

Diagnostics as a Reasoning Process: From Logic Structure to Software Design

Matteo Cristani¹, Francesco Olivieri², Claudio Tomazzoli¹, Luca Viganò³
and Margherita Zorzi¹

¹Department of Computer Science, University of Verona, Italy

²Data61, CSIRO, Australia

³Department of Informatics, King's College London, UK

Diagnostic tests are used to determine anomalies in complex systems such as organisms or built structures. Once a set of tests is performed, the experts *interpret* their results and make decisions based on them. This process is named *diagnostic reasoning*. In diagnostic reasoning a decision is established by using both rules and general knowledge on the tests and the domain. The artificial intelligence community has focused on devising and automating different methods of diagnosis for medicine and engineering, but, to the best of our knowledge, the decision process in logical terms has not yet been investigated thoroughly. The automation of the diagnostic process would be helpful in a number of contexts, in particular when the number of test sets to make decision is too wide to be dealt with manually. To tackle such challenges, we shall study logical frameworks for diagnostic reasoning, automation methods and their computational properties and technologies implementing these methods. In this paper, we present the formalization of a hybrid reasoning framework $T^{\mathcal{L}}$ that hosts tests and deduction rules on tests, and an algorithm that transforms a $T^{\mathcal{L}}$ theory into defeasible logic, for which an implemented automated deduction technology (called Spindle) exists. We evaluate the methodology by means of a real-world example related to the Open Web Application Security Project requisites. The full diagnostic process is driven from the definition of the issue to the decision.

ACM CCS (2012) Classification: Computing methodologies → Artificial intelligence → Distributed artificial intelligence → Multi-agent systems

Theory of computation → Logic → Modal and temporal logics

Keywords: tests, experiments, hybrid reasoning, labelled logic, temporal logic, defeasible logic, SPINDle engine

1. Introduction

Diagnostic reasoning is the process of evaluating the results of operations (questions or practical actions) in order to establish which specific conditions hold on an individual or, generically, a sample. This class of operations are usually called *tests*. A number of scientific fields exploit test-based knowledge acquisition: computer science, engineering, earth sciences, biology, medicine and many others.

A test commonly reveals a *property*, usually in search of anomalies, with a margin of error, and provides information about causes of the anomaly. As a consequence, by establishing "cause and effect" relationships, tests provide information about possible solutions. Consider medical diagnosis: specific symptoms suggest which tests are to be done, and the results of such tests help the specialists in identifying which disease is currently on and, consequently, which therapies are needed. The steps of this process are *test-driven* and *knowledge-driven* decisions and, therefore, hybrid reasoning processes.

The result of a test is not exact. Tests are prone to errors, for they reveal a property without proving it in a logical way. This is in contrast to what happens in deductive systems, where a reasoning process starts from premises considered true and, through derivation rules, infers consequences. In other words, tests reveal truth on tested conditions in a provisional way.

The diagnostic reasoning processes that we mentioned above are commonly executed on a huge number of data for several specific diagnostic processes, for instance:

1. In information security, when vast system logs and security data are issued and tested.
2. In geology, where the process of testing data for decisions related to anti-earthquake protections are named with the portmanteau word *geognostic* investigations.
3. In medicine, especially in epidemic control population tests, where the number of tests executed and controlled is wide.

In the situations listed above, the decisions to be made are complex and the diagnostic processes per se can be time-consuming. It would be therefore worth to assist the specialist (information engineer, earth scientist, medical scientist, *etc.*) in the process with a computer-assisted diagnostic reasoning system.

In particular, when a diagnostic test is performed, there are numerous configurations of the results that may be complex to treat simultaneously, for a medical doctor or another specialist, and basic automated learning techniques, such as data mining ones, cannot be used for this purpose satisfactorily. Being able to employ reasoning techniques for computer-assisted diagnosis has been one of the main goals of the research efforts on AI in medicine (see [19, 29, 34, 6] for references and specific approaches). In these cases, it would be useful to assist the medical doctor or another specialist in decision making by providing an automated tool able to decide about the logical consistency and the logical consequences of a set of test results, on top of general issues of those tests themselves, including statistical behaviors, temporal relationships and revealing capabilities.

To tackle these challenges, we started a research program that comprises the development of a mechanized reasoning technology, able to decide the validity of test sets, the assessment of the technology on real-world cases and the comparison with human behaviors in these cases. In this paper, we move the first steps in the research cycle, providing both the definition of a logical framework for the diagnostic reasoning, the $T^{\mathcal{L}}$ logic, and the development of an

architecture that applies the reasoning process of $T^{\mathcal{L}}$ to a real-world case study in information security.

We use the expressiveness of Labelled Modal Logic [10], [11], [38], [43], with temporal and statistical information added to a basic propositional language. Experiments are modeled in terms of tests viewed as Bayesian classifiers, which reveal one or more properties of a sample.

We define the syntax of formulae and relational rules between labels in $T^{\mathcal{L}}$ and sketch ideas about a full deduction systems la Prawitz, by presenting the deduction rules; however, we do not provide soundness and completeness results as these are beyond the scope of this paper. We propose examples of how $T^{\mathcal{L}}$ works and provide technical issues in the construction of the mentioned experimental technology; we also show how to build an architecture to host the developed mechanization of reasoning.

The remainder of the paper is organized as follows. Section 2 discusses some background of this research. Section 3 reviews relevant related literature. Section 4 introduces the logic $T^{\mathcal{L}}$, defining the basic alphabets, the syntax of formulae, test labels (procedures applied during an experiment), and labelled formulae. In particular, Section 4.2 formalizes central notions of the diagnostic-based reasoning, provides a specific analysis of the non-monotonic aspects of the logic itself and focuses also on the relations between tests. In Section 5, we investigate the structure of an architecture for diagnostic reasoning that we further discuss in a real world setting in Section 6. We conclude with Section 7 by summarizing research, discussing some open problems, and sketching future lines of research.

2. Background

In this section, we briefly recall some basic notions from statistical information retrieval and learning [25], and we discuss concrete examples of diagnostic procedure.

From a mathematical perspective, a test is naturally interpreted as a statistical classifier, *i.e.*, a function f that, fed with an input a , is able to predict a probability distribution over a set of

classes. Oversimplifying, f assigns to a a label y that represents the answer (the classification of a). This classification is not exact and therefore the answer given by the classifier can be wrong. For example, if f encodes the problem "Does x enjoy property P ?", the answer "Yes" to this question classifies a as an element of the set of objects that enjoy P , and this can be described by an assertion such as $f: P(x)$. There is an implicit epistemic meaning of this assertion, corresponding to the ability of f to assert $P(x)$ as happens, for instance, in announcement logics or in Agent Communication Languages, where agents make assertions, or in the pure epistemic interpretation of the classical modal logic K, where agents know (or believe) assertions. Also in those systems, truth of sentences may not be guaranteed by the assertion, belief or knowledge of the sentences. Someone may assert, believe or know something, but this something might actually be false.

A large taxonomy of probabilistic classifiers has been developed. In this paper, we focus on the simplest type of classifiers, called (Naive) Bayes (or Bayesian) classifiers, which exploit some strong statistical assumptions [35]. Bayesian classifiers work well in many complex real-world situations and thus represent the execution of tests in an acceptable way.

Classifiers are prone to error. In this context, errors are described either as false positive results, or false negative results⁴. In the remainder of this paper, we omit the word result(s) whenever it is clear from context; we also speak of true positive and true negative for those answers that coincide with the answers given by a logical formula. Scientific research in this area aims to reduce errors in Bayesian classifiers, obtaining better methods to derive knowledge from experiments. The following example both provides a concrete instance of diagnostic reasoning and permits us to introduce the notions of error, and their taxonomy, based on the relation between properties and the revelation of them.

Example 1. Western-Blot is a technique used in biology to confirm the existence of antibodies against a particular pathogenic factor. This is determined by the application of the test in a manner that can be considered without false negatives. Western-Blot, however, has a number of false positives. In contrast, the Elisa test (or, simply, Elisa) analogously lacks false negatives but it exhibits a larger number of false positives than Western-Blot when applied to the same pathogenic factor.

Usually, the sequence of tests depends upon their cost more than their reliability. For instance, Elisa is a cheaper procedure than Western-Blot, and thus Elisa is typically applied before than Western-Blot.

To illustrate this, assume that Elisa answers positively on a given sample. We cannot conclude with certainty that the pathogenic factor is present in the tested organism, due to the high number of false positives exhibited by Elisa. Thus, we apply the Western-Blot test to confirm the validity of Elisa's result. We now derive a negative answer. Since it is assumed that Western-Blot is without false negatives, we can conclude that the pathogenic factor is not present in the organism, against the evidence provided by Elisa.

Example 1 shows a way of deriving truth from tests that is common in those systems. It is straightforward to see that tests with no false negatives that give a negative answer, as well as tests with no false positives that give a positive answer, are always truthful.

3. Related Work

The notion of assisted diagnosis and the usage of intelligent systems in medicine for diagnostic purposes have been a mainstream research topic in artificial intelligence in medicine. Since the pioneering works of Reiter [34] and Davis [19]

⁴False positives and false negatives are concepts analogous to type I and type II errors in statistical hypothesis testing, where a positive result corresponds to rejecting the null hypothesis and a negative result corresponds to not rejecting the null hypothesis. Roughly speaking, a false positive, commonly called a "false alarm", is a result that indicates that a given condition exists while in fact it does not, whereas a false negative is a test result that indicates that a condition does not hold while in fact it does. In principle, tests can be considered without false negatives when the number of false negative results is irrelevant to the decision process as happens, for instance, for those tests that present 1 case of false negative in 1 million. For the purpose of our logical framework, we can assume that this means that there are no false negatives.

these studies have been focusing on two methods: case-based reasoning (see [29] for several recent references to this approach) and statistical methods applied to reasoning (inspired by the original work of Johnson *et al.* [21]; see [23, 28] for recent investigations). The nature of errors in tests for diagnostic reasoning has been studied to support the idea that a test has some intrinsic probability of revealing the property it has been devised for. Therefore, the majority of these investigations have focused upon the ideas that a test can be erroneous in making decisions and that, based on a potentially erroneous decision, we have a tree of possible decisions that have a degree of validity, depending strictly on the validity of the starting decision. There is a long stream of investigations based upon fuzzy and probabilistic reasoning methods of computer science applied to medicine, started by Pis *et al.* in [32] and followed by many other investigations, notably [7] in rheumatology. There have also been some comparative studies, such as [6] and [31]. These methodological studies have given rise to a series of architectural proposals making use of probabilistic methods (see, *e.g.*, [8] and [24]).

4. Focusing on Experimental Knowledge: The Logic $T^{\mathcal{L}}$

We introduce logic $T^{\mathcal{L}}$ that is devised to perform approximate reasoning on tests. Informally, a (well-formed) formula of $T^{\mathcal{L}}$ represents a property of a sample (or individual) that can be revealed, with a margin of error, by a suitable experiment, built out from a sequence of tests. Information about tests is represented by labels, which are metalinguistic logical objects that "adorn" the pure syntactical level of formulae.

4.1. Syntax of $T^{\mathcal{L}}$

The alphabet of $T^{\mathcal{L}}$ is built out of the variable symbol x , a denumerable set of symbols for constants each denoted by lowercase Latin letters (possibly indexed), and a denumerable set of unary predicates, denoted by capital Latin letters \mathcal{P} , \mathcal{Q} ..., possibly indexed.

Predicates represent properties. When applied to an individual constant, a predicate returns an

element of a given domain. Properties are revealed by tests, which are not included in the syntax of formulae; rather, we introduce tests in the syntax of labels that we give below.

A *ground atomic formula* (ground formula, hereafter) is an atomic formula of the form $\mathcal{P}(c)$, where c is a constant. We write $g\mathcal{F}$ to denote the set of ground formulae.

Formulae in $T^{\mathcal{L}}$ are built from the set of atomic formulae by means of the usual logical connectives: \perp , \neg , \wedge , \rightarrow .

Formally, the set $a\mathcal{F}$ of well-formed assertion formulae is the smallest set such that:

- (i) $g\mathcal{F} \subseteq a\mathcal{F}$,
- (ii) $\perp \in a\mathcal{F}$,
- (iii) if $A \in a\mathcal{F}$ then $\neg A \in a\mathcal{F}$,
- (iv) if $A, B \in a\mathcal{F}$ then $(A \wedge B) \in a\mathcal{F}$, and
- (v) if $A, B \in a\mathcal{F}$ then $(A \rightarrow B) \in a\mathcal{F}$.

We denote well-formed assertion formulae by \mathcal{A} , \mathcal{B} , \mathcal{C} ..., possibly indexed, and call them formulae or assertions for short.

Basic literals are formed by letters or negations of letters, applied to constants, *i.e.*, $\mathcal{P}(c)$ and $\neg\mathcal{P}(c)$. For example, if *Fever* is a predicate and *John* is a constant, then *Fever(John)* is a literal.

Following the tradition of *labelled deduction systems* [38, 26, 27], we extend the syntax above by introducing a class of labels that represent experiments, *i.e.*, instants of time in which tests of properties are performed on a sample, under some environmental conditions. *Labels* are built from a set \mathcal{R} of symbols for tests denoted by variables and possibly indexed. Tests in label symbols carry information about the *execution time* (the instant in which the test is performed) and the *experimental condition* (condition, for short), which is the history of actions performed during the experiment and (possibly) additional information provided/known during the diagnostic process. This reflects the fact that a particular test can be conditioned by a specific situation (like the environment, a medical condition, *etc.*). For instance, when a geologist conducts a forecast of the position of underground water, among other examinations there is an extraction of a vertical cylinder of ground: if the terrain is very humid, then the stratification of the underground can be

different than usual, leading to a change in the forecast itself.

To formalize these ideas, we introduce the set \mathcal{T} of symbols for time instants t , possibly indexed, and the set \mathcal{A} of experimental conditions denoted by φ , possibly indexed. In this paper, which provides a first investigation, we define \mathcal{A} simply as the set that contains finite compositional sequences of tests $\tau_1 \dots \tau_k$, where we assume that $\tau_{i+1} \in \mathcal{R}$ has been applied after $\tau_i \in \mathcal{R}$ on the same sample. Clearly, we can have $\varphi = \emptyset$.

Given a fixed denumerable set $Lab_{\mathcal{T}}$ of labels of the form $\tau^{(t;\varphi)}$, where τ is a test able to reveal one or more properties, t represents a time instant (of a given timeline) and φ is the experimental condition. Labels are denoted by l and r , possibly indexed. A test label is a construct that is more expressive than a test symbol: a test label represents a test put into a context, *i.e.*, equipped with additional information such as its time (when it is applied) and the history of the experiment, *i.e.*, the trace of previously applied tests (in the same experiment).

In this paper, we focus on diagnostic reasoning about ground formulae and leave the extension to propositional or first-order logic to future work (see Section 7).

We define labelled formulae as follows. A labeled (well-formed) formula is a formula of the form $\tau^{(t;\varphi)}: A$, where $A \in g\mathcal{F}$.

Intuitively, $\tau^{(t;\varphi)}: \mathcal{P}(c)$ denotes the assertion " τ reveals \mathcal{P} at time t on the sample c , under conditions φ ". For instance, we can write $Elisa^{(Monday;Fever)}: Ebola(John)$ to express that we execute the Elisa test on a sample on Monday, with the patient John having a Fever, to reveal the existence of an infection of Ebola.

Ground facts are ground formulae without labels. We need to introduce one epistemic negation to denote the fact that a formula is not revealed by a test, which is conceptually different than stating that a test reveals the negation of a formula. We thus introduce the negation $?$ that ranges over labelled formulae, in contrast to the logical connective \neg that we already introduced above. Note that neither $\tau^{(t;\varphi)}: A$ implies $\tau^{(t;\varphi)}: \neg A$, nor $\tau^{(t;\varphi)}: \neg A$ implies $\tau^{(t;\varphi)}: A$.

4.2. Orders and Relation for Tests and Observable Properties

We now discuss the mechanization of experimental reasoning and how to provide a logical foundation of test-based knowledge. In this paper, we mainly focus on test labels and on the reasoning processes performed during a procedure that aims at extracting experimental knowledge from some resources (typically, a sample). We can define a partial order between two test labels, both related to temporal information and statistical measures for test performances. We start by defining some temporal orders between labels. We write $t_1 < t_2$ to denote the usual temporal order between time instants, and $\varphi_1 \triangleleft \varphi_2$ to denote the order between conditions. We state that $\varphi_1 \triangleleft \varphi_2$ indicates that φ_1 is a prefix of φ_2 . Following the tradition of labelled deduction systems [26], [27], [38], we define relational formulae by lifting the orders to labels.

Definition 1. (Temporal Relational Formulae).

- $\tau_1^{(t_1, \varphi_1)} \ll \tau_2^{(t_2, \varphi_2)}$ iff $t_1 < t_2$ and $\varphi_1 \triangleleft \varphi_2$
- $\tau_1^{(t_1, \varphi_1)} \mapsto \tau_2^{(t_2, \varphi_2)}$ iff $t_1 < t_2$, $\varphi_2 = \varphi_1 \cdot \tau_1$ and there is no \bar{t} such that $t_1 < \bar{t} < t_2$,

where " $\varphi \cdot \tau_1$ " denotes the condition obtained by performing τ_1 after the events described in φ_1 .

Note that we are modeling the notion of temporal composition of tests. In particular, \ll represents a general temporal application sequence, whereas $\tau_1 \mapsto \tau_2$ represents the execution of the test τ_2 immediately after the execution of the test τ_1 . Note also that the above formula requires the introduction of a logic with branching future time (see Section 7).

With a slight abuse of notation, we write $\tau_1 \mapsto \tau_2$ to denote the test obtained by composing τ_1 and τ_2 ; we treat $\tau_1 \mapsto \tau_2$ as a symbol in \mathcal{R} , and we then use it as a label. We now introduce three orders based on test metrics for elements in $Lab_{\mathcal{T}}$.

Definition 2. (Metric-based Relational Formulae). We write

- $(\tau_1^{(t_1, \varphi_1)} >_a \tau_2^{(t_2, \varphi_2)})[A]$ if τ_1 at time t_1 and under condition φ_1 is more accurate in revealing A than τ_2 at time t_2 under condition φ_2 .
- $(\tau_1^{(t_1, \varphi_1)} >_p \tau_2^{(t_2, \varphi_2)})[A]$ if τ_1 at time t_1 and under condition φ_1 is more precise in revealing A than τ_2 at time t_2 under condition φ_2 .

- $(\tau_1^{(t_1, \varphi_1)} >_r \tau_2^{(t_2, \varphi_2)})[A]$ if τ_1 at time t_1 and under condition φ_1 has greater recall in revealing A than τ_2 at time t_2 under condition φ_2 .

The base of empirical reasoning about tests is the deduction of truth on tests that are correct (with no false positives) or complete (with no false negatives). We introduce the modal operator \square_+ to denote the fact that a test has no false positives, and the modal operator \square_- to denote that it has no false negatives. The modal operators \square_+ and \square_- relate to accuracy, precision and recall. We use these terms in the usual meaning they have in machine learning and specifically in the theory of Bayesian classifiers. Accuracy is the probabilistic complement of error rate of a test, precision is the probabilistic complement of negative error rate (namely the probability of the test giving a correct positive answer), and recall is the probabilistic complement of positive error rate (thus the probability of a test giving a correct negative answer). If a test is both correct and complete, then so is the property it reveals. This can be expressed by means of logical rules. For example, when two tests are differently accurate, and both lack false positives, then they are also ordered in the same way by precision. Analogously, when they lack false negatives, they are also ordered in the same way with respect to recall. We formalize these concepts as follows, where MAR stands for Map Accuracy to Recall, and MAP stands for Map Accuracy to Precision:

$$\frac{\tau_1^{(t, \varphi)} <_a \tau_2^{(t, \varphi)} \square_+ \tau_1 \square_+ \tau_2}{\tau_1^{(t, \varphi)} <_p \tau_2^{(t, \varphi)}} \text{MAP}$$

$$\frac{\tau_1^{(t, \varphi)} <_a \tau_2^{(t, \varphi)} \square_- \tau_1 \square_- \tau_2}{\tau_1^{(t, \varphi)} <_r \tau_2^{(t, \varphi)}} \text{MAR}$$

Interference between $(\tau_1^{(t, \varphi)} >_a \tau_2^{(t, \varphi)})[A]$ assertions and $(\tau_1^{(t, \varphi)} >_p \tau_2^{(t, \varphi)})[A]$ or $(\tau_1^{(t, \varphi)} >_r \tau_2^{(t, \varphi)})[A]$ is managed by means of rules like the following one:

$$\frac{(\tau_1^{(t, \varphi)} >_p \tau_2^{(t, \varphi)})[A] (\tau_1^{(t, \varphi)} >_r \tau_2^{(t, \varphi)})[A]}{(\tau_1^{(t, \varphi)} >_a \tau_2^{(t, \varphi)})[A]} P-R$$

This rule can be reproduced, analogously, for the accuracy as related to recall and to precision.

$$\frac{(\tau_1^{(t, \varphi)} >_p \tau_2^{(t, \varphi)})[A] (\tau_1^{(t, \varphi)} >_r \tau_2^{(t, \varphi)})[A]}{(\tau_1^{(t, \varphi)} >_a \tau_2^{(t, \varphi)})[A]} A-R$$

$$\frac{(\tau_1^{(t, \varphi)} >_a \tau_2^{(t, \varphi)})[A] (\tau_1^{(t, \varphi)} <_r \tau_2^{(t, \varphi)})[A]}{(\tau_1^{(t, \varphi)} >_p \tau_2^{(t, \varphi)})[A]} A-P$$

The modal interplay between different metrics is an interesting problem from both the proof theoretical viewpoint and the practical one (related to the software design). We leave to future work the implementation of the interplay between different metrics. In Example 2 we introduce the key idea that we will exploit in the following: the result of a test is measured by the accuracy hypothesis we assume for the test. For instance, when a test is valued 0.8 accurate, we mean that we believe the test result is true in 80% of the cases, whilst we think that the test gives a wrong answer in 20% of the cases.

Example 2. Assume that we execute Elisa (Eli) on sample John (J) to test for HIV. We execute the test on Monday (Mon), under the history of no previous test. The test results are positive. Now, since Elisa has no false negatives but has false positives (and is not particularly accurate), we execute Western-Blot (WB) on Tuesday (Tue) to confirm/refute Elis result. Western-Blot, obviously is executed with the history of Elisa, which does not interfere with it. The test results are negative. Now, since Western-Blot has not false negatives, we conclude that the sample is HIV-free.

$$\frac{Eli^{(Mon, \emptyset)} : HIV(J)WB^{(Tue, Eli)} : \neg HIV(J)Mon < Tue \square_- WB^{(Tue, Eli)} \square_- Eli^{(Mon, \emptyset)}}{\neg HIV(J)}$$

Western-Blot is more accurate than Elisa, so $WB^{(Tue, Eli)} >_a Eli^{(Mon, \emptyset)}$.

The example shows that the best accuracy of a test τ_i with respect to a test τ_j induces a first, in-

tuitive notion of prevalence: if revealed formulae are contradictory, we trust the more reliable experiment. This will become central in Section 3.3, when we move toward defeasible theories.

It is well known that, when using tests for revealing properties employed in empirical sciences, a given test can interfere with the result of other tests. For instance, certain therapeutic tests (such as the attempt at solving a dangerous potential bacterial infection by the prophylaxis with antibiotics) can make the results of other tests unreliable.

We say that test τ_1 obfuscates test τ_2 if performing τ_1 on a sample before τ_2 diminishes τ_2 's ability to reveal a given property.

On the other hand, τ_1 gifts a property on a test τ_2 when its application extends the ability of τ_2 to reveal the property itself.

This reasoning is based on the application of tests in sequence, which is the reason why we have introduced an implicit notion of time. We assume that time is discrete, and that tests are executed at a given instant of time. We introduce a notion of absolute time and associate directly temporal instants to test execution only. Partial obfuscation and partial gift can be intuitively described as follows:

- We say that a test τ_1 (for a property A) *a-obfuscates* (\searrow_a) the test τ_2 of a property B if, when τ_1 is executed before τ_2 , then the accuracy of $\tau_2 : B$ is less than it would have been if the test τ_1 on A was not executed.
- We say that that a test τ_1 (for a property A) *a-gifts* (\nearrow_a) a property B if $\tau_1 : B$, when, contrary to *a-obfuscation*, the accuracy of the test for B increases.

We can similarly define p-obfuscation, p-gift, r-obfuscation and r-gift referring to obfuscation and gift for precision and recall, instead of accuracy. More formally, we can provide the following relation, exploiting metric-based relational formulae (Definition 2).

Definition 3. Obfuscation and Gift

- $(\tau_1^{(t_1, \varphi_1)} \searrow_a \tau_2^{(t_2, \varphi_2)})[B]$ iff $t_1 < t_2$ and $\tau_2^{(t, \varphi)} >_a \tau_2^{(t_2, \varphi_2)}$ for $t < t_1$ or for φ s.t. $\tau_1 \notin \varphi$.
- $(\tau_1^{(t_1, \varphi_1)} \nearrow_a \tau_2^{(t_2, \varphi_2)})[B]$ iff $t_1 < t_2$ and $\tau_2^{(t_2, \varphi_2)} >_a \tau_2^{(t, \varphi)}$ for $t < t_1$ or for φ s.t. $\tau_1 \notin \varphi$.

Similar rules for recall and precision can be ob-

tained by replacing relations \nearrow_a and \searrow_a with the counterparts \nearrow_p , \nearrow_r , \searrow_p and \searrow_r .

Total obfuscation and *total gift* have a specific logical interpretation. A test τ_1 (revealing a property A) totally obfuscates another test τ_2 (revealing a property B). After the execution of τ_1 it is no longer possible to reveal B by means of τ_2 :

$$\frac{\tau_1^{(t_1, \varphi_1)} : A(c)(\tau_1^{(t_1, \varphi_1)} \searrow_a \tau_2^{(t_2, \varphi_2)})[B] t_1 < t_2}{\sim \tau_2^{(t_2, \varphi_2)} : B(c)} \text{totalObf}$$

Dually, a test τ_1 (revealing a property A) totally gifts another test τ_2 (revealing a property B). After the execution of τ_1 it is no longer necessary to reveal B , since the information that B holds for the sample is obtained as a side effect of the execution of τ_1 . Since B does not require to be revealed but, after the execution of τ_1 , it became a ground knowledge, we can classify B as a fact.

$$\frac{\tau_1^{(t_1, \varphi_1)} : A(c)(\tau_1^{(t_1, \varphi_1)} \nearrow_a \tau_2^{(t_2, \varphi_2)})[B] t_1 < t_2}{\tau_1^{(t_1, \varphi_1)} : B(c)} \text{totalGift}$$

From now on, for the sake of simplicity, in total gift and total obfuscation symbols we omit to specify the "a" symbol so that \searrow_a and \nearrow_a will simply be \searrow and \nearrow .

Clinical diagnostic is a useful setting to show what kind of hybrid knowledge we are modeling, but it is not the only context in which this knowledge can be found.

We discuss here an example related to information security. In particular, we consider the Open Web Application Security Project (OWASP)3, Top Ten Most Critical Web Application Security Risks. OWASP's Top Ten is updated regularly and the latest edition includes A_1 -Injection, A_2 -Broken Authentication, A_3 -Sensitive Data Exposure, A_4 -XML External Entities (XXE), A_5 -Broken Access Control, A_6 -Security Misconfiguration, A_7 -Cross-Site Scripting (XSS), A_8 -Insecure Deserialization, A_9 -Using Components with Known Vulnerabilities, A_{10} -Insufficient Logging and Monitoring. These risks are not independent from each other; being exposed to one risk sometimes entails being also exposed to another one in the list. Being exposed to Injection (A_1) means that

untrusted data is sent to an interpreter as part of a command or query and this data can trick the interpreter into executing unintended commands or accessing data without proper authorization; this can imply also to be exposed to the risk of Broken Authentication (A_2). In the following example, we interpret vulnerability scanning tools as a test to reveal a given risk of the OWASP Top Ten.

We write $\tau_{(Ai)}^{(t,\varphi)}$ to say that the security risk A_i is revealed by a test (a suitable scanning tool) $\tau_{(Ai)}$. To classify test/scanning tools, *i.e.*, to measure software performances, we adopt the standard binary classification of algorithm behaviour.

Example 3. (Total obfuscation).

Let be C a web application. Risk A_9 (Using Components with Known Vulnerabilities) obfuscates risk A_{10} (Insufficient Logging and Monitoring).

$$\frac{\tau_{A_9}^{(Mon,\emptyset)} : A_9(C) \left(\tau_{A_9}^{(Mon,\emptyset)} \downarrow \tau_{A_{10}}^{(Tue,\tau_{A_{10}})} \right) [A_{10}]}{\sim \tau_{A_{10}}^{(Tue,\tau_{A_9})} : A_{10}(C)}$$

Example 4. (Total gift).

Let C be a web application. Risk A_1 (Injection) totally gifts risk A_2 (Broken Authentication).

$$\frac{\tau_{A_1}^{(Mon,\emptyset)} : A_1(C) \left(\tau_{A_1}^{(Mon,\emptyset)} \uparrow \tau_{A_2}^{(Tue,\tau_{A_2})} \right) [A_2]}{\tau_{A_1}^{(Mon,\emptyset)} : A_1(C)}$$

An interference between tests τ_1 and τ_2 may occur. We write $\tau_1 \perp \tau_2$ if τ_1 and τ_2 are *non-interfering*, *i.e.*, if they do not obfuscate or gift each other in either direction.

4.3. Defeasible Logic and Diagnostic Reasoning

One of the most characterising aspects of the experiment-based reasoning is the possibility that a property revealed positively by a test is revealed negatively by another test. Generally speaking, we want to devise a method of reasoning that allows us to accommodate contra-

dictory assertions, in a non-monotonic fashion. In [14], [15], [16], [17], [18], [37], some of us have investigated the use of Defeasible Logic as a means for managing data coming from external sources and validated by means of data mining methods. Non-monotonic reasoning accommodates conclusions when dealing with potential conflicts. When derivations may lead to potentially contradictory conclusions, we may typically have two strategies to avoid inconsistencies. In a credulous approach, we branch by creating two distinct sets of conclusions: one for each of the contradictory conclusions. Opposite, with a skeptical approach, we need a (preference) mechanism to establish whether one conclusion is preferred to the other one (in literature, this is typically referred to as a superiority, or preference relation; see [20] for a systematic analysis). If such a mechanism is not able to solve the conflict, no conclusion is derived unless exceptions are given. Exceptions can be seen as particular conditions preventing to draw a specific conclusion⁵. In this paper, we shall not consider credulous settings. The formalism that we employ here to convert intrinsic non-monotonic aspects of $T^{\mathcal{L}}$ logic is Defeasible Logic (DL), a skeptical non-monotonic reasoning framework that accommodates assertions, priorities and negative exceptions as introduced above.

There are three distinct sources of non-monotonicity when reasoning about tests:

- Two different tests may give out different results on the same sample.
- One test cannot be used to conclude a diagnosis, because another test has modified the sample or created a condition that prevents the use of the sample.
- One test can be used to conclude diagnosis, because another test has modified the sample or created a condition that allows the use of the sample for concluding on the diagnosis without performing the test at all.

We emphasize these aspects in Section 5, where we introduce a rewriting algorithm that transforms a set of labelled $T^{\mathcal{L}}$ rules into a defeasible theory that can be processed, in turn, by a de-

⁵Such exceptions are known as negative exceptions. In credulous settings, another type of exceptions is possible: positive exceptions, whose purpose is to force a particular derivation.

feasible engine.

In this perspective, DL can be viewed as a meta-logic: its rules habilitate the expression of diagnostic reasoning in a natural way and, thanks to the rewriting algorithm, a defeasible theory is produced. Once we have produced such a defeasible theory, we can process the theory by means of the reasoning technology SPINdle. In Defeasible Logic [3], [30], we indeed have rules for opposite derivations, although not all concluded. In the situation where rules for opposite literals are activated the logic does not produce any inconsistency but does not draw any conclusion unless a preference (or superiority) relation states that one rule prevails over the other. A defeasible theory D is defined as a structure $(F, R, >)$, where:

- F is the set of facts, a set of atomic assertions (literals) considered to be always true (e.g., a fact is that "the stove is ON", formally "stove ON"),
- R is the set of rules, which in turn contains three finite sets of rules: strict rules (denoted by \Rightarrow), defeasible rules (denoted by symbol \rightarrow), and defeaters (denoted by symbol \rightsquigarrow).
- $>$ is a binary relation over R , restricted on defeasible rules with opposite conclusions.

A defeasible rule can be defeated by contrary evidence; defeaters are special rules whose only purpose is to defeat defeasible rules by producing contrary evidence. Our framework does not use strict rules or defeaters, but only defeasible rules. The superiority relation establishes that some rules override the conclusion of another one with the opposite conclusion. Like in [3], we consider only a version of this logic that can be reduced to a propositional theory and does not contain defeaters. In DL, a proof P of length n is a finite sequence $P(1), \dots, P(n)$ of tagged literals of the type $\pm \Delta p$ and $\pm \partial p$. The idea is that, at every step of the derivation, a literal is either proven or disproven. The set of positive and negative conclusions is called extension. The meanings of tagged literals is as follows

- $+\Delta q$, which means that there is a definite proof for q in D ; such a proof uses strict rules and facts only.
- $-\Delta q$ which means that q is not definitely provable in D .

- $+\partial q$ which means that q is defeasibly provable in D .
- $-\partial q$ which means that q is refuted, or not defeasibly provable in D .

Formalisation of the proof tags is out of the scope of the present paper. An idea of how the derivation mechanism works is proposed in the following Example 5.

Example 5.

Let $D = (a, b, \mathcal{R}, >)$ be a defeasible theory such that

$$\begin{aligned} \mathcal{R} = \{ & r_1 : a \rightarrow c \\ & r_2 : b \Rightarrow d \\ & r_3 : c \rightarrow -d \\ & > = \{(r_2, r_3)\} \\ & \} \end{aligned}$$

Then we derive $+\Delta c$ via r_1 and $+\partial d$ since r_2 is stronger than r_3 .

Note that it does not matter whether the antecedents of a rule have been proven as strict conclusions: if such literals are used to allow a defeasible rule to fire (as for r_2), the conclusion will be defeasible.

5. An Architecture for Diagnostic Reasoning

5.1. An Algorithm for Transforming $T^{\mathcal{L}}$ Assertions onto Defeasible Theories

In this section, we define an algorithm that translates diagnostic results in a DL theory. We employ three kinds of objects: (a) facts, (b) rules and (c) priorities, namely superiority relations that establish which rule prevails when conflicting conclusions arise. We treat $T^{\mathcal{L}}$ as formed by two layers, the first formed by the defeasible objects and the second formed by meta-rules, establishing how to deal with the rules themselves. Essentially, we consider the facts as known truths, incontrovertible data, such as the results of tests without errors, or direct anamnestic data, for instance the age of a patient. Rules are instead the central part of the logical structure and relate the tests to the diag-

nosis, providing defeasible derivations.

The relations between tests, both temporal and evaluation ones, including gift and obfuscation, are treated as meta-rules, namely rules providing room for derived priorities. In the application of the translation algorithm we show that the meta-rules can be synchronized, translated into defeasible rules, and then used to make a decision about the meaning of a $T^{\mathcal{L}}$ theory in linear time. First of all, we introduce the intended meaning of the elements of the system. This architecture of the solution is specified in detail in Section 5.2.

Meta-rules are written with two constraints: time and experimental evaluation. In particular, these rules are transformed into defeasible rules, extended with a temporal label expressing the initial time instant t of the (open) interval in which the rule is available to be put before the literals appearing in the rule itself and criteria based on measures, again mapped onto labels in the form p_+ or p_- above the derivation operation sign. The transformation algorithm *SincroCutII* is introduced below, Algorithm 1. The algorithm takes as input a set of meta-rules and a set of evaluations of an experiment and transforms them into a defeasible theory by checking the temporal constraints, interference, and modal relations among tests. The result of the algorithm is a defeasible theory. The model of these meta-rules is inspired by studies of one of the authors [12]. We now describe how the algorithm works. It takes as input a finite set of $T^{\mathcal{L}}$ assertions, and gives, as output, a defeasible theory. The first cycle initializes basic data structures, used to host the converted tokens. The second cycle of the algorithm computes the facts in the theory, and therefore determines the base for the subsequent derivations by the Spindle reasoner. The third cycle reads meta-rules and priorities and translates them into the defeasible theory under construction.

Notation $r : t$ extracts the temporal information of a rule. The procedure *evalExperiment* extracts the result of an experiment. The procedure *createNewRule* creates a new empty rule.

Observe, moreover, that obfuscation and gift, the relations between rules that are providing room for temporal re-processing of rules themselves, are treated as modals in the second cycle of the algorithm itself.

Algorithm 1. Transformation algorithm SincroCutII.

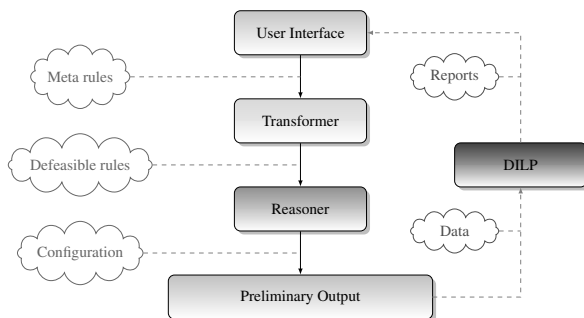
Input: an ordered set of *Metarules* Θ , current time t and a set of *evaluations* of the experiments Ψ ;
Output: a defeasible theory $T = \langle F, R, P \rangle$ (facts, rules, priorities);
 $\Psi' \leftarrow \Psi$
repeat
 $m \leftarrow \text{pop}(\Psi')$; $l \leftarrow \text{evalExperiment}(m)$;
 // l is in the form $A(l) : C(l) = \tau_i^{(t_i, \varphi_i)} : P_k(C)$
 $L \leftarrow \text{push}(l, L)$;
until $\Psi' = \emptyset$;
 $F \leftarrow \emptyset$; $L' \leftarrow L$;
repeat
 if $l.t_1 \leq t$ **then**
 $f \leftarrow A(l).\varphi$; $F \leftarrow \text{push}(f, F)$;
 $r' \leftarrow ' \Rightarrow ' + C(l)$; $R \leftarrow \text{push}(r', R)$;
 end if
until $L' = \emptyset$
repeat
 $r \leftarrow \text{pop}(\Theta)$;
 if $\text{type}(r) = \text{priority}$ **then**
 $P \leftarrow \text{push}(r, P)$;
 else
 if $r.t_1 \leq t$ **then**
 $\eta \leftarrow \emptyset$; $L' \leftarrow L$;
 repeat
 $l' \leftarrow \text{pop}(F')$;
 if $l' \in A(r)$ **then**
 $\eta \leftarrow (A(r) - l')$;
 end if
 until $(L' = \emptyset \text{ or } \eta \neq \emptyset)$;
 if $\eta \neq \emptyset$ **then**
 $r' \leftarrow \text{createNewRule}()$;
 if $\eta \simeq \downarrow$ **then**
 $r' \leftarrow \Rightarrow P_k(C)$; $r'' \leftarrow \Rightarrow \sim P_k(C)$;
 $R \leftarrow \text{push}(r', R)$; $R \leftarrow \text{push}(r'', R)$;
 else if $\eta \simeq \uparrow$ **then**
 $r' \leftarrow \Rightarrow P_k(C)$;
 $R \leftarrow \text{push}(r', R)$;
 else
 $A(r') \leftarrow \eta$; $C(r') \leftarrow C(r)$; $R \leftarrow \text{push}(r', R)$;
 end if
 end if
 end if
 end if
 end if
 until $\Theta = \emptyset$;
return T ;

A certain rule is a candidate for rewriting only if the synchronizer has acknowledged that its clock time falls within the validity interval of the rule when the rule has a validity interval explicitly specified, or at the exact instant of the rule if the rule is not tagged so and therefore is considered instantaneous.

The algorithm executes the translation. At this stage of our research, we do not provide proofs of soundness and completeness of the deduction system introduced in Section 4, and, consequently we will not discuss properties of correctness and completeness of the implemented solution. All these investigations are left to future work, along with the discussion of the semantics of the logical framework, and the corresponding canonical models. Being the framework based on DL, clearly the semantics of the former depends upon the semantics of the latter into which the $T^{\mathcal{L}}$ assertions are translated.

What we can prove here is the complexity of the method, which is independent of the issues discussed above (soundness of the $T^{\mathcal{L}}$ deduction rules, semantics of $T^{\mathcal{L}}$, completeness, canonical models, correctness of the algorithm, completeness of the algorithm). In fact, the algorithm is linear in the number of literals appearing in the $T^{\mathcal{L}}$ set of assertions given as input to the algorithm, since the limited number of cycles that can be executed is the number of literals.

5.2. Architecture of a System Implementing $T^{\mathcal{L}}$ Transformation



In this section, we briefly introduce an architecture for diagnostic reasoning that is based on four modules, some documented in this paper, some yet to come. The architecture is described in terms of functions of the modules. We introduce here the DILP module, a module used to perform *recommendations* on the rules to introduce that is based upon the Machine Learning methods of Inductive Logic Programming, an approach that is also applied to DL.

- **User Interface:** allows the user to input "meta rules" and provides visualization of all data coming from the DILP module;
- **Transformer:** takes as input a set of "meta rules" and gives as output a set of defeasible rules to be used by the Reasoner, in accordance with the time given by its internal clock mechanism and an evaluation algorithm;
- **Reasoner:** uses the rules and determines the "should be" conclusions; Preliminary Output: is responsible for the delivery to the user;
- **DILP:** gathers data and makes analysis delivering summaries and possible rules to be displayed by the User Interface (future extension).

The output of the User Interface is an ordered set of "Meta rules" which are one of the input of the Transformer that runs continuously and at given times uses the algorithm *SincroCutII* to produce a set of defeasible rules.

These rules are given to the Reasoner whose output is using the $+@$ conclusion given the set of rules coming from the Transformer. The Defeasible Logic rule engine used is a Prolog-like engine called SPINdle [22]. At different times different conclusions are possible, due to the work of the Transformer. We now show how the algorithm works by means of a detailed example.

6. Case Study

We use a concrete example to show how our model can fit a real-life scenario. We consider a web application of a bank, located in Italy, which had to be tested against the OWASP Top Ten risks that we listed above.

The analysis is split between two different contractors, namely subjects in charge of analyzing a subset of the list. Contractor α uses a combination of automated testing and human validation and can cover risk $\{A_1, A_3, A_4, A_5, A_{10}\}$ while contractor β performs only automated testing and can cover risk $\{A_2, A_4, A_6, A_7, A_8, A_9, A_{10}\}$; α is considered to perform tests with a higher accuracy than the ones delivered by β . Both execute the tests sequentially one per day and α

has been engaged after β has completed its task. For the sake of space in the rest of this example we will write A_i instead of $A_i(\text{BankApp})$ to describe the ground formulas we are revealing. There is an obfuscation on A_{10} given a test on A_9 and a gift on A_2 given a test on A_1

$$\frac{\tau_{A_9}^{(t_1, \emptyset)} : A_9 \left(\tau_{A_9}^{(t_1, \emptyset)} \downarrow \tau_{A_{10}}^{(t_2, \tau_{A_{10}})} \right) [A_{10}]}{\sim \tau_{A_{10}}^{(t_2, \tau_{A_9})} : A_{10}}$$

$$\frac{\tau_{A_1}^{(t_1, \emptyset)} : A_1 \left(\tau_{A_1}^{(t_1, \emptyset)} \uparrow \tau_{A_2}^{(t_2, \tau_{A_{10}})} \right) [A_2]}{\tau_{A_1}^{(t_1, \emptyset)} : A_2}$$

We also state that $\alpha : A_1 > \beta : A_1$, $\alpha : A_4 > \beta : A_4$ and $\alpha : A_{10} > \beta : A_{10}$ by knowledge on the accuracy of tests and $\alpha : A_2 > \beta : A_2$ because of the gift specified above. The results are:

$$\left\{ \alpha_{A_1}^{(1, \phi_1)} : A_1, \alpha_{A_3}^{(2, \phi_2)} : \neg A_3, \alpha_{A_4}^{(3, \phi_3)} : A_4, \right. \\ \left. \alpha_{A_5}^{(4, \phi_4)} : \neg A_5, \alpha_{A_{10}}^{(5, \phi_5)} : A_{10} \right\}$$

$$\left\{ \beta_{A_2}^{(6, \phi_6)} : \neg A_2, \beta_{A_4}^{(7, \phi_7)} : \neg A_4, \beta_{A_6}^{(8, \phi_8)} : A_6, \beta_{A_7}^{(9, \phi_9)} : \neg A_7, \right. \\ \left. \beta_{A_8}^{(10, \phi_{10})} : A_8, \beta_{A_9}^{(11, \phi_{11})} : A_9, \beta_{A_{10}}^{(12, \phi_{12})} : \neg A_{10} \right\}$$

This is therefore the set of meta-rules. Once the Translator has performed Algorithm *SincroCutIII* we have, at any time after all the tests have been executed ($t > 12$):

$$\rightarrow \phi_1, \rightarrow \phi_2, \rightarrow \phi_3, \rightarrow \phi_4, \rightarrow \phi_5, \rightarrow \phi_6, \rightarrow \phi_7, \rightarrow \phi_8, \\ \rightarrow \phi_9, \rightarrow \phi_{10}, \rightarrow \phi_{11}, \rightarrow \phi_{12}$$

$$\begin{aligned} r_{\alpha 1} : \Rightarrow A_1 & \quad r_{\beta 1} : \Rightarrow \neg A_2 \\ r_{\alpha 2} : \Rightarrow \neg A_3 & \quad r_{\beta 2} : \Rightarrow A_4 \\ r_{\alpha 3} : \Rightarrow \neg A_4 & \quad r_{\beta 3} : \Rightarrow A_6 \\ r_{\alpha 4} : \Rightarrow \neg A_5 & \quad r_{\beta 4} : \Rightarrow \neg A_7 \\ r_{\alpha 5} : \Rightarrow A_{10} & \quad r_{\beta 5} : \Rightarrow A_8 \\ & \quad r_{\beta 6} : \Rightarrow A_9 \\ & \quad r_{\beta 7} : \Rightarrow \neg A_{10} \end{aligned}$$

$$\begin{aligned} o_{\beta 1} : \Rightarrow A_{10} & \quad g_{\alpha 1} \succ r_{\beta 1} \\ o_{\beta 2} : \Rightarrow \neg A_{10} & \quad r_{\alpha 3} : \succ r_{\beta 2} \\ g_{\alpha 1} : \Rightarrow A_2 & \quad r_{\alpha 5} : \succ r_{\beta 7} \\ & \quad r_{\alpha 5} : \succ o_{\beta 2} \end{aligned}$$

Given that theory, the Reasoner concludes $+\partial A_1, +\partial A_2, +\partial \neg A_3, +\partial \neg A_4, +\partial \neg A_5, +\partial \neg A_6, +\partial \neg A_7, +\partial \neg A_8, +\partial \neg A_9, +\partial \neg A_{10}$, as shown in Appendix A.

We can therefore conclude that the application is subject to risks A_1 -Injection, A_2 -Broken Authentication, A_6 -Security Misconfiguration, A_8 -Insecure Deserialization, A_9 -Using Components with Known Vulnerabilities, A_{10} -Insufficient Logging and Monitoring, but not to A_3 -Sensitive Data Exposure, A_4 -XML External Entities (XXE), A_5 -Broken Access Control, A_7 Cross-Site Scripting (XSS).

7. Discussion and Conclusions

In this paper, we have developed the logic $T^{\mathcal{L}}$ [17], which is able to formalize a form of diagnostic reasoning based both on deduction and on experimental knowledge. We introduced some notions about experiment-based deduction, following a perspective clearly oriented to reasoning mechanization. In comparison with [17], we also focused on the (natural) defeasible aspects of diagnostic knowledge. To this end, we introduced a rewriting algorithm *SincroCutII*, that takes as input $T^{\mathcal{L}}$ formulas and transforms them into a defeasible theory by checking temporal, accuracy and interference constraints between tests. The result of the *SincroCutII* is a defeasible theory. By means of an example from a real-life scenario, we carried out a case study using the defeasible engine SPINdle. We are currently working in three directions. First, the system $T^{\mathcal{L}}$ can be improved as a (stand-alone) labelled temporal logic framework, and its proof theory seems to be a challenging and interesting task. On the semantic side, we observed that the natural interpretation for $T^{\mathcal{L}}$ is related to some interpretations of the branching time logic UB [9]. The most suitable style is Prawitz' natural deduction [4], [5], [33]. Following [13], [39], [40], we are developing a labelled, non-monotonic natural deduction system. Second, the defeasible flavor we pointed out in this paper

seems to be the right perspective to move toward a more expressive automatic reasoner. In particular, we aim to extend the deduction system both to include more refined quantitative information about tests and to address more complex diagnostic based deduction, including multi-level defeasible mechanisms.

Finally, there is a strong relation between tests and resources. In a more refined framework, a test could reasonably consume a resource in revealing a property. This reflects what effectively happens in a number of laboratory experiments and we plan to investigate whether this could be captured by means of an approach inspired by Linear Logic [1].

References

- [1] J. Y. Girard, "Linear Logic", *Theoretical Computer Science*, vol. 50, no. 1, pp. 1–101, 1987.
[http://dx.doi.org/10.1016/0304-3975\(87\)90045-4](http://dx.doi.org/10.1016/0304-3975(87)90045-4)
- [2] G. Antoniou *et al.*, "A Family of Defeasible Reasoning Logics and its Implementation", *Proceedings of the 14th European Conference on Artificial Intelligence, ECAI'00*, Amsterdam, The Netherlands, pp. 459–463, 2000.
- [3] G. Antoniou *et al.*, "Representation Results for Defeasible Logic", *ACM Trans. Comput. Log.*, vol. 2, no. 2, pp. 255–287, 2001.
<http://dx.doi.org/10.1145/371316.371517>
- [4] F. Aschieri and M. Zorzi, "Non-determinism, Non-termination and the Strong Normalization of System T", in: M. Hasegawa (Ed.), *Typed Lambda Calculi and Applications, Proceedings of the 11th International Conference, TLCA, 2013*, Eindhoven, The Netherlands, vol. 7941 of *Lecture Notes in Computer Science*, pp. 31–47, Springer, 2013.
http://dx.doi.org/10.1007/978-3-642-38946-7_5
- [5] F. Aschieri and M. Zorzi, "On Natural Deduction in Classical First-order Logic: Curry-Howard, Correspondence, Strong Normalization and Herbrand's Theorem", *Theor. Comput. Sci.*, vol. 625, pp. 125–146, 2016.
<https://doi.org/10.1016/j.tcs.2016.02.028>
- [6] J. Bellamy, "Medical Diagnosis, Diagnostic Spaces, and Fuzzy Systems", *Journal of the American Veterinary Medical Association*, vol. 210, no. 3, pp. 390–396, 1997.
<https://www.ncbi.nlm.nih.gov/pubmed/9057925>
- [7] M. Belmonte-Serrano *et al.*, "Renoir: An Expert System Using Fuzzy Logic for Rheumatology Diagnosis", *International Journal of Intelligent Systems*, vol. 9, no. 11, pp. 985–1000, 1994.
<http://dx.doi.org/10.1002/int.4550091102>
- [8] K. Boegl *et al.*, "Knowledge Acquisition in the Fuzzy Knowledge Representation Framework of a Medical Consultation System", *Artif. Intell. Med.*, vol. 30, no. 1, pp. 1–26, 2004.
[http://dx.doi.org/10.1016/S0933-3657\(02\)00073-8](http://dx.doi.org/10.1016/S0933-3657(02)00073-8)
- [9] C. Caleiro *et al.*, "A labeled deduction system for the logic UB", in: C. Sánchez, K. B. Venable, and E. Zimanyi (Eds.), *2013 20th International Symposium on Temporal Representation and Reasoning*, Pensacola, FL, USA, 26–28, 2013, pp. 45–53 IEEE Computer Society, 2013.
<http://dx.doi.org/10.1109/TIME.2013.14>
- [10] C. Combi *et al.*, "A logical framework for XML reference specification", in: Q. Chen, A. Hameurlain, F. Toumani, R. R. Wagner, and H. Decker (Eds.), *Database and Expert Systems Applications - Proceedings of 26th International Conference, DEXA 2015, Valencia, Spain, Part II*, vol. 9262 of *Lecture Notes in Computer Science*, pp. 258–267, Springer, 2015.
http://dx.doi.org/10.1007/978-3-319-22852-5_22
- [11] C. Combi *et al.*, "A Hybrid Logic for XML Reference Constraints", *Data Knowl. Eng.*, vol. 115, pp. 94–115, 2018.
<http://dx.doi.org/10.1016/j.datak.2018.02.005>
- [12] M. Cristani *et al.*, "Ontology-driven compression of temporal series: A case study in SCADA technologies", in *19th International Workshop on Database and Expert Systems Applications DEXA 2008*, Turin, Italy, pp. 734–738, IEEE Computer Society, 2008.
<http://dx.doi.org/10.1109/DEXA.2008.72>
- [13] M. Cristani *et al.*, "It could rain: weather forecasting as a reasoning process", in: R. J. Howlett, L. C. Jain, Z. Popovic, D. B. Popovic, S. N. Vukosavic, C. Toro, and Y. Hicks (Eds.), *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 22nd International Conference KES-2018*, Belgrade, Serbia, 2018, *Procedia Computer Science*, vol. 126, pp. 850–859, Elsevier, 2018.
<http://dx.doi.org/10.1016/j.procs.2018.08.019>
- [14] M. Cristani *et al.*, "Energy saving by ambient intelligence techniques", in: L. Barolli, F. Xhafa, M. Takizawa, T. Enokido, A. Castiglione, and A. De Santis (Eds.), *17th International Conference on Network-Based Information Systems, NBiS 2014, Salerno, Italy*, pp. 157–164, IEEE Computer Society, 2014.
<http://dx.doi.org/10.1109/NBiS.2014.39>
- [15] M. Cristani *et al.*, "Improving energy saving techniques by ambient intelligence scheduling", in: L. Barolli, M. Takizawa, F. Xhafa, T. Enokido, and J. H. Park (Eds.), *29th IEEE International Conference on Advanced Information Networking and Applications, AINA 2015, Gwangju, South Korea*, pp. 324–331, IEEE Computer Society, 2015.

- <http://dx.doi.org/10.1109/AINA.2015.202>
- [16] M. Cristani *et al.*, "Automatic synthesis of best practices for energy consumptions", in *10th International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing, IMIS 2016, Fukuoka, Japan*, pp. 154–161, IEEE Computer Society, 2016.
<http://dx.doi.org/10.1109/IMIS.2016.79>
- [17] M. Cristani *et al.*, "Towards a logical framework for diagnostic reasoning", in: G. Jezic, Y. J. Chen-Burger, R. J. Howlett, L. C. Jain, Lj. Vlacic, and R. Šperka (Eds.), *12th KES Conference on Agent and Multi-Agent Systems: Technologies and Applications, KES-AMSTA 2018, Smart Innovation, Systems and Technologies*, vol. 96, pp. 144–155, Cham, 2019. Springer International Publishing.
http://dx.doi.org/10.1007/978-3-319-92031-3_14
- [18] M. Cristani *et al.*, "Defeasible reasoning about electric consumptions", in: L. Barolli, M. Takizawa, T. Enokido, A. J. Jara, and Y. Bocchi (Eds.), *30th IEEE International Conference on Advanced Information Networking and Applications, AINA 2016, Crans-Montana, Switzerland, 23-25 March, 2016*, pp. 885–892, IEEE Computer Society, 2016.
<http://dx.doi.org/10.1109/AINA.2016.62>
- [19] R. Davis, "Diagnostic reasoning based on structure and behavior", *Artificial Intelligence*, vol. 24, no. 1, pp. 347–410, 1984.
<https://dblp.org/rec/bib/journals/ai/Davis84>
- [20] P. M. Dung *et al.*, "Argumentation-based proof procedures for credulous and sceptical non-monotonic reasoning", in: A. C. Kakas and F. Sadri (Eds.), *Computational Logic: Logic Programming and Beyond, Essays in Honour of Robert A. Kowalski, Part II, Lecture Notes in Computer Science*, vol. 2408, pp. 289–310, Springer, 2002.
http://dx.doi.org/10.1007/3-540-45632-5_12
- [21] P. E. Johnson *et al.*, "Expertise and error in diagnostic reasoning", *Cognitive Science*, vol. 5, no. 3, pp. 235–283, 1981.
http://dx.doi.org/10.1207/s15516709cog0503_3
- [22] H. Lam and G. Governatori, "The making of spindle", in: G. Governatori, J. Hall, and A. Paschke (Eds.), *Rule Interchange and Applications, Proceedings of International Symposium, RuleML 2009, Las Vegas, Nevada, USA*, vol. 5858 of Lecture Notes in Computer Science, pp. 315–322, Springer, 2009.
http://dx.doi.org/10.1007/978-3-642-04985-9_29
- [23] A. E. Lawson and E. S. Daniel, "Inferences of clinical diagnostic reasoning and diagnostic error", *Journal of Biomedical Informatics*, vol. 44, no. 3, pp. 402–412, 2011, Biomedical Complexity and Error.
<http://dx.doi.org/10.1016/j.jbi.2010.01.003>
- [24] Q. Liu *et al.*, "Design and implement for diagnosis systems of hemorheology on blood viscosity syndrome based on GRC", in *Proceedings of the 9th International Conference on Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, RSFDGrC'03*, pp. 413–420, Berlin, Heidelberg, 2003. Springer-Verlag.
http://dx.doi.org/10.1007/3-540-39205-X_69
- [25] C. D. Manning *et al.*, "Introduction to Information Retrieval", Cambridge University Press, New York, NY, USA, 2008.
- [26] A. Masini *et al.*, "A qualitative modal representation of quantum register transformations", in *38th IEEE International Symposium on Multiple-Valued Logic (ISMVL 2008)*, Dallas, Texas, USA, pp. 131–137, IEEE Computer Society, 2008.
<http://dx.doi.org/10.1109/ISMVL.2008.36>
- [27] A. Masini *et al.*, "Modal deduction systems for quantum state transformations", *Multiple-Valued Logic and Soft Computing*, vol. 17, no. 5–6, pp. 475–519, 2011.
<http://www.oldcitypublishing.com/journals/mvlsc-home/mvlsc-issue-contents/mvlsc-volume-17-number-5-6-2011/mvlsc-17-5-6-p-475-519/>
- [28] M. McShane *et al.*, "Inconsistency as a diagnostic tool in a society of intelligent agents", *Artificial Intelligence in Medicine*, vol. 55, no. 3, pp. 137–148, 2012.
- [29] D. McSherry, "Conversational case-based reasoning in medical decision making", *Artificial Intelligence in Medicine*, vol. 52, no. 2, pp. 59–66, 2011., *Artificial Intelligence in Medicine AIME 2009*.
- [30] D. Nute, "Handbook of Logic in Artificial Intelligence and Logic Programming", in: D. M. Gabbay, C. J. Hogger, and J. A. Robinson (Eds.), vol. 3, pp. 353–395, Oxford University Press, Inc., New York, NY, USA, 1994.
- [31] N. H. Phuong and V. Kreinovich, "Fuzzy logic and its applications in medicine", *International Journal of Medical Informatics*, vol. 62, no. 2, pp. 165–173, 2001.
- [32] P. Pis and R. Mesiar, "Fuzzy model of inexact reasoning in medicine", *Computer Methods and Programs in Biomedicine*, vol. 30, no. 1, pp. 1–8, 1989.
- [33] D. Prawitz, "Natural Deduction: A Proof-Theoretical Study", Dover Publications, 1965.
- [34] R. Reiter, "A theory of diagnosis from first principles", *Artificial Intelligence*, vol. 32, no. 1, pp. 57–95, 1987.
- [35] I. Rish, "An empirical study of the naive bayes classifier", in *IJCAI 2001 workshop on empirical methods in artificial intelligence*, vol. 3, pp. 41–46, IBM New York, 2001.
- [36] R. Seising, "From vagueness in medical thought

to the foundations of fuzzy reasoning in medical diagnosis", *Artificial Intelligence in Medicine*, vol. 38, no. 3, pp. 237–256, 2006.

- [37] C. Tomazzoli *et al.*, "Non-monotonic reasoning rules for energy efficiency", *JAISE*, vol. 9, no. 3, pp. 345–360, 2017.
- [38] L. Viganò, "Labelled non-classical logics", Kluwer, 2000.
<https://dblp.org/rec/bib/books/daglib/0003059>
- [39] L. Viganò *et al.*, "Quantum state transformations and branching distributed temporal logic", in: U. Kohlenbach, P. Barceló, and R. J. G. B. de Queiroz (Eds.), *Logic, Language, Information, and Computation - 21st International Workshop, WoLLIC 2014, Valparaíso, Chile, 2014*. Proceedings, Lecture Notes in Computer Science, vol. 8652, pp. 1–19, Springer, 2014.
http://dx.doi.org/10.1007/978-3-662-44145-9_1
- [40] L. Viganò *et al.*, "A branching distributed temporal logic for reasoning about entanglement-free quantum state transformations", *Inf. Comput.*, vol. 255, pp. 311–333, 2017.
<http://dx.doi.org/10.1016/j.ic.2017.01.007>
- [41] P. Vineis, "Methodological insights: fuzzy sets in medicine", *Journal of Epidemiology & Community Health*, vol. 62, no. 3, pp. 273–278, 2008.
<http://dx.doi.org/10.1136/jech.2007.063644>
- [42] A. Yardimci, "Soft computing in medicine", *Applied Soft Computing*, vol. 9, no. 3, pp. 1029–1043, 2009.
<http://dx.doi.org/10.1016/j.asoc.2009.02.003>
- [43] A. Masini and M. Zorzi, "A logic for quantum register measurements", *Axioms*, vol. 8, no. 1, pp. 25, 2019.
<http://dx.doi.org/10.3390/axioms8010025>

Received: September 2018

Revised: January 2019

Accepted: January 2019

Contact address:

Matteo Cristiani
Department of Computer Science
University of Verona,
Italy

Francesco Olivieri
Data61
CSIRO
Australia

Claudio Tomazzoli
Department of Computer Science
University of Verona,
Italy

Luca Viganò
Department of Informatics
King's College London
UK

Margherita Zorzi
Department of Computer Science
University of Verona,
Italy

MATTEO CRISTIANI is Senior Researcher at the Department of Computer Science of the University of Verona. He obtained his Ph.D. in Electronic and Information Engineering from the University of Padova in 1995. He has been a principal investigator in several industrial projects and is the author of more than seventy journal and conference papers in the area of artificial intelligence. His research interests include machine learning, automated reasoning, and natural language processing.

FRANCESCO OLIVIERI completed his Master degree in theoretical Computer Science with Magna cum Laude at the University of Verona in 2009. Francesco consequently obtained two doctorates, one from the University of Verona in Multi-Agents Systems (2014), and one from the Griffith University, in Non-Classical Logics (2015). Francesco is currently working at Data61 in the Software & Computational Systems Group, CSIRO, where he also covers the role of co-chair for the CSIRO Postdoctoral Fellows Committee. His research areas vary from legal informatics, logics, natural language processing, business process visualization, verification and compliance. He is the inventor and promotor of the "Socrates Peripatetic Scuola" to: 1) provide an opportunity to disadvantaged or overlooked students, to nurture their learning passion and talent for science that they could not receive otherwise, and 2) to provide university scholarships.

CLAUDIO TOMAZZOLI obtained his M.S. in Electronic Engineering (1997) from the University of Padova, Italy, and his Ph.D. in Computer Science (2014) from the University of Verona, Italy. He was an IT professional before going back to studies, and is currently a research fellow at the University of Verona. His research interests cover energy saving techniques, machine learning, automatic document processing, and image retrieval.

LUCA VIGANÒ received his Ph.D. in Computer Science from the University of Saarland, Germany, in 1997, and his Habilitation in Computer Science from the University of Freiburg, Germany, in 2003. He is Professor of Computer Science and Head of the Cybersecurity Group at the Department of Informatics of King's College London, UK. His research focuses on formal methods and tools for the specification, verification, and construction of secure systems, and on the theory and applications of non-classical and security logics. On these topics, he has taught several classes, tutorials, and industrial courses, and has published extensively. He has served as PC-chair and PC-member in over 120 international conferences and workshops, some of which he founded or helped found.

MARGHERITA ZORZI received a PhD in Computer Science from the University of Verona in 2009. She was a post-doc researcher in Verona and Paris (Université Paris Nord) before becoming Assistant Professor at the University of Verona in 2018. Her principal research interests are logical methods in computer science, quantum computing foundations and statistical natural language processing. Zorzi has published in international peer-reviewed conferences and journals (e.g., *Information & Computation*, *Mathematical Structure in Computer Science*, and *Theoretical Computer Science*). She has been reviewer for top-level conferences and journals, and PC member of international conferences and workshops.
