
Patrick Rodler

Don't Treat the Symptom, Find the Cause!
Efficient Artificial-Intelligence Methods
for (Interactive) Debugging

HABILITATION THESIS

Subject Area: Computer Science

Specialization: Artificial Intelligence

Alpen-Adria-Universität Klagenfurt

Klagenfurt, December 2022

Abstract

In the modern world, we are permanently using, leveraging, interacting with, and relying upon systems of ever higher sophistication, ranging from our cars, recommender systems in e-commerce, and networks when we go online, to integrated circuits when using our PCs and smartphones, the power grid to ensure our energy supply, security-critical software when accessing our bank accounts, and spreadsheets for financial planning and decision making. The complexity of these systems coupled with our high dependency on them implies both a non-negligible likelihood of system failures, and a high potential that such failures have significant negative effects on our everyday life. For that reason, it is a vital requirement to keep the harm of emerging failures to a minimum, which means minimizing the system downtime as well as the cost of system repair. This is where model-based diagnosis comes into play.

Model-based diagnosis is a principled, domain-independent approach that can be generally applied to troubleshoot systems of a wide variety of types, including all the ones mentioned above, and many more. It exploits and orchestrates i.a. techniques for knowledge representation, automated reasoning, heuristic problem solving, intelligent search, optimization, stochastics, statistics, decision making under uncertainty, machine learning, as well as calculus, combinatorics and set theory to detect, localize, and fix faults in abnormally behaving systems.

In this thesis, we will give an introduction to the topic of model-based diagnosis, point out the major challenges in the field, and discuss a selection of approaches from our research addressing these issues. For instance, we will

- present methods for the optimization of the time and memory performance of diagnosis systems,
- tackle the improvement of the usefulness and usability of debuggers,
- introduce efficient techniques for a semi-automatic debugging by interacting with a user or expert,
- reveal results of profound theoretical and empirical analyses,
- address how to conquer particularly challenging diagnostic use cases,
- make recommendations as to the best use, configuration, and combination of diagnostic algorithms,
- provide assistance for other researchers and practitioners by a survey and didactic treatises,
- critically call into question and scrutinize techniques and practices adopted in the field,
- describe our full-fledged debugging tool that incorporates more than a decade of research, and
- demonstrate how our methods can be effectively harnessed in important application domains such as optimal scheduling, spreadsheet debugging, or the quality assurance of the Semantic Web.

Contents

List of Tables	1
List of Figures	1
1 Introduction and Preliminaries	3
1.1 Model-Based Diagnosis	3
1.2 Sequential Diagnosis	8
1.3 Diagnostic Tasks	10
1.4 Application Areas	11
1.5 Related Research Fields	12
1.6 A Generic (Interactive) Diagnosis System	14
1.7 Research Challenges and Goals	15
2 Our Research: Overview	19
2.1 From the Diagnosis System Perspective	19
2.2 From the Type of Research Perspective	19
2.3 Works Included in this Thesis: In a Nutshell	20
3 Our Research: Details	23
3.1 Diagnosis Computation	23
3.1.1 Memory-Limited Diagnosis Computation	23
3.1.2 DynamicHS: Optimizing Reiter’s Hitting-Set Tree for Sequential Diagnosis	25
3.1.3 StaticHS: Reducing Costs by Solving a Generalized Sequential Diagnosis Problem	26
3.1.4 On the Impact of Diagnosis Computation Strategies on Diagnostic Decision Making	28
3.1.5 A Taxonomy and Classification of Diagnosis Computation Algorithms	30
3.2 Query Computation and Selection	31
3.2.1 Proposal and Theoretical Analysis of Query Selection Heuristics	31
3.2.2 Empirical Analysis of Query Selection Heuristics	33
3.2.3 Efficient Query Computation and Selection by Systematic Search	34
3.3 User Studies and Debugging Tool	36
3.3.1 User Study: Are Query-Based Debuggers Really Helping Users?	36
3.3.2 On User Types, Assumptions, Optimization Criteria, and How to Make Query- Based Debugging Simpler and More Efficient	37
3.3.3 OntoDebug: A Full-Fledged Ontology Debugging Tool	39
3.4 Diagnostic Reasoning	41
3.4.1 Optimization by Randomization for Hard Diagnostic Problem Cases	41
3.4.2 A Formal Proof and Simple Explanation of the Seminal QuickXplain Algorithm	43
3.5 Diagnostic Modeling	44
3.5.1 On Modeling Techniques for Spreadsheet Debugging: A Theoretical and Empirical Analysis	44

3.5.2 Reducing Model-Based Diagnosis to Knowledge Base Debugging	46
Bibliography	49
A Papers Included in the Habilitation Thesis	55

List of Tables

2.1	Research works grouped based on the addressed aspect(s) of a generic diagnosis system	20
2.2	Research works grouped by type of paper	20
2.3	Overview of the research works included in this thesis	21

List of Figures

1.1	Full adder example: Components, system description, observations, predictions	4
1.2	Full adder example: Minimal conflicts	5
1.3	Full adder example: Minimal diagnoses	6
1.4	Full adder example: Diagnosis computation using HS-Tree	7
1.5	Full adder example: Sequential diagnosis, situation after the first measurement	9
1.6	Model-based diagnosis application domains	11
1.7	Related fields that influence model-based diagnosis	12
1.8	Related fields influenced by model-based diagnosis	13
1.9	Generic (sequential) diagnosis system: Modules, external entities, functionality	14
1.10	Generic (sequential) diagnosis system: Problems, research challenges, goals	16
3.1	Meaning of the accentuations used in Figures 3.2–3.16.	23
3.2	Memory-limited diagnosis computation	24
3.3	DynamicHS—Streamlining Reiter’s hitting-set tree for sequential diagnosis	26
3.4	StaticHS—Reducing costs by solving a generalized sequential diagnosis problem	27
3.5	On the impact of diagnosis computation strategies on diagnostic decision making	29
3.6	A taxonomy and classification of diagnosis computation algorithms	31
3.7	Proposal and theoretical analysis of query selection heuristics	32
3.8	Empirical analysis of query selection heuristics	34
3.9	Efficient query computation and selection by systematic search	35
3.10	User study	37
3.11	How to make query-based debugging simpler and more efficient	38
3.12	OntoDebug—A full-fledged ontology debugging tool	39
3.13	Optimization by randomization for hard diagnostic problem cases	41
3.14	A formal proof and simple explanation of the QuickXplain algorithm	44
3.15	On modeling techniques for spreadsheet debugging	45
3.16	Reducing model-based diagnosis to knowledge base debugging	46

Chapter 1

Introduction and Preliminaries

We first provide a short and example-driven introduction to model-based diagnosis, the overarching topic of this thesis. Then, we outline the main ideas behind sequential diagnosis, a research area dedicated to efficient interactive diagnosis. Next, we describe the principal diagnostic tasks of interest, shed light on application areas, and discuss links to related research domains. Finally, we expound the structure, modules and way of functioning of a generic (interactive) diagnosis system, based on which we explain the main research challenges and goals in the field tackled by our works described in later chapters.

Please note that we intentionally keep the discussion of the preliminaries on a higher level of abstraction and abstain from a very formal and detailed treatment in this chapter, whose main aim is to give the reader an intuition and overview of the topic. More in-depth accounts of technicalities can be found in the papers included in this thesis (cf. Appendix A).

1.1 Model-Based Diagnosis

Model-based diagnosis [Rei87, dKW87] is a popular, well-understood, domain-independent and principled paradigm for troubleshooting malfunctioning systems of a wide variety of types, such as physical devices, software, knowledge bases, or vehicles. The theory of model-based diagnosis [Rei87] assumes a *system* (e.g., a circuit) that is composed of a set of *components* (e.g., gates), and a *system description* (e.g., a logical knowledge base characterizing the circuit). The latter is expressed in a formal (logical) language and can be used to derive the expected behavior of the system by means of automated deduction systems. If the predicted system behavior, under the assumption that all components are functioning normally, is not in line with *observations* made or *measurements* conducted about the system, the goal is to locate the abnormal system components that are responsible for this discrepancy.

Example 1.1 (System description, components, observations, predictions) Consider the faulty full adder (cf. [Rei87]) depicted in Fig. 1.1. We have five system components, comprising two xor-gates (X_1, X_2), two and-gates (A_1, A_2), and one or-gate (O_1). Moreover, we have observations of the circuit's inputs (1,0,1) and outputs (sum bit equal to 1 and carry bit equal to 0). No measurements are given. The system description could, e.g., be formulated as a logical knowledge base, specifying the (normal) behavior of all gates, which gates are connected by wires, as well as relevant general domain knowledge. E.g., for the xor-gate X_1 , we would have the logical axiom $ok(X_1) \rightarrow out(X_1) = xor(in_1(X_1), in_2(X_1))$ where $out(G)$ describes the value at the output terminal of gate G , $in_1(G), in_2(G)$ denote the two input values of gate G , xor is the standard xor-function, and ok is a distinguished predicate used to state that a component is functioning normally. Further, the connection between, e.g., the xor-gates X_1 and X_2 is specified by the axiom $out(X_1) = in_1(X_2)$. An example of a general domain axiom is $\neg(1 = 0)$, which expresses that the constants 1 and 0 are different, i.e., if we derive that the value at some gate terminal is both 1 and 0, then we obtain a contradiction. Now, using a theorem

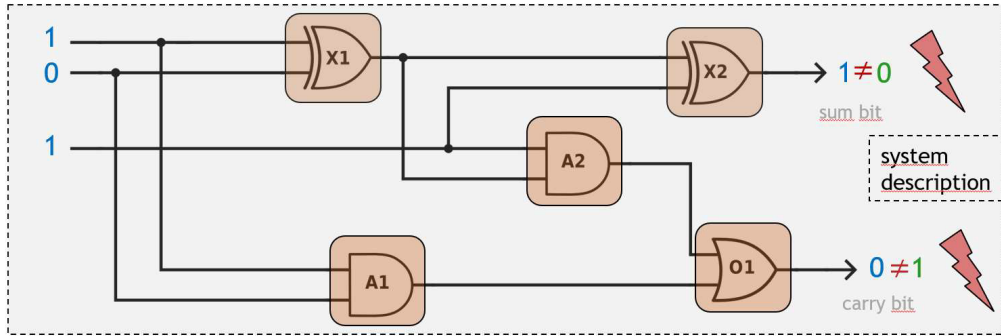


Figure 1.1: Full adder example: Components (orange boxes), observations (blue numbers), predictions under the assumption of normally behaving components (green numbers), and consistency check between observed and predicted values (red inequality signs; red flashes indicate inconsistency). The system description is a formal characterization of the depicted circuit, e.g., a set of logical sentences expressing the (normal) behavior of the components, their interrelations, as well as other relevant domain knowledge.

prover to predict the circuit’s outputs given its inputs, while assuming all five gates to be normal, i.e., $\{ok(X_1), ok(X_2), ok(A_1), ok(A_2), ok(O_1)\}$, we obtain a contradiction. In particular, we deduce that the sum bit must be 0, and the carry bit must be 1, which however is inconsistent with the given observations of the two outputs. As a conclusion, we know that some component must be abnormal, i.e., the $ok(G)$ assumption for some gate(s) G must be retracted to restore consistency between predictions and observations. \square

Different Interpretations of the Diagnosis Problem. In general, there are two main interpretations of the system description on the one hand, and the observations and measurements on the other. First, as illustrated in Example 1.1, the system description might constitute the correct specification of a system (e.g., a physical device) in which case the observations and measurements represent evidence about the actual system (the real device) implementing the specification. We term this interpretation of system model vs. evidence as *correct vs. actual*. Second, the system description might correspond to a model of a faulty artifact (e.g., a buggy program) in which case the observations and measurements are specifications of the intended system behavior (e.g., giving evidence what the correct program should do), and the “normal” behavior of components (e.g., lines of code) is the description of the components just as they are given in the faulty artifact. We refer to this interpretation as *actual vs. intended*.

Diagnosis Problem Instance and Diagnoses. Given a *diagnosis problem instance (DPI)*—consisting of the system description, the system components, the observations and the measurements—a (minimal) *diagnosis* is a (minimal¹) set of components that, when assumed abnormal, while all other components are assumed normal, leads to consistency between the system description (predicted behavior) and the observations and measurements (real/intended behavior). Simply put, a diagnosis is an explanation for the observed or measured system misbehavior, and a minimal diagnosis is an explanation that assumes a minimal set of components to be abnormal. In practice, for reason of computational efficiency and obeying the principle of parsimony [Rei87], the focus of diagnosis research and tools is often restricted to minimal diagnoses. The set of minimal diagnoses is representative of all diagnoses given the common *weak fault model* [dKMR92] setting where the system description characterizes only the normal, and no abnormal, behavior of components. This is also the diagnostic setting we consider throughout this work.

Conflicts and Hitting-Set Property. Computing (minimal) diagnoses requires checking the consistency of different component (ab)normality assumptions with the observations and measurements. Basically, the relevant search space spans all subsets of the set of system components, which makes an exhaustive exploration infeasible in all but the simplest cases. However, a more focused generation of diagnoses is

¹Whenever we speak of *minimality*, we mean minimality wrt. set inclusion.




Figure 1.2: Full adder example: Minimal conflicts.

enabled by the notion of a conflict. Intuitively, when the windshield wipers of our car do not work, the attention of the diagnostic analysis can concentrate on a small subset of the car’s components, e.g., the motor of the wipers, or the connecting cables, but not, say, on parts of the car’s engine or the braking system, because they do not influence the observed behavior of the car [vLP08, Chap. 10]. That is, observations are usually not only inconsistent with the system description under the assumption that *all* components are normal, but already when an (often relatively) small subset of the components is assumed fault-free. In more technical terms, a (*minimal*) *conflict* is a (minimal) set of system components that, when assumed normal, leads to an inconsistency between the system description and the observations and measurements. So, at least one component in each (minimal) conflict must be assumed abnormal in order to explain the system misbehavior. This directly leads us to the hitting-set relationship between conflicts and diagnoses [Rei87]: A (minimal) diagnosis is a (minimal) hitting set of all minimal conflict sets. Given a collection of sets \mathbf{S} , X is a (*minimal*) *hitting set* of \mathbf{S} iff (X is a minimal set such that) X includes only elements occurring in sets of \mathbf{S} and X has a non-empty intersection with each set in \mathbf{S} .

Example 1.2 (Conflicts and diagnoses) For our full-adder problem in Fig. 1.1, we find two minimal conflicts, i.e., $\mathbf{C} = \{C_1, C_2\} = \{\langle X_1, X_2 \rangle, \langle X_1, A_2, O_1 \rangle\}$, as illustrated in Fig. 1.2.² E.g., C_1 is a conflict because, assuming the two xor-gates X_1, X_2 to be functioning normally, (1) we can derive from the two inputs $(in_1(X_1), in_2(X_1)) = (1, 0)$ of X_1 that the output $out(X_1)$ of X_1 must be equal to 1, and (2) from $out(X_1) = in_1(X_2) = 1$ as well as $in_2(X_2) = 1$, we can deduce that $out(X_2) = 0$, which means a contradiction to the observed value $out(X_2) = 1$. The minimality of the conflicts follows from the fact that, once we omit one of the normality assumptions for X_1 or X_2 , we cannot derive $out(X_2) = 0$ anymore. An analogous argumentation is applicable for the second conflict C_2 .

By computing all minimal hitting sets for \mathbf{C} , we can determine all minimal diagnoses $\mathbf{D} = \{D_1, D_2, D_3\} = \{[X_1], [A_2, X_2], [O_1, X_2]\}$, as shown by Fig. 1.3. E.g., diagnosis D_1 corresponds to the assumption $\{\neg ok(X_1), ok(X_2), ok(A_1), ok(A_2), ok(O_1)\}$. This assumption explains the faulty behavior of the full-adder because, when assuming X_1 to be abnormal (i.e., $\neg ok(X_1)$), which is equivalent to dropping the restriction $out(X_1) = xor(in_1(X_1), in_2(X_1))$ imposed on the behavior of X_1 (cf. Example 1.1), none of the two predictions $out(X_2) = 0, out(O_1) = 1$ which contradict the observations can be inferred. In other words, the system description under this assumption is consistent with the observations. This is exactly the definition of a diagnosis. Minimality of the diagnosis D_1 can be directly seen by observing that the only proper subset $[\]$ (the empty diagnosis which assumes none of the components faulty) of $[X_1]$ is not a diagnosis (cf. the argumentation in Example 1.1). \square

Classical Algorithms for Diagnosis and Conflict Computation. As a consequence of the hitting-set property discussed above, conflicts can help to largely restrict the search space for diagnoses, and also permit the computation of diagnoses in a particular order of preference or plausibility. The two most

²Throughout this work, we denote conflicts by angle brackets and diagnoses by squared brackets.




Figure 1.3: Full adder example: Minimal conflicts and minimal diagnoses.

common such notions of diagnostic preferability are minimum-cardinality (i.e., those diagnoses assuming the fewest faulty components are favored) and maximal probability (i.e., the most likely diagnoses based on, e.g., known component failure rates are deemed most plausible). Notably, it is often not necessary to generate *all* (minimal) conflicts if the goal is to find only a set of the most-preferred minimal diagnoses.

One prominent and commonly adopted algorithm for diagnosis computation based on conflicts is Reiter's HS-Tree [Rei87]. It constructs a node- and edge-labeled tree in breadth-first manner (cf. Fig. 1.4). Each internal tree node n is labeled by a minimal conflict which has an empty intersection with $path(n)$, the set of edge labels along the path from the root to n (where $path(n) = \emptyset$ for the root node n). If there is no such conflict, $path(n)$ is a *diagnosis* (as it must already hit all minimal conflicts) and is labeled by

- (a) $\times_{(\supset \mathcal{D}_i)}$ if there is some already found minimal diagnosis \mathcal{D}_i with $\mathcal{D}_i \subset path(n)$ (i.e., $path(n)$ is a *non-minimal diagnosis* and node n is closed, i.e., becomes a leaf node), or
- (b) $\checkmark_{(\mathcal{D}_i)}$ else, where the index (\mathcal{D}_i) indicates that $path(n)$ is the i -th found minimal diagnosis denoted by \mathcal{D}_i (i.e., $path(n)$ is a *minimal diagnosis*³ and node n is closed, i.e., becomes a leaf node).

Each node labeled by a minimal conflict $\mathcal{C} = \{c_1, \dots, c_k\}$ has k outgoing edges, each labeled by a different element c_i from \mathcal{C} and pointing to a newly generated node. Before labeling a node n , the algorithm tests if there is another node n' with $path(n') = path(n)$, and if so, n is labeled by \times , i.e., node n is closed and becomes a leaf node (because in terms of its path it is a duplicate and hence can be discarded). Finally, when the tree T is completed, i.e., there are no more unlabeled nodes, the set of minimal diagnoses is exactly given by $Sol_{HS} := \{path(n) \mid n \text{ is labeled by } \checkmark\}$.

In fact, the execution of HS-Tree can be stopped at any point in time; in this case, Sol_{HS} will comprise a set of most preferred diagnoses (those with minimal cardinality). Since conflict computation, taking into account the involved theorem prover calls, is usually the most expensive operation in the HS-Tree algorithm, all minimal conflicts computed throughout the tree construction process are stored, and before a conflict is freshly computed for a node, the collection of stored conflicts is scanned for a suitable conflict with an empty intersection with the current node's path.⁴

Literature offers a variety of algorithms for conflict computation (cf. e.g., [Rod22c]). One of the most commonly used in model-based diagnosis is the seminal QuickXplain algorithm [Jun04, Rod22c] (cf. Example 1.4 and Sec. 3.4.2). It performs a recursive divide-and-conquer analysis of the set of system components and requires per computed minimal conflict a number of calls to a consistency checker that is linear in the number of system components. Hence, HS-Tree requires $O(|Sol_{HS}| + |COMPS| * n_C)$ consistency checks, one for each diagnosis \mathcal{D}_i (\checkmark label) to verify that it hits every conflict, plus $O(|COMPS|)$ calls per computed minimal conflict, where $COMPS$ denotes the set of system components and n_C the number of computed minimal conflicts throughout the execution of HS-Tree. Note, no theorem prover

³Observe that $path(n)$ must indeed be a *minimal* diagnosis in this case since the check in (a) was negative and due to the breadth-first tree construction, which guarantees that all nodes n' with $path(n') \subset path(n)$ must already have been processed (and all diagnoses of lower cardinality must already have been found) before n is labeled.

⁴Please refer to, e.g., [Rod22b] for a more elaborate discussion and a precise algorithmic description of HS-Tree.




Figure 1.4: Full adder example: Diagnosis computation using HS-Tree. Blue font...minimal conflicts; green shaded areas...paths that are minimal diagnoses; green font...node labels for minimal diagnoses; red font...node labels for non-minimal diagnoses; yellow circles...order of node processing (breadth-first); orange font...label indicating if conflict is freshly computed (C) or reused from stored conflicts (R).

calls are necessary to determine non-minimal diagnoses or duplicate paths (\times labels), cf. [Rod22b].

Example 1.3 (Diagnosis computation using HS-Tree) Reconsider our full-adder example from Fig. 1.1. The tree built by the HS-Tree algorithm for this problem is displayed by Fig. 1.4. The set of computed diagnoses is indeed equal to all minimal diagnoses (cf. the set \mathbf{D} discussed in Example 1.2). The algorithm requires two conflict computations (eight plus five consistency checks, cf. Example 1.4) and three additional consistency checks (to prove that the branches marked with \checkmark are diagnoses). For more complex systems with a large set of components, this conflict-based computation can make a huge difference as compared to a brute-force exploration of the diagnosis search space, as the size and number of minimal conflicts is often relatively small. Note, to compute only one preferred diagnosis, say, a minimum-cardinality diagnosis assuming the fewest abnormal components (in this case $[X_1]$), the computation of a single minimal conflict would suffice in this particular example. \square

Example 1.4 (Conflict computation using QuickXplain) The idea underlying conflict computation by means of QuickXplain can be illustrated as follows. Consider, e.g., the node labeled at time point ③ in HS-Tree (cf. Fig. 1.4). To determine the minimal conflict $\langle X_1, X_2 \rangle$ to label this node, QuickXplain will execute the Steps (1)–(5) shown below. Each step involves one consistency check using a suitable (sound and complete) theorem prover. First, note that the goal when labeling the node with the number ③ in the HS-Tree is finding a minimal conflict that has an empty intersection with the node’s path $\{A_2\}$. In other words, we need to determine a minimal conflict that is a subset of the set S including all components minus the node’s path, i.e., $S = \{A_1, O_1, X_1, X_2\} = \{A_1, A_2, O_1, X_1, X_2\} \setminus \{A_2\}$. So, in Step (1), QuickXplain checks whether a conflict is contained in S by testing whether the assumption that all the four components in S are normal is consistent with the observations (below we highlight the components assumed normal and taking part in the consistency check by a single underlining). The result is negative (i.e., “inconsistent”, shown by a \ominus) which means that there is at least one (minimal) conflict among the subsets of S . Let us refer to the minimal conflict that will be found by QuickXplain by \mathcal{C} . In Step (2), the algorithm partitions S into two equal-sized subsets and tests whether the left subset, $\{A_1, O_1\}$, is already a conflict. Since the normality assumption of these two components does not contradict the observations (i.e., we obtain “consistent”, signified by the \oplus), it can be inferred that some component from $\{X_1, X_2\}$ must be in \mathcal{C} . Hence, QuickXplain starts to analyze the right subset $\{X_1, X_2\}$ and recursively splits this set in two equally-sized subsets, $\{X_1\}$ and $\{X_2\}$, and adds the left set $\{X_1\}$ to the normally assumed components for the next consistency check performed in Step (3). We find that this check is positive (\oplus),

which implies that X_2 is in \mathcal{C} (denoted by a double underlining in the illustration below). Intuitively, X_2 must be in \mathcal{C} as it is known from Step (1) that $\{A_1, O_1, X_1, X_2\}$ contains a conflict, and from Step (3) that, once X_2 is removed, the set does not comprise any conflicts. In Step (4), a consistency test for (the normality assumption of) $\{A_1, O_1, X_2\}$ is executed to figure out whether the left subset, $\{X_1\}$, is not needed to build a conflict given the knowledge that the right subset, $\{X_2\}$, is in the conflict. However, we observe that consistency is given (\oplus), which entails that X_1 must be in \mathcal{C} as well (due to an analogue argumentation as given for X_2). Now, as it has been derived that both X_1 and X_2 are in \mathcal{C} , the algorithm tests in Step (5) if the left subset of S , i.e., $\{A_1, O_1\}$, is still relevant for the conflict in the light of this knowledge. Since the consistency check returns negatively (\ominus), QuickXplain deduces that both components A_1, O_1 are irrelevant to \mathcal{C} and returns the minimal conflict $\mathcal{C} = \langle X_1, X_2 \rangle$ as a final result.

We observe that the computation of \mathcal{C} takes five consistency checks. Likewise, we find that the computation of the conflict $\mathcal{C}' := \langle X_1, A_2, O_1 \rangle$ at node ① in Fig. 1.4, given the list of components $[A_1, A_2, O_1, X_1, X_2]$, requires eight theorem prover calls. \square

- (1) $[A_1, O_1, X_1, X_2] \ominus \rightarrow$ the set $\{A_1, O_1, X_1, X_2\}$ includes a conflict
- (2) $[A_1, O_1, X_1, X_2] \oplus \rightarrow$ some element from $\{X_1, X_2\}$ is in \mathcal{C}
- (3) $[A_1, O_1, X_1, X_2] \oplus \rightarrow X_2$ is in \mathcal{C}
- (4) $[A_1, O_1, X_1, \underline{X_2}] \oplus \rightarrow X_1$ is in \mathcal{C}
- (5) $[A_1, O_1, \underline{X_1}, \underline{X_2}] \ominus \rightarrow A_1, O_1$ are not in \mathcal{C}

1.2 Sequential Diagnosis

When a system malfunction is detected, there is in many cases not enough information available to locate the real cause unambiguously. Rather, one often faces a substantial number of different (minimal) diagnoses given the initial system observations and measurements. However, only one of the diagnoses, which we refer to as the *actual diagnosis*, pinpoints the (unknown) actually faulty components. All other diagnoses are spurious in that they assume at least one actually normal component to be faulty, or at least one actually faulty component to be normal, or both.

One strategy to handle multiple diagnoses is to adopt techniques to rank or restrict the computed diagnoses based on some informative criterion such as maximal probability or minimal cardinality [dK91]. Whereas such techniques can be very powerful if the actual diagnosis appears (early) in the solution list, they cannot give the user a guarantee that the actually faulty components will be located (efficiently).

This is where sequential diagnosis techniques [dKW87] enter the stage. Their aim is to collect additional information about the system in order to gradually rule out spurious diagnoses and thereby reduce the diagnostic uncertainty. More specifically, sequential diagnosis approaches iteratively pose *queries* to an oracle (e.g., an electrical engineer in case of a malfunctioning circuit, or a domain expert in case of a faulty knowledge base). At this, queries are selected in a way that each query answer elicits some new information about the diagnosed system that allows to discriminate between diagnoses. Depending on the system type, queries can capture, e.g., system tests (given specific inputs of the system, determine the resulting outputs), probes (make a measurement at a particular location in the system), or expert questions (answer a question about the intended behavior of the system) [Rod22i].

The basic technical idea behind queries is to exploit the fact that different diagnoses predict different system behaviors or properties (e.g., intermediate values or outputs). Observing a system aspect for which the predictions of multiple diagnoses disagree then leads to the invalidation of those diagnoses whose predictions are inconsistent with the query answer. Since the determination of optimal queries (which minimize the overall effort for the interacting oracle) is NP-hard [PA90], sequential diagnosis methods often have to rely on myopic strategies to estimate the favorability of different query candidates. One approach is to evaluate the quality of queries in terms of the utility of the expected situation after

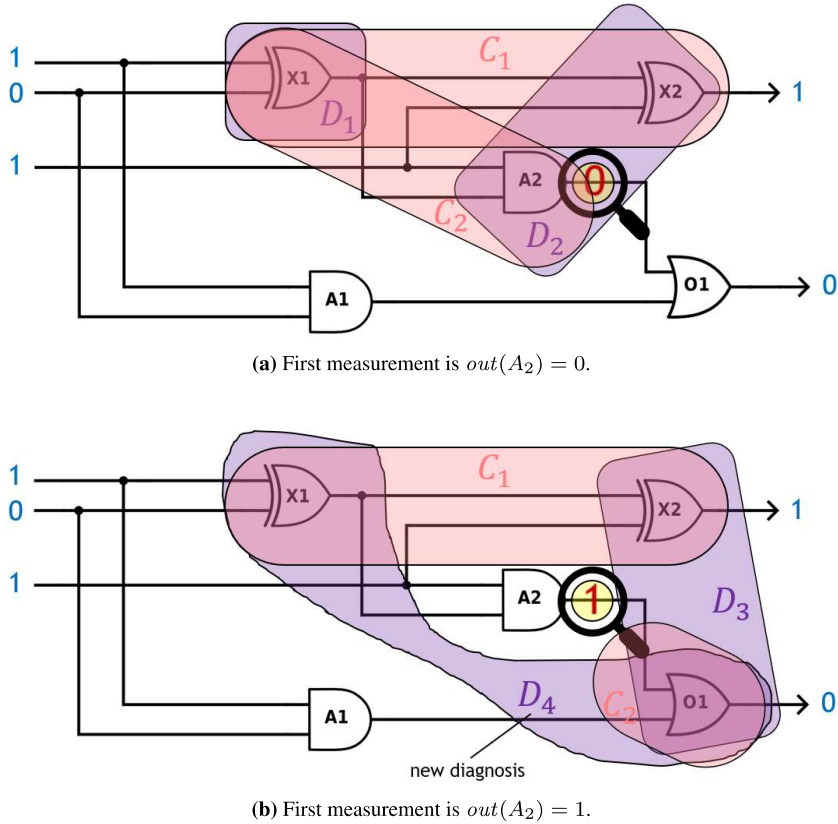


Figure 1.5: Full adder example: Sequential diagnosis, situation after the first measurement.

knowing their answer, taking into account features such as the number of invalidated diagnoses, or the estimated additional oracle effort. Such *one-step-lookahead analysis* [dKRS92] based on query selection heuristics has proven to constitute a particularly favorable trade-off between computational efficiency and diagnostic effectivity, and is today state-of-the-art in sequential diagnosis. Accordingly, a range of *heuristics* that quantify the goodness of queries based on various information-theoretic considerations have been proposed in literature [Mor82, dKW87, PA90, SFFR12, RSFF13, Rod16b, Rod17, RS18c].

A generic *sequential diagnosis process* can be roughly thought of as a recurring execution of

- (1) the computation of a set of (preferred) minimal diagnoses,
- (2) the selection of the most informative query according to some query selection heuristic based on the computed diagnoses,
- (3) the presentation of the computed query to the oracle which provides a query answer, and
- (4) the exploitation of the query outcome to refine the system knowledge.

This iterative process continues until sufficient diagnostic certainty is obtained, e.g., only one minimal diagnosis (with overwhelming probability) remains.

Example 1.5 (Sequential diagnosis) Reconsider our full-adder from Fig. 1.1. As expounded in Example 1.2, the initial system knowledge in terms of observations and measurements leaves the user (e.g., an electrical engineer) with three minimal diagnoses. Let us assume that in the past the failure rate of the or-gates has been registered to be higher than for other types of gates. In the light of this meta information, diagnosis $\mathcal{D}_3 = [X_2, O_1]$ would appear to be the most likely and plausible. However, let us also suppose that the intention is to avoid the potential cost involved in the unnecessary replacement of healthy gates as well as the cost of missing the replacement of faulty gates as a consequence of choosing a diagnosis

that turns out not to be the actual fault explanation. To this end, the user might decide to use a sequential diagnosis engine as an aid to efficiently locate the real fault with certainty.

Initiating the sequential diagnosis process given the known diagnoses $\mathbf{D} = \{\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3\}$ (cf. Example 1.2), a query selection engine might propose a query asking the engineer to measure the voltage at the wire $out(A_2)$ that connects the gates A_2 and O_1 . The rationale behind this query could be to differentiate between the most probable diagnosis \mathcal{D}_3 (being the only diagnosis including the or-gate) and the remaining diagnoses $\{\mathcal{D}_1, \mathcal{D}_2\}$. Note that this query has the property of eliminating one of the known diagnoses irrespective of the oracle's answer; that is, already before the query outcome is known, it is guaranteed that the query will bring relevant information for diagnosis discrimination. In various diagnostic scenarios, it is demanded of query selection engines to generate only such discriminating queries.

Let us now analyze the effects the two possible query answers, $out(A_2) = 0$ and $out(A_2) = 1$, have on the minimal conflicts and diagnoses. Given $out(A_2) = 0$, we learn that at least one of $\{X_1, A_2\}$ must be at fault, since assuming both of them to work nominally, the output of X_1 (inputs 1 and 0) would be 1, and the output of A_2 (inputs 1 and 1) would be 1 as well. Fig. 1.5a shows the implication of this reasoning, i.e., the “shrinkage” of the minimal conflict \mathcal{C}_2 to $\langle X_1, A_2 \rangle$, which causes the invalidation of diagnosis \mathcal{D}_2 (as can be easily seen by means of the hitting-set relationship between diagnoses and conflicts). If, on the other hand, the oracle measures $out(A_2) = 1$, the conflict \mathcal{C}_2 reduces to $\langle O_1 \rangle$ because we now know that one of its inputs is 1 (query answer) and its output is 0 (observation), which entails that O_1 must be abnormal. This means that, of the original minimal diagnoses, only \mathcal{D}_3 remains valid. However, notably, in this case we detect that a new diagnosis, $\mathcal{D}_4 = [X_1, O_1]$, has emerged as a superset of \mathcal{D}_1 . In general, it holds that each minimal diagnosis after the query answer is incorporated into the system knowledge is either one of the original minimal diagnoses or a proper superset of some original minimal diagnosis.

For any outcome we might have got for the first query, the debugger would suggest a measurement at $out(X_1)$ as a second query. Regardless of its outcome, diagnostic certainty will be obtained. To see this, let us consider all four cases: First, for the situation $out(A_2) = 0$ (Fig. 1.5a), we have:

- If $out(X_1) = 0$, the new minimal conflicts are given by $\mathbf{C} = \{\langle X_1 \rangle\}$. I.e., $\mathcal{D}_1 = [X_1]$ is the only remaining minimal diagnosis.
- If $out(X_1) = 1$, the new minimal conflicts are given by $\mathbf{C} = \{\langle X_2 \rangle, \langle A_2 \rangle\}$. I.e., $\mathcal{D}_2 = [X_2, A_2]$ is the only remaining minimal diagnosis.

Second, for the situation $out(A_2) = 1$ (Fig. 1.5b), we get:

- If $out(X_1) = 0$, the new minimal conflicts are given by $\mathbf{C} = \{\langle X_1 \rangle, \langle O_1 \rangle\}$. I.e., $\mathcal{D}_4 = [X_1, O_1]$ is the only remaining minimal diagnosis.
- If $out(X_1) = 1$, the new minimal conflicts are given by $\mathbf{C} = \{\langle X_2 \rangle, \langle O_1 \rangle\}$. I.e., $\mathcal{D}_3 = [X_2, O_1]$ is the only remaining minimal diagnosis.

Hence, the user is overall required to make two measurements in the circuit in order to determine the actual diagnosis with full certainty. \square

1.3 Diagnostic Tasks

In model-based diagnosis, there are three main tasks [vLP08, Chap. 10]:

- *fault detection* (is there a fault? / is some component faulty?),
- *fault localization* (where is the fault? / which components are faulty?), and
- *fault identification and repair* (what is faulty with each component? / how to fix it?).

Except for the case where faulty components are simply replaced by new ones, all these three steps are usually necessary for the successful debugging of an abnormal system, and they are executed in the order they are listed. Fault localization can be seen as a core phase in model-based diagnosis as it incorporates a fault detection step (localization is only relevant once it has been ascertained that faults are indeed present in the system), and its outcome represents one solution to the repair problem (in terms of a simple, but often still reasonable and efficient component replacement). Given a strong fault model, where the system description, besides the normal behavior, also specifies abnormal component behavior(s) called fault

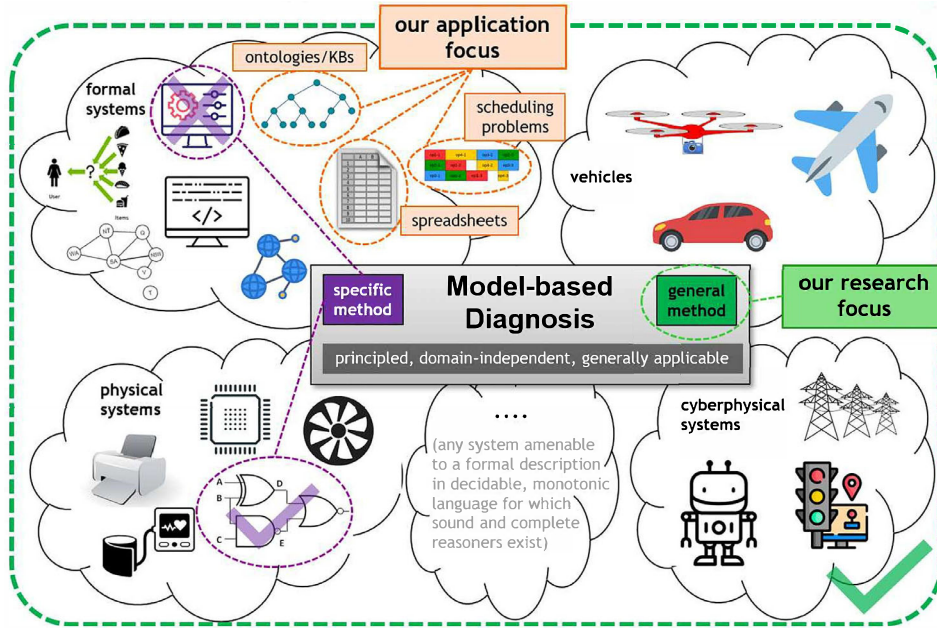


Figure 1.6: Model-based diagnosis application domains grouped into categories (*clouds*), differentiation between specific (*violet*) and general (*green*) methods, and illustration of research and application focus (*light green* and *light orange*) of the works included in this thesis. The *bold green dashed line* and *checkmark sign* indicate that general methods aim at preserving maximal generality in order to be (at least in principle) applicable to any model-based diagnosis application, which is not the case for specific methods (as exemplarily visualized by the *violet dashed line*).

modes, the fault identification task (find a fault mode assignment for each component that is consistent with the observations and measurements) might be directly tackled as a second step in the process, without first locating the fault. A more detailed discussion of this is beyond the scope of this work and can be found in [vLP08, Chap. 10]. In any case, the fault identification and repair step will usually require domain expertise in terms of a more detailed insight into the functioning, structure, construction, etc. of the components. The works included in this thesis cover all three above-mentioned tasks, however with a focus on the fault localization problem.

1.4 Application Areas

Model-based diagnosis can be and has been applied to a myriad of different application areas. As suggested by Fig. 1.6, we might categorize some of the prominent use cases into

- *formal systems* (e.g., software, spreadsheets, knowledge bases, ontologies, networks, scheduling problems, recommender systems, constraint satisfaction problems, or configuration systems),
- *vehicles* (e.g., drones, aircrafts, or cars),
- *cyberphysical systems* (e.g., power grids, traffic control systems, or robots), and
- *physical systems* (e.g., office machines, medical devices, circuits, hardware systems, or turbines).

While the provision of a suitable system model might sometimes be one of the most challenging steps towards diagnosing (e.g., physical) systems when relying on the “correct vs. actual” interpretation of the diagnosis problem (cf. Sec. 1.1), a readily usable system model can often be directly extracted from the buggy system in an automatic way when considering the “actual vs. intended” perspective on the problem. The latter strategy is often applicable to formal systems, which constitute the primary focus of our empirical evaluations conducted in the course of the works included in this thesis.

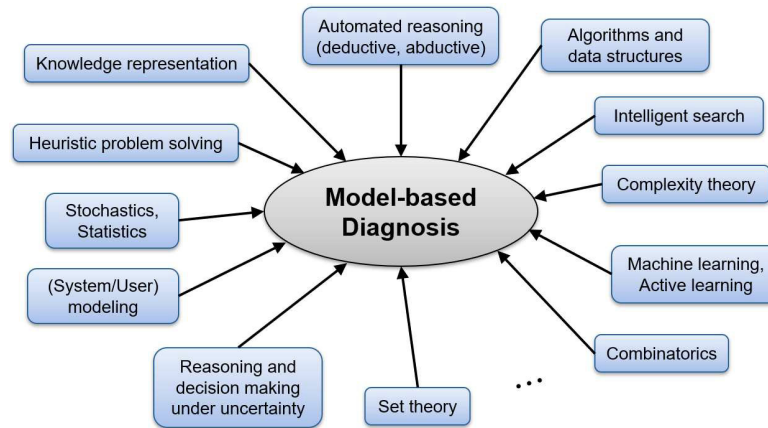


Figure 1.7: (Some) related subjects and fields that influence model-based diagnosis research.

In general, any system that is decomposable into a set of components and amenable to a formal description in a decidable monotonic knowledge representation language, for which sound and complete theorem provers (for consistency checking) exist, is a potential use case for model-based diagnosis.

Whereas model-based diagnosis is a generally applicable and domain-independent framework, this does not automatically mean that all implementations of the framework, i.e., model-based debugging approaches and tools, are built so as to maintain the maximal generality. As suggested in [Rod22d], we might discern specific and general diagnosis methods (cf. Fig. 1.6). The former are designed to harness specifics of a particular (class of) system type(s) in order to streamline the diagnosis process while accepting a somewhat narrowed down application scope of the approach. The latter, by contrast, aim at maintaining a maximally broad applicability by making no assumptions about the (structure, type, domain, etc. of the) diagnosed system, the language used to describe the system, or the inference mechanism used. The largest part of the works included in this thesis belong to the latter category and can in principle be generally applied to any model-based diagnosis problem as per Reiter's theory [Rei87]. As regards applications, our research works mainly target ontology and knowledge base debugging, spreadsheet diagnosis, as well as over-constrained scheduling problems (cf. Fig. 1.6).

1.5 Related Research Fields

Model-based diagnosis is a core discipline in model-based problem solving, which is embedded in the field of knowledge representation and reasoning, a large and important sub-area of artificial intelligence. By reason of its profound mathematical and theoretical foundations, the great variety of its applications, the consideration of hard combinatorial and set theoretic problems, the broad range of proposed algorithms, the diversity of possible diagnostic scenarios, the partial observability and uncertainty usually involved in the process, the usefulness of machine learning techniques to boost diagnostic performance, and many others, model-based diagnosis has intersections with a multitude of other research areas. It is both influenced by and does itself influence a range of other fields, where Fig. 1.7 and Fig. 1.8 give an overview of some of them.

Prominent influencing fields are, by the very nature of model-based diagnosis, of course knowledge representation, reasoning, and modeling techniques, as these fields deal with languages and ways to express system descriptions and with efficient automated deduction techniques on top of these formalisms for model-based consistency checking. Discrete-mathematical subjects such as combinatorics and set theory, as well as complexity theory, stochastics, and algorithms and data structures are central to theoretical analyses in the field. Moreover, given the NP-hardness and thus general intractability of the problems addressed in model-based diagnosis (e.g., diagnosis computation or query optimization), heuristic problem

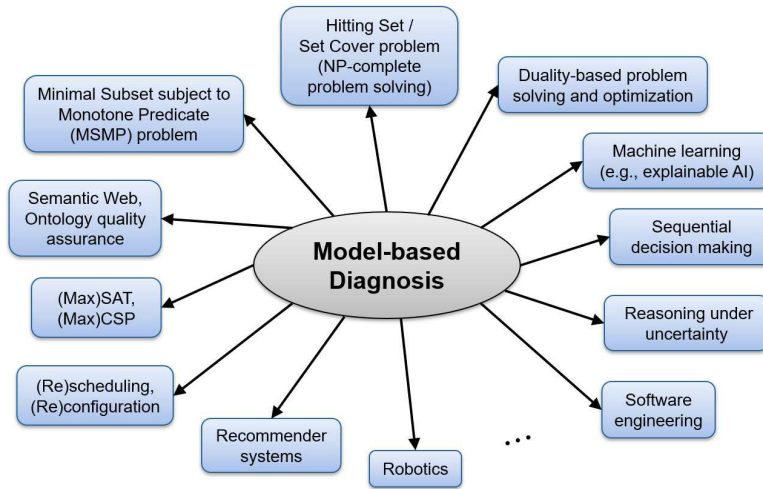


Figure 1.8: (Some) related subjects and fields that are influenced by model-based diagnosis research.

solving and intelligent search techniques as well as machine learning approaches can in some settings be vital to the development of practical and successful diagnosis methods. Finally, statistics and uncertain reasoning play a pivotal role in dealing with the problem of intelligent and reasonable decision making given only partial information.

On the other hand, also model-based diagnosis research has contributed to advancements in several fields. First, there are the various domains it has been successfully applied in, such as robotics, recommender systems, software engineering, the Semantic Web, ontology engineering, as well as scheduling and configuration problems. Apart from its application-related impact, model-based diagnosis has brought up substantial theoretical findings and algorithmic solutions that (can) provide valuable ideas and insights for, or propel the state-of-the-art of important research domains. Examples are the hitting-set (and the closely related set-cover) problem, which is one of the key problems studied in model-based diagnosis due to the hitting-set relationship between diagnoses and conflicts (cf. Sec. 1.1). Worth mentioning in this context is also the field of duality-based problem solving and optimization, as characterized in [Sla14]. Here, the problem of finding a (cost-)optimal set from a collection $\mathbf{X} = \{S_1, \dots, S_k\}$ of subsets S_i of some universe U is investigated, and suggested to be tackled by using best-first hitting set computation (e.g., [Rei87, Rod15]) over the dual collection \mathbf{X}^* of \mathbf{X} . At this, \mathbf{X}^* is defined as the collection including all subsets S_j^* of U whose complement $U \setminus S_j^*$ is not in \mathbf{X} . Since the collection of diagnoses is dual to the collection of conflicts, and vice versa, progress in conflict-based diagnosis computation can often be directly exploited for arbitrary duality-based computation or optimization problems. Other prominent notions linked by such a duality relationship are, e.g., minimal unsatisfiable subsets vs. minimal correction subsets of clauses in a CNF [BK15], which are highly relevant to SAT research [BHMW09], or conflicts vs. relaxations in the context of constraint satisfaction problems [Jun04]. Another interesting observation is that the computation of both minimal diagnoses and minimal conflicts is an instance of the so-called Minimal Subset subject to a Monotone Predicate (MSMP) problem [MJB13]. As a consequence, efficient methods for diagnosis and conflict calculation, if sufficiently general (cf. Sec. 1.4 and [Rod22c]), also constitute solutions to this more general MSMP problem. Instances of the latter include, e.g., minimal unsatisfiable subsets, minimal unsatisfiable cores, minimal correction subsets, prime implicants and implicates, as well as justifications and most concise optimal queries to an oracle [dKW87, Mar95, BL03, DHN06, Hor11, Rod22i]. Furthermore, techniques from model-based diagnosis are relevant to sequential decision making, reasoning under uncertainty, and machine learning. Particularly noteworthy in this regard is the more recent field of explainable AI which has been gaining increasing attention and interest throughout the last years and which is devoted to supply-

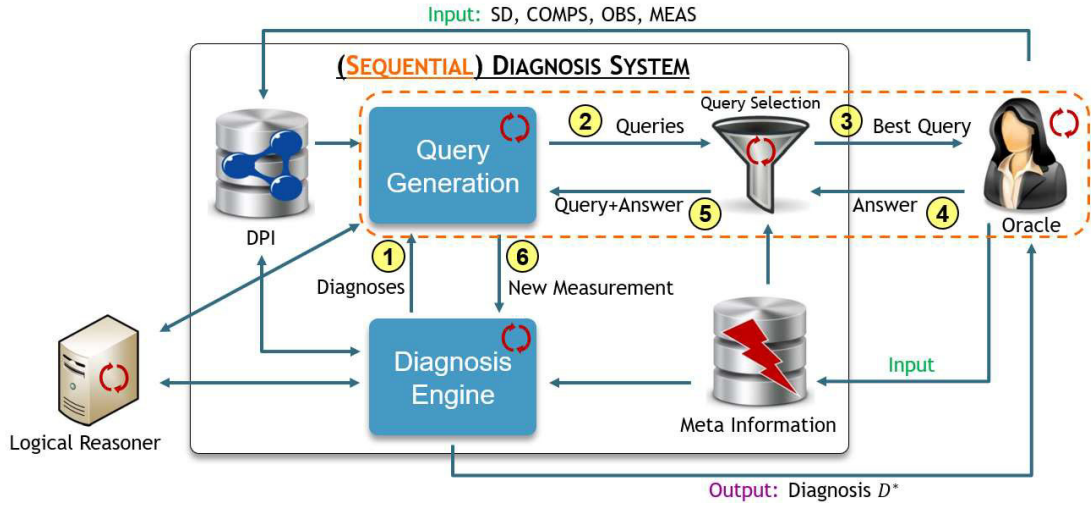


Figure 1.9: Generic (sequential) diagnosis system: Inputs, outputs, modules, stored information, external entities, and functionality. The orange dashed line shows the elements additionally present if the debugging system targets sequential diagnosis. A circle of two red arrows indicates that the operation executed by the respective entity is potentially expensive (e.g., wrt. time, memory, cost, or effort). The yellow circled numbers visualize the process flow; the depicted loop is reiterated until a predefined stop criterion is satisfied.

ing users of artificial intelligence and especially machine learning tools with understandable and useful explanations for generated outputs. Finally, the principles underlying model-based diagnosis approaches share several crucial parallels with other model-based problem solving techniques such as system design, (re)configuration, or failure-modes-and-effects analysis [vLP08, Chap. 10]. Thus, ideas and approaches for one area can be fruitful and beneficial also for the others.

1.6 A Generic (Interactive) Diagnosis System

In order to get a feel for the problems, challenges, and research goals in the context of (sequential) model-based diagnosis (discussed in Sec. 1.7), it is instructive to examine, on a sufficiently abstract level, the structure, modules, properties, and other aspects shared by most, if not all, (interactive) diagnosis tools and approaches. To this end, let us consider a generic diagnosis system, as depicted in Fig. 1.9:

Structure and external entities: On the structural level, we have the diagnosis system and two entities outside of the system, viz the oracle (or: user) and the logical reasoner. Within the system, we discern two databases (in the most abstract conceptual sense), one storing the current diagnosis problem instance (DPI), and the other storing any meta information relevant to the modules in the system or to the present diagnostic scenario or application. Often, the meta information encompasses, e.g., fault information (such as component failure rates), parameter settings for algorithms (e.g., stopping or optimization criteria), heuristics (e.g., for query selection), or general preferences or configurations of the diagnosis system (e.g., which type of queries to ask [Rod22f], or which type of problem to solve [RH18b]). In addition, the system incorporates three modules, viz a diagnosis engine, a query generation unit, as well as a query selection component. Note that, in case the diagnosis system is not intended for the purpose of sequential diagnosis, it does not need to include the query generation and query selection modules, and no communication with an oracle is necessary during the process (see the orange dashed line in Fig. 1.9).

Inputs and outputs (cf. green and violet font in Fig. 1.9): The inputs to a diagnosis system are, first, a diagnosis problem instance (DPI) capturing the initial state of knowledge about the diagnosed system in terms of the system description (SD), the components (COMPS), as well as the observations (OBS) and

measurements (MEAS), cf. Sec. 1.1. Technically, SD, OBS, and MEAS will each be a set of sentences in a specific formal (often: logical) language, such as propositional logic, predicate calculus, description logics, finite-state machines, differential equations, or constraints, and COMPS a set of the system components of interest to the diagnostic task. Second, the user might (often: optionally) provide available meta information as described above. The output of a diagnosis system can in general take manifold shapes, depending on the type, application context, purpose, goal, way of deployment, etc. of the diagnosis system as well as of the user's preferences (cf. [Rod22d]). Usually, the output will be either one (often: minimal or preferred) diagnosis, or an (often: small) set of diagnoses (often: sorted by some preference criterion).

Process flow (cf. yellow circled numbers in Fig. 1.9, and enumeration (1)–(4) in Sec. 1.2): Given the DPI and potentially some relevant meta information such as fault probabilities, a preference criterion, a number of diagnoses to compute, a timeout, or other algorithm-specific parameters, the first step, accomplished by the diagnosis engine, is the computation of a set of so-called *leading* minimal diagnoses [dKW89], which are often the most preferred diagnoses, i.e., the most probable or minimum-cardinality ones. If the system is not sequential, then these diagnoses might be directly presented to the user (perhaps in the form of a ranking) for further inspection or processing. Otherwise, in a second step, the leading diagnoses are forwarded to the query generation unit, which usually generates a set of query candidates. The complexity of this step can range from trivial (e.g., if the possible queries, such as probing points in the form of wires in a circuit, are immediately extractable from the system description [dKW87]) to very sophisticated (e.g., if queries, such as possible expert questions in ontology debugging, are implicit and computationally hard to extract [Rod22i]). In a third step, the query selection component filters the query candidates for the best (or a sufficiently good) query according to some (heuristic) measure given as a part of the meta information. The best query is then, in a fourth step, posed to the oracle, which performs according measurements, probes, tests, etc. in case of a tangible system under inspection, or thinks about the correct answer to a given expert question in case of a formal system to be debugged. The oracle is usually assumed to be capable or knowledgeable enough to answer an adequate fraction of the possible queries such that the sequential diagnosis session can reasonably proceed (after an acceptable number of queries are perhaps skipped and replaced by alternative ones). The provided answer of the oracle along with the asked query is then used to formulate a (logical) sentence to be added to the DPI in terms of a new measurement in the course of the fifth and sixth steps. This loop is continued until some stop criterion is met, e.g., an adequate degree of diagnostic certainty is achieved (cf. Sec. 1.2).

1.7 Research Challenges and Goals

Let us now take a look at the generic diagnosis system discussed in the previous section through the eyes of a researcher, by evaluating components and other aspects of the system regarding their influence on the *diagnostic performance* in terms of time, memory or other costs, and on the *diagnostic effectivity* in terms of the quality of the produced outputs. The result of this analysis is displayed in Fig. 1.10.

External entities: Both the calls to a logical reasoner and the consultations of an interacting oracle are usually among the most expensive “operations” throughout a diagnosis session. Since both these entities

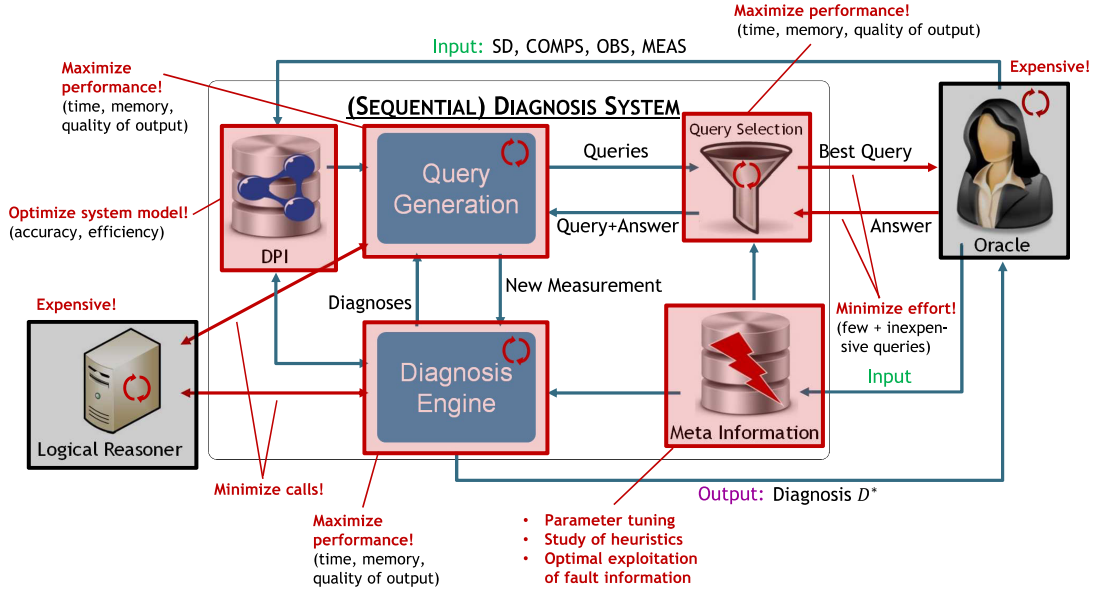


Figure 1.10: Generic (sequential) diagnosis system: Problems as well as research challenges and goals. *Black boxes* illustrate that the respective entity can (often) not be directly influenced by the diagnosis system, i.e., is considered a black-box with a defined output given a defined input, but unknown internals. The meaning of *red boxes* is that the respective entities can be actively shaped in that they are the result of design choices when building the diagnosis system, i.e., their internals can basically be implemented, organized and configured at will. *Thick red arrows* indicate that the communication described by this arrow is generally expensive. The research challenges and goals are depicted by *red-colored text*.

are frequently (and also in the works included in this thesis) viewed as black-boxes⁵, i.e., as functions with unknown internals that yield a defined output given a defined input, an important goal towards efficient model-based diagnosis is to minimize the communication with both the reasoner and the query-answering oracle (cf. the thick red arrows in Fig. 1.10).

Internal components: For the three modules performing computations, i.e., the diagnosis engine, the query generation, and the query selection, the objective is to maximize their time and memory performance as well as the the quality of their output, where the definition of quality might depend on, e.g., the user’s preferences, the meta information, the given diagnostic scenario, or the used algorithms. In addition, diagnosis-specific gains can also be achieved by leveraging modifications of the stored data, i.e., the meta information and DPI. Concerning the latter, since there are always multiple possible ways of modeling one and the same thing, one goal could be to find a representation of the diagnosed system that optimally balances diagnostic accuracy (reasonable number of generated diagnoses) with diagnostic efficiency (reasonable diagnosis computation time). Regarding the meta information, possible goals are,

⁵Regarding the logical reasoner, [Rod22d] gives an overview of some prominent diagnosis approaches, while also showing that most of them view the reasoner as a black-box. Alternatively, one can consider a reasoner as a so-called glass-box, where the idea is to streamline the reasoner’s internals towards more efficient computations relevant to the diagnostic task (e.g., by extracting conflicts as a byproduct of negative consistency checks). Studies contrasting black-box and glass-box techniques can be found, e.g., in [Hor11, Kal06]. Roughly, their conclusion is that the latter usually outperform the former wrt. computation time. On the other hand, advantages of black-box methods are their (1) robustness (no sophisticated, error-prone modifications of complex reasoning algorithms), (2) simplicity (internals of reasoner irrelevant), (3) flexibility (e.g., one can use a portfolio reasoning approach by switching to the most efficient reasoner in a simple plug-in fashion depending on the language used to describe the diagnosed system [RGH12]), and (4) up-to-dateness (which allows users to directly benefit from advances in general research on automated reasoning). Concerning the interacting oracle, it was shown that making suggestions to the user wrt. their answering behavior (e.g., by soliciting additional inputs or explanations) can bring a significant boost in terms of fault localization efficiency [Rod22f]. Similar to the case of the logical reasoner, this could be seen as an interpretation of the oracle as a “glass-box”. Nevertheless, presumably for simplicity, it appears that most methods in the literature regard the oracle as a black-box.

e.g., to find optimal configurations of the algorithms involved in the diagnosis process, optimal settings that suit a particular (group of) user(s) best, optimal heuristics (e.g., for query selection) to employ based on the given diagnosis scenario, or the optimal usage of the given (and possibly misleading [RSFF13]) fault information in the process.

Summary: These considerations boil down to the following research objectives (cf. red text in Fig. 1.10):

- *Diagnosis engine:* Maximize the performance wrt. time, memory and output quality.
- *Query generation:* Maximize the performance wrt. time, memory and output quality.
- *Query selection:* Maximize the performance wrt. time, memory and output quality.
- *DPI:* Optimize the system model wrt. diagnostic accuracy and efficiency.
- *Meta information:* Optimize the provided data for best diagnostic performance.
- *Logical reasoner:* Minimize the expense of communications with the reasoner.
- *Oracle:* Minimize the expense for the oracle wrt. time, effort, and other costs.

The research works included in this thesis consider all these research objectives, i.e., they cover all modules and aspects of the discussed generic diagnosis system. The contributions of these works will be explored in more detail in the upcoming chapters.

Chapter 2

Our Research: Overview

The aim of this chapter is to give the reader a concise overview of the conducted research in which the author of this thesis was involved. First, we approach this task in two orthogonal ways, by synopsising the research works (*i*) in Sec. 2.1, from the perspective of the aspects of a generic diagnosis system that are addressed, as well as (*ii*) in Sec. 2.2, regarding their type (e.g., whether the respective paper is theoretical, empirical, application-related, etc.). Finally, in Sec. 2.3, we focus our attention on the works included in this thesis and outline them from the two perspectives (i) and (ii), along with other relevant information.

2.1 From the Diagnosis System Perspective

Recalling the generic diagnosis system explicated in Sec. 1.6 and the associated research problems discussed in Sec. 1.7, Tab. 2.1 enumerates our research works grouped by the diagnosis system aspect (and related research objective) that was addressed in the respective paper. Note that several papers consider multiple system aspects and thus occur in different groups. As the table shows, we have been active along various lines of research, and we devoted our research to all research objectives identified above. This is a result of our holistic perspective on the model-based diagnosis field and of our overarching goal to advance the state-of-the-art of diagnosis systems along different dimensions.

2.2 From the Type of Research Perspective

To be able to better categorize and assess the conducted research, it can be helpful to view the works from another different angle. To this end, Tab. 2.2 lists the research works grouped by the type of the respective paper. At this, we discern

- papers proposing novel methodologies or algorithms (*Methodology/Algorithm*),
- application papers which are—withstanding the generality of the suggested approaches, cf. Sec.1.4—devoted to one particular application of model-based diagnosis, (*Application*),
- explanatory papers of didactic or survey nature (*Explanatory/Didactic/Survey*),
- empirical papers including comprehensive experimental evaluations or user studies (*Empirical*),
- papers including substantial theoretical results such as algorithm derivations, correctness proofs or complexity analyses (*Theoretical*),
- critical papers that fundamentally challenge the state-of-the-art (*Critical*), and
- tool papers describing, analyzing, evaluating, and differentiating debugging tools (*Tool*).

The table demonstrates that our research activities not only consider different research avenues, but have also produced a heterogeneous collection of works in terms of the nature of the publications.

Aspect of Diagnosis System	Research Works where Aspect is Addressed
Diagnosis Engine	[HJN ⁺ 22, Rod22a, Rod22b, Rod22c, Rod22d, Rod22e, Rod22g, Rod21a, Rod21b, RTJ21, Rod20a, Rod20c, Rod20d, Rod20e, RE20, RT20, Rod19, RJSF19, RH18a, RH18b, SRS18a, Rod16a, Rod15, SFRF14a, SFRF14c, SFRF12]
Query Generation	[Rod22a, Rod22f, Rod22i, Rod20b, RE19a, RE19b, RE19c, Rod17, RSS17a, RSS17b, Rod16a, Rod16b, Rod15, SFFR12]
Query Selection	[Rod22a, Rod22f, Rod22g, Rod22i, RE20, RE19b, RS18a, RS18c, Rod17, RSS17a, RSS17b, Rod16b, Rod15, SFRF14b, SFRF14c, RSFF13, RSFF12, SFFR12, SFRF12]
DPI	[HJN ⁺ 22, Rod22b, Rod21a, Rod20c, RJSF19, RH18b, SRS18a, SRS ⁺ 18b, SRS ⁺ 18c, RS17]
Meta Information	[Rod22a, Rod22b, Rod22d, Rod22e, Rod22f, Rod22g, Rod22h, Rod22i, Rod21a, RE20, RJSF19, RE19c, RS18a, RS18c, SRS18a, Rod17, Rod15, SFRF14b, SFRF14c, RSFF13, RSFF12, SFFR12, SFRF12, RSFF11]
Logical Reasoner	[Rod22b, Rod22c, Rod22d, Rod21a, RTJ21, Rod20c, RT20]
Oracle	[Rod22f, RE19a, RE19b, RE19c, RJSF19, SRS18a, SRS ⁺ 18b, SRS ⁺ 18c]

Table 2.1: Research works grouped based on the addressed aspect(s) of a generic diagnosis system. Per group (i.e., table row), works are arranged chronologically and sorted by author name(s). All works with year tokens larger or equal to “16” have been carried out during the author’s habilitation period; we also include earlier works in the lists in order to give a more complete, reasonable and illustrative picture of the author’s research.

Type of Work	Research Works of this Type
Methodology/Algorithm	[HJN ⁺ 22, Rod22a, Rod22b, Rod22d, Rod22e, Rod22f, Rod22g, Rod22h, Rod22i, Rod21a, Rod21b, RTJ21, Rod20a, Rod20b, Rod20c, Rod20d, Rod20e, RE20, RT20, Rod19, RE19a, RE19c, RJSF19, RH18a, RH18b, Rod17, RSS17a, RSS17b, Rod16b, RSFF13, RSFF12, SFRF12, RSFF11]
Application	[HJN ⁺ 22, Rod22f, RTJ21, RT20, RE19a, RE19c, RJSF19, RS18c, SRS18a, SRS ⁺ 18b, SRS ⁺ 18c, RSFF12, SFFR12, SFRF12, RSFF11],
Explanatory/Didactic/Survey	[Rod22c, Rod22d, SRS18a, SRS ⁺ 18c, Rod17]
Empirical	[HJN ⁺ 22, Rod22a, Rod22b, Rod22e, Rod22f, Rod22g, Rod22h, Rod22i, Rod21a, Rod21b, RTJ21, Rod20c, Rod20d, Rod20e, RE20, RT20, RE19a, RE19c, RJSF19, RS18a, RH18a, RH18b, RS18c, Rod17, SFRF14c, RSFF13, SFFR12, SFRF12, RSFF11]
Theoretical	[HJN ⁺ 22, Rod22a, Rod22b, Rod22c, Rod22d, Rod22e, Rod22f, Rod22h, Rod22i, Rod21a, RTJ21, Rod20c, RT20, RH18b, Rod17, RS17, RSS17a, RSS17b, Rod16a, Rod16b, Rod15, SFRF14c, RSFF13, SFFR12, SFRF12, RSFF11]
Critical	[Rod22f, Rod22g, RJSF19]
Tool	[Rod22d, RJSF19, SRS18a, SRS ⁺ 18b, SRS ⁺ 18c]

Table 2.2: Research works grouped by type of paper. Per group (i.e., table row), works are arranged chronologically and sorted by author name(s). All works with year tokens larger or equal to “16” have been carried out during the author’s habilitation period; we also include earlier works in the lists in order to give a more complete, reasonable and illustrative picture of the author’s research.

2.3 Works Included in this Thesis: In a Nutshell

Let us now have a look at the summary of the research works included in this thesis given by Tab. 2.3. The table can be read row-wise to learn about the main focus and type of a particular paper, and column-wise to find the papers that address a particular system aspect or are of a particular type. Note that many of the papers consider more than the system aspects mentioned in the table, but for clarity the table indicates only the topics that constitute the papers’ emphasis (please see Tab. 2.1 and for an overview that considers also secondary matters investigated by the papers).

Research Work Included in Thesis					Diagnosis System Aspects in Focus							Type of Work						
Paper	Cat.	Venue	CA	SA	Diagnosis Engine	Query Generation	Query Selection	DPI	Meta Information	Logical Reasoner	Oracle	Methodology/Algo.	Application	Explan./Didact./Survey	Empirical	Theoretical	Critical	Tool
[Rod22c]	A	Artificial Intelligence	✓	✓	✓							✓			✓	✓		
[Rod22f]	A	Knowledge-Based Systems	✓	✓		✓	✓				✓	✓	✓		✓	✓	✓	✓
[Rod22b]	A	Information Sciences	✓	✓	✓			✓		✓		✓			✓	✓		
[Rod22c]	A	Artificial Intelligence Review	✓	✓	✓					✓				✓		✓		
[Rod22i]	(A)	Artificial Intelligence (*)	✓	✓		✓	✓					✓			✓	✓		
[HJN ⁺ 22]	(A)	Artificial Intelligence (**)	✓		✓			✓				✓	✓		✓	✓		
[Rod22g]	A	AAAI	✓	✓	✓				✓			✓			✓		✓	
[Rod22d]	C	DX	✓	✓	✓							✓		✓	✓	✓		✓
[RTJ21]	A	KR	✓		✓					✓		✓	✓		✓	✓		
[Rod20c]	A	ECAI	✓	✓	✓			✓		✓		✓			✓	✓		
[RJSF19]	A	Knowledge-Based Systems	✓		✓	✓	✓	✓	✓	✓		✓	✓		✓		✓	✓
[RH18b]	A	SoCS	✓		✓			✓				✓			✓	✓		
[RS18c]	A	RuleML+RR	✓				✓		✓						✓			
[SRS18a]	B	FoKS	✓		✓	✓	✓	✓	✓	✓			✓	✓				✓
[Rod17]	C	DX	✓	✓				✓	✓			✓		✓	✓	✓		
[RS17]	C	DX	✓					✓										✓

Table 2.3: Overview of the research works included in this thesis (*left table sector*), based on the aspect of a generic diagnosis system *mainly* addressed in the paper (*middle table sector*) and the paper type (*right table sector*). Meaning of the columns:

(*Venue*) quotes the publication venue of the respective paper, where journal names are given in full while conference names are abbreviated to keep the table size manageable (the reader is kindly asked to refer to the bibliography for the full names); a single asterisk (*) means that the work is currently under revision, and the double asterisk (**) signifies that the work will be submitted shortly.

(*Cat.*) gives the rating of the publication venue of the paper (according to its classification in the expert-curated public research database “FoDok” of the University of Klagenfurt), where “A” refers to top-tier journals (Q1 on Scimago Journal & Country Rank) or top-tier international conferences (acceptance rate below 0.33), “B” means good journals (Q2 on Scimago Journal & Country Rank) or good international conferences (acceptance rate below 0.66), while all works not falling in categories A or B are classified by “C”; the rating is parenthesized for not yet published works.

(*CA*) indicates whether the author of this thesis is a corresponding author of the respective paper.

(*SA*) marks whether the author of this thesis is the single author of the respective work.

Chapter 3

Our Research: Details

In this chapter, we give some more details about the research papers included in this thesis. We thematically subdivide the discussion into five topics, each discussed in one of the Sections 3.1–3.5. Each of these sections is organized into subsections, one for each publication. The description of each paper is structured into three parts. First, we provide one figure for each publication which illustrates the aspects of the generic diagnosis system addressed as well as the research goals aimed at (cf. Figures 3.2–3.16). The legend pertinent to these figures is shown in Fig. 3.1. Second, we treat its contents, where we briefly review the paper’s highlights by sketching the paper’s motivation, contributions as well as results.⁶ Third, we provide some meta information about each work including the publication venue, its scientific classification⁷ (cf. Tab. 2.3), the contributions of the author of this thesis to the paper as well as other publications on the topic where the author of this thesis was involved. All papers covered in this chapter can be found in Appendix A.



Figure 3.1: Meaning of the accentuations used in Figures 3.2–3.16.

3.1 Diagnosis Computation

3.1.1 Memory-Limited Diagnosis Computation

Publication: System Aspects and Goals Addressed See Fig. 3.2.

Publication: Contents in a Nutshell

Motivation: Various model-based diagnosis scenarios require the computation of the most preferred minimal diagnoses. These diagnoses can be helpful, e.g., for a reasonable decision making and a well-founded early termination in the course of a sequential diagnosis process [Rod22g]. Existing algorithms that are sound (i.e., output only minimal diagnoses) and complete (i.e., can return all minimal diagnoses), however, require exponential space to achieve this task. This can prevent the application of these techniques to diagnosis settings which, e.g., pose substantial memory requirements on the diagnostic methods, or suffer from too little memory.

Contributions: To enable successful diagnosis both on memory-restricted devices and for memory-intensive problem cases, we propose the RBF-HS algorithm for diagnosis computation which combines two seminal search methods, Reiter’s hitting set computation strategy HS-Tree [Rei87] and Korf’s path-finding technique Recursive Best-First Search (RBFS) [Kor93]. RBF-HS is the *first* diagnosis search

⁶Please find animated presentations of some of the discussed works on <http://isbi.aau.at/ontodebug/publications>.

⁷For each journal publication, we provide the journal’s quartile in the Scimago Journal & Country Rank as well as its Scimago Journal Rank Indicator (cf. <https://www.scimagojr.com/>). To generate the ranking (for the year 2021), we used “Computer Science” as subject area, and “Artificial Intelligence” as subject category.

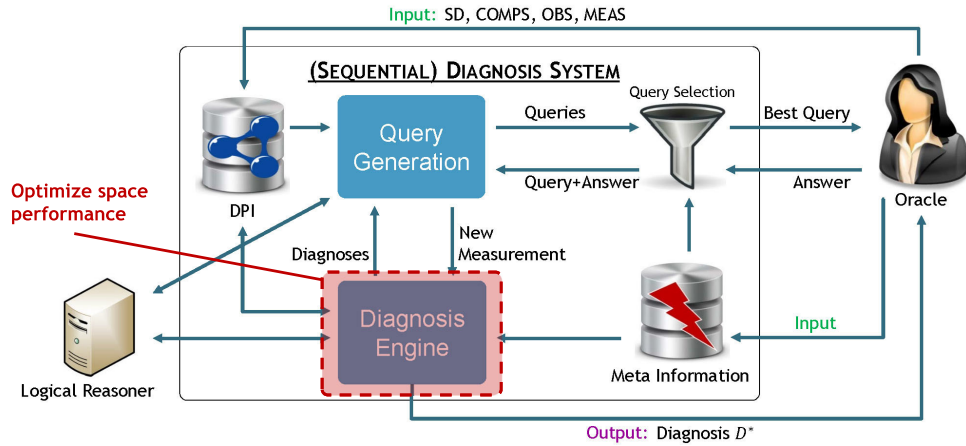


Figure 3.2: Memory-limited diagnosis computation: Addressed system aspects and goals.

that is sound, complete, best-first, generally applicable (to any diagnosis problem as per Reiter’s theory [Rei87]), and *linear-space*. In addition, we present HBF-HS, a hybrid between RBF-HS and HS-Tree, the idea behind which is to find a trade-off between runtime optimization and a restricted space consumption that does not exceed the available memory. Notably, both suggested approaches are not restricted to diagnosis problems, but applicable to best-first hitting set computation in general, which is relevant to various research and application domains beyond the frontiers of model-based diagnosis (cf., e.g., [Sla14]).

Results: We conducted comprehensive experiments on real-world benchmarks where we compared RBF-HS and HBF-HS to HS-Tree, the most popular algorithm with the same properties (soundness, completeness, best-first property, general applicability). First, for the task of computing minimum-cardinality diagnoses, we found that

- in most cases, RBF-HS substantially reduces memory requirements by up to several orders of magnitude while exhibiting no more than marginal runtime losses, if any.
- in more than a third of the cases, RBF-HS even achieved *both* runtime *and* memory improvements. E.g., in one case we observed 65% runtime savings coupled with a 99.9% memory reduction; in another case, HS-Tree ran out of memory (32 GB) while RBF-HS never required the storage of more than the comparably negligible number of 125 tree nodes.
- whenever the runtime of RBF-HS was significantly higher than that of HS-Tree, the use of HBF-HS could level this overhead while still reasonably limiting the used memory.
- RBF-HS scales to large numbers of computed diagnoses and to problems involving high-cardinality minimal diagnoses.

Second, when computing minimal diagnoses in descending order of probability, we observe that RBF-HS tends to trade memory savings more or less one-to-one for runtime overheads (which has well-understood theoretical reasons we discuss in the paper). Again, HBF-HS proves to be a reasonable remedy to cut down the runtime while complying with practicable memory bounds.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[Rod22e]
<i>Authors:</i>	Patrick Rodler
<i>Title:</i>	Memory-limited model-based diagnosis
<i>Year:</i>	2022
<i>Publication Venue (Type):</i>	Artificial Intelligence (Journal)
<i>Venue Metrics:</i>	A (Scimago: Q1, Journal Rank Indicator 1.673)
<i>Own Contributions:</i>	Single-author publication
<i>Other Publications on the Topic:</i>	[Rod21b] (Int'l Symposium on Combinatorial Search) [Rod20d] (Int'l Workshop on Principles of Diagnosis) [Rod22h] (Technical report)

3.1.2 DynamicHS: Optimizing Reiter's Hitting-Set Tree for Sequential Diagnosis

Publication: System Aspects and Goals Addressed See Fig. 3.3.

Publication: Contents in a Nutshell

Motivation: To reason about the best next query, sequential diagnosis methods usually require a sample of diagnoses at each step of the iterative diagnostic process. The computation of this sample can be accomplished by various diagnostic search algorithms [Rod22d]. Among those, Reiter's HS-Tree [Rei87] is one of the most popular due to its desirable properties and general applicability. However, HS-Tree *per se* does not encompass any specific provisions for being employed in an iterative way. Usually, it is used in a stateless fashion throughout the diagnosis process to (re)compute a sample of diagnoses per iteration, each time given the latest (updated) system knowledge including all so-far collected measurements. At this, the built search tree is discarded between two iterations, albeit often large parts of the tree have to be rebuilt in the next iteration, involving redundant operations and calls to costly reasoning services. Already Raymond Reiter, in his seminal⁸ paper [Rei87] from 1987, pointed to the research problem of devising a reasonable strategy to reuse an existing hitting-set search tree to compute diagnoses after new system information is obtained. This very issue remained unanswered for decades.

Contributions: Addressing Reiter's longstanding open question, we propose DynamicHS, a variant of HS-Tree that maintains state throughout the diagnostic session and embraces special strategies to minimize the number of expensive reasoner invocations. The main objective of this new algorithm is to allow for more efficient computations in sequential diagnosis than HS-Tree while maintaining *all* the advantages (general applicability, soundness, completeness, best-first property) of the latter.

Results: In extensive experiments, we used a corpus of real-world benchmarks to compare DynamicHS with HS-Tree in various sequential diagnosis scenarios (considering different numbers of diagnoses to be computed per iteration, and various employed query selection heuristics). The main insights were:

- DynamicHS is superior to HS-Tree in terms of computation time in 99.4% of the investigated diagnostic scenarios. Roughly, these savings are achieved by trading less time (fewer redundant operations and reasoner calls) for more space (statefulness), where the additional memory required by DynamicHS was reasonable in the vast majority of the scenarios; and, whenever HS-Tree was applicable in our experiments in terms of memory requirements, DynamicHS was so as well.
- The average runtime savings over HS-Tree achieved by DynamicHS are substantial and statistically significant in most scenarios, and reach median and maximal values of 52% and 75%. That is, HS-Tree requires up to an average of four times the computation time of DynamicHS in the tested diagnosis scenarios.
- In single diagnosis sessions, we observed that it took HS-Tree up to more than nine times as much time as DynamicHS—notably, while both methods always compute the same diagnostic solutions.

⁸On Google Scholar, the paper boasts almost 4500 citations as of December 2022.

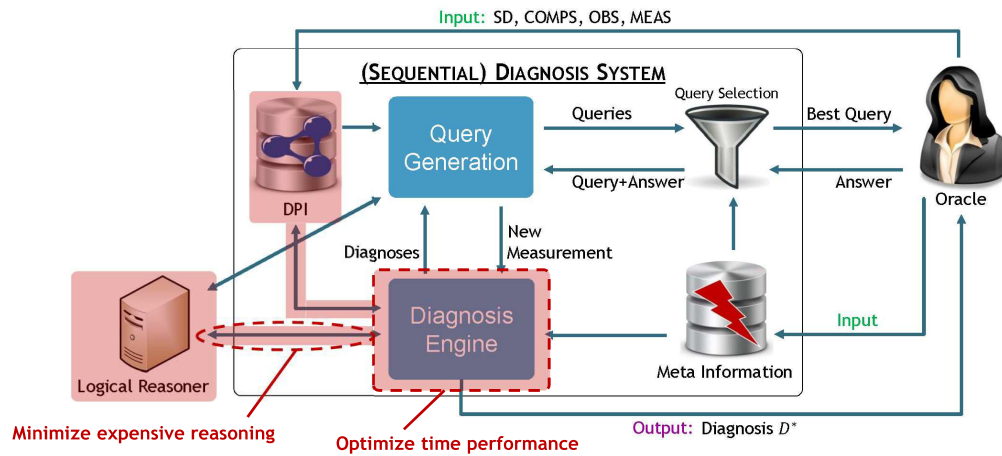


Figure 3.3: DynamicHS—Streamlining Reiter’s hitting-set tree for sequential diagnosis: Addressed system aspects and goals.

- The median runtime savings of DynamicHS per scenario appear to be neither dependent on the number of diagnoses computed nor on the measurement selection heuristic used.
- Considering the hardest cases per diagnosis scenario, which were up to one order of magnitude harder than the average cases, the time savings obtained by means of DynamicHS are even more substantial than on average, and reach median and maximal values of 64 % and 89 %, respectively.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[Rod22b]
<i>Authors:</i>	Patrick Rodler
<i>Title:</i>	DynamicHS: Streamlining Reiter’s hitting-set tree for sequential diagnosis
<i>Year:</i>	2022
<i>Publication Venue (Type):</i>	Information Sciences (Journal)
<i>Venue Metrics:</i>	A (Scimago: Q1, Journal Rank Indicator 2.290)
<i>Own Contributions:</i>	Single-author publication
<i>Other Publications on the Topic:</i>	[Rod20c] (European Conference on Artificial Intelligence) [Rod20e] (Int’l Symposium on Combinatorial Search) [Rod20a] (Int’l Workshop on Principles of Diagnosis) [Rod21a] (Technical report)

3.1.3 StaticHS: Reducing User Interaction Costs by Solving a Generalized Sequential Diagnosis Problem

Publication: System Aspects and Goals Addressed See Fig. 3.4.

Publication: Contents in a Nutshell

Motivation: In sequential diagnosis, new measurements are successively added to an initially given DPI. Each such acquisition of new information leads to a change of the solution space of minimal diagnoses. In fact, as we touched upon in Example 1.5, each minimal diagnosis for the newly obtained DPI is either equal to or a proper superset of some minimal diagnosis for the original DPI. In other words, any “new” minimal diagnosis emerging throughout the sequential diagnosis process assumes faulty strictly more components than some initial minimal diagnosis. In many real-world applications such as physical

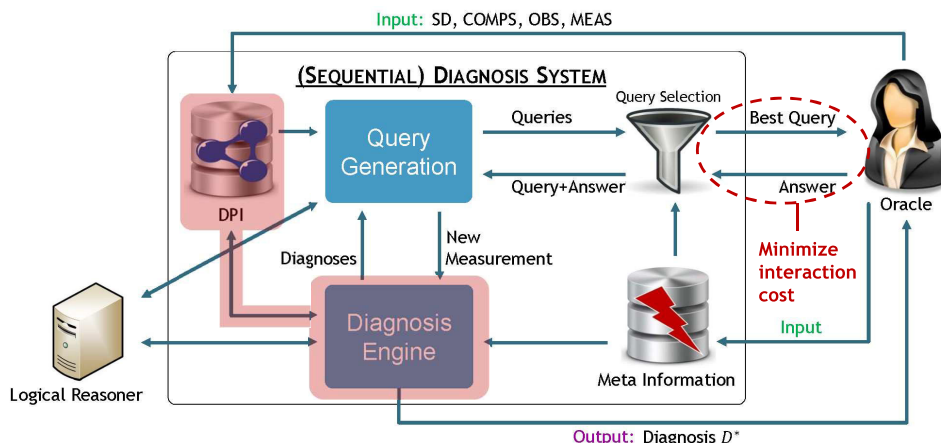


Figure 3.4: StaticHS—Reducing user interaction costs by solving a generalized sequential diagnosis problem: Addressed system aspects and goals.

devices, however, components are usually much more likely to be normal than at fault (at a given point in time). Thus, there is a high chance of the actual diagnosis (pinpointing the actually faulty components) being among the minimal diagnoses for the original input DPI. In such scenarios, one might—as a first step—only want to explore the initial diagnosis solution space, and neglect all “new” minimal diagnoses arising after DPI transitions.

Contributions: We suggest an alternative interpretation of the sequential diagnosis problem, called StatSD, where new measurements are not used to formulate a new DPI (as in the standard interpretation that we refer to as DynSD for disambiguation in this section), but only as constraints on the initial solution space of minimal diagnoses. The goal is to continue gathering information until only one (sufficiently likely) minimal diagnosis remains among the solutions of the initial DPI (and not the current DPI including all collected measurements as in DynSD). As a study of the literature revealed, existing sequential diagnosis methods focus only on DynSD.

As a consequence, we propose StaticHS, a novel sound and complete diagnosis search that can generate minimal diagnoses in best-first order, is suitable to solve both StatSD and DynSD, and is as generally applicable as Reiter’s HS-Tree [Rei87]. We theoretically prove that, if the actual diagnosis is among the minimal diagnoses for the originally given DPI, then for any set of oracle queries required to solve the DynSD problem (e.g., by means of HS-Tree), a subset of these queries already suffices to locate the actual diagnosis when the StatSD problem is considered and solved (e.g., by means of StaticHS) instead. Hence, if we could ensure that the same queries are asked regardless of whether StatSD or DynSD is solved (which we generally cannot since the diagnoses available to guide query selection may be different for both problems), then solving StatSD would be proven more efficient than DynSD given the preconditions mentioned above.

To solve DynSD, StaticHS will solve a series of StatSD problems to obtain the same solution diagnosis as any standard sequential diagnosis algorithm geared towards DynSD. At this, different strategies (influencing performance, not correctness) can be taken defining which specific StatSD problems are tackled. One such strategy is to use a new DPI whenever a new measurement is obtained, in which case StaticHS will act as any standard sequential diagnosis technique. Consequently, (1) StaticHS is a generalization of other techniques, such as HS-Tree, in the context of sequential diagnosis, and (2) solving multiple StatSD problems is a more general strategy than solving one DynSD problem.

Results: As somewhat expected based on the theoretical findings, empirical examinations using real-world problems revealed that StaticHS, compared to a standard sequential diagnosis algorithm addressing DynSD (in our tests w.l.o.g. implemented by HS-Tree), reduces the required number of queries substan-

tially both when (a) tackling the StatSD problem suffices (actual diagnosis is among the initial minimal diagnoses) and when (b) the consideration of the DynSD problem is necessary (actual diagnosis is not among the initial minimal diagnoses). In numbers, the avg. / max. saved measurement effort for the oracle in scenarios (a) and (b) amounted to 20 % / 65 % and 21 % / 65 %, respectively. Considering the individual executed sequential diagnosis sessions, the significant superiority of StaticHS as opposed to a method attacking DynSD is manifested in terms of a lower / same / higher number of queries in 76 % / 21 % / 3 % of the cases. These savings are a result of StaticHS's ability to combine a search space reduction (StatSD) with completeness (DynSD). As to the average algorithm reaction time (oracle waiting time), StaticHS outperformed HS-Tree for a half of the benchmarks, while it was slower for the other half. However, whenever StaticHS exhibited a higher reaction time, which in absolute terms was always less than 5 sec, this was compensated by very few queries StaticHS required to locate the actual diagnosis.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[RH18b]
<i>Authors:</i>	Patrick Rodler, Manuel Herold
<i>Title:</i>	StaticHS: A variant of Reiter's hitting set tree for efficient sequential diagnosis
<i>Year:</i>	2018
<i>Publication Venue (Type):</i>	Int'l Symposium on Combinatorial Search (Conference)
<i>Venue Metrics:</i>	A
<i>Own Contributions (est.):</i>	Concept/Idea: 100 % Theory/Algorithms: 100 % Development/Implementation: 30 % Evaluation: 80 % Manuscript: 100 %
<i>Other Publications on the Topic:</i>	[RH18a] (Int'l Workshop on Principles of Diagnosis)

3.1.4 On the Impact of Diagnosis Computation Strategies on Diagnostic Decision Making

Publication: System Aspects and Goals Addressed See Fig. 3.5.

Publication: Contents in a Nutshell

Motivation: Statistical samples, in order to be representative, have to be drawn from a population in a random and unbiased way. Nevertheless, it is common practice in the field of model-based diagnosis to make estimations from (biased) best-first samples. One example is the computation of a few most probable diagnoses for a defective system and the use of these to assess which aspect of the system, if measured, would bring the highest information gain.

Contributions: In this work, we scrutinize whether these statistically not well-founded conventions, that both diagnosis researchers and practitioners have adhered to for decades, are indeed reasonable. To this end, we empirically analyze six diagnosis computation methods for generating samples of minimal diagnoses: one best-first strategy yielding the most probable diagnoses, one outputting a random sample of diagnoses, and, as a baseline, a worst-first technique which returns the least probable diagnoses; moreover, we suggest an efficient heuristic method based on [SFRF14c] that can be configured to produce approximations of the best-first, the random, and the worst-first samples. The particular focus of our evaluations is on the comparison between (statistically unfounded, but commonly used) best-first and (statistically well-founded, but not commonly adopted) random sampling approaches. Specifically, we

- study the samples' (a) "theoretical" representativeness ("accuracy") in terms of how well they allow to estimate features (eliminated diagnoses, answer probabilities) of queries taken into account

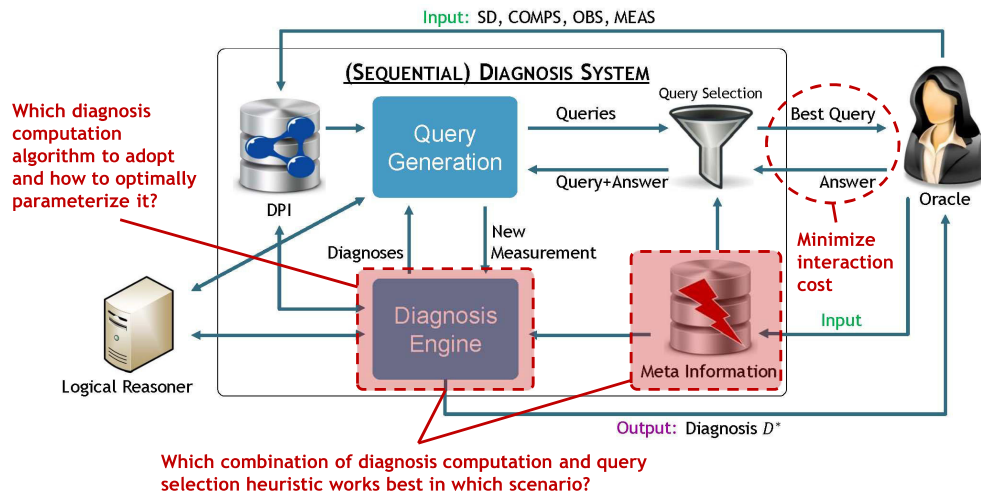


Figure 3.5: On the impact of diagnosis computation strategies on diagnostic decision making: Addressed system aspects and goals.

by query selection heuristics to suggest the best next query, as well as their (b) “practical” representativeness (“efficiency”) in terms of how many queries and how much time are needed to locate the fault in case the particular sample type is used,

- investigate the impact of the used sample size (number of computed diagnoses),
- examine the optimal trade-off between sampling efficiency (diagnosis computation time) and effectivity (effort until the actual diagnosis is located), and
- explore how the approximate sampling techniques (which do not guarantee a specific set of diagnoses to be computed) compare to the exact ones (which provably generate a precisely defined subset of all diagnoses).

Results: The bottom line of our investigations is: Somewhat surprisingly, the best-first samples including the most probable diagnoses commonly used in the field proved to be the best choice in a large fraction of the investigated cases. Yet, we also find that, for certain configurations of a diagnosis system, best-first samples imply drastic overheads compared to other sample types. Random samples, though enabling highly reliable estimations, often led to a worse diagnostic efficiency than biased ones. We discuss reasons for this phenomenon and make recommendations which configurations wrt. diagnosis computation algorithm, query selection heuristic and sample size users of diagnosis systems should adopt for best diagnostic performance. E.g., best-first samples are favorable for small sample sizes or when the information gain [dKW87] or split-in-half [SFFR12] heuristics are used, whereas random ones are best for larger sample sizes or when adopting the risk optimization [RSFF13] or most probable singleton [Rod17] techniques for query selection. Further, our results suggest a time-information trade-off in diagnosis sampling, i.e., more efficient sampling tends to imply less effective queries. Finally, we find that the suggested approximate, and often efficient, sampling technique based on the Inv-HS-Tree algorithm [SFRF14c] in many cases provides a good balance between sampling efficiency and diagnostic effectivity.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[Rod22g]
<i>Authors:</i>	Patrick Rodler
<i>Title:</i>	Random vs. best-first: Impact of sampling strategies on decision making in model-based diagnosis
<i>Year:</i>	2022
<i>Publication Venue (Type):</i>	AAAI Conference on Artificial Intelligence (Conference)
<i>Venue Metrics:</i>	A
<i>Own Contributions:</i>	Single-author publication
<i>Other Publications on the Topic:</i>	[RE20] (Int’l Workshop on Principles of Diagnosis)

3.1.5 A Taxonomy and Classification of Diagnosis Computation Algorithms

Publication: System Aspects and Goals Addressed See Fig. 3.6.

Publication: Contents in a Nutshell

Motivation: Diagnosis computation is one of the most integral tasks in model-based diagnosis as it allows to generate fault hypotheses, which are essential for both fault localization and repair. Due to its generality, the model-based diagnosis formalism has been used to express and tackle debugging problems in a wide diversity of application areas, as expounded in Sec. 1.4. This has led to a remarkable multitude and heterogeneity of the diagnosis computation methods proposed in the literature, which are often motivated by and tailored for application-specific requirements and problem cases. As a result, it is a hard task for both researchers and practitioners to

- get an overview of existing approaches,
- identify the crucial properties of diagnostic techniques,
- assess the methods based on these properties, and
- choose the appropriate approach for a research- or application-related diagnostic task at hand.

Contributions: To account for this, we present a taxonomy for diagnosis computation algorithms (under the commonly used weak fault model setting, cf. Sec. 1.1). Specifically, we introduce and formally define a range of features which are arguably vital for a proper understanding, comparison, selection, and use of diagnostic techniques. We explain the influence of each feature on the proper selection of a diagnosis algorithm for a diagnostic task, discuss the potential impact of different feature manifestations on the performance of diagnosis algorithms, and examine relationships among the features.

Results: To demonstrate the value and application of the proposed taxonomy, we provide a multi-dimensional assessment and categorization of over 30 important diagnostic methods in the literature.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[Rod22d]
<i>Authors:</i>	Patrick Rodler
<i>Title:</i>	How should I compute my candidates? A taxonomy and classification of diagnosis computation algorithms
<i>Year:</i>	2022
<i>Publication Venue (Type):</i>	Int’l Workshop on Principles of Diagnosis (Conference)
<i>Venue Metrics:</i>	C
<i>Own Contributions:</i>	Single-author publication
<i>Other Publications on the Topic:</i>	[Rod22b] (Information Sciences)

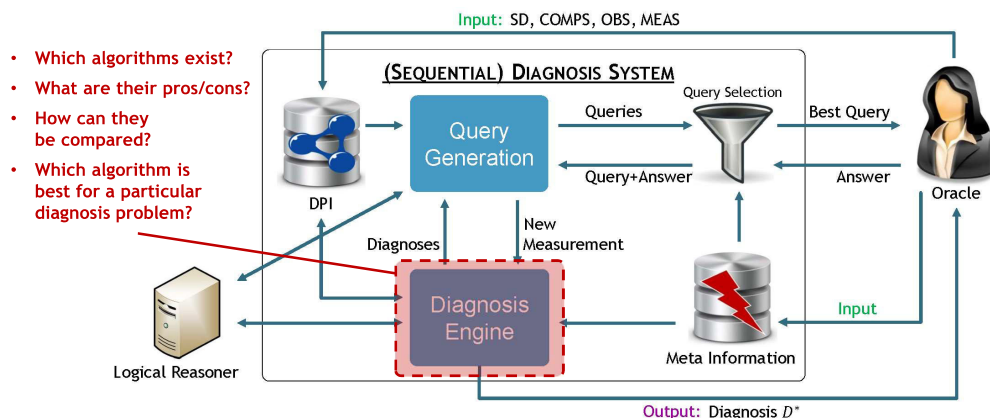


Figure 3.6: A taxonomy and classification of diagnosis computation algorithms: Addressed system aspects and goals.

3.2 Query Computation and Selection

3.2.1 Proposal and Theoretical Analysis of Query Selection Heuristics

Publication: System Aspects and Goals Addressed See Fig. 3.7.

Publication: Contents in a Nutshell

Motivation: Research in the field of Active Learning (AL) [Set10] provides a range of diverse general heuristics targeting the optimization of hypotheses discrimination tasks. While traditionally and very fruitfully exploited in machine learning, e.g., for efficient text classification, image retrieval, concept learning, machine translation, or natural language processing, the key idea behind AL is that a learner can achieve greater accuracy with less newly collected information if the used training data can be adaptively chosen based on its current state of knowledge. At each iteration of the learning process, the active learner can consult an oracle, e.g. a human expert, to label any query from some predefined query space. The new information in terms of the query's label is then taken into account to update the learner's current knowledge state. As this bears close resemblance to the generic information acquisition process pursued by sequential diagnosis systems (cf. Sections 1.2 and 1.6), many AL strategies, termed Query Selection Measures (QSMs), are basically tailored for being used in the context of sequential diagnosis. QSMs are real-valued functions quantifying the quality of queries. However, only few of these AL QSMs, such as information entropy [dKW87], have so far been carried over to or employed in interactive diagnosis.

Contributions: Motivated by these observations, the contributions of this work are the following: We

- investigate the relationship between AL and sequential diagnosis and discuss commonalities of and differences between both fields,
- survey all popular AL QSMs and adapt them for being used as heuristics in sequential diagnosis,
- define a plausible general discrimination preference order (DPO) on queries (formalizing the notion of the "diagnoses discrimination power" of queries),
- figure out superiority relationships between QSMs based on their compliance with the DPO, which suggests a preference order on QSMs,
- harness this preference order to make recommendations regarding the selection of a suitable QSM for a given diagnostic scenario,
- derive improved (parameterized) versions from some QSMs to overcome unveiled deficits (e.g., objectively suboptimal decisions) arising from idiosyncrasies of model-based diagnosis,
- give equivalence classes of QSMs under various conditions (query spaces, QSM parametrizations),

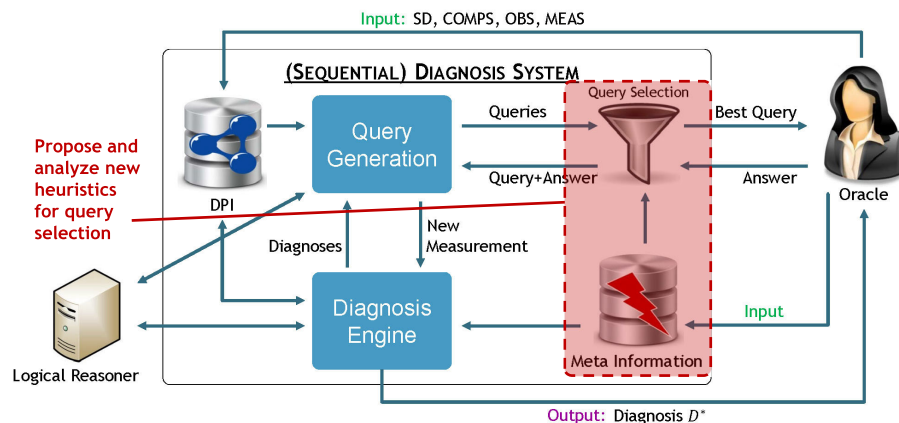


Figure 3.7: Proposal and theoretical analysis of query selection heuristics: Addressed system aspects and goals.

- theoretically analyze QSM functions regarding their global optima and determine qualitative properties of optimal queries,
- show how these properties can be used to formulate heuristics and pruning techniques that allow us to design efficient systematic search procedures for optimal queries wrt. any given QSM,
- based on the systematic search technique devised, suggest a paradigm shift in sequential diagnosis from the commonly adopted *pool-based query selection* strategy (where the best query is filtered out from a potentially large set of candidates) towards a *query synthesis* approach (whose aim is to find a sufficiently good query while exploring a minimal fraction of the query space in a goal-directed way based on heuristics). This novel approach can be especially promising for diagnostic problem cases (e.g., in knowledge base debugging [Rod15]) where a pool of queries is neither explicitly given nor efficiently computable.

Results: In experiments using a benchmark dataset of real-world diagnosis problems, we compared the suggested query synthesis (QS) against the commonly used pool-based query selection (PB) approach. The main insights (from preliminary results) were:

- On average, over all cases, it took QS for all investigated QSMs less than 1 sec to achieve a query optimality (compared to the theoretically optimal QSM-value) of more than 99.999% while exploring only minor fractions of less than 4% the complete query space.
- PB could not handle scenarios where much diagnostic evidence (i.e., a large sample of diagnoses, cf. Sec. 3.1.4) is available for query selection. More precisely, while QS could easily and efficiently deal with up to 80 given minimal diagnoses, PB could in most cases only process no more than 10 ones within a timeout of an hour.
- PB consumed substantially more time than QS. Specifically, the former required minimally / on average / maximally 27 / 787 / 2528 (!) times the query computation time of the latter. By absolute numbers, the minimal / average / maximal PB execution time amounted to 6 / 137 / 566 sec whereas QS never required more than 0.6 sec for the same cases.

Overall, these observations indicate both a very high efficiency and a very high query quality achieved by the proposed novel query computation paradigm, and let QS appear to be the method of choice (at least) for model-based diagnosis problems with query spaces of large size or implicit nature (where query candidates need to be expensively computed, e.g., by means of logical reasoning [SFFR12]).

Publication: Meta Information

<i>Publication included in Thesis:</i>	[Rod17]
<i>Authors:</i>	Patrick Rodler
<i>Title:</i>	On active learning strategies for sequential diagnosis
<i>Year:</i>	2017
<i>Publication Venue (Type):</i>	Int'l Workshop on Principles of Diagnosis (Conference)
<i>Venue Metrics:</i>	C
<i>Own Contributions:</i>	Single-author publication
<i>Other Publications on the Topic:</i>	[Rod16b] (Technical report)

3.2.2 Empirical Analysis of Query Selection Heuristics

Publication: System Aspects and Goals Addressed See Fig. 3.8.

Publication: Contents in a Nutshell

Motivation: A variety of query selection heuristics have been suggested for sequential diagnosis (cf. Sec. 3.2.1), but it is unclear which heuristic to opt for given a particular diagnostic scenario. Questions of interest are, e.g.: Is some heuristic always superior to all others? On which factors does the performance of the particular heuristics depend? Under which circumstances should one use which heuristic?

Contributions: To give well-founded guidance to users of debugging systems and to bring light to these and other questions, we conduct comprehensive empirical evaluations of the query selection heuristics that have been proposed in model-based diagnosis literature. Using real-world benchmark diagnosis problems, we investigate the heuristics under varying conditions regarding the (a) diagnosis probability distributions, (b) quality (meaningfulness) of the probabilities, (c) available diagnostic evidence (size of the computed diagnosis sample) for query computation, and (d) diagnostic structure (i.e., number of system components; number and cardinality of diagnoses; reasoning complexity).

Results: Our main findings based on experiments using real-world problem cases are:

- Using an appropriate heuristic is essential. For inadequate choices of heuristics, we observe overheads in terms of oracle effort of avg. / max. 100 % / 250 %.
- The one and only (generally) best heuristic does not exist, or has not yet been found.
- The main factors affecting the number of queries required for fault localization are the number and size of the given diagnoses as well as the degree of bias in and quality of the given fault information.
- In different diagnosis scenarios, higher or lower numbers of computed diagnoses can be appropriate in combination with a used heuristic. We make recommendations in this regard.
- Different heuristics prevail in different scenarios. Specifically, for
 - (a) misleading fault information (e.g., derived from historical data which does not apply to the current problem), the random query selection strategy is overall preferable to all others,
 - (b) vague fault information (e.g., where the probabilities are partially reasonable, and partially not), the MPS heuristic [Rod17], which selects a query that maximizes the probability of an elimination of a maximal number of diagnoses, turns out to be best approach, and for
 - (c) reasonable fault information, the BME heuristic [Rod17], which proposes a query that maximizes the number of diagnoses that can be eliminated with a probability of more than 50 %, is most favorable (closely followed by the well-known information entropy heuristic [dKW87] and the EMCb heuristic [Rod17], which suggests a query that maximizes the expected diagnosis elimination rate).

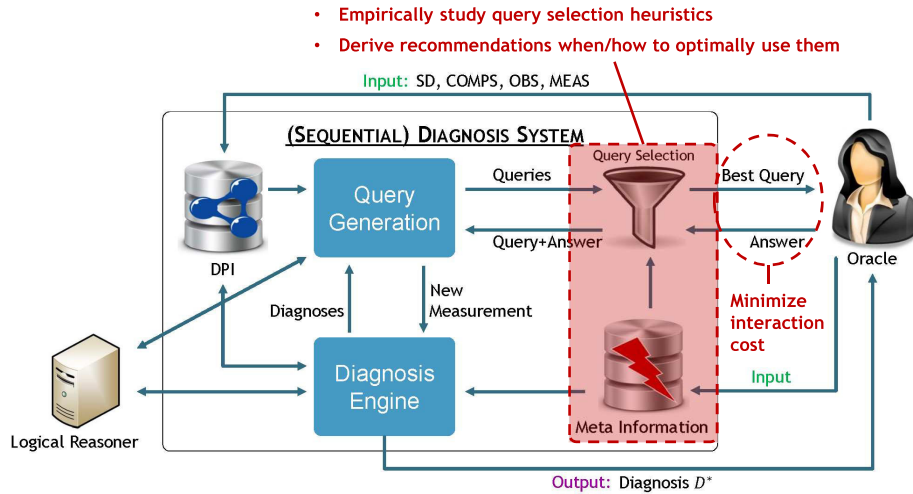


Figure 3.8: Empirical analysis of query selection heuristics: Addressed system aspects and goals.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[RS18c]
<i>Authors:</i>	Patrick Rodler, Wolfgang Schmid
<i>Title:</i>	On the impact and proper use of heuristics in test-driven ontology debugging
<i>Year:</i>	2018
<i>Publication Venue (Type):</i>	Int'l Joint Conference on Rules and Reasoning (Conference)
<i>Venue Metrics:</i>	A
<i>Own Contributions (est.):</i>	Concept/Idea: 100 % Theory/Algorithms: 100 % Development/Implementation: 40 % Evaluation: 80 % Manuscript: 100 %
<i>Other Publications on the Topic:</i>	[RS18a] (Int'l Workshop on Principles of Diagnosis) [RS18b] (Technical report)

3.2.3 Efficient Query Computation and Selection by Systematic Search

Publication: System Aspects and Goals Addressed See Fig. 3.9.

Publication: Contents in a Nutshell

Motivation: As the determination of optimal queries is NP-hard [PA90], sequential diagnosis methods often have to rely on myopic strategies to gauge the gain of different query candidates (cf. Sec. 1.2). One approach is to evaluate the quality of queries in terms of the utility of the expected situation after knowing their answer. Such one-step-lookahead analysis has proven to constitute a particularly favorable trade-off between computational efficiency and diagnostic effectivity, and is today state-of-the-art in sequential diagnosis. However, this solves only a part of the problem, as various other sources of complexity remain. Examples include the strong reliance of sequential diagnosis tools on costly reasoning services, highly expressive knowledge representation languages used for system modeling, large numbers of query candidates, and that queries might neither be explicitly given nor efficiently computable (cf., e.g., [Rod15]).

Existing sequential diagnosis methods often assume a particular *(i)* type of diagnosed system, *(ii)* formalism to describe the system, *(iii)* inference engine, *(iv)* type of query to be of interest, *(v)* query quality

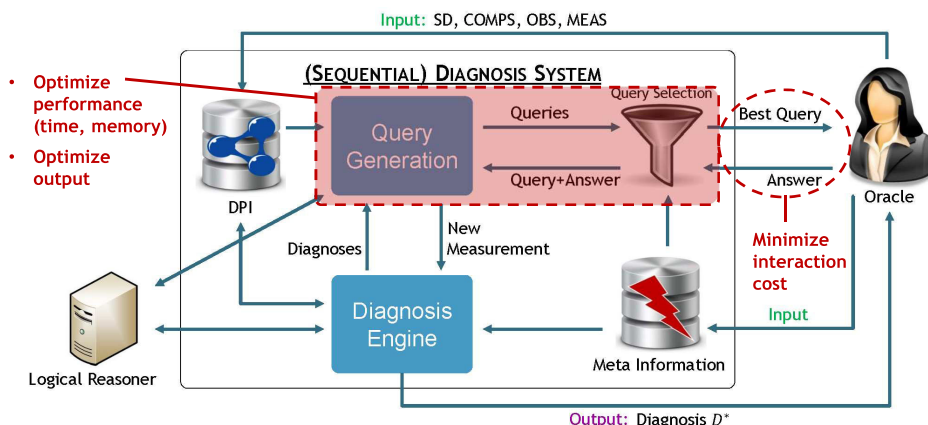


Figure 3.9: Efficient query computation and selection by systematic search: Addressed system aspects and goals.

criterion to be adopted, or (vi) diagnosis computation algorithm to be employed. Further, they often cannot deal with large or implicit query spaces or with expressive logics, or require inputs (e.g., certain inferences) that cannot always be provided.

Contributions: As a remedy, we propose a novel one-step lookahead query computation technique for sequential diagnosis that overcomes the mentioned issues of existing methods. More specifically, our approach (1) is well-founded and based on a solid theory, (2) involves a systematic search for optimal queries, (3) can operate on implicit and huge query spaces, (4) allows for a two-stage optimization of queries (wrt. both their number and their cost), (5) is designed to reduce the required logical inferences to a minimum, and (6) is generally applicable in that it can deal with any type of diagnosis problem as per Reiter’s theory [Rei87], is applicable with any monotonic knowledge representation language, can interact with a multitude of diagnosis engines and logical reasoners, and allows for a quality optimization of queries based on any of the common heuristics in literature (cf. Sections 3.2.1 and 3.2.2).

Results: Using a benchmark of real-world diagnosis problems, we extensively studied the performance of the novel technique (N) and compared it with the state-of-the-art method (S). The gained insights are:

- N is complete, i.e., it can explore all queries (as opposed to S),
- N is non-redundant, i.e., it considers each query at most once (as opposed to S),
- N can compute queries without any expensive reasoner calls (as opposed to S),
- N gives theoretical guarantees about the achieved query quality (as opposed to S),
- N is by orders of magnitude faster than S,
- N *always* returns as good or better queries than S, and
- N scales to input sizes of up to hundreds of diagnoses while S can handle only single-digit numbers.

As a conclusion, given this new approach, query computation and selection is no longer a bottleneck in sequential diagnosis, even in particularly challenging scenarios where queries are neither explicitly given nor efficiently extractable from the system model.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[Rod22i]
<i>Authors:</i>	Patrick Rodler
<i>Title:</i>	Sequential diagnosis by systematic search
<i>Year:</i>	2022
<i>Publication Venue (Type):</i>	Artificial Intelligence (under revision) (Journal)
<i>Venue Metrics:</i>	A (Scimago: Q1, Journal Rank Indicator 1.673)
<i>Own Contributions:</i>	Single-author publication
<i>Other Publications on the Topic:</i>	[RSS17b] (Int’l Workshop on Principles of Diagnosis) [Rod16b] (Technical report) [RSS17a] (Technical report) [Rod22a] (Technical report)

3.3 User Studies and Debugging Tool**3.3.1 User Study: Are Query-Based Debuggers Really Helping Users?**

Publication: System Aspects and Goals Addressed See Fig. 3.10.

Publication: Contents in a Nutshell

Motivation: Interactive debugging based on queries that are automatically generated by the diagnosis system is an attractive methodology as tools implementing this technique can interactively guide their users to the true cause of the observed problem. This query-based approach has been successfully adopted, e.g., for ontology debugging problems, an equally challenging as topical application domain in the era of the Semantic Web [BHL⁺01]. It has proven to be a feasible and efficient approach in various evaluations based on computational experiments (e.g., [SFFR12, RS18c]). While such simulation-based evaluations can assess and compare the performance of debugging algorithms in terms of various interesting aspects such as the time and space complexity, the number of reasoner calls, the theoretical number of required user interactions, or the precision of the fault localization process, they also have certain limitations and often cannot fully inform us about a method’s true usefulness. E.g., they cannot be used to determine if certain assumptions made by an evaluated debugging method hold for actual users, or to gauge the user acceptance. This can be accomplished by an evaluation approach based on user studies where the performance and behavior of experts while using a debugger is observed and analyzed. No user studies with a focus on the investigation of query-based ontology debugging have been published so far.

Contributions: We conducted user studies in the form of testing and debugging exercises that were specifically designed to evaluate if query-based debugging is truly favorable over a previous debugging approach based on test cases, where users specify and add measurements (logical sentences) to the DPI manually in order to prune the diagnosis space. The studies were performed in the context of ontology debugging, where we used a version of the OntoDebug tool [SRS18a] (cf. Sec. 3.3.3) for the experiments, which i.a. involved fault localization tasks given ontologies with injected faults. The study participants were computer science students who had a certain level of education in the development and debugging of ontologies and who received some initial training with the tool. The examined research questions were related to (i) the efficiency and effectiveness of query-based debugging (i.e., do experts need less time, do they find more faults?), (ii) the cognitive ability of users to find out which of the identified diagnoses is the correct one, and (iii) the difficulty of answering system-generated queries for experts.

Results: Our studies revealed the following:

- A query-based approach can make the debugging process more efficient (wrt. user time and effort), without leading to a loss in effectiveness (wrt. the rate of identified faulty ontology axioms).
- The measured overhead using the manual approach amounts to an average of 37% wrt. time and 117% wrt. effort (captured by the number of required mouse clicks).

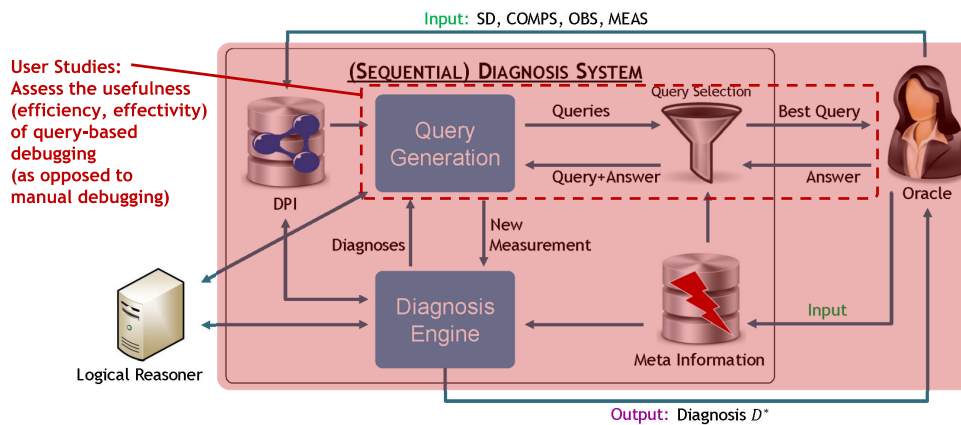


Figure 3.10: User study: Addressed system aspects and goals.

- The capability of a debugger to rank the actual diagnosis higher in a list of candidates visible to the users does not translate into a more effective debugging process. That is, users generally do not have a “perfect bug understanding”, a result in line with similar studies conducted in the software debugging domain [PO11].
- The experiments revealed that “oracle errors”, where users provide faulty inputs, are relatively frequent. Indeed, at least one fault occurred to a quarter of the study participants.

As unveiled by a survey of the literature, oracle errors are a largely open issue, as algorithmic testing or debugging methods usually do not assume, handle or examine oracle errors. As a first step towards approaching this topic, we thus propose a prediction model for oracle errors based on insights from our studies. An assessment of the model evinces that queries it estimates to be hard in fact (1) lead to a higher failure rate, (2) are perceived to be harder, and (3) result in a lower confidence of users in their answers, and vice versa.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[RJSF19]
<i>Authors:</i>	Patrick Rodler, Dietmar Jannach, Konstantin Schekotihin, Philipp Fleiss
<i>Title:</i>	Are query-based ontology debuggers really helping knowledge engineers?
<i>Year:</i>	2019
<i>Publication Venue (Type):</i>	Knowledge-Based Systems (Journal)
<i>Venue Metrics:</i>	A (Scimago: Q1, Journal Rank Indicator 2.192)
<i>Own Contributions (est.):</i>	Concept/Idea: 40 % Theory/Algorithms: 50 % Development/Implementation: 50 % Evaluation: 70 % Manuscript: 60 %
<i>Other Publications on the Topic:</i>	None

3.3.2 On User Types, Assumptions, Optimization Criteria, and How to Make Query-Based Debugging Simpler and More Efficient

Publication: System Aspects and Goals Addressed See Fig. 3.11.

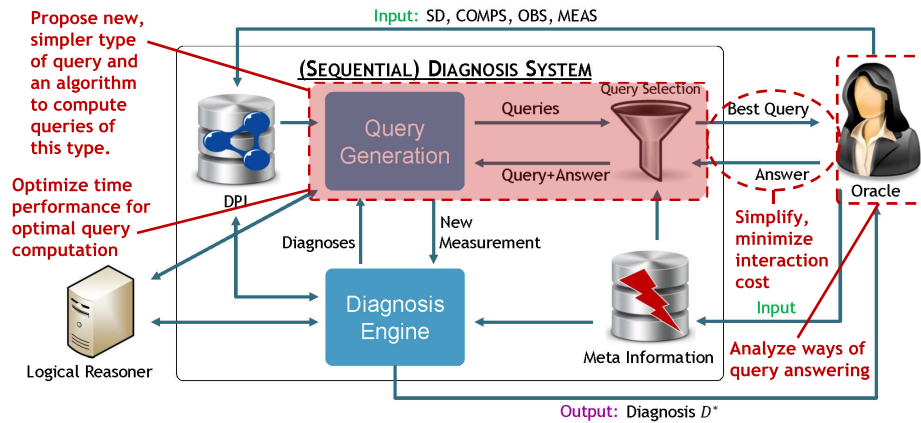


Figure 3.11: On user types, assumptions, optimization criteria, and how to make query-based debugging simpler and more efficient: Addressed system aspects and goals.

Publication: Contents in a Nutshell

Motivation: To suggest as informative queries as possible, existing sequential diagnosis methods draw on various algorithmic optimizations as well as heuristics. However, these computations are often based on certain assumptions about the interacting expert and its behavior, which often have not been fully validated.

Contributions: In this work, we study the reasonability of assumptions about users (experts), optimization techniques and criteria made by state-of-the-art interactive debugging systems in the application context of ontology debugging. We demonstrate that the made assumptions might not always be adequate and discuss consequences of their violations. In particular, we characterize a range of expert types with different query answering behavior and show that existing approaches are far from achieving optimal efficiency for all of them. In addition, we find that the cost metric adopted by state-of-the-art techniques might not always be realistic and that a change of metric can have a decisive impact on the best choice of query answering strategy. As a remedy, based on the insights of our investigations, we suggest a new—and simpler—type of expert question that leads to a stable fault localization performance for all analyzed expert types and effort metrics, and has a range of further advantages over existing techniques, e.g., smaller query search spaces. Moreover, we present a polynomial-time algorithm for computing optimal queries of the suggested type which is fully compatible with existing concepts (e.g., query selection heuristics [Rod17]) and infrastructure (e.g., debugging user interfaces [SRS18a]) in the field.

Results: Comprehensive experiments encompassing more than 23 000 fault localization sessions on faulty real-world problem cases attest that the new querying method is substantially and statistically significantly superior to existing techniques both in terms of the number of necessary expert interactions and in terms of the query computation time. We find that relying on the new querying method is better than drawing on existing techniques in 98 % of the investigated diagnosis scenarios, and can save an interacting expert more than 80 % of their work. In addition, the new approach can reduce the expert’s waiting time for the next query by more than three orders of magnitude, or, in other words, avoid an overhead of more than 120 000 %. E.g., for one problem case we recorded computation times for the optimal normal query by the state-of-the-art approach of avg. / max. 150 sec / 430 sec, whereas an optimal query of the new type was generated in 0.49 sec / 1.06 sec by the novel algorithm, which never required more than 1.1 sec to find a best next query throughout all conducted experiments. Beside these findings, we demonstrate that the efficiency of existing query-based tools can be significantly boosted by suggesting an appropriate query answering strategy to an expert; we also make recommendations in this regard. Further, we suggest optimal configurations of a debugger for situations where the new type of query is used.

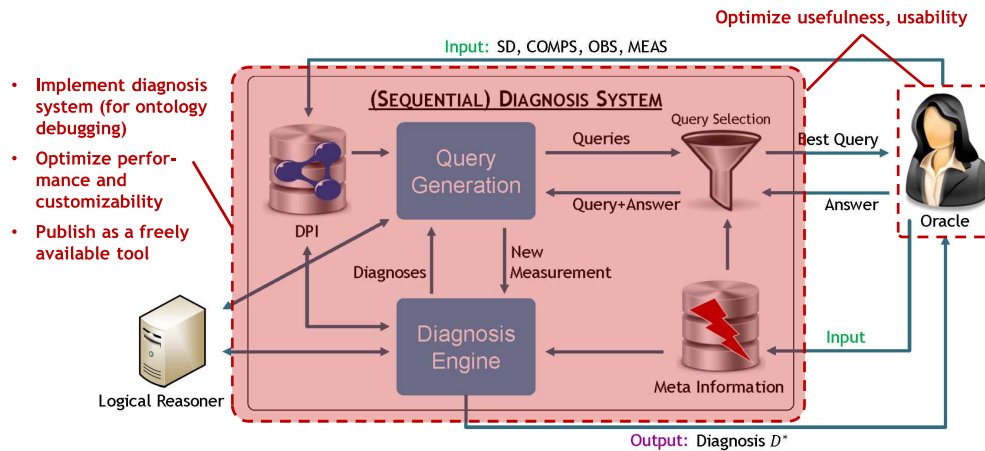


Figure 3.12: OntoDebug—A full-fledged ontology debugging tool: Addressed system aspects and goals.

Publication: Meta Information

Publication included in Thesis: [Rod22f]
Authors: Patrick Rodler
Title: One step at a time: An efficient approach to query-based ontology debugging
Year: 2022
Publication Venue (Type): Knowledge-Based Systems (Journal)
Venue Metrics: A (Scimago: Q1, Journal Rank Indicator 2.192)
Own Contributions: Single-author publication
Other Publications on the Topic: [RE19c] (Int'l Conference on Industrial Engineering and Other Applications of Applied Intelligent Systems)
 [RE19a] (Joint Ontology Workshops)
 [RE19b] (Technical report)
 [Rod20b] (Technical report)

3.3.3 OntoDebug: A Full-Fledged Ontology Debugging Tool

Publication: System Aspects and Goals Addressed See Fig. 3.12.

Publication: Contents in a Nutshell

Motivation: Applications of semantic systems require their users to design ontologies that correctly formalize knowledge about a domain. One example of a vital field extensively adopting ontologies for highly critical applications is biomedicine, cf., e.g., the OBO project⁹ or the NCI-Thesaurus¹⁰. In many cases, factors such as an insufficient understanding of a knowledge representation language, problems concerning modeling techniques and granularity, collaborative development processes, large and complex domains to be described, or the inability to foresee all implications of formulated axioms result in faulty ontologies. Possible deficiencies of ontologies are, e.g., logical contradictions or wrong entailments, which can have severe consequences, e.g., when a wrong therapy is suggested for a patient in health-related applications. What additionally exacerbates the problem is that the root cause of such errors can be extremely hard to identify for humans.

⁹<http://obo.sourceforge.net>

¹⁰<http://ncit.nci.nih.gov>

Debugging tools help to localize faults in ontologies by finding explanations of discrepancies between the actual ontology and the intended one. Existing ontology development environments, such as Protégé [Mus15], Swoop [KPS⁺06], or ORE [LB10], implement various expedient methodologies to assist users in debugging tasks. However, they are mainly based on the principle of presenting collections of conflicts (also referred to as justifications in this context) to the interacting user, who may then need to execute manual analyses in order to figure out the actually faulty axioms and how to appropriately fix them.

Since, in practice, it is often the case that the initially given information is insufficient to unambiguously localize the faulty axioms, there may exist a possibly large number of minimal conflicts and minimal diagnoses, which can render a manual analysis tedious at best, and error-prone or sheerly infeasible at worst. In such scenarios, tool support in the form of interactive debuggers, which semi-automatically guide users to the actual fault by asking a series of questions (whether certain axioms must or must not be entailments of the intended ontology), can greatly alleviate the load on the user, both cognitively and in terms of time and effort (cf., Sec. 3.3.1 and [RJSF19]). However, to date no such interactive debugging tools have been proposed for ontology debugging.

Contributions: We present OntoDebug, an interactive ontology debugger integrated in Protégé, the most popular open-source ontology editor in the world that is used for the maintenance, development, and quality assurance of OBO ontologies and the NCI-Thesaurus (see above), and for numerous other critical applications. OntoDebug is free, open-source, and publicly available in the standard plug-ins repository of Protégé. Its development and implementation was a yearslong endeavor, underpinned and considerably shaped by well-founded research results of our group over more than a decade.

Results: OntoDebug builds upon the ideas of the existing tools discussed above and improves on them i.a. as follows:

- The tool aids the user wrt. all three principal model-based diagnosis tasks (cf. Sec. 1.3), i.e., in terms of fault detection, localization, and repair.
- The plug-in supports an interactive query-based ontology debugging. By using the techniques suggested in [SFFR12, Rod15] and further improved, e.g., in [Rod17, RS18c, RSS17b], the tool is user-centric and based on a query-answer dialogue where the user can rely on their domain knowledge to answer questions about axioms of the intended ontology in order to successively restrict the set of diagnoses. In this vein, the system guides the user throughout the fault localization process and, as opposed to existing approaches, the user is not required to manually specify, analyze or compare various fault explanations in terms of diagnoses and conflicts.
- The plug-in enables a test-driven ontology development, which is closely related and based on the same principles as the familiar test-driven software development [Bec03]. It supports the specification of test cases, which capture requirements on the intended ontology. In particular, the user can define axioms which must be entailments or non-entailments of the intended ontology. These are then automatically verified and a debugging session is started if at least one of them is violated.
- The tool offers a repair interface enabling non-intrusive modifications of the ontology. In the process towards defining a correct repair, a user can introduce and test various modifications of faulty axioms on a virtual copy of the ontology. All modifications are only applied to the productive ontology after the user is satisfied with the obtained result.
- OntoDebug provides the opportunity of a versatile parametrization, configuration, individualization, and combination of various algorithms. Such a proper selection and fine-tuning of used techniques, e.g., based on research results, can help to streamline the performance of OntoDebug based on a given diagnosis problem setting.

As of now, the tool boasts a tally of more than 50K downloads.

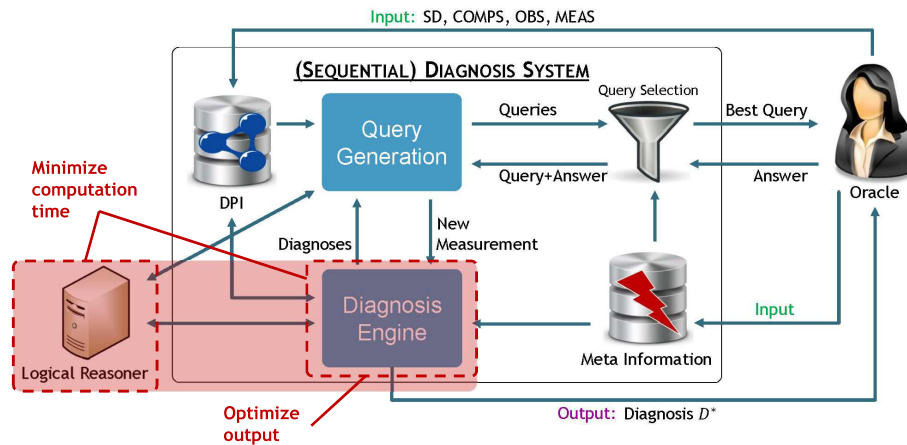


Figure 3.13: Optimization by randomization for hard diagnostic problem cases: Addressed system aspects and goals.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[SRS18a]
<i>Authors:</i>	Konstantin Schekotihin, Patrick Rodler, Wolfgang Schmid
<i>Title:</i>	OntoDebug: Interactive ontology debugging plug-in for Protégé
<i>Year:</i>	2018
<i>Publication Venue (Type):</i>	Int'l Symposium on Foundations of Information and Knowledge Systems (Conference)
<i>Venue Metrics:</i>	B
<i>Own Contributions (est.):</i>	Concept/Idea: 50 % Theory/Algorithms: 60 % Development/Implementation: 25 % Evaluation: n.a. Manuscript: 50 %
<i>Other Publications on the Topic:</i>	[SRS ⁺ 18b] (Int'l Conference on Biological Ontology) [SRS ⁺ 18c] (Int'l Conference on Biological Ontology)

3.4 Diagnostic Reasoning

3.4.1 Optimization by Randomization for Hard Diagnostic Problem Cases

Publication: System Aspects and Goals Addressed See Fig. 3.13.

Publication: Contents in a Nutshell

Motivation: As discussed in Sec. 1.1, the logical reasoning is in many cases the most computationally costly operation in the course of a model-based diagnosis task. Thus, the (expected) time required per theorem prover call also has an influence on the selection of the algorithms employed for debugging. In particular, given theorem proving is cheap (taking time, say, in the range of milliseconds per call of an inference engine), a wider range of algorithms will be applicable to a problem instance than in case reasoning operations are expensive. One particularly challenging domain in this regard is given by, e.g., over-constrained scheduling problems involving too many jobs (products) that cannot all be accomplished (produced) until a given deadline, e.g., due to seasonal order fluctuations, unforeseen machine breakdowns, or incoming high-priority orders. In this diagnosis use case, times per reasoner call,

even if executed by a highly performant cutting-edge solver, can reach up to several minutes. When the objective in such a scenario is the computation of a preferred diagnosis, e.g., a minimal-cost or minimum-cardinality set of jobs to be dropped from the schedule in a way the remaining jobs can be serviced in time, we often measure times in the range of multiple hours per diagnosis computation. Such long waiting times however might not be practical in today's highly dynamical production regimes. For such demanding diagnostic conditions, a more efficient diagnosis computation and optimization approach is required to achieve a maximal customer satisfaction (e.g., by postponing the fewest products as a result of computing a preferred diagnosis) while maintaining a higher decision flexibility (through a reasonable diagnosis computation time), which can be an essential factor towards giving the company a competitive edge.

Contributions: Catering for the described diagnosis use cases where the inefficiency of logical inference operations becomes too severe a bottleneck for standard diagnosis techniques to be put into action, we suggest a novel algorithm for the computation of preferred (e.g., minimum-cardinality or most probable) diagnoses. The method is based on the principle of optimization by randomization, where random minimal diagnoses are successively generated until a (sufficiently) optimal diagnosis is found. The underlying rationale is to trade one hard optimization problem for a range of easier decision problems by explicitly solving (multiple instances of) the subset-minimality problem implicit in the optimization problem. Applied to the mentioned production industry domain, given an over-constrained scheduling task stated as a constraint optimization problem, this means to forgo the (*very inefficient*) direct determination of a preferred diagnosis by means of a state-of-the-art constraint solver. Instead, our approach involves calling an MSMP algorithm such as QuickXplain [Jun04, Rod22c] (cf. Sections 1.1, 1.5 and 3.4.2) multiple times with a randomly modified input (cf. [Rod22g]) to generate a random sample of minimal diagnoses, where each such execution of the MSMP algorithm requires a linear number of (*comparably much more efficient*) consistency checks. We note that this suggested separation between the optimization problem to be solved and its inherent relaxed subset-minimality problem is usually not supported by modern constraint solvers, which is a confirmation that alternative approaches like the one introduced can be useful.

The proposed algorithm is efficiently parallelizable since no information exchange is necessary between successive diagnosis computations, and it can use and benefit from the most suitable and performant algorithms for the involved sub-problems (random number generation, theorem proving, and MSMP problem) as all modules of the method are viewed as black-boxes.

Results: In extensive tests on popular benchmark instances, we compared the performance of the novel optimization-by-randomization technique to a direct approach to solving the optimization problem drawing on a world's leading solver for scheduling problems. The results showed that, first, the new method manifested significant time improvements. In fact, it was *in all cases* able to find better diagnoses in less than half the time required by the direct approach. Second, the proposed strategy could achieve substantial quality enhancements of the computed diagnoses in that over 25 % / 60 % fewer product orders needed to be canceled or postponed on average / maximally.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[RTJ21]
<i>Authors:</i>	Patrick Rodler, Erich Teppan, Dietmar Jannach
<i>Title:</i>	Randomized problem-relaxation solving for over-constrained schedules
<i>Year:</i>	2021
<i>Publication Venue (Type):</i>	Int’l Conference on Principles of Knowledge Representation and Reasoning (Conference)
<i>Venue Metrics:</i>	A
<i>Own Contributions (est.):</i>	Concept/Idea: 50 % Theory/Algorithms: 60 % Development/Implementation: 0 % Evaluation: 50 % Manuscript: 70 %
<i>Other Publications on the Topic:</i>	[RT20] (Int’l Workshop on Principles of Diagnosis)

3.4.2 A Formal Proof and Simple Explanation of the QuickXplain Algorithm

Publication: System Aspects and Goals Addressed See Fig. 3.14.

Publication: Contents in a Nutshell

Motivation: In his seminal¹¹ paper [Jun04] of 2004, Ulrich Junker proposed the QuickXplain algorithm, which provides a divide-and-conquer computation strategy to find within a given set an irreducible subset with a particular (monotone) property. Beside its original application in the domain of constraint satisfaction problems, the algorithm has since then found widespread adoption in areas as different as model-based diagnosis, recommender systems, verification, or the Semantic Web. This popularity is due to the frequent occurrence of the problem of finding irreducible subsets (cf. Sec. 1.5) on the one hand, and to QuickXplain’s general applicability and favorable computational complexity on the other hand. However, although (we regularly experience and have conducted a study to confirm that) people are having a hard time understanding QuickXplain and seeing why it works correctly, a proof of correctness of the algorithm has never been published.

Contributions: This is what we account for in this work, by explaining QuickXplain in a novel tried and tested way, based on a simple, accessible “flat” notation instead of a tree (cf. Example 1.4), and by presenting a first formal correctness proof of it. The latter is expounded as a “proof to explain”, where a particular focus is put on its intelligibility by constructing and explaining it in a way to illuminate the workings and to foster the understandability and intuition of the algorithm.

Results: Apart from showing the correctness of the algorithm and excluding the later detection of errors (*proof and trust effect*), the added value of the availability of a formal proof is, e.g., (i) that the workings of the algorithm often become completely clear only after studying, verifying and comprehending the proof (*didactic effect*), (ii) that the shown proof methodology can be used as a guidance for proving other recursive algorithms (*transfer effect*), and (iii) the possibility of providing “gapless” correctness proofs of systems that rely on (results computed by) QuickXplain, such as numerous model-based debuggers (*completeness effect*).

¹¹On Google Scholar, the paper boasts more than 550 citations as of December 2022.

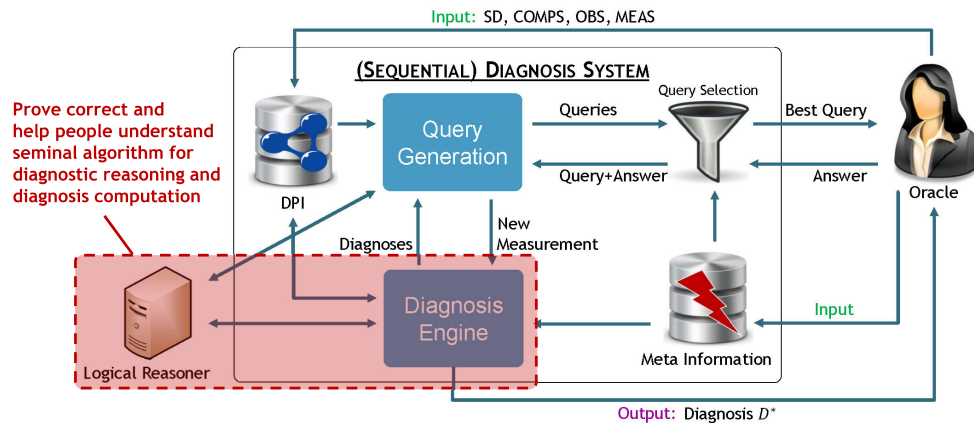


Figure 3.14: A formal proof and simple explanation of the QuickXplain algorithm: Addressed system aspects and goals.

Publication: Meta Information

Publication included in Thesis:	[Rod22c]
Authors:	Patrick Rodler
Title:	A formal proof and simple explanation of the QuickXplain algorithm
Year:	2022
Publication Venue (Type):	Artificial Intelligence Review (Journal)
Venue Metrics:	A (Scimago: Q1, Journal Rank Indicator 2.180)
Own Contributions:	Single-author publication
Other Publications on the Topic:	[Rod15] (PhD thesis) [Rod20f] (Technical report)

3.5 Diagnostic Modeling

3.5.1 On Modeling Techniques for Spreadsheet Debugging: A Theoretical and Empirical Analysis

Publication: System Aspects and Goals Addressed See Fig. 3.15.

Publication: Contents in a Nutshell

Motivation: Based on a formal description of the system, model-based diagnosis allows to precisely and deterministically reason about potential faults responsible for an observed system misbehavior. Under certain circumstances such a formal model can even be extracted from the buggy system in a fully automatic way, such as often in case of software. Since logical reasoning is a core building block of the diagnosis process, the performance of model-based debuggers is largely affected by the reasoning efficiency. The latter in turn depends on the complexity and expressivity of the system description. As highly detailed system descriptions capturing the exact semantics of the system are often beyond the reach of state-of-the-art reasoning techniques, researchers have proposed to reduce the level of detail by introducing more abstract ways to describe the system. When relying on such non-exact system models, it is of interest in how far the reduced detail in terms of modeling translates to an increased efficiency of computations based on the model, and which effects it has on the informativeness of diagnostic inference.

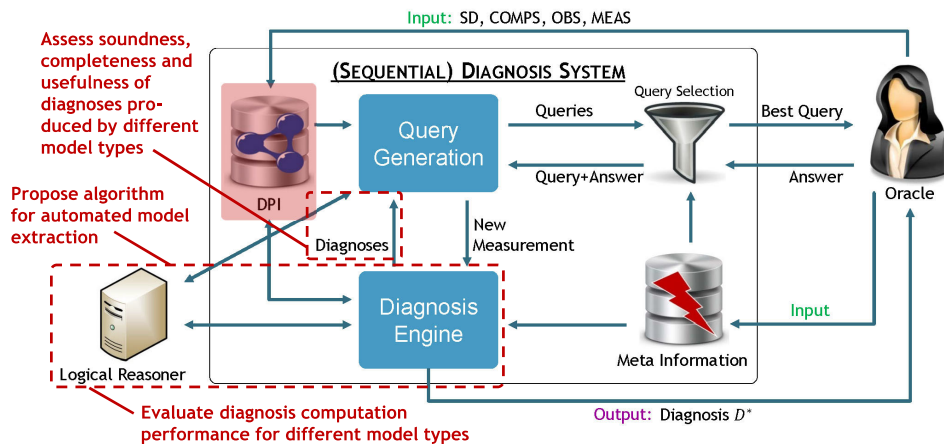


Figure 3.15: On modeling techniques for spreadsheet debugging: Addressed system aspects and goals.

Contributions: In this work, we thoroughly analyze different system modeling techniques with a focus on fault localization in spreadsheets, one of the most popular end-user programming paradigms used by millions in a wide range of areas. Specifically, we present three constraint model types to characterize a spreadsheet at various levels of abstraction, show how they can be extracted from a buggy spreadsheet automatically, and provide theoretical and empirical investigations of the implications of the abstraction level on the computed diagnostic solutions as well as on the computation performance. The model types under study are (1) an exact value-based model that captures the precise semantics of the spreadsheet, (2) a functional dependency model that abstracts from the concrete numerical values and allows to reason only in terms of whether spreadsheet cells are correct or faulty, and (3) a newly proposed qualitative deviation model whose level of abstractness lies between the two other model types in that it allows to express whether (numerical) cell values are too low, exact, or too high.

Results: The main theoretical findings are that

- all model types are diagnostically complete in that no diagnosis can be missed, i.e., the localization of the actually faulty cells is facilitated by each model type,
- the non-exact abstract models may be unsound in that they may generate spurious diagnoses that are in fact no valid fault explanations, and
- in general, a higher degree of abstraction implies an equal or lower diagnostic accuracy (equally many or more generated diagnoses) and equally many or more spurious diagnoses.

Empirically, based on comprehensive experiments using real-world and synthetic benchmarks of faulty spreadsheets, the key lessons learned are that

- the exact model produces significantly fewer diagnoses than the abstract models on average, albeit the number of returned diagnoses is always within one order of magnitude for all model types,
- the exact model is often not (efficiently) applicable, thus abstract models are well motivated,
- diagnosis computation performance is very similar for both abstract models and by at least one and up to more than four orders of magnitude faster on average than for the value-based model,
- the proposed qualitative deviation model is as good regarding diagnostic accuracy as the exact model in a third of all examined cases, and up to more than half of the cases per diagnostic scenario,
- in general, when the exact model exhibits poor performance, then abstract models can be a powerful and efficient surrogate.

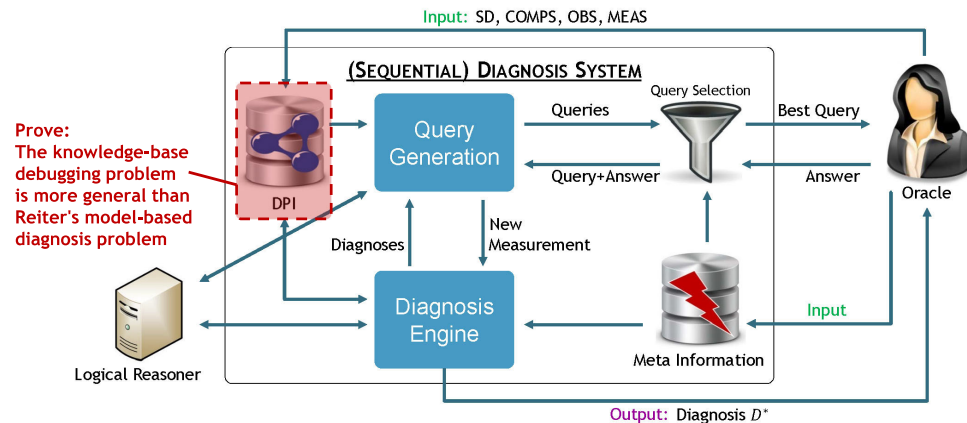


Figure 3.16: Reducing model-based diagnosis to knowledge base debugging: Addressed system aspects and goals.

Publication: Meta Information

<i>Publication included in Thesis:</i>	[HJN ⁺ 22]
<i>Authors:</i>	Birgit Hofer, Dietmar Jannach, Iulia Nica, Patrick Rodler, Franz Wotawa
<i>Title:</i>	On modeling techniques for spreadsheet debugging: A theoretical and empirical analysis
<i>Year:</i>	2022
<i>Publication Venue (Type):</i>	Artificial Intelligence (to be submitted shortly) (Journal)
<i>Venue Metrics:</i>	A (Scimago: Q1, Journal Rank Indicator 1.673)
<i>Own Contributions (est.):</i>	Concept/Idea: 40 % Theory/Algorithms: 65 % Development/Implementation: 0 % Evaluation: 70 % Manuscript: 85 %
<i>Other Publications on the Topic:</i>	None

3.5.2 Reducing Model-Based Diagnosis to Knowledge Base Debugging

Publication: System Aspects and Goals Addressed See Fig. 3.16.

Publication: Contents in a Nutshell

Motivation: In the literature, we have witnessed a range of ways to formalize the model-based diagnosis (MBD) problem. This variety originates, e.g., from different application areas of interest to be addressed, different notational conventions or preferences, or the desire to represent aspects not captured by Reiter's original formulation of the problem [Rei87]. In particular, we have in our research intensively dealt with the knowledge base debugging (KBD) problem [Rod15], which draws on concepts from model-based diagnosis to find faults in a monotonic knowledge base. However, on a very precise theoretical level, the exact relationship between Reiter's account and ours was for long not fully clear and subject of recurring discussions between researchers.

Contributions: This paper brings light into the darkness regarding this matter. To this end, we theoretically analyze and compare the two views on the debugging problem.

Results: We find that KBD is a generalization of MBD in that any MBD problem can be reduced to a KBD problem and solutions of the former can be directly extracted from solutions of the latter. Moreover, our studies reveal that the sequential MBD problem (cf. Sec. 1.2) is a special case of the sequential KBD problem (cf. [Rod15]) in that the latter allows a user to provide more types of measurements. As a consequence of these results, KBD approaches can be applied to all systems amenable to MBD. So, research on KBD is relevant also to MBD. The particular consequence of this is that our line of research, which has often considered the KBD formulation, is, without any restrictions, generally applicable to the model-based diagnosis problem (cf. Sec. 1.4).

Publication: Meta Information

Publication included in Thesis: [RS17]
Authors: Patrick Rodler, Konstantin Schekotihin
Title: Reducing model-based diagnosis to knowledge base debugging
Year: 2017
Publication Venue (Type): Int'l Workshop on Principles of Diagnosis (Conference)
Venue Metrics: C
Own Contributions (est.): Concept/Idea: 30 %
Theory/Algorithms: 70 %
Development/Implementation: n.a.
Evaluation: n.a.
Manuscript: 95 %
Other Publications on the Topic: None

Bibliography

- [Bec03] Kent Beck. *Test-Driven Development: By Example*. Addison-Wesley Professional, 2003.
- [BHL⁺01] Tim Berners-Lee, James Hendler, Ora Lassila, et al. The Semantic Web. 2001.
- [BHMW09] Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors. *Handbook of Satisfiability*. IOS Press, 2009.
- [BK15] Fahiem Bacchus and George Katsirelos. Using Minimal Correction Sets to More Efficiently Compute Minimal Unsatisfiable Sets. In *Int'l Conference on Computer Aided Verification*, pages 70–86, 2015.
- [BL03] Elazar Birnbaum and Eliezer L. Lozinskii. Consistent subsets of inconsistent systems: Structure and behaviour. *Journal of Experimental & Theoretical Artificial Intelligence*, 15(1):25–46, 2003.
- [DHN06] Nachum Dershowitz, Ziyad Hanna, and Alexander Nadel. A scalable algorithm for minimal unsatisfiable core extraction. In *International Conference on Theory and Applications of Satisfiability Testing*, pages 36–41, 2006.
- [dK91] Johan de Kleer. Focusing on Probable Diagnoses. In *AAAI Conference on Artificial Intelligence*, pages 842–848, 1991.
- [dKW89] Johan de Kleer and Brian C. Williams. Diagnosis with behavioral modes. In *Int'l Joint Conference on Artificial Intelligence*, volume 89, pages 1324–1330, 1989.
- [dKMR92] Johan de Kleer, Alan K. Mackworth, and Raymond Reiter. Characterizing diagnoses and systems. *Artificial Intelligence*, 56, 1992.
- [dKRS92] Johan de Kleer, Olivier Raiman, and Mark Shirley. One step lookahead is pretty good. In *Readings in Model-Based Diagnosis*, pages 138–142, 1992.
- [dKW87] Johan de Kleer and Brian C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97–130, 1987.
- [HJN⁺22] Birgit Hofer, Dietmar Jannach, Julia Nica, Patrick Rodler, and Franz Wotawa. On Modeling Techniques for Spreadsheet Debugging: A Theoretical and Empirical Analysis. *Artificial Intelligence (to be submitted shortly)*, 2022.
- [Hor11] Matthew Horridge. *Justification Based Explanation in Ontologies*. PhD thesis, Univ. of Manchester, 2011.
- [Jun04] Ulrich Junker. QuickXplain: Preferred Explanations and Relaxations for Over-Constrained Problems. In *AAAI Conference on Artificial Intelligence*, volume 3, pages 167–172, 2004.
- [Kal06] Aditya Kalyanpur. *Debugging and Repair of OWL Ontologies*. PhD thesis, Univ. of Maryland, College Park, 2006.

- [Kor93] Richard E. Korf. Linear-space best-first search. *Artificial intelligence*, 62(1):41–78, 1993.
- [KPS⁺06] Aditya Kalyanpur, Bijan Parsia, Evren Sirin, Bernardo Cuenca Grau, and James Hendler. Swoop: A Web Ontology Editing Browser. *Journal of Web Semantics*, 4(2):144–153, 2006.
- [LB10] Jens Lehmann and Lorenz Bühmann. ORE - A Tool for Repairing and Enriching Knowledge Bases. In *Int'l Semantic Web Conference*, pages 177–193, 2010.
- [Mar95] Pierre Marquis. Knowledge compilation using theory prime implicates. In *Int'l Joint Conference on Artificial Intelligence*, pages 837–845, 1995.
- [MJB13] Joao Marques-Silva, Mikoláš Janota, and Anton Belov. Minimal sets over monotone predicates in boolean formulae. In *Int'l Conference on Computer Aided Verification*, pages 592–607, 2013.
- [Mor82] Bernard M.E. Moret. Decision trees and diagrams. *ACM Computing Surveys*, 14(4):593–623, 1982.
- [Mus15] Mark A. Musen. The Protégé project: A look back and a look forward. *AI matters*, 1(4):4–12, 2015.
- [PA90] Krishna R. Pattipati and Mark G. Alexandridis. Application of heuristic search and information theory to sequential fault diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(4):872–887, 1990.
- [PO11] Chris Parnin and Alessandro Orso. Are automated debugging techniques actually helping programmers? In *Int'l Symposium on Software Testing and Analysis*, pages 199–209, 2011.
- [RE19a] Patrick Rodler and Michael Eichholzer. How You Ask Matters: A Simple Expert Questioning Approach for Efficient Ontology Fault Localization. In *Joint Ontology Workshops*, 2019.
- [RE19b] Patrick Rodler and Michael Eichholzer. A new expert questioning approach to more efficient fault localization in ontologies. Technical Report, Univ. of Klagenfurt, 2019. (arXiv:1904.00317)
- [RE19c] Patrick Rodler and Michael Eichholzer. On the usefulness of different expert question types for fault localization in ontologies. In *Int'l Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems*, pages 360–375, 2019.
- [RE20] Patrick Rodler and Fatima Elichanova. Do We Really Sample Right In Model-Based Diagnosis? In *Int'l Workshop on Principles of Diagnosis*, 2020.
- [Rei87] Raymond Reiter. A Theory of Diagnosis from First Principles. *Artificial Intelligence*, 32(1):57–95, 1987.
- [RGH12] Ana Armas Romero, Bernardo Cuenca Grau, and Ian Horrocks. MORE: Modular combination of OWL reasoners for ontology classification. In *Int'l Semantic Web Conference*, 2012.
- [RH18a] Patrick Rodler and Manuel Herold. Reducing Sequential Diagnosis Costs by Modifying Reiter's Hitting Set Tree. In *Int'l Workshop on Principles of Diagnosis*, 2018.
- [RH18b] Patrick Rodler and Manuel Herold. StaticHS: A Variant of Reiter's Hitting Set Tree for Efficient Sequential Diagnosis. In *Int'l Symposium on Combinatorial Search*, 2018.

- [RJSF19] Patrick Rodler, Dietmar Jannach, Konstantin Schekotihin, and Philipp Fleiss. Are query-based ontology debuggers really helping knowledge engineers? *Knowledge-Based Systems*, 179:92–107, 2019.
- [Rod15] Patrick Rodler. *Interactive Debugging of Knowledge Bases*. PhD thesis, Univ. of Klagenfurt, 2015.
- [Rod16a] Patrick Rodler. A Theory of Interactive Debugging of Knowledge Bases in Monotonic Logics. Master’s thesis, Univ. of Klagenfurt, 2016.
- [Rod16b] Patrick Rodler. Towards Better Response Times and Higher-Quality Queries in Interactive Knowledge Base Debugging. Technical report, Univ. of Klagenfurt, 2016. (arXiv:1609.02584)
- [Rod17] Patrick Rodler. On Active Learning Strategies for Sequential Diagnosis. In *Int’l Workshop on Principles of Diagnosis*, pages 264–283, 2017.
- [Rod19] Patrick Rodler. Towards Optimizing Reiter’s HS-Tree for Sequential Diagnosis. Technical report, Univ. of Klagenfurt, 2019. (arXiv:1907.12130)
- [Rod20a] Patrick Rodler. DynamicHS: Optimizing Reiter’s HS-Tree for Sequential Diagnosis. In *Int’l Workshop on Principles of Diagnosis*, 2020.
- [Rod20b] Patrick Rodler. On Expert Behaviors and Question Types for Efficient Query-Based Ontology Fault Localization. Technical report, Univ. of Klagenfurt, 2020. (arXiv:2001.05952)
- [Rod20c] Patrick Rodler. Reuse, Reduce and Recycle: Optimizing Reiter’s HS-Tree for Sequential Diagnosis. In *European Conference on Artificial Intelligence*, pages 873–880, 2020.
- [Rod20d] Patrick Rodler. Sound, Complete, Linear-Space, Best-First Diagnosis Search. In *Int’l Workshop on Principles of Diagnosis*, 2020.
- [Rod20e] Patrick Rodler. Too Good to Throw Away: A Powerful Reuse Strategy for Reiter’s Hitting Set Tree. In *Int’l Symposium on Combinatorial Search*, pages 135–136, 2020.
- [Rod20f] Patrick Rodler. Understanding the QuickXplain Algorithm: Simple Explanation and Formal Proof. Technical report, Univ. of Klagenfurt, 2020. (arXiv:2001.01835)
- [Rod21a] Patrick Rodler. Appendix to the Paper: DynamicHS: Streamlining Reiter’s Hitting-Set Tree for Sequential Diagnosis. Technical report, Univ. of Klagenfurt, 2021.
- [Rod21b] Patrick Rodler. Linear-Space Best-First Diagnosis Search. In *Int’l Symposium on Combinatorial Search*, pages 188–190, 2021.
- [Rod22a] Patrick Rodler. Appendix to the Paper: Sequential Model-Based Diagnosis by Systematic Search. Technical report, Univ. of Klagenfurt, 2022.
- [Rod22b] Patrick Rodler. DynamicHS: Streamlining Reiter’s hitting-set tree for sequential diagnosis. *Information Sciences*, 2022.
- [Rod22c] Patrick Rodler. A formal proof and simple explanation of the QuickXplain algorithm. *Artificial Intelligence Review*, 55(8):6185–6206, 2022.
- [Rod22d] Patrick Rodler. How should I compute my candidates? A taxonomy and classification of diagnosis computation algorithms. In *Int’l Workshop on Principles of Diagnosis*, 2022.
- [Rod22e] Patrick Rodler. Memory-limited model-based diagnosis. *Artificial Intelligence*, 305:103681, 2022.

- [Rod22f] Patrick Rodler. One step at a time: An efficient approach to query-based ontology debugging. *Knowledge-Based Systems*, 251:108987, 2022.
- [Rod22g] Patrick Rodler. Random vs. Best-First: Impact of Sampling Strategies on Decision Making in Model-Based Diagnosis. In *AAAI Conference on Artificial Intelligence*, 2022.
- [Rod22h] Patrick Rodler. RBF-HS: Recursive Best-First Hitting Set Search. Technical report, Univ. of Klagenfurt, 2022. (arXiv:2010.04282)
- [Rod22i] Patrick Rodler. Sequential Diagnosis by Systematic Search. *Artificial Intelligence (under revision)*, 2022.
- [RS17] Patrick Rodler and Konstantin Schekotihin. Reducing Model-Based Diagnosis to Knowledge Base Debugging. In *Int'l Workshop on Principles of Diagnosis*, pages 284–296, 2017.
- [RS18a] Patrick Rodler and Wolfgang Schmid. Comparing the Performance of Traditional and Novel Heuristics for Sequential Diagnosis. In *Int'l Workshop on Principles of Diagnosis*, 2018.
- [RS18b] Patrick Rodler and Wolfgang Schmid. Evaluating Active Learning Heuristics for Sequential Diagnosis. Technical report, Univ. of Klagenfurt, 2018. (arXiv:1807.03083)
- [RS18c] Patrick Rodler and Wolfgang Schmid. On the Impact and Proper Use of Heuristics in Test-Driven Ontology Debugging. In *Int'l Joint Conference Rules and Reasoning*, pages 164–184, 2018.
- [RSFF11] Patrick Rodler, Kostyantyn Shchekotykhin, Philipp Fleiss, and Gerhard Friedrich. Balancing Brave and Cautious Query Strategies in Ontology Debugging. In *Joint Workshop on Knowledge Evolution and Ontology Dynamics*, 2011.
- [RSFF12] Patrick Rodler, Kostyantyn Shchekotykhin, Philipp Fleiss, and Gerhard Friedrich. RIO: Minimizing User Interaction in Debugging of Aligned Ontologies. In *Int'l Workshop on Ontology Matching*, pages 49–60, 2012.
- [RSFF13] Patrick Rodler, Kostyantyn Shchekotykhin, Philipp Fleiss, and Gerhard Friedrich. RIO: Minimizing User Interaction in Ontology Debugging. In *Int'l Conference on Web Reasoning and Rule Systems*, pages 153–167, 2013.
- [RSS17a] Patrick Rodler, Wolfgang Schmid, and Konstantin Schekotihin. A Generally Applicable, Highly Scalable Measurement Computation and Optimization Approach to Sequential Model-Based Diagnosis. Technical report, Univ. of Klagenfurt, 2017. (erXiv:1711.05508)
- [RSS17b] Patrick Rodler, Wolfgang Schmid, and Konstantin Schekotihin. Inexpensive Cost-Optimized Measurement Proposal for Sequential Model-Based Diagnosis. In *Int'l Workshop on Principles of Diagnosis*, pages 200–218, 2017.
- [RT20] Patrick Rodler and Erich Teppan. The Scheduling Job-Set Optimization Problem: A Model-based Diagnosis Approach. In *Int'l Workshop on Principles of Diagnosis*, 2020.
- [RTJ21] Patrick Rodler, Erich Teppan, and Dietmar Jannach. Randomized Problem-Relaxation Solving for Over-Constrained Schedules. In *Int'l Conference on Principles of Knowledge Representation and Reasoning*, pages 696–701, 2021.
- [Set10] Burr Settles. Active Learning Literature Survey. Technical report, Univ. of Wisconsin-Madison, 2010.
- [SFFR12] Kostyantyn Shchekotykhin, Gerhard Friedrich, Philipp Fleiss, and Patrick Rodler. Interactive Ontology Debugging: Two Query Strategies for Efficient Fault Localization. *Journal of Web Semantics*, 12–13:88–103, 2012.

- [SFRF12] Kostyantyn Shchekotykhin, Philipp Fleiss, Patrick Rodler, and Gerhard Friedrich. Direct computation of diagnoses for ontology alignment. In *Int'l Workshop on Ontology Matching*, pages 244–245, 2012.
- [SFRF14a] Kostyantyn Shchekotykhin, Gerhard Friedrich, Patrick Rodler, and Philipp Fleiss. A direct approach to sequential diagnosis of high cardinality faults in knowledge bases. In *Int'l Workshop on Principles of Diagnosis*, 2014.
- [SFRF14b] Kostyantyn Shchekotykhin, Gerhard Friedrich, Patrick Rodler, and Philipp Fleiss. Interactive Ontology Debugging using Direct Diagnosis. In *Int'l Workshop on Debugging Ontologies and Ontology Mappings*, 2014.
- [SFRF14c] Kostyantyn Shchekotykhin, Gerhard Friedrich, Patrick Rodler, and Philipp Fleiss. Sequential diagnosis of high cardinality faults in knowledge-bases by direct diagnosis generation. In *European Conference on Artificial Intelligence*, pages 813–818, 2014.
- [Sla14] John Slaney. Set-theoretic duality: A fundamental feature of combinatorial optimisation. In *European Conference on Artificial Intelligence*, pages 843–848, 2014.
- [SRS18a] Konstantin Schekotihin, Patrick Rodler, and Wolfgang Schmid. OntoDebug: Interactive Ontology Debugging Plug-in for Protégé. In *Int'l Symposium on Foundations of Information and Knowledge Systems*, pages 340–359, 2018.
- [SRS⁺18b] Konstantin Schekotihin, Patrick Rodler, Wolfgang Schmid, Matthew Horridge, and Tania Tudorache. A Protégé Plug-In for Test-Driven Ontology Development. In *Int'l Conference on Biological Ontology*, 2018.
- [SRS⁺18c] Konstantin Schekotihin, Patrick Rodler, Wolfgang Schmid, Matthew Horridge, and Tania Tudorache. Test-Driven Ontology Development in Protégé. In *Int'l Conference on Biological Ontology*, 2018.
- [vLP08] Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter. *Handbook of Knowledge Representation*. Elsevier, 2008.

Appendix A

Papers Included in the Habilitation Thesis

Please find in the attachment the papers included in this thesis. The order of the papers is the same as adhered to in the discussion of the papers in Chap. 3.

**Appendix omitted in
online version for
licensing reasons.**