



Vienna University of Technology
Information Systems Institute
Distributed Systems Group

Identifying Incompatible Implementations of Industry Standard Service Interfaces for Dependable Service-Based Applications

Under Review for Publication in -

C. Inzinger, W. Hummer, B. Satzger, P. Leitner, S. Dustdar
inzinger@infosys.tuwien.ac.at

TUV-1841-2012-1

6/7/12

In this paper we study fault localization techniques for identification of incompatible configurations and implementations in service-based applications (SBAs). We consider SBAs with abstract service interfaces that integrate multiple concrete service implementations from various providers. Practice has shown that standardized interfaces alone do not guarantee compatibility of services originating from different partners. Hence, dynamic runtime instantiations of such SBAs pose a great challenge to reliability and dependability. The aim of this work is to monitor and analyze successful and faulty executions in SBAs, in order to proactively detect incompatible configurations at runtime. Our approach is based on well-established machine learning techniques, and extends state-of-the-art fault localization by explicitly addressing temporary and changing fault conditions. Moreover, the presented fault localization technique works on a per-request basis and is able to take individual service inputs into account. Considering not only the service configuration but also the service input data as a parameter for the fault localization algorithm increases the computational complexity by an order of magnitude. Hence, our extensive performance evaluation is targeted at large-scale SBAs and illustrates the feasibility and decent scalability of the approach.

Keywords: Fault Localization, Dependable Systems, Service-Based Applications

Identifying Incompatible Implementations of Industry Standard Service Interfaces for Dependable Service-Based Applications

Christian Inzinger, Waldemar Hummer, Benjamin Satzger, Philipp Leitner and Schahram Dustdar
Distributed Systems Group, Vienna University of Technology
Argentinierstrasse 8/184-1, A-1040 Vienna, Austria
{lastname}@dsg.tuwien.ac.at
<http://dsg.tuwien.ac.at/>

Abstract—

In this paper we study fault localization techniques for identification of incompatible configurations and implementations in service-based applications (SBAs). We consider SBAs with abstract service interfaces that integrate multiple concrete service implementations from various providers. Practice has shown that standardized interfaces alone do not guarantee compatibility of services originating from different partners. Hence, dynamic runtime instantiations of such SBAs pose a great challenge to reliability and dependability. The aim of this work is to monitor and analyze successful and faulty executions in SBAs, in order to proactively detect incompatible configurations at runtime. Our approach is based on well-established machine learning techniques, and extends state-of-the-art fault localization by explicitly addressing temporary and changing fault conditions. Moreover, the presented fault localization technique works on a per-request basis and is able to take individual service inputs into account. Considering not only the service configuration but also the service input data as a parameter for the fault localization algorithm increases the computational complexity by an order of magnitude. Hence, our extensive performance evaluation is targeted at large-scale SBAs and illustrates the feasibility and decent scalability of the approach.

I. INTRODUCTION

Distributed and mission-critical enterprise applications are becoming more and more reliant on external services, provided by suppliers, customers or other members of service value networks [1] (SVNs). In many industries, the technical interfaces of these services are nowadays governed by industry standards, specified by bodies such as the TM Forum¹ (TMF), the Association for Retail Technology Standards² (ARTS) or the International Air Transport Association³ (IATA). Hence, integration of services provided by different business partners into a single service-based application (SBA) becomes feasible. Additionally, as oftentimes a multitude of potential partners are providing implementations of the same standardized interfaces, SBAs are enabled to dynamically switch providers at runtime, i.e., dynamically select the most suitable

implementation of a given standardized interface based on fluid business requirements.

Unfortunately, practice has shown that standardized interfaces alone do not guarantee compatibility of services originating from different partners. Many industry standards are prone to underspecification, while others simply allow multiple alternative (and incompatible) implementations to co-exist. Additionally, and particularly for younger specifications, not every vendor can be trusted to interpret each standard text in the same way. Consequently, there are practical cases, where SBAs, which should work correctly in the abstract, fail to function because of unexpected incompatibilities of service implementations chosen at runtime. Note that this does not necessarily mean that any single one of the chosen service implementations is faulty in itself – it merely means that two or more chosen service implementations do not work in conjunction (even though both may work perfectly in combination with other services).

In this paper, we present a machine learning driven approach to identify such incompatibilities of industry standard implementations. We analyze runtime event logs emitted by the SBA using decision tree techniques and principal component analysis, with the goal of suggesting combinations of service implementations that should not be used in conjunction. Our approach takes into account not only the actual service implementations themselves, but also the received input and the produced output data of implementations. We discuss our approach based on an example industry standard from the TMF, the Next Generation Operation Systems and Software (NGOSS) [2] standard. Furthermore, we quantify the benefits of our approach based on a numerical evaluation.

The remainder of the paper is structured as follows. In Section II we introduce an illustrative scenario from the telecommunications domain which highlights the characteristics of SBAs as studied in this work. Section III discusses related work in the field of reliable distributed systems and fault localization in SBAs. The core part of the paper is Section IV, where we establish a model for fault localization in SBAs and describe our approach in detail. Section VII

¹<http://www.tmforum.org/browse.aspx>

²<http://www.nrf-arts.org/>

³<http://www.iata.org/Pages/default.aspx>

covers a comprehensive experimental evaluation and discusses strengths and limitations of the approach. Finally, Section VIII concludes the paper and points to future research directions.

II. SCENARIO

The motivation for this paper is based on a scenario from the telecommunications services domain. The enhanced Telecom Operations Map (eTOM)⁴, which forms part of the NGOSS program, is a widely adopted industry standard for implementation of business processes promoted by the TMF. MTOSI [3] is an XML-based technology stack defined for NGOSS, which consists of a set of unified, transport-independent interfaces for network and service management. Our scenario is condensed from the TMF’s Case Study Handbook [4] as well as two eTOM-related IBM publications on practical application of SOA in such systems [5], [6].

A. Service Delivery the eTOM Way

Figure 1 depicts the service delivery process in Business Process Modeling Notation (BPMN). The process consists of six activities (denoted i_1, \dots, i_6). We refer to these activities as *interfaces* or *abstract services*. Each abstract service activity has a set of sub-activities which we denote as *concrete service implementations* (denoted c_1, \dots, c_{14} in the figure). At runtime the process selects and executes one concrete service for each service interface. The data flow between the service interfaces of the scenario process is illustrated in Figure 2.

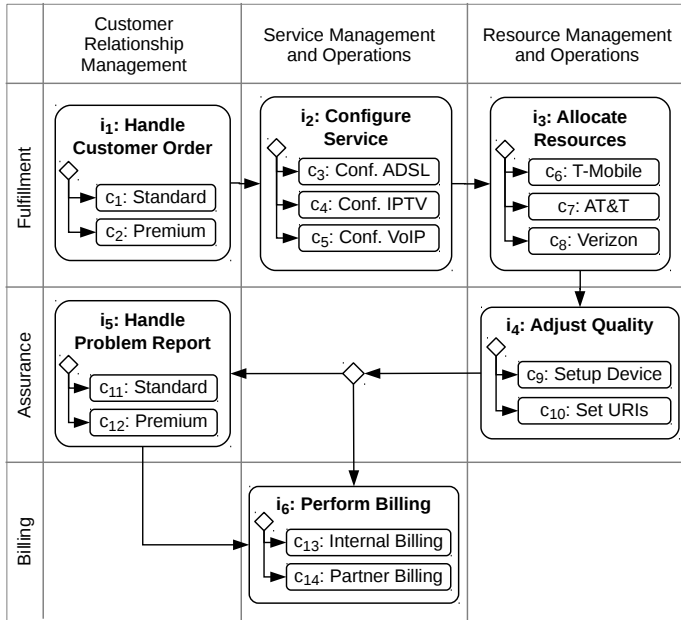


Fig. 1. Service-Based Scenario Application, based on [5] and [6]

The process is initiated by the abstract service i_1 (Handle Customer Order) which is offered in two variants for standard and premium users. Depending on the order input, the process

then configures a particular service (ADSL, IPTV or VoIP). The third abstract service selects one among three partner providers to allocate resources required for service delivery. Telecommunication services are typically associated with Quality of Service (QoS) attributes, which are fine-tuned by abstract service i_4 . For instance, this activity configures parameters in the ADSL device or sets the location URI (Uniform Resource Identifier) of IPTV endpoints, in correspondence with QoS requirements. If a problem is detected at runtime, the optional reporting service is executed in activity i_5 . Finally, the process terminates after storing billing information, either for paying partner providers or for internal accounting if the service was delivered in-house. Besides regular termination, the process may also be interrupted by exceptions at any stage of execution (not depicted in Figure 1). We assume that the information whether the execution has terminated regularly or exceptionally is available for each instance of the process.

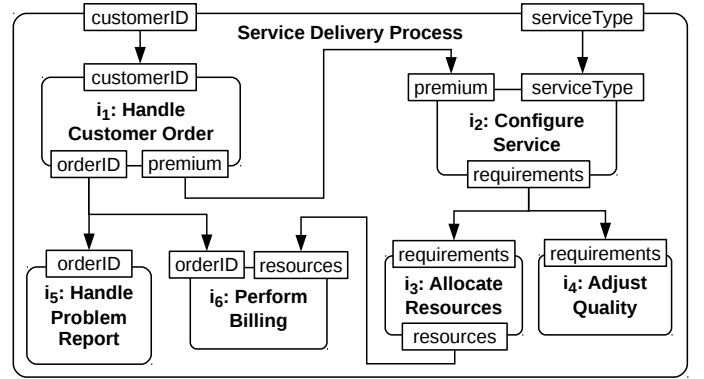


Fig. 2. Data Flow in the Scenario Process

One defining characteristic of eTOM and the presented scenario is process decomposition, which means that business processes are modeled at different levels of abstraction, from the high-level business goals view down to the technical implementation level. In our scenario this is illustrated by the distinction between abstract and concrete services, though in fact the number of abstraction levels can be higher than two.

B. Challenges for Reliable Service Delivery

The scenario outlined in Section II entails challenges to reliability that are typically encountered in service-based applications. Interface standardization (such as MTOSI in the case of our scenario) per se does not guarantee compatibility of services originating from different partners. The interactions among services contain complex dependencies and data flows. The number of variations, i.e., possible instantiations of the process, grows exponentially with the combination of concrete services as well as the provided user input. Hence, comprehensive upfront verification and validation in terms of integration testing is not always feasible and can only cover a certain percentage of the possible instantiations. Therefore, in addition to rigorous testing methods, reliable operation of business-

⁴<http://www.tmfforum.org/BusinessProcessFramework/1647/home.html>

critical SBAs requires proactive monitoring to analyze and avoid incompatible configurations at runtime.

III. RELATED WORK

In this section we discuss existing approaches related to reliability, fault detection, and fault localization in SBAs and distributed systems in general.

A. Software Testing

Our work is related to the broad field of software testing where a plethora of approaches for fault localization have been proposed. Generally, software testing stands for the process of executing a program or systems with the intent of finding errors [7]. Testing approaches are often divided into white- and black-box testing. In white-box (or logic-driven) testing the internals of the software under test are visible to the tester. Black-box (input/output-driven) testing has to get along with no information about internal structure. In our problem formulation, we are facing a black-box model in which we can observe the system behavior but have no details about the internals. In addition to testing the pure correctness of software, important further aspects may be tested, such as performance, reliability, and security. Formal verification of software is an alternative to testing that is often employed in highly safety-critical environments.

Canfora et al. [8] provide an extensive overview of testing services and SBAs. The seminal work by Narayanan and McIlraith [9] was among the first to perform automated simulation and verification based on a semantic model of Web services. Another related approach has been presented by Hummer et al. [10], which performs upfront integration testing with different combinations of concrete service implementations. Due to the huge search space, even in medium sized SBAs, their test case generation approach is not able to consider the service input and output data, whereas the efficient fault localization algorithms used in this work allow us to do so. Concluding, in software testing a system is actively executed to find problems; in this work, however, we do not control the software but we monitor its execution to localize faults and fault reasons at runtime.

B. Software Fault Localization

Software fault localization helps to identify bugs in software on the source code level. Oftentimes a two-phase procedure is applied: 1) finding suspicious code that may contain bugs and 2) examining the code and deciding whether it contains bugs with the goal of fixing them. Research mainly focused on the former, the identification of suspicious code parts with prioritization based on its likelihood of containing bugs [11]. The seminal paper by Hutchins et al. [12] introduces an evaluation environment for fault localization (often referred to as the *Siemens suite*), consisting of seven base programs (in different versions) that have been seeded with faults on the source code level. Renieres et al. [13] present a fault localization technique for identifying suspicious lines of a

program's source code. Based on the existence of a faulty run of the program and many correct runs they select the correct run that is most similar to the faulty one. Proximity is defined based on the program dependence graph. Then, they compare the two runs and produce a report of suspicious program lines. This general functionality is very common in software fault localization. Guo et al. [14] propose a different similarity metric based on control flow. The metric takes into account the sequence of statement instances rather than just the according set. Our work differs from traditional software fault localization in that we do not analyze program code but assume to only be able to observe the external behavior of services. We also assume that the environment or service implementations may change during runtime, in contrast to the analysis of static code.

C. Monitoring and Fault Detection in Distributed Systems

Monitoring and fault detection are key challenges for implementing reliable distributed systems, including SBAs. Fault detectors are a general concept in distributed systems and aim at identify faulty components. In asynchronous systems it is in fact impossible to implement a perfect fault detector [15], because faults cannot be distinguished with certainty from lost or delayed messages. Heartbeat messages can be used for probabilistic detection of faulty components; in this case a monitored component or service has the responsibility to send heartbeats to a remote entity. The fault detector presented in [16] considers the heartbeat inter-arrival times and allows for a computation of a component's faulty behavior probability based on past behavior. Lin et al. [17] describes a middleware architecture called *Llama* that advocates a service bus that users can install on existing service-based infrastructures. It collects and monitors service execution data which enable to incorporate fault detection mechanisms using the data. Such a service bus can be used to collect the data necessary for our analysis. The major body of research in the area of monitoring and fault detection in SBAs deals with topics like SLAs (service-level agreements) and service compositions rather than compatibility issues [18].

D. Fault Analysis and Adaptation in Distributed Systems

Fault analysis derives knowledge from faults that have been experienced. Adaptation tries to leverage this knowledge to reconfigure the system to overcome faults. Zhou et al. [19] have proposed GAUL, an problem analysis technique for unstructured system logs. Their approach is based on enterprise storage systems, whereas we focus on dynamic service-based applications. At its core, GAUL uses a fuzzy match algorithm based on string similarity metrics to associate problem occurrences with log output lines. The aim of GAUL differs from our approach since we assume the existence of structured log files and focus on the localization of faulty configuration parameters. Control of SOAs mostly relies on static approaches, such as predefined policies [20]. Techniques

from artificial intelligence can be used to improve management policies for SBAs during runtime. Markov decision processes, for instance, represent a possible way for modeling the decision-making problems that arise in controlling SBAs. Markov decision processes and algorithms to solve them have been shown effective in reducing the impact of defects in service implementations by adapting the SBA at runtime [21]. In this work we focus on fault localization rather than on how to react in the face of faults.

IV. BASIC FAULT LOCALIZATION APPROACH

This section discusses our novel fault localization technique. In Section IV-A we establish a notion for the model of service-based systems. Sections IV-B and IV-C discuss preprocessing and machine learning techniques used to learn rules which describe the reasons for faults based on the data contained in the model.

A. System Model

We establish a generalized model which forms the basis for the concepts presented in the paper. The core model artifacts are summarized in Table I and briefly discussed in the following. Where applicable, the table also contains examples which refer back to the scenario in Section II.

A SBA consists of a set of industry standard service interfaces I and a set of implementations (C). The mapping between interface and implementation is defined by the function $c : I \rightarrow \mathcal{P}(C)$, where $\mathcal{P}(C)$ denotes the power set of C . P denotes the domain of possible input parameters, each defined by name and value domain. Function p returns all inputs required by an interface. The set F defines data flows as pairs of interfaces (i_x, i_y) , where the output of i_x becomes the input of i_y . Transitive data flows spanning more than two services can be derived from F . Moreover, we define T as the sequence of logged execution traces (in chronological order). Finally, the function r is used to express the result of a trace, i.e., whether the trace represents a successful or failed execution of the SBA.

Symbol	Description
$I = \{i_1, \dots, i_n\}$	Set of industry standard interfaces defined by the SBA. Example: $I = \{i_1, \dots, i_6\}$
$C = \{c_1, \dots, c_m\}$	Set of available concrete implementations to interfaces. Example: $C = \{c_1, \dots, c_{14}\}$
$c : I \rightarrow \mathcal{P}(C)$	Function that returns all concrete candidate implementations for an interface. Example: $c(i_2) = \{i_3, i_4, i_5\}$
$P = [N \times D]$	Domain of service input parameters. Each input parameter is defined by a name (N) and a domain of possible data values (D). Example: $P = \{('premium', \{true, false\}), ('serviceType', String), \dots\}$

$p : I \rightarrow \mathcal{P}(P)$	Function that returns all input parameters for an interface. Example: $p(i_1) = \{('customerID', String)\}$
$F \subseteq I \times I$	Set of direct data flows (dependencies) between two services. Example: $F = \{(i_1, i_2), (i_2, i_3), (i_2, i_4), \dots\}$
$t_x : K \rightarrow V,$ $K = I \cup (I \times N),$ $V = S \cup D,$ $x \in \{1, \dots, k\}$	Log trace representing one execution of the SBA. The function maps from a set of keys (K) to values (V). In particular, interfaces (I) map to implementations (S), and parameter names ($I \times N$) map to parameter values (D). Example: $t_1: a_1 \mapsto c_2, i_2 \mapsto c_3, \dots, (i_1, 'customerID') \mapsto 'joe123', (i_2, 'premium') \mapsto true, \dots$
$T = \langle t_1, \dots, t_k \rangle$	Sequence of logged execution traces.
$r : \{1, \dots, k\} \mapsto \{success, fault\}$	Function that determines for an integer $x \in \{1, \dots, k\}$ whether the execution represented by the trace t_x was successful or has failed.
$E_S \subseteq \mathcal{P}(\mathcal{P}(I \rightarrow C))$	Incompatible assignment. If the implementations in E are used in combination, a fault occurs at runtime. Example: $E_C = \{\{(i_1 \mapsto c_2)\}, \{(i_2 \mapsto c_4), (i_3 \mapsto c_8)\}\}$.
$E_P \subseteq \mathcal{P}(\mathcal{P}(I \rightarrow C) \cup ((I \times N) \mapsto D))$	Incompatible assignment with specific input data. Example: $E_P = \{\{(i_1 \mapsto c_2)\}, \{((i_2, 'premium') \mapsto false), (i_3 \mapsto c_8)\}\}$

TABLE I. Description of Variables

Summarizing the model, the core idea of our approach is to analyze log traces of SBA executions for fault localization. We consider two classes of properties as part of the traces: 1) runtime binding of interfaces to concrete implementations, and 2) service input parameters, i.e., data provided by the user to the application as well as data flowing between services.

B. Trace Data Preparation

Table II lists an excerpt of six exemplary traces for the scenario application. We follow the terminology typically used in machine learning and denote the column titles as *attributes* and the rows starting from the second row as *instances*. The first attribute (t_x) is the instance identifier attribute, the last attribute ($r(x)$) is denoted *class attribute*.

Evidently, the number of attributes and combinations of attribute values can grow very large. To estimate the number of possible traces for a medium sized application, let us consider an imaginary SBA using 10 interfaces ($|I| = 10$), 3 candidate implementations per interface ($|c(i_x)| = 3 \forall i_x \in I$), 3 input parameters per service ($|p(i_x)| = 3 \forall i_x \in I$), and 100 possible data values per parameters ($|d| = 100 \forall i_x \in I, (n, d) \in p(i_x)$). The theoretical number of possible executions in this SBA is $3^{10} * 100^{3^{10}} = 5.9049 * 10^{64}$. Efficient localization of faults in such large problem spaces evidently poses a huge algorithmic challenge. Even more problematically, the problem

t_x	i_1	i_2	i_3	..	$t_x(i_1, 'customerID')$	$t_x(i_2, 'premium')$..	$r(x)$
t_1	c_1	c_3	c_7	..	'joe123'	false	..	success
t_2	c_2	c_4	c_6	..	'aliceXY'	true	..	success
t_3	c_1	c_5	c_8	..	'joe123'	false	..	fault
t_4	c_2	c_5	c_8	..	'bob456'	true	..	success
t_5	c_2	c_4	c_7	..	'aliceXY'	true	..	success
t_6	c_1	c_4	c_8	..	'lindaABC'	false	..	fault
..								

TABLE II
EXAMPLE TRACES FOR SCENARIO APPLICATION

space becomes infinite if the service parameters use non-finite data domains (e.g., *String*).

The first step towards feasible fault analysis is to reduce the problem space to the most relevant information. We propose a two-step approach to achieve this:

- 1) Identifying (ir)relevant attributes: The first manual preprocessing step is to decide, based on domain knowledge about the SBA, which attributes are relevant for fault localization. For instance, in our scenario we can safely say that the unique *orderID* attribute does not have a direct influence on whether the execution succeeds or fails. On the other hand, the parameter *serviceType* can indeed have a direct influence on the result, namely if one of the services ADSL, IPTV, or VoIP is faulty. Per default, all attributes are deemed relevant, but removing part of the attributes from the execution traces helps to reduce the search space.
- 2) Partitioning of data domains: Research on software testing and dependability has shown that faults in programs are often not solely incurred by a single input value, but usually depend on a range of values with common characteristics [22]. Partition testing strategies therefore divide the domain of values into multiple sub-domains and treat all values within a sub-domain as equal. As a simple example, considering that a service has a parameter with type *Integer* (i.e., $\{-2^{31}, \dots, +2^{31} - 1\}$), a valid partitioning would be to treat negative/positive values and zero as separate sub-domains: $\{\{-2^{31}, \dots, -1\}, \{0\}, \{1, \dots, +2^{31} - 1\}\}$. If explicit knowledge about suitable partitioning is available, input value domains can be partitioned manually as part of the preprocessing. However, efficient methods have been proposed to automatize this procedure (e.g., [23]).

C. Learning Rules from Decision Trees

Using the preprocessed trace data, we strive to identify the attribute values or combinations of attribute values that are likely responsible for faults in the application. For this purpose, we utilize decision trees [24], a popular technique in

machine learning. Note that decision trees are usually used for classification, which means to learn rules from a set of training instances with the aim of predicting the class attribute of a new instance. However, our purpose is not classification because in our problem formulation the value of the class attribute (*success* or *fault*) is known for each trace instance; instead, we are interested in learning a decision tree and obtaining the rules which apply to a particular value of the class attribute (*fault*).

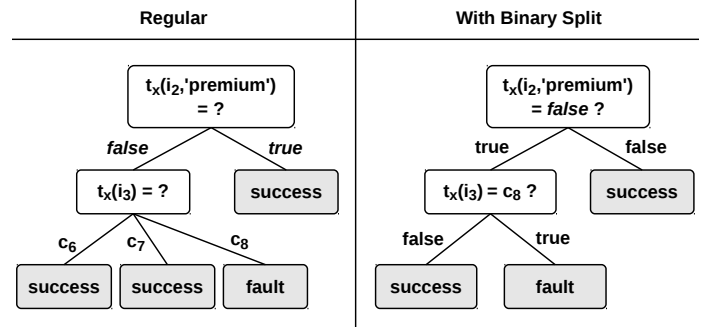


Fig. 3. Exemplary Decision Tree in Two Variants

Figure 3 illustrates decision trees based on our scenario and the example traces in Table II. The figure shows two variants of the same tree which classifies non-premium services from Provider 3 ($t_x(i_3) = c_8$, see Figure 1). The inner nodes are decision nodes which divide the traces search space, and the leaf nodes indicate the trace results. The left-hand side of the figure shows a regular decision tree where each decision node splits according to the possible values of an attribute. The right-hand side shows the same tree with binary split (i.e., each decision node has two outgoing edges).

Algorithm 1 Obtain Incompatibility Rules from Decision Tree

- 1: $E_I \leftarrow \emptyset$
 - 2: **for all** *fault* leaf nodes **as** n **do**
 - 3: $path \leftarrow$ path of nodes from n to root node
 - 4: $E_{temp} \leftarrow \emptyset$
 - 5: **for all** decision node along $path$ **as** d **do**
 - 6: **if** condition of d is *true* along $path$ **then**
 - 7: $E_{temp} \leftarrow E_{temp} \cup$
 - 8: **end if**
 - 9: **end for**
 - 10: $E_I \leftarrow E_I \cup E_{temp}$
 - 11: **end for**
 - 12: **for all** $E_x, E_y \in E_I$ **do**
 - 13: **if** E_x is covered by E_y **then**
 - 14: $E_I \leftarrow E_I \setminus E_x$
 - 15: **end if**
 - 16: **end for**
-

The decision tree with binary split is used to automatically derive incompatible attribute values. Basically, the procedure is to loop over all *fault* leaf nodes and to create a combination

of attribute assignments along the path from the leaf to the root node. The detailed algorithm is presented in Algorithm 1.

V. COPING WITH TEMPORARY AND CHANGING FAULTS

So far, we have shown how trace data can be collected, transformed into a decision tree, and used for obtaining rules which describe which configurations have led to a fault. The assumption so far was that faults are deterministic and static. However, in real-life systems which are influenced by various external factors, we have to be able to cope with temporary and changing faults. Our approach is hence tailored to react to such irregularities in dynamically changing environments.

A temporary fault manifests itself in the log data as a trace $t \in T$ whose result $r(t)$ is supposed to be *success*, but the actual result is $r(t) = \textit{fault}$. Such temporary faults can lead to a situation of contradicting instances in the data set. Two trace instances $t_1, t_2 \in T$ contradict each other if all attributes are equal except for the class attribute:

$$\{(k, v) \mid (k, v) \in t_1\} = \{(k, v) \mid (k, v) \in t_2\}, \\ r(t_1) \neq r(t_2).$$

Fortunately, state-of-the-art decision tree induction algorithms are able to cope with such temporary faults which can be considered as noise in the training data (e.g., [25]).

If the reasons for faults within an SBA change permanently, we need a mechanism to let the machine learning algorithms forget old traces and train new decision trees based on fresh data. Before discussing strategies for maintaining multiple decision trees, we first briefly discuss in Section V-A how the accuracy of an existing classification model is tested over time.

A. Assessing the Accuracy of Decision Trees

Let D be the set of decision trees used for obtaining fault combination rules. We use the function $rc : (D \times \{1, \dots, k\}) \rightarrow \{\textit{success}, \textit{fault}\}$, where k is the highest trace index (cf. Table I), to express how a decision tree classifies a certain trace. Over a subset $T_d \subseteq T$ of the traces classified by a decision tree d , we determine four measures typically used for assessing accuracy in information retrieval and machine learning [26]:

- True Positives: $TP(T_d) = \{t_x \in T_d \mid rc(d, x) = \textit{fault} \wedge r(x) = \textit{fault}\}$
- True Negatives: $TN(T_d) = \{t_x \in T_d \mid rc(d, x) = \textit{success} \wedge r(x) = \textit{success}\}$
- False Positives: $FP(T_d) = \{t_x \in T_d \mid rc(d, x) = \textit{fault} \wedge r(x) = \textit{success}\}$
- False Negatives: $FN(T_d) = \{t_x \in T_d \mid rc(d, x) = \textit{success} \wedge r(x) = \textit{fault}\}$

From the four basic measures we obtain further metrics to assess the quality of a decision tree. The *precision* expresses how many of the traces identified as faults were actually faults ($TP/(TP + FP)$). *Recall* expresses how many of the faults were actually identified as such ($TP/(TP + FN)$). Finally,

the *F1 score* [27] integrates precision and recall into a single value (harmonic mean):

$$F1(d) = 2 * \frac{\textit{precision} * \textit{recall}}{\textit{precision} + \textit{recall}}$$

B. Maintaining a Pool of Decision Trees

In the following we discuss our approach to cope with changing fault conditions over time, based on a sample execution of the scenario application introduced in Section II. Figure 4 illustrates a representative sequence of execution traces ($\{t_1, t_2, t_3, \dots\}$); time progresses from the left-hand side to the right-hand side of the figure. In the top of the figure the trace results ($r(t_x)$) are printed, where ‘‘S’’ represents *success* and ‘‘F’’ represents *fault*. As the traces arrive with progressing time we utilize deduction algorithms to learn decision trees from the data. At time point 1, the decision tree d_1 is initialized and starts the training phase. The learning algorithm has an initial training phase which is required to collect a sufficient amount of data to generate rules that pass the required statistical confidence level. After the initial training phase the quality of the decision tree rules is assessed by classifying new incoming traces. In Figure 4 correct classifications are printed in normal text, while incorrect classifications are printed in bold underlined font.

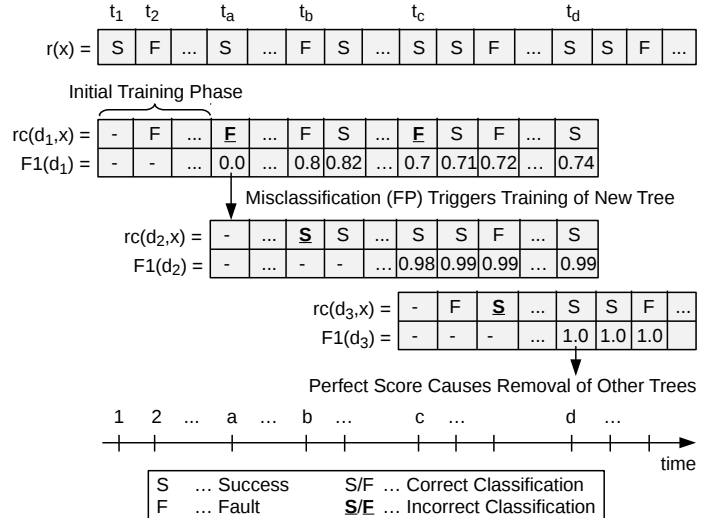


Fig. 4. Maintaining Multiple Trees to Cope with Changing Faults

We have marked four particularly interesting time points (a, b, c, d) in Figure 4, which we discuss in the following.

- In time point a the tree d_1 misclassifies the trace t_a as a false positive. This misclassification triggers the parallel training of a new decision tree d_2 based on the traces starting with t_a .
- A false negative misclassification by d_2 happens in time point b . However, since this happens during the initial training phase of d_2 , we simply regard the trace t_b as useful information for the learner and add it to the training set. No further action is required.

- Time point c contains another false positive misclassification of d_1 . In the meantime, $F1(d_1)$ had risen due to some correct classifications, but now the score is pushed down to 0.7. Again, as in time point a , the generation of a new tree d_3 is triggered.
- At time d the changing environment seems to have stabilized and decision tree d_3 reached a state with perfect classification ($F1(d_3) = 1$). At this point, the remaining decision trees are rejected. The old trees are still stored for reference, but are not trained with further data to save computing power.

VI. IMPLEMENTATION

Our prototype implementation of the presented fault localization approach is implemented in Java. We utilize the open-source machine learning framework *Weka*⁵. Weka contains an implementation of the popular *C4.5* decision tree deduction algorithm [28], denoted *J48 classifier* in Weka. *C4.5* has been applied successfully in many application areas and is known for its good performance characteristics.

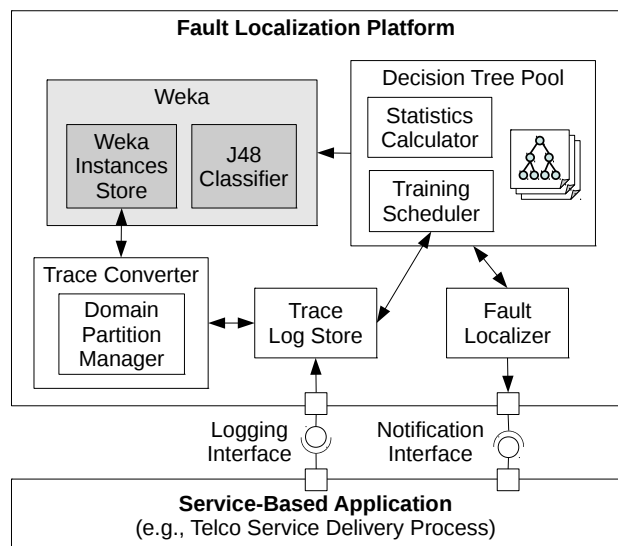


Fig. 5. Prototype Implementation Architecture

Figure 5 outlines the architecture of the Fault Localization Platform with the core components. Third-party components (Weka) are depicted with light grey background color. The service-based application submits its log traces (service bindings plus input messages) to the Logging Interface and provides a Notification Interface to receive fault localization updates. The Trace Log Store receives trace data and forwards them to the Trace Converter. The Domain Partition Manager maintains the customizable value partitions for input messages. For instance, if a trace contains an integer input parameter $x = -173$ and the chosen domain partition for x is $\{negative, zero, positive\}$ then the Trace Converter

⁵<http://www.cs.waikato.ac.nz/ml/weka/>

transforms the input to $x = negative$. The transformed traces are put to the Weka Instances Store. The Decision Tree Pool utilizes the Weka J48 Classifier to maintain the set of trees. The Statistics Calculator determines quality measures for the learned classifiers, and the Training Scheduler triggers the adaptation of the tree pool to changing environments.

VII. EVALUATION

In the following we evaluate different aspects of our proposed fault localization approach. We have set up a comprehensive evaluation framework as part of *Indenica*⁶, a research project aiming at developing a virtual platform for service computing. The framework generates realistic traces of large service compositions, against which we run our fault detection algorithms.

A. Evaluation Setup

The test composition traces are generated randomly, with assumed uniform distribution of the underlying random generator. Table III shows six different SBA instances with corresponding parameter settings which are considered for evaluation. The table also lists for each setting the probability that a fault occurs in a random execution of the SBA.

ID	$ I ,$ $i \in I$	$ c(i) ,$ $i \in I$	$ p(i) ,$ $(n, d) \in p(i)$	$ d ,$ $i \in I$	$\{ e , e \in E_I\}$	Fault Probability
S_1	5	5	10	20	$\{1\}$	$4 * 10^{-2}$
S_2	5	5	10	20	$\{2\}$	$2 * 10^{-3}$
S_3	5	5	10	20	$\{3\}$	$1 * 10^{-4}$
S_4	5	5	10	20	$\{3, 3, 3\}$	$3 * 10^{-4}$
S_5	10	10	10	100	$\{3, 4\}$	$1.001 * 10^{-6}$
S_6	10	10	10	100	$\{4\}$	$1 * 10^{-12}$

TABLE III
FAULT PROBABILITIES FOR EXEMPLARY SBA MODEL SIZES

All tests have been performed on a machine with two Intel Xeon E5620 quad-core CPUs, 32 GB RAM, and running Ubuntu Linux 11.10 with kernel version 3.0.0-16.

B. Basic Properties of the Presented Approach

1) *Trace Limits*: First, we evaluate how many fault traces are required by the J48 classifier to pass the threshold for reliable fault detection. The scenario SBAs S_1, S_2, S_3 (cf. Table III) were used in Figure 6, 20 iterations of the test were executed, and the figure contains three boxes representing the range of minimum and maximum values. As shown in Figure 6, the number of traces required to successfully detect a faulty configuration depends mostly on the complexity (i.e., probability) of the fault with regard to the total scenario size.

A single fault configuration in the configuration S_1 was on average detected after observing between 90 and 190 traces.

⁶<http://www.indenica.eu/>




Fig. 6. Number of Traces Required to Detect Faults of Different Probabilities

If we multiply these values with the fault probability of $4 * 10^{-2}$, we get a range of 4 to 8 fault traces required for the localization. Also with more complex (and hence unlikely) faults the relative figures do not appear to change considerably. With a fault probability of $2 * 10^{-3}$ and $1 * 10^{-4}$ the faults are detected after observing 3/16 and 4/7 minimum/maximum fault traces, respectively. The data suggest that there is a strong relationship between the number of required fault traces and the fault probability.

2) *Noise Resilience*: As discussed in Section V, we anticipate the existence of temporary faults in the system. Temporary faults create noise in the trace logs. Therefore, we evaluate the performance of our approach using different noise levels. In Figure 7 we analyze how the F1 score develops with increasing noise ratio. The figure contains four lines, one each for the scenario settings $S1 - S4$. To ensure that the algorithm actually obtained enough traces for fault localization (see limits in Section VII-B1), we executed the localization run after 200000 observed traces.




Fig. 7. Noise resilience of the fault localization mechanism. Our approach is able to maintain reasonable accuracy in the presence of noisy data.

The results look quite promising and our fault localization mechanism is able to maintain reasonable accuracy even in very noisy environments, due to the high noise resilience of the underlying C4.5 decision tree algorithm. Note that starting

from noise ratio 0.01 there are more spurious execution traces than actual errors observable. Nevertheless, acceptable localization accuracy is maintained, except for $p = 0.0001$, where the environment noise affects ten times more traces than the fault.

C. Changing Fault Conditions

As discussed in Section II, our fault localization approach is able to cope with changing environments, which we show in the following. Figure 8 shows the performance of our approach in the presence of changing faults. The evaluation is set up as follows: Initially a fault combination $FC1$ is active. At trace 33000, the implementation that causes the fault $FC1$ is repaired, but the fix introduces a new fault $FC2$ that is fixed at trace 66000. At trace 66000, another fault $FC3$ occurs, and an attempted fix at trace 88000 introduces an additional fault $FC4$, while $FC3$ remains active. At trace 121000, both $FC3$ and $FC4$ are fixed, but two new faults $FC5$ and $FC6$ are introduced to the system. The occurrence probability for each of the fault combinations ($FC1 - FC6$) is set to $2 * 10^{-3}$ (corresponding to scenario setting $S2$ in Table III).

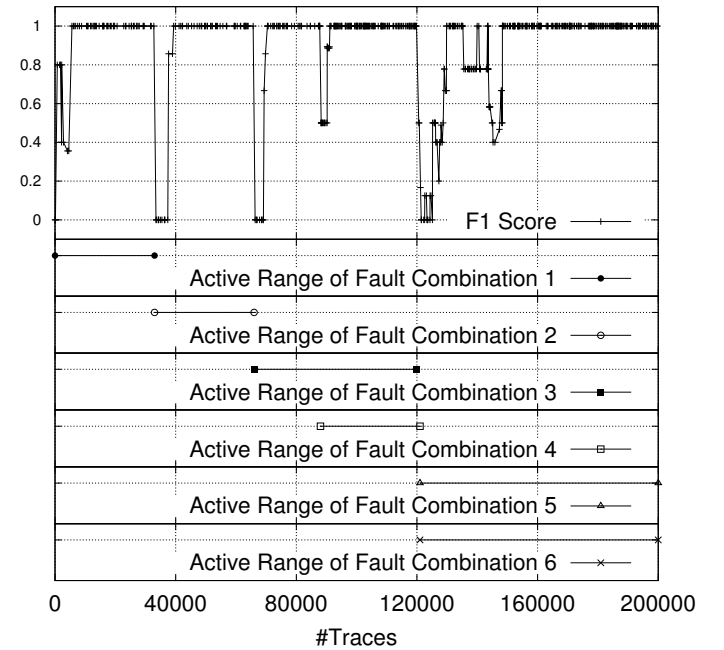


Fig. 8. Fault Localization Accuracy for Dynamic Environment with Changing Faults

This scenario is designed to mimic a realistic situation, but serves mainly to highlight several aspects of our solution. After about 4000 observed execution traces the localizer provides a first guess as to the cause of the fault, but the classification is not yet correct. After around 5200 observed execution traces, the localizer was able to analyze enough error traces to provide an accurate localization result. Note that at that time, only about 6 error traces have been observed, yet the algorithm already produces a correct result. At trace 33000, the previously detected fault $FC1$ disappears and is replaced by $FC2$. Due to

the pool of decision trees maintained by our localizer, $FC2$ can again be accurately localized roughly 6000 traces later. Similarly, after $FC2$ disappears, $FC3$ is localized roughly 5000 traces after its introduction. The decision tree pool allows for the effective localization of new faults introduced to the system at any time. At trace 88000, $FC4$ is introduced, and can again be accurately localized after observing around 5000 traces. $FC3$ and $FC4$ disappear at trace 121000 and are replaced by simultaneously occurring errors $FC5$ and $FC6$. This situation is more challenging for our approach, as seen in the rightmost 80000 traces in Figure 8. The spikes between trace 121000 and 150000 represent different localization attempts that are later invalidated by contradicting execution traces. Finally, however, the localization stabilizes and both faults $FC5$ and $FC6$ are accurately detected.

D. Runtime Considerations

In the following we provide insights into the runtime performance in different configurations and discuss strategies for fine-tuning the performance depending on the target machine.

Due to the nature of the tackled problem, as well as the usage of C4.5 decision trees to generate rules, there are some practical limitations on the number of traces and scenario sizes that can be analyzed using our approach within a reasonable time. Figures 9 and 10 show the time needed to localize faults for various trace window sizes using different exemplary base scenarios on our evaluation machine.




Fig. 9. Localization Time for different trace window sizes in the scenario $S1$ for input sizes $|d| = \{5, 10, 50, 100\}$.

Figure 9 shows the time needed for localization runs with the base scenario $S1$ for input sizes $|d| = \{5, 10, 50, 100\}$. We observe satisfactory computational scaling properties for our approach, showing an approximately linear increase of localization time with rising number of traces analyzed.

Figure 10 shows the required localization time for the base scenario $S5$ for input sizes $|d| = \{5, 10, 50, 100\}$. The figures shown above illustrate that the time needed for a single localization run increases roughly linearly with increasing window sizes. Larger trace windows allow the algorithm to find more complex faults. If fast localization results are needed, the




Fig. 10. Localization Time for different trace window sizes in the scenario $S5$ for input sizes $|d| = \{5, 10, 50, 100\}$.

window size must be kept adequately small, at the cost of the system not being able to localize faults above a certain complexity.

Furthermore, the frequency of localization runs must be considered when implementing our approach in systems with very frequent incoming traces (in the area of hundreds or thousands of traces per second). Evidently, there is a natural limit to the number of traces that can be processed per time unit. Figure 11 shows the localization speed as number of traces processed per second compared to different fault localization intervals (i.e., number of traces after which fault localization is triggered periodically) for different window sizes ($|T|$, i.e., number of considered traces).




Fig. 11. Localization performance in traces per second for different fault localization intervals and window sizes, using scenario $S6$

The data in Figure 11 can be seen as a performance benchmark for the machine(s) on which the fault localization is executed. Executing this test on different machines will result in different performance footprints, which serves as a decision support for configuring window size and localization interval. For instance, if our application produces 1500 traces per second (i.e., processes 1500 requests per second), a localization interval greater than 400 should be used. Currently, the selection happens manually, but as part of our future work we investigate means to fine-tune this configuration automatically.

VIII. CONCLUSION

In this paper we describe a fault localization technique that is able to identify which combinations of service bindings and input data cause problems in SBAs. The analysis is based on log traces, which accumulate during runtime of the SBA. A decision tree learning algorithm is employed to construct a tree from which we extract rules, describing which configurations are likely to lead to faults. For providing a fine-grained analysis we do not only consider the service bindings but also data on message level. This allows to find incompatibilities that go beyond “service A has incompatibility issues with service B” leading to rules of the form “service A has incompatibility issues with service B for messages of type C”. Such rules can help to safely use partial functionality of services. We present extensions to our basic approach that help to cope with dynamic environments and changing fault patterns. We have conducted experiments based on a real-world industry scenario of realistic size. The results provide evidence that the employed approach leads to successful fault localization for dynamically changing conditions, and is able to cope with the large amounts of data that accumulate by considering fine-grained data on message level.

As future work we plan to extend our approach beyond the pure fault localization aspects; in particular, we will use the extracted rules for guiding automated reconfiguration when a fault occurs. Furthermore, we intend to integrate test coverage mechanisms that help to actively investigate faults. This can be used for systematic test execution of insightful configurations and input requests which further narrow down the search space of possible fault reasons.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Commission’s Seventh Framework Programme [FP7/2007-2013] under grant agreement 257483 (Indenica).

REFERENCES

- [1] B. Blau, J. Kramer, T. Conte, and C. van Dinther, “Service value networks,” in *Commerce and Enterprise Computing, 2009. CEC '09. IEEE Conference on*, 2009, pp. 194–201.
- [2] M. J. Creaner and J. P. Reilly, *NGOSS Distilled: The Essential Guide to Next Generation Telecoms Management*. TeleManagement Forum, 2005.
- [3] F. Caruso, D. Milham, and S. Orobec, “Emerging industry standard for managing next generation transport networks: TMF MTOSI,” in *Network Operations and Management Symposium, 2006. NOMS 2006. 10th IEEE/IFIP*, april 2006, pp. 1–15.
- [4] T. Forum, “Case study handbook,” December 2009.
- [5] M. Fiammante, “Dynamic soa and bpm: From simplified integration to dynamic processes,” *Dynamic SOA and BPM: Best Practices for Business Process Management and SOA Agility*, 2009.
- [6] S. M. Glen and J. Andexer, “A practical application of soa,” <http://www.ibm.com/developerworks/webservices/library/ws-soa-practical/>, October 2007.
- [7] G. J. Myers, C. Sandler, and T. Badgett, *The Art of Software Testing, Third Edition*, 3rd ed. Wiley, 2011.
- [8] G. Canfora and M. Di Penta, “Testing services and service-centric systems: challenges and opportunities,” *IT Professional*, vol. 8, no. 2, pp. 10–17, march-april 2006.
- [9] S. Narayanan and S. A. McIlraith, “Simulation, verification and automated composition of web services,” in *11th International Conference on World Wide Web (WWW)*. ACM, 2002, pp. 77–88.
- [10] W. Hummer, O. Raz, O. Shehory, P. Leitner, and S. Dustdar, “Test coverage of data-centric dynamic compositions in service-based systems,” in *Proceedings of the 2011 Fourth IEEE International Conference on Software Testing, Verification and Validation*, ser. ICST '11. Washington, DC, USA: IEEE Computer Society, 2011, pp. 40–49. [Online]. Available: <http://dx.doi.org/10.1109/ICST.2011.55>
- [11] W. E. Wong and V. Debroy, “Software fault localization,” *Part of the IEEE Reliability Society 2009 Annual Technology Report.*, 2009.
- [12] M. Hutchins, H. Foster, T. Goradia, and T. Ostrand, “Experiments of the effectiveness of dataflow- and controlflow-based test adequacy criteria,” in *16th International Conference on Software Engineering*, 1994, pp. 191–200.
- [13] M. Renieres and S. Reiss, “Fault localization with nearest neighbor queries,” in *18th IEEE International Conference on Automated Software Engineering*, 2003, pp. 30–39.
- [14] L. Guo, A. Roychoudhury, and T. Wang, “Accurately choosing execution runs for software fault localization,” in *CC*, 2006, pp. 80–95.
- [15] T. D. Chandra and S. Toueg, “Unreliable failure detectors for reliable distributed systems,” *J. ACM*, vol. 43, no. 2, pp. 225–267, 1996.
- [16] B. Satzger, A. Pietzowski, W. Trumler, and T. Ungerer, “A new adaptive accrual failure detector for dependable distributed systems,” in *SAC*, Y. Cho, R. L. Wainwright, H. Haddad, S. Y. Shin, and Y. W. Koo, Eds. ACM, 2007, pp. 551–555.
- [17] K.-J. Lin, M. Panahi, Y. Zhang, J. Zhang, and S.-H. Chang, “Building accountability middleware to support dependable soa,” *Internet Computing, IEEE*, vol. 13, no. 2, pp. 16–25, march-april 2009.
- [18] M. Papazoglou, P. Traverso, S. Dustdar, and F. Leymann, “Service-oriented computing: State of the art and research challenges,” *Computer*, vol. 40, no. 11, pp. 38–45, nov. 2007.
- [19] P. Zhou, B. Gill, W. Belluomini, and A. Wildani, “Gaul: Gestalt analysis of unstructured logs for diagnosing recurring problems in large enterprise storage systems,” in *29th IEEE Symposium on Reliable Distributed Systems (SRDS)*, 2010, pp. 148–159.
- [20] T. Phan, J. Han, J.-G. Schneider, T. Ebringer, and T. Rogers, “A survey of policy-based management approaches for service oriented systems,” in *19th Australian Conference on Software Engineering*, 2008, pp. 392–401.
- [21] C. Inzinger, B. Satzger, W. Hummer, P. Leitner, and S. Dustdar, “Non-intrusive policy optimization for dependable and adaptive service-oriented systems,” in *Proceedings of the 2012 ACM Symposium on Applied Computing (SAC'12)*, Trento, Italy, 2012, p. (to appear).
- [22] E. Weyuker and B. Jeng, “Analyzing partition testing strategies,” *IEEE Transactions on Software Engineering (TSE)*, vol. 17, no. 7, pp. 703–711, 1991.
- [23] M. R. Chmielewski and J. W. Grzymala-Busse, “Global discretization of continuous attributes as preprocessing for machine learning,” *International Journal of Approximate Reasoning*, vol. 15, no. 4, pp. 319–331, 1996.
- [24] J. R. Quinlan, “Induction of decision trees,” *Machine Learning*, vol. 1, pp. 81–106, 1986.
- [25] D. W. Aha, “Tolerating noisy, irrelevant and novel attributes in instance-based learning algorithms,” *Int. J. Man-Mach. Stud.*, vol. 36, no. 2, pp. 267–287, Feb. 1992.
- [26] R. Baeza-Yates and R.-N. Berthier, *Modern information retrieval*. ACM Press, Addison-Wesley, 1999.
- [27] G. Hripesak and A. S. Rothschild, “Technical brief: Agreement, the f-measure, and reliability in information retrieval,” *JAMIA*, vol. 12, no. 3, pp. 296–298, 2005.
- [28] J. R. Quinlan, *C4.5: programs for machine learning*. Morgan Kaufmann Publishers Inc., 1993.