Deep NN (5%). 23% of approaches adopted support vector machines. One also adopted decision trees (5%), and used adaptive boosting (5%). 5% did not specify a technique.

RQ5 (Evaluation of Approach): Results were most often evaluated using the mutation score (55%), followed by number of correct classifications (18%), classification accuracy (18%), and ROC (5%). One study did not perform evaluation.

4.3 Limitations and Open Challenges

The sampled studies show great promise. They illustrate the potential for solving the oracle problem. However, we have observed multiple limitations and challenges that must be overcome to transition research into use in real-world software development.

Volume, Contents, and Collection of Training Data: Supervised ML approaches, even semi-supervised approaches, require training data to create the predictive model that serves as the test oracle. There are multiple challenges related to the *required volume* of training data, the *required contents* of the training data, and the *human effort* required to produce that training data.

Regardless of the specific type of test oracle, the volume of training data that is needed can be vast. This data is generally attained from labeled system execution logs, which means that the SUT needs to be executed *many* times to gather the information needed for training the model. Approaches based on deep learning could produce highly accurate test oracles, but may require thousands of executions to gather the required training data. Some approaches also must

preprocess the collected data before training. The time required to produce the training data can be high and must be considered.

This is particularly true for expected value oracles. Even if the output is abstracted into a small pool of representative values, predicting one of several values is a more difficult task than a boolean classification, and requires significant training data for *each* of the values that can result to make accurate classifications. In addition, the training data for expected value oracles must come from passing test cases—i.e., the output must be what was expected—or labels must be hand-applied by humans. A small number of cases based on failing output may be acceptable if the algorithm is resilient to noise in the training data, but training on faulty code can easily result in an inaccurate model. This introduces a significant barrier to automating training by, e.g., generating test input and simply recording the output that results.

Oracles that produce a direct test verdict model a simpler classification problem—is the result a pass or a fail? However, the requirements on the contents of the underlying data are significant. Each entry in the dataset must be assigned a verdict in order to train the model. This requires either existing test oracles—reducing the need for a ML-based oracle in the first place—or human labeling of test results. Humans are limited in their ability to serve as an oracle, as judging test results is time-consuming and can be erroneous as tester becomes fatigued [4, 22]. This makes it difficult to produce a significant volume of training data. Further complicating this problem is the fact that training a test verdict oracle requires the training data to contain a large number of *failing test cases*. This implies that faults have already been discovered in the system and, presumably, fixed before the oracle is trained. This also will reduce the potential effectiveness of a ML-based oracle.

Metamorphic relation oracles face a similar dilemma. In many of the proposed approaches, the training data consists of source code features labeled with a classification representing whether a particular type of metamorphic relation holds over that feature. This training data must be hand-labeled by a human tester with knowledge of whether these relations hold or not. This requires significant up-front knowledge and effort to establish the ground truth.

Regardless of the oracle type, generating oracles for complex systems will require ML techniques that can extrapolate from limited training data and that can tolerate noise in the training data. Means of generating synthetic training data, like in the work of Nair et al. [26], demonstrate the potential for *data augmentation* to help in overcoming this limitation.

Retraining and Feedback: After training, models generated by supervised learning techniques have a fixed error rate and do not learn from new mistakes made after training. In other words, if the training data is insufficient or inaccurate, the generated oracle will remain inaccurate as long as it remains in use. The ability to improve the oracle based on additional feedback after training could help account for limitations in the initial training data.

There are two primary means to overcome this limitation—either retraining the model using an enriched training dataset, or adopting a reinforcement learning approach that can adapt its expectations based on attained feedback on the accuracy of its decisions. Both means carry challenges. Retraining requires (a) establishing a schedule for when to train the updated model, and (b), an active effort on

the part of human testers to enrich and curate the training dataset. Enriching this dataset—as well as the use of RL—requires some kind of feedback mechanism to judge the accuracy of the oracle. This is likely to require human feedback on, at least, a subset of the decisions made, reducing the potential cost savings.

RQ6 (Challenges): Oracle generation is limited by the required quantity, quality, and content of training data. Assembling training data may require significant human effort. Models should be retrained over time.

Complexity of Modeled Functionality: Many approaches are demonstrated on highly simplistic functions, with only a few lines of code and a small number of possible outputs. While it is intuitive to *start* with highly simplistic examples to examine the viability of an approach, application of such techniques in the field would require oracle generation for far more complex system functions. If a function is simple, there is likely little need for oracle generation in the first place. It remains to be seen whether generated oracles can predict the output of real-world production code, or even simple code with an unconstrained or lightly constrained output space.

Generation of an expected output oracle that can model any arbitrary function with unconstrained output may be prohibitively difficult for even the most effective ML techniques available today. Some abstraction should be expected. One possibility to consider is a variable level of abstraction—e.g., a training-time decision to cluster the output into an adjustable number of representative values (i.e., the centroid of each cluster). Training could take place over different settings for this parameter, and an acceptable balance between quality and level-of-detail could be explored.

Variety, Complexity, and Tuning of ML Techniques: Many of the proposed approaches—especially the earlier ones—are based on simple neural networks with only a few hidden layers. These techniques have strict limitations in the complexity of the functions they can model, and have been superseded by newer ML techniques. Deep learning techniques, which may utilize a high number of hidden layers, may be key in building models of more complex functions. One approach to date has utilized deep learning [24], and we would expect more to explore these techniques in the coming years. However, deep learning also introduces steep requirements on the training data that may limit its applicability [36].

Almost all of the proposed approaches are based on a single ML technique. An approach explored in other domains is the use of *ensembles* [20]. In such approaches, models are trained on the same data using a variety of techniques. Each model is asked for a prediction, then the final prediction is based on the consensus of the ensemble. Ensembles are often able to reach stable, accurate conclusions in situations where a single model may be inaccurate. Ensembles may be a way to overcome the fragility of current oracle generation approaches.

Many ML techniques have a number of *hyperparameters* that can be tuned (e.g., the learning rate, number of hidden units, or activation function) [23]. Hyperparameter tuning can have a major impact on model accuracy, and can enable significant improvements in the results of even simple ML techniques. The proposed approaches do not explore the impact of hyperparameter tuning on the trained models. This is an oversight that should be corrected in future work.

RQ6 (Challenges): Applied techniques may be insufficient for modeling complex functions with many possible outputs. Varying degrees of output abstraction should be explored. Deep learning and ensemble techniques, as well as hyperparameter tuning, should be explored.

Lack of a Standard Benchmark: The emergence of bug benchmarks (e.g., [9, 30]) has enabled sophisticated analyses and comparison of approaches to automated input generation and program repair. To date, oracle generation has often been evaluated on case examples—often over-simplistic examples—where code or metadata is unavailable. This makes comparison and replication difficult.

The creation of a benchmark for oracle generation research could advance the state-of-the-art in the field, spur new research advances, and enable replication and extension of proposed approaches. Such a benchmark should contain a variety of code examples from multiple domains and of varying levels of complexity, allowing the field to move beyond over-simplistic examples. Code examples should be paired with the metadata needed to support oracle generation. This would include sample test cases and human-created test oracles, at minimum. Such a benchmark could also include sample training data that could be augmented over time by researchers.

Lack of Replication Package or Open Code: A common dilemma in software engineering research is lack of access to the code built by researchers or the data used to draw conclusions. Often, the paper itself is not sufficient to allow replication or application of the technique in a new context. This applies to research in oracle generation as well. Some studies make use of open-source ML frameworks (e.g., scikit-learn). This is positive, in that the tools are trustworthy and available. However, without the authors’ code and data, there may not be enough information to enable replication. Further, these frameworks themselves evolve over time, and the attained results may differ because the underlying ML technique has changed since the original study was published.

New approaches should include a replication package with the source code written by the authors, execution scripts, and the versions of external dependencies that were used at the time that the study was performed. This should also include data used by the authors in their analyses.

RQ6 (Challenges): Research is limited by overuse of simplistic examples, the lack of common benchmarks, and the unavailability of code and data. A benchmark should be created for evaluating oracle research, and researchers should be encouraged to provide replication packages and open code.

5 THREATS TO VALIDITY

External and Internal Validity: Our conclusions are based on the studies sampled. It is possible that we may have omitted important studies or sampled an inadequate number of studies. This can affect internal validity—the evidence we use to make conclusions—and external validity—the generalizability of our findings. SLRs are not required to reflect all studies from a research field. Rather, their selection protocol (search string, inclusion and exclusion criteria) should be sufficient to ensure an adequate sample of the field. We

believe that our selection strategy was appropriate. We tested different search strings, and performed a validation exercise to test the robustness of our string. We have used four databases, covering the majority of relevant software engineering venues. Our final set of studies includes 22 primary studies, which we believe is sufficient to make informed conclusions.

Conclusion Validity: The analyses performed are qualitative, and require inference from the authors. This could introduce bias into our conclusions. For example, subjective judgements are required as part of article selection, data extraction, and coding (e.g., categorizing studies based on the oracle type). To control for bias, protocols were discussed and agreed upon by both authors, and independent verification took place on—at least—a sample of all decisions made by either author.

Construct Validity: We used a set of properties to guide data extraction. These properties may have been incomplete or misleading. However, we have tried to establish properties that were appropriate and directly informed by our research questions. These properties were iteratively refined using a selection of papers, and we believe they have allowed us to thoroughly answer the research questions.

6 CONCLUSIONS

Machine learning has the potential to solve the “*test oracle problem*”—the challenge of automatically generating oracles for a function. We have characterized emerging research in this area through a systematic literature review examining oracle types, researcher goals, the ML techniques applied, how the generation process was assessed, and the open research challenges in this emerging field.

Based on a sample of 22 relevant studies, we observed that ML algorithms have been used to generate test verdict, metamorphic relation, and—most commonly—expected output oracles. The ML algorithms train predictive models that serve either as a stand-in for an existing test oracle—predicting a test verdict—or as a way to learn information about a function—either the expected output or metamorphic relations—that can be used as part of issuing a verdict.

Almost all studies employed a supervised or semi-supervised approach, trained on labeled system executions or source code metadata. Of these approaches, many used some type of neural network—including Backpropagation NNs, Multilayer Perceptrons, RBF NNs, probabilistic NNs, and Deep NNs. Others applied include support vector machines, decision trees, and adaptive boosting. Results were most often evaluated using the mutation score, number of correct classifications, classification accuracy, and ROC.

The studies show great promise, but there are significant open challenges. Generation is limited by the required quantity, quality, and content of training data. Models should be retrained over time. Applied techniques may be insufficient for modeling complex functions with many possible outputs. Varying degrees of output abstraction, deep learning and ensemble techniques, and hyperparameter tuning should be explored. In addition, research is limited by overuse of simplistic examples, lack of common benchmarks, and unavailability of code and data. A robust open benchmark should be created, and researchers should provide replication packages.

ACKNOWLEDGMENTS

This research was supported by Vetenskapsrådet grant 2019-05275.

REFERENCES

- [1] K. K. Aggarwal, Yogesh Singh, A. Kaur, and O. P. Sangwan. 2004. A Neural Net Based Approach to Test Oracle. *SIGSOFT Softw. Eng. Notes* 29, 3 (May 2004), 1–6. <https://doi.org/10.1145/986710.986725>
- [2] M. Moein Almasi, Hadi Hemmati, Gordon Fraser, Andrea Arcuri, and Janis Benefelds. 2017. An Industrial Evaluation of Unit Test Generation: Finding Real Faults in a Financial Application. In *Proceedings of the 39th IEEE/ACM International Conference on Software Engineering (ICSE)—Software Engineering in Practice Track (SEIP)* (Buenos Aires, Argentina) (ICSE 2017). ACM, New York, NY, USA.
- [3] Ethem Alpaydin. 2020. *Introduction to machine learning*. MIT press.
- [4] E.T. Barr, M. Harman, P. McMinn, M. Shahbaz, and Shin Yoo. 2015. The Oracle Problem in Software Testing: A Survey. *IEEE Transactions on Software Engineering* 41, 5 (May 2015), 507–525. <https://doi.org/10.1109/TSE.2014.2372785>
- [5] Ronyérison Braga, Pedro Santos Neto, Ricardo Rabêlo, José Santiago, and Matheus Souza. 2018. A Machine Learning Approach to Generate Test Oracles. In *Proceedings of the XXXII Brazilian Symposium on Software Engineering (Sao Carlos, Brazil) (SBES '18)*. Association for Computing Machinery, New York, NY, USA, 142–151. <https://doi.org/10.1145/3266237.3266273>
- [6] J. Ding and D. Zhang. 2016. A machine learning approach for developing test oracles for testing scientific software. *Proceedings of the International Conference on Software Engineering and Knowledge Engineering, SEKE 2016-January* (2016), 390–395. <https://doi.org/10.18293/SEKE2016-137> cited By 4.
- [7] Vedran Dunjko and Hans J Briegel. 2018. Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Reports on Progress in Physics* 81, 7 (2018), 074001.
- [8] Vinicius HS Durelli, Rafael S Durelli, Simone S Borges, Andre T Endo, Marcelo M Eler, Diego RC Dias, and Marcelo P Guimarães. 2019. Machine learning applied to software testing: A systematic mapping study. *IEEE Transactions on Reliability* 68, 3 (2019), 1189–1212.
- [9] Gregory Gay and René Just. 2020. Defects4J as a Challenge Case for the Search-Based Software Engineering Community. In *Search-Based Software Engineering*, Aldeida Aleti and Annibale Panichella (Eds.). Springer International Publishing, Cham, 255–261.
- [10] G. Gay, M. Staats, M. Whalen, and M. Heimdahl. 2015. Automated Oracle Data Selection Support. *Software Engineering, IEEE Transactions on PP*, 99 (2015), 1–1. <https://doi.org/10.1109/TSE.2015.2436920>
- [11] G. Gay, M. Staats, M. Whalen, and M.P.E. Heimdahl. 2015. The Risks of Coverage-Directed Test Case Generation. *Software Engineering, IEEE Transactions on PP*, 99 (2015). <https://doi.org/10.1109/TSE.2015.2421011>
- [12] F. Gholami, N. Attar, H. Haghighi, M. V. Asl, M. Valueian, and S. Mohamadyari. 2018. A classifier-based test oracle for embedded software. In *2018 Real-Time and Embedded Systems and Technologies (RTEST)*. 104–111. <https://doi.org/10.1109/RTEST.2018.8397165>
- [13] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. 2016. *Deep learning*. Vol. 1. MIT press Cambridge.
- [14] Bonnie Hardin and Upulee Kanewala. 2018. Using Semi-Supervised Learning for Predicting Metamorphic Relations. In *Proceedings of the 3rd International Workshop on Metamorphic Testing* (Gothenburg, Sweden) (MET '18). Association for Computing Machinery, New York, NY, USA, 14–17. <https://doi.org/10.1145/3193977.3193985>
- [15] D. J. Hiremath, M. Claus, W. Hasselbring, and W. Rath. 2020. Automated identification of metamorphic test scenarios for an ocean-modeling application. In *2020 IEEE International Conference on Artificial Intelligence Testing (AITest)*. 62–63. <https://doi.org/10.1109/AITest49225.2020.00016>
- [16] H. Jin, Y. Wang, N. Chen, Z. Gou, and S. Wang. 2008. Artificial Neural Network for Automatic Test Oracles Generation. In *2008 International Conference on Computer Science and Software Engineering*, Vol. 2. 727–730. <https://doi.org/10.1109/CSSE.2008.774>
- [17] U. Kanewala and J.M. Bieman. 2013. Using machine learning techniques to detect metamorphic relations for programs without test oracles. *2013 IEEE 24th International Symposium on Software Reliability Engineering, ISSRE 2013* (2013), 1–10. <https://doi.org/10.1109/ISSRE.2013.6698899> cited By 30.
- [18] U. Kanewala, J.M. Bieman, and A. Ben-Hur. 2016. Predicting metamorphic relations for testing scientific software: A machine learning approach using graph kernels. *Software Testing Verification and Reliability* 26, 3 (2016), 245–269. <https://doi.org/10.1002/stvr.1594> cited By 35.
- [19] B. Kitchenham and S. Charters. 2007. Guidelines for performing Systematic Literature Reviews in Software Engineering.
- [20] Ekrem Kocaguneli, Tim Menzies, and Jacky W. Keung. 2012. On the Value of Ensemble Effort Estimation. *IEEE Transactions on Software Engineering* 38, 6 (2012), 1403–1416. <https://doi.org/10.1109/TSE.2011.111>
- [21] W. Makondo, R. Nallanthighal, I. Mapanga, and P. Kadehu. 2016. Exploratory test oracle using multi-layer perceptron neural network. In *2016 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*. 1166–1171. <https://doi.org/10.1109/ICACCI.2016.7732202>
- [22] Phil McMinn, Mark Stevenson, and Mark Harman. 2010. Reducing Qualitative Human Oracle Costs Associated with Automatically Generated Test Data. In *Proceedings of the First International Workshop on Software Test Output Validation* (Trento, Italy) (STOV '10). ACM, New York, NY, USA, 1–4. <https://doi.org/10.1145/1868048.1868049>
- [23] Leandro L. Minku. 2019. A novel online supervised hyperparameter tuning procedure applied to cross-company software effort estimation. *Empirical Software Engineering* 24, 5 (2019), 3153–3204. <https://doi.org/10.1007/s10664-019-09686-w>
- [24] A.K. Monsefi, B. Zakeri, S. Samsam, and M. Khashehchi. 2019. Performing Software Test Oracle Based on Deep Neural Network with Fuzzy Inference System. *Communications in Computer and Information Science* 891 (2019), 406–417. https://doi.org/10.1007/978-3-030-33495-6_31 cited By 0.
- [25] Kshirasagar Naik and Priyadarshi Tripathy. 2011. *Software testing and quality assurance: theory and practice*. John Wiley & Sons.
- [26] Aravind Nair, Karl Meinke, and Sigrid Eldh. 2019. Leveraging Mutants for Automatic Prediction of Metamorphic Relations Using Machine Learning. In *Proceedings of the 3rd ACM SIGSOFT International Workshop on Machine Learning Techniques for Software Quality Evaluation* (Tallinn, Estonia) (MaL-TeSQuE 2019). Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3340482.3342741>
- [27] Alessandro Orso and Gregg Rothermel. 2014. Software Testing: A Research Travelogue (2000–2014). In *Proceedings of the on Future of Software Engineering* (Hyderabad, India) (FOSE 2014). ACM, New York, NY, USA, 117–132. <https://doi.org/10.1145/2593882.2593885>
- [28] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning, Second Edition An Introduction*. 550 pages.
- [29] D. J. Richardson, S. L. Aha, and T. O'Malley. 1992. Specification-based Test Oracles for Reactive Systems. In *Proc. of the 14th Int'l Conf. on Software Engineering*. Springer, 105–118.
- [30] Ripon K. Saha, Yingjun Lyu, Wing Lam, Hiroaki Yoshida, and Mukul R. Prasad. 2018. Bugs.Jar: A Large-Scale, Diverse Dataset of Real-World Java Bugs. In *Proceedings of the 15th International Conference on Mining Software Repositories* (Gothenburg, Sweden) (MSR '18). Association for Computing Machinery, New York, NY, USA, 10–13. <https://doi.org/10.1145/3196398.3196473>
- [31] Om Prakash Sangwan, Pradeep Kumar Bhatia, and Yogesh Singh. 2011. Radial Basis Function Neural Network Based Approach to Test Oracle. *SIGSOFT Softw. Eng. Notes* 36, 5 (Sept. 2011), 1–5. <https://doi.org/10.1145/2020976.2020992>
- [32] S.R. Shahamiri, W.M.N.W. Kadir, S. Ibrahim, and S.Z.M. Hashim. 2011. An automated framework for software test oracle. *Information and Software Technology* 53, 7 (2011), 774–788. <https://doi.org/10.1016/j.infsof.2011.02.006> cited By 31.
- [33] S.R. Shahamiri, W.M.N. Wan Kadir, and S. Bin Ibrahim. 2010. An automated oracle approach to test decision-making structures. *Proceedings - 2010 3rd IEEE International Conference on Computer Science and Information Technology, ICCSIT 2010* 5 (2010), 30–34. <https://doi.org/10.1109/ICCSIT.2010.5563989> cited By 10.
- [34] S.R. Shahamiri, W.M.N. Wan-Kadir, S. Ibrahim, and S.Z.M. Hashim. 2012. Artificial Neural Networks as multi-networks automated test oracle. *Automated Software Engineering* 19, 3 (2012), 303–334. <https://doi.org/10.1007/s10515-011-0094-z> cited By 23.
- [35] A. Singhal, A. Bansal, and A. Kumar. 2016. An approach to design test oracle for aspect oriented software systems using soft computing approach. *International Journal of Systems Assurance Engineering and Management* 7, 1 (2016), 1–5. <https://doi.org/10.1007/s13198-015-0402-2> cited By 1.
- [36] Luke Taylor and Geoff Nitschke. 2017. Improving Deep Learning using Generic Data Augmentation. arXiv:1708.06020 [cs.LG]
- [37] M. Vanmali, M. Last, and A. Kandel. 2002. Using a neural network in the software testing process. *International Journal of Intelligent Systems* 17, 1 (2002), 45–62. <https://doi.org/10.1002/int.1002> cited By 63.
- [38] Vineeta, A. Singhal, and A. Bansal. 2014. Generation of test oracles using neural network and decision tree model. In *2014 5th International Conference - Confluence The Next Generation Information Technology Summit (Confluence)*. 313–318. <https://doi.org/10.1109/CONFLUENCE.2014.6949311>
- [39] M. Ye, B. Feng, L. Zhu, and Y. Lin. 2006. Neural networks based automated test oracle for software testing. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 4234 LNCS - III (2006), 498–507. <https://www.scopus.com/inward/record.uri?eid=s2.0-33750708220&partnerID=40&md5=858633bdc3cbbdac0195bd004eb4141e> cited By 9.
- [40] P. Zhang, X. Zhou, P. Pelliccione, and H. Leung. 2017. RBF-MLMR: A Multi-Label Metamorphic Relation Prediction Approach Using RBF Neural Network. *IEEE Access* 5 (2017), 21791–21805. <https://doi.org/10.1109/ACCESS.2017.2758790>
- [41] R. Zhang, Y.-W. Wang, and M.-Z. Zhang. 2019. Automatic test oracle based on probabilistic neural networks. *Advances in Intelligent Systems and Computing* 752 (2019), 437–445. https://doi.org/10.1007/978-981-10-8944-2_50 cited By 0.