

Sound, Complete, Linear-Space, Best-First Diagnosis Search*

Patrick Rodler¹

¹University of Klagenfurt, Austria,
patrick.rodler@aau.at

Abstract

Various model-based diagnosis scenarios require the computation of the most preferred fault explanations. Existing algorithms that are sound (i.e., output *only* actual fault explanations) and complete (i.e., can return *all* explanations), however, require exponential space to achieve this task. As a remedy, to enable successful diagnosis on memory-restricted devices and for memory-intensive problem cases, we propose RBF-HS, a diagnostic search based on Korf’s well-known RBFS algorithm. RBF-HS can enumerate an arbitrary fixed number of fault explanations in best-first order within linear space bounds, without sacrificing the desirable soundness or completeness properties. Evaluations using real-world diagnosis cases show that RBF-HS, when used to compute minimum-cardinality fault explanations, in most cases saves substantial space (up to 98 %) while requiring only reasonably more or even less time than Reiter’s HS-Tree, a commonly used and as generally applicable sound, complete and best-first diagnosis search.

1 Introduction

In model-based diagnosis, heuristic search techniques have proven to be a powerful tool for the computation of parsimonious fault explanations (*minimal diagnoses*) [1; 2; 3; 4; 5; 6; 7; 8]. Due to the (NP-)hardness of the diagnosis computation problem,¹ it is in most applications infeasible to determine all (minimal) diagnoses.

For this reason—besides the often stipulated guarantee that only minimal diagnoses are generated (*soundness*) and no minimal diagnosis is missed (*completeness*)—the focus of search techniques is usually laid on the *best* minimal diagnoses, e.g., the most probable or minimum-cardinality

*Due to the paper length restrictions, some parts of our analyses had to be omitted in this version. For the full version (including a thorough discussion of related work, complexity analyses, all proofs, an extended evaluation section, more examples, and a more comprehensive description of the algorithm) see <http://isbi.aau.at/ontodebug/publications>.

¹Given a non-empty set of minimal diagnoses, computing another minimal diagnosis is NP-hard, even if theorem proving is in P [9].

ones. In particular, the set of the best diagnoses appears to be more appropriate than just *any* sample of diagnoses, i.e.,

- if all components have a very low failure probability (actual diagnosis among minimum-cardinality diagnoses), as is the case for many physical devices,
- if the given probabilistic information is trustworthy and well-founded (actual diagnosis among most probable diagnoses), like in projects with a long history of bug-fixes, or
- in sequential diagnosis² scenarios where an early termination appears to be reasonable only if the best remaining solution is always known.

Like for many algorithms, there is a trade-off between time and space complexity for diagnosis search methods. Among the factors time and space, the latter can be the more critical criterion. Because, if the memory consumption of an algorithm exceeds the amount of available memory, the problem becomes intractable, whereas, with a higher time demand, an algorithm does still work in principle and will deliver a solution (albeit with a potentially longer waiting time). In model-based diagnosis, there is a range of scenarios which (a) pose substantial memory requirements on diagnostic search methods or (b) suffer from too little memory. One example for (a) are problems involving high-cardinality diagnoses, e.g., when two systems are integrated and a multitude of errors emerge at once [5; 11]. Manifestations of (b) are frequently found in today’s era of the Internet of Things (IoT), distributed or autonomous systems, and ubiquitous computing, where low-end microprocessors, often with only a small amount of RAM, are built into almost any device. Whenever such devices should perform (self-)diagnosing actions [12; 13], memory-aware diagnosis algorithms are a must [14; 15].

Existing (sound and complete) best-first diagnosis search methods require an exponential amount of memory in that all paths in the search tree must be stored in order to guarantee that the best one is expanded in each iteration. Hence, they often disqualify for scenarios like (a) and (b) above. As a remedy, we have devised a diagnosis search called *Recursive Best-First Hitting Set Search (RBF-HS)*, which is based

²*Sequential diagnosis* [10] aims at the gradual elimination of spurious diagnoses through the suggestion of informative system measurements, until some stop criterion is met, e.g., only one diagnosis remains. We say that sequential diagnosis relies on *early termination* if it might already stop although multiple diagnoses are still possible, e.g., if the probability of some diagnosis exceeds some predefined threshold.

on Korf’s well-known *Recursive Best-First Search (RBFS)* algorithm [16]. RBF-HS features all the desirable properties of diagnostic searches, i.e., *soundness*, *completeness*, and the *best-first property*, and is able to return an arbitrary fixed number of the best existing solutions within *linear memory* bounds. Moreover, RBF-HS is *generally applicable* regardless of the (monotonic) system description language or the particular logical theorem prover used.

In extensive evaluations on minimum-cardinality diagnosis computation tasks over real-world diagnosis cases, we compare the performance of RBF-HS with that of Reiter’s HS-Tree, a state-of-the-art sound, complete and best-first diagnosis search that is as generally applicable as RBF-HS. The results evince that RBF-HS is comparable with HS-Tree in terms of runtime even though the latter requires significantly (up to orders of magnitude) more memory.

2 Preliminaries

We first briefly characterize MBD concepts used throughout this work, based on the framework of [17; 18] which is (slightly) more general [19] than Reiter’s theory [20].³

Diagnosis Problem. We assume that the diagnosed system, consisting of a set of components $\{c_1, \dots, c_k\}$, is described by a finite set of logical sentences $\mathcal{K} \cup \mathcal{B}$, where \mathcal{K} (possibly faulty sentences) includes knowledge about the behavior of the system components, and \mathcal{B} (correct background knowledge) comprises any additional available system knowledge and system observations. More precisely, there is a one-to-one relationship between sentences $ax_i \in \mathcal{K}$ and components c_i , where ax_i describes the nominal behavior of c_i (*weak fault model*). E.g., if c_i is an AND-gate in a circuit, then $ax_i := out(c_i) = and(in1(c_i), in2(c_i))$; \mathcal{B} in this case might contain sentences stating, e.g., which components are connected by wires, or observed circuit outputs. The inclusion of a sentence ax_i in \mathcal{K} corresponds to the (implicit) assumption that c_i is healthy. Evidence about the system behavior is captured by sets of positive (P) and negative (N) measurements [10; 20; 21]. Each measurement is a logical sentence; positive ones $p \in P$ must be true and negative ones $n \in N$ must not be true. The former can be, depending on the context, e.g., observations about the system, probes or required system properties. The latter model properties that must not hold for the system, e.g., if \mathcal{K} is a biological knowledge base to be debugged, a negative test case might be $\forall X bird(X) \rightarrow flies(X)$ (“every bird flies”). We call $\langle \mathcal{K}, \mathcal{B}, P, N \rangle$ a *diagnosis problem instance (DPI)*.

Example 1 (Diagnosis Problem) Tab. 1 depicts a DPI stated in propositional logic. The “system” (which is the knowledge base itself in this case) comprises five “components” c_1, \dots, c_5 , and the “nominal behavior” of c_i is given by the respective axiom $ax_i \in \mathcal{K}$. There is neither any background knowledge ($\mathcal{B} = \emptyset$) nor any positive measurements ($P = \emptyset$) available from the start. But, there is one negative measurement (i.e., $N = \{\neg A\}$), which postulates that $\neg A$ must not be an entailment of the correct system (knowledge base). Note, however, that \mathcal{K} (i.e., the assumption that all “components” work nominally) in this case does entail $\neg A$ (e.g.,

³The main reason for using this more general framework is its ability to handle negative measurements (things that must *not* be true for the diagnosed system) which are helpful, e.g., for diagnosing knowledge bases [21; 17].

$\mathcal{K} =$	$\{ax_1 : A \rightarrow \neg B$	$ax_2 : A \rightarrow B$	$ax_3 : A \rightarrow \neg C$
	$ax_4 : B \rightarrow C$	$ax_5 : A \rightarrow B \vee C$	$\}$
$\mathcal{B} = \emptyset$	$P = \emptyset$	$N = \{\neg A\}$	

Table 1: Example DPI stated in propositional logic.

due to the axioms ax_1, ax_2) and therefore some axiom in \mathcal{K} must be faulty (i.e., some “component” is not healthy). \square

Diagnoses. If the system description along with the positive measurements (under the assumption \mathcal{K} that all components are healthy) is inconsistent, i.e., $\mathcal{K} \cup \mathcal{B} \cup P \models \perp$, or some negative measurement is entailed, i.e., $\mathcal{K} \cup \mathcal{B} \cup P \models n$ for some $n \in N$, some assumption(s) about the healthiness of components, i.e., some sentences in \mathcal{K} , must be retracted. We call such a set of sentences $\mathcal{D} \subseteq \mathcal{K}$ a *diagnosis* for the DPI $\langle \mathcal{K}, \mathcal{B}, P, N \rangle$ iff $(\mathcal{K} \setminus \mathcal{D}) \cup \mathcal{B} \cup P \not\models x$ for all $x \in N \cup \{\perp\}$. We say that \mathcal{D} is a *minimal diagnosis* for dpi iff there is no diagnosis $\mathcal{D}' \subset \mathcal{D}$ for dpi . The set of minimal diagnoses is representative of all diagnoses under the weak fault model [22], i.e., the set of all diagnoses is equal to the set of all supersets of minimal diagnoses. Thus, diagnosis approaches usually restrict their focus to only minimal diagnoses. We furthermore denote by \mathcal{D}^* the *actual diagnosis* which pinpoints the actually faulty axioms, i.e., all elements of \mathcal{D}^* are in fact faulty and all elements of $\mathcal{K} \setminus \mathcal{D}^*$ are in fact correct.

Example 2 (Diagnoses) For our DPI in Tab. 1 we have four minimal diagnoses, given by $\mathcal{D}_1 := [ax_1, ax_3]$, $\mathcal{D}_2 := [ax_1, ax_4]$, $\mathcal{D}_3 := [ax_2, ax_3]$, and $\mathcal{D}_4 := [ax_2, ax_5]$. For instance, \mathcal{D}_1 is a minimal diagnosis as $(\mathcal{K} \setminus \mathcal{D}_1) \cup \mathcal{B} \cup P = \{ax_2, ax_4, ax_5\}$ is both consistent and does not entail the given negative measurement $\neg A$. \square

Diagnosis Probability Model. In case useful meta information is available that allows to assess the likeliness of failure for system components, the probability of diagnoses (of being the actual diagnosis) can be derived. Specifically, given a function pr that maps each sentence (system component) $ax \in \mathcal{K}$ to its failure probability $pr(ax) \in (0, 1)$, the probability $pr(X)$ of a diagnosis candidate⁴ $X \subseteq \mathcal{K}$ (under the common assumption of independent component failure) is computed as the probability that all sentences in X are faulty, and all others are correct, i.e., $pr(X) := \prod_{ax \in X} pr(ax) \prod_{ax \in \mathcal{K} \setminus X} (1 - pr(ax))$.

Example 3 (Diagnosis Probabilities) Let the component probabilities for the DPI in Tab.1 be $\langle pr(ax_1), \dots, pr(ax_5) \rangle = \langle .1, .05, .1, .05, .15 \rangle$. Then, we can compute the probabilities of all minimal diagnoses from Example 2 as $\langle pr(\mathcal{D}_1), \dots, pr(\mathcal{D}_4) \rangle = \langle .0077, .0036, .0036, .0058 \rangle$. For instance, $pr(\mathcal{D}_1)$ is calculated as $.1 * (1 - .05) * .1 * (1 - .05) * (1 - .15)$. The normalized diagnoses probabilities would then be $\langle .37, .175, .175, .28 \rangle$. Note, this normalization makes sense if not all diagnoses, but only *minimal* diagnoses are of interest, which is usually the case in model-based di-

⁴Note, the probability (of being equal to the actual diagnosis) of some $X \subseteq \mathcal{K}$ which is not a diagnosis is trivially zero. Still, it is reasonable to define the probability pr for such sets as well. The reason is that diagnosis searches like the one discussed in this work grow diagnoses stepwise, starting from the empty set, and it can make a substantial difference (in terms of performance), which of those partial diagnoses are further explored when. To this end, the probabilities can provide a valuable guidance.

agnosis applications for complexity reasons. \square

Conflicts. Useful for diagnosis computation is the notion of a conflict [10; 20]. A conflict is a set of healthiness assumptions for components c_i that cannot all hold given the current knowledge about the system. More formally, $C \subseteq \mathcal{K}$ is a *conflict* for the DPI $\langle \mathcal{K}, \mathcal{B}, P, N \rangle$ iff $C \cup \mathcal{B} \cup P \models x$ for some $x \in N \cup \{\perp\}$. We call C a *minimal conflict* for dpi iff there is no conflict $C' \subset C$ for dpi .

Example 4 (Conflicts) For our running example, dpi , in Tab. 1, there are four minimal conflicts, given by $\mathcal{C}_1 := \langle ax_1, ax_2 \rangle$, $\mathcal{C}_2 := \langle ax_2, ax_3, ax_4 \rangle$, $\mathcal{C}_3 := \langle ax_1, ax_3, ax_5 \rangle$, and $\mathcal{C}_4 := \langle ax_3, ax_4, ax_5 \rangle$. For instance, \mathcal{C}_4 , in CNF equal to $(\neg A \vee \neg C) \wedge (\neg B \vee C) \wedge (\neg A \vee B \vee C)$, is a conflict because adding the unit clause (A) to this CNF yields a contradiction, which is why the negative test case $\neg A$ is an entailment of \mathcal{C}_4 . The minimality of the conflict \mathcal{C}_4 can be verified by rotationally removing from \mathcal{C}_4 a single axiom at the time and controlling for each so obtained subset that this subset is consistent and does not entail $\neg A$. \square

Relationship between Conflicts and Diagnoses. Conflicts and diagnoses are closely related in terms of a hitting set and a duality property [20]:

Hitting Set Property A (minimal) diagnosis for dpi is a (minimal) hitting set of all minimal conflicts for dpi .

(X is a *hitting set* of a collection of sets \mathbf{S} iff $X \subseteq \bigcup_{S_i \in \mathbf{S}} S_i$ and $X \cap S_i \neq \emptyset$ for all $S_i \in \mathbf{S}$)

Duality Property Given a DPI $dpi = \langle \mathcal{K}, \mathcal{B}, P, N \rangle$, X is a diagnosis (or: contains a minimal diagnosis) for dpi iff $\mathcal{K} \setminus X$ is not a conflict (or: does not contain a minimal conflict) for dpi .

3 RBF-HS Algorithm

The Idea. Korf’s heuristic search algorithm RBFS [16] provides the inspiration for RBF-HS. Historically, the main motivation that led to the engineering of RBFS was the problem that best-first searches by that time required exponential space. The goal of RBFS is to trade (more) time for (much less) space by means of a “(re)explore-current-best & backtrack & forget-most & remember-essential & update-cost” cycle. In this vein, RBFS works within linear-space bounds while maintaining completeness and the best-first property.

Notation. RBF-HS is depicted by Alg. 1. It deals with nodes, where each node n is a subset of or equal to a diagnosis and corresponds to a set of edge labels along one branch from the root of the constructed hitting set tree (cf. [20]). Nodes can be unlabeled (initially, after being generated) or labeled by either *valid* (node is a minimal diagnosis), *closed* (node is a non-minimal or already computed minimal diagnosis), or by a minimal conflict (node is a subset of a diagnosis). There are two functions, f and F , that assign a cost to each node, where f defines initial costs ($f := pr$) and remains constant throughout the execution of RBF-HS, and F specifies backed-up (or: learned) costs and is subject to change while RBF-HS runs.

Inputs and Output. RBF-HS accepts a DPI $dpi = \langle \mathcal{K}, \mathcal{B}, P, N \rangle$, a probability measure pr (see Sec. 2), and a stipulated number ld of minimal diagnoses to be returned as input arguments. It outputs the ld (if existent) most probable (wrt. pr) minimal diagnoses for dpi .

Note, for correctness reasons [18], the probability model pr must be *cost-adjusted*, i.e., $pr(ax) < 0.5$ for all

Algorithm 1 RBF-HS

Input: tuple $\langle dpi, pr, ld \rangle$ comprising

- a DPI $dpi = \langle \mathcal{K}, \mathcal{B}, P, N \rangle$
- a probability measure pr that assigns a failure probability $pr(ax) \in (0, 1)$ to each $ax \in \mathcal{K}$ (cf. Sec. 2), where pr is cost-adjusted; *note:* the cost function $f(n) := pr(n)$ for all tree nodes n
- the number ld of leading minimal diagnoses to be computed

Output: list \mathbf{D} where \mathbf{D} is the list of the ld (if existent) most probable (as per pr) minimal diagnoses wrt. dpi , sorted by probability in descending order

```

1: procedure RBF-HS( $dpi, pr, ld$ )
2:    $\mathbf{D} \leftarrow []$ ,  $\mathbf{C} \leftarrow []$ 
3:    $\mathcal{C} \leftarrow \text{FINDMINCONFLICT}(dpi)$ 
4:   if  $\mathcal{C} = \emptyset$  then
5:     return  $\mathbf{D}$ 
6:   if  $\mathcal{C} = \text{'no conflict'}$  then
7:     return  $\emptyset$ 
8:    $\mathbf{C} \leftarrow \text{ADD}(\mathcal{C}, \mathbf{C})$ 
9:   RBF-HS'( $\emptyset, f(\emptyset), -\infty$ ) ▷  $\emptyset$  is the root node
10:  return  $\mathbf{D}$ 

11: procedure RBF-HS'( $n, F(n), bound$ )
12:   $L \leftarrow \text{LABEL}(n)$ 
13:  if  $L = \text{'closed'}$  then
14:    return  $-\infty$ 
15:  if  $L = \text{'valid'}$  then ▷ new minimal diagnosis found
16:     $\mathbf{D} \leftarrow \text{ADD}(n, \mathbf{D})$ 
17:    if  $|\mathbf{D}| \geq ld$  then
18:      exit procedure
19:    return  $-\infty$ 
20:  Child_Nodes  $\leftarrow \text{EXPAND}(n, L)$ 
21:  for  $n_i \in \text{Child\_Nodes}$  do
22:    if  $f(n_i) > F(n)$  then ▷  $f$  true,  $n$  was already expanded before
23:       $F(n_i) \leftarrow \min(F(n), f(n_i))$ 
24:    else
25:       $F(n_i) \leftarrow f(n_i)$ 
26:  if  $|\text{Child\_Nodes}| = 1$  then ▷ add dummy node  $n_d$  with  $F(n_d) = -\infty$ 
27:    Child_Nodes  $\leftarrow \text{ADDDUMMYNODE}(\text{Child\_Nodes})$ 
28:  Child_Nodes  $\leftarrow \text{SORTDECREASINGBYF}(\text{Child\_Nodes})$ 
29:   $n_1 \leftarrow \text{GETANDDELETEFIRSTNODE}(\text{Child\_Nodes})$  ▷  $n_1 \dots$  best child
30:   $n_2 \leftarrow \text{GETFIRSTNODE}(\text{Child\_Nodes})$  ▷  $n_2 \dots$  2nd-best child
31:  while  $F(n_1) \geq bound \wedge F(n_1) > -\infty$  do
32:     $F(n_1) \leftarrow \text{RBF-HS}'(n_1, F(n_1), \max(bound, F(n_2)))$ 
33:    Child_Nodes  $\leftarrow \text{INSERTSORTEDBYF}(n_1, \text{Child\_Nodes})$ 
34:     $n_1 \leftarrow \text{GETANDDELETEFIRSTNODE}(\text{Child\_Nodes})$  ▷  $n_1 \dots$  best child
35:     $n_2 \leftarrow \text{GETFIRSTNODE}(\text{Child\_Nodes})$  ▷  $n_2 \dots$  2nd-best child
36:  return  $F(n_1)$ 

37: procedure LABEL( $n$ )
38:  for  $n_i \in \mathbf{D}$  do
39:    if  $n \supseteq n_i$  then ▷ goal test, part 1 (is  $n$  non-minimal?)
40:      return closed ▷  $n$  is a non-minimal diagnosis
41:  for  $\mathcal{C} \in \mathbf{C}$  do
42:    if  $\mathcal{C} \cap n = \emptyset$  then ▷ cheap non-goal test (is  $n$  not a diagnosis?)
43:      return  $\mathcal{C}$  ▷  $n$  is not a diagnosis; reuse  $\mathcal{C}$  to label  $n$ 
44:   $L \leftarrow \text{FINDMINCONFLICT}(\langle \mathcal{K} \setminus n, \mathcal{B}, P, N \rangle)$ 
45:  if  $L = \text{'no conflict'}$  then ▷ goal test, part 2 (is  $n$  diagnosis?)
46:    return valid ▷  $n$  is a minimal diagnosis
47:  else ▷  $n$  is not a diagnosis
48:     $\mathbf{C} \leftarrow \text{ADD}(L, \mathbf{C})$  ▷  $L$  is a new minimal conflict ( $\notin \mathbf{C}$ )
49:  return  $L$ 

50: procedure EXPAND( $n, \mathcal{C}$ )
51:  Succ_Nodes  $\leftarrow []$ 
52:  for  $e \in \mathcal{C}$  do
53:    Succ_Nodes  $\leftarrow \text{ADD}(n \cup \{e\}, \text{Succ\_Nodes})$ 
54:  return Succ_Nodes

```

$ax \in \mathcal{K}$ must hold. This can be accomplished for any given pr by choosing an arbitrary fixed $c \in (0, 0.5)$ and by setting $pr_{adj}(ax) := c * pr(ax)$ for all $ax \in \mathcal{K}$. Observe that this adjustment does not affect the relative probabilities in that $pr_{adj}(ax)/pr_{adj}(ax') = k$ whenever $pr(ax)/pr(ax') = k$, i.e., no information is lost in the sense that the fault probability order of components will remain invariant.

To effect that diagnoses of minimum cardinality (instead of maximal probability) are preferred by RBF-HS, the probability model must satisfy $pr(ax) := c$ for all $ax \in \mathcal{K}$ for some arbitrary fixed $c \in (0, 0.5)$. Note, this is equivalent to

defining $pr(n) := 1/|n|$ for all nodes n .

Trivial Cases. At the beginning (line 2), RBF-HS initializes the solution list of found minimal diagnoses \mathbf{D} and the list of already computed minimal conflicts \mathbf{C} . Then, two trivial cases are checked, i.e., if no diagnoses exist (lines 4–5) or the empty set is the only diagnosis (lines 6–7) for dpi . Note, the former case applies iff \emptyset is a conflict for dpi , which implies that $\mathcal{K} \setminus \emptyset = \mathcal{K}$ is not a diagnosis for dpi by the Duality Property (cf. Sec. 2), which in turn means that no diagnosis can exist since diagnoses are subsets of \mathcal{K} and each superset of a diagnosis must be a diagnosis as well (weak fault model, cf. Sec. 2). The latter case holds iff there is no conflict at all for dpi , i.e., in particular, \mathcal{K} is not a conflict, which is why $\mathcal{K} \setminus \mathcal{K} = \emptyset$ is a diagnosis by the Duality Property, and consequently no other *minimal* diagnosis can exist.

If none of these trivial cases is given, the call of FINDMINCONFLICT (line 3) returns a non-empty minimal conflict \mathcal{C} (line 8 is reached), which entails by the Hitting Set Property (cf. Sec. 2) that a non-empty (minimal) diagnosis will exist. For later reuse (note: conflict computation is an expensive operation), \mathcal{C} is added to the computed conflicts \mathbf{C} , and then the recursive sub-procedure RBF-HS' is called (line 9). The arguments passed to RBF-HS' are the root node \emptyset , its f -value, and the initial bound set to $-\infty$.

Recursion: Abstract View. For better understanding, it is instructive to look upon RBF-HS' as a succession of the following blocks:

- node labeling (line 12),
- node elimination or addition to solutions (lines 13–19),
- node expansion (line 20),
- node cost inheritance (lines 21–25),
- child node preparation (lines 26–28), and
- recursive child node exploration (lines 29–36).

Recursion: Principle. The basic principle of the recursion (RBF-HS') is to always explore the open node (initially, only the root node \emptyset) with highest F -value (initially, F -values are f -values) in a depth-first manner, until the best node in the currently explored subtree has a lower F -value than the globally best alternative node (whose F -value is always stored by *bound*). Then backtrack and propagate the best F -value among all child nodes up at each backtracking step. Based on their latest known F -value, the child nodes at each tree level are re-sorted in best-first order of F -value. When re-exploring an already explored, but later forgotten, subtree, the F -value of nodes in this subtree is, if necessary, updated through an inheritance from parent to children (cf. [16]). In this vein, a re-learning of already learned backed-up F -values, and thus repeated and redundant work, is avoided. Exploring a node n in RBF-HS means labeling n and assigning it to the set of computed minimal diagnoses (collection \mathbf{D}) if the label is *valid*, and to discard n (no assignment to any collection) in case it is labeled *closed*. In both these cases, the backed up F -value of n is set to $-\infty$, which prevents the algorithm to be misled in prospective iterations by good F -values of these already explored nodes. If n 's label is a minimal conflict \mathcal{C} , then $|\mathcal{C}|$ child nodes $\{n \cup \{c\} | c \in \mathcal{C}\}$ are generated and recursively explored. This recursive backtracking search is executed until either \mathbf{D} comprises the desired number ld of minimal diagnoses or the hitting set tree has been explored in its entirety.

Sub-Procedures. The workings of the sub-procedures called throughout RBF-HS are:

- FINDMINCONFLICT(dpi) receives a DPI $dpi =$

$\langle \mathcal{K}, \mathcal{B}, P, N \rangle$ and outputs a minimal conflict $\mathcal{C} \subseteq \mathcal{K}$ if one exists, and 'no conflict' else. A well-known algorithm that can be used to implement this function is QUICKXPLAIN [23; 24].

- ADD(x, L) takes an object x and a list of objects L as inputs, and returns the list obtained by appending the element x to the end of the list L .
- ADDDUMMYNODE(L) takes a list of nodes L , appends an artificial node n with $f(n) := -\infty$ to L , and returns the result.
- GETANDDELETEFIRSTNODE(L) accepts a sorted list L , deletes the first element from L and returns this deleted element.
- GETFIRSTNODE(L) accepts a sorted list L and returns L 's first element.
- SORTDECREASINGBYF(L) accepts a list of nodes L , sorts L in decending order of F -value, and returns the resulting sorted list.
- INSERTSORTEDBYF(n, L) accepts a node n and a list of nodes L sorted by F -value, and inserts n into L in a way the sorting of L by F -value is preserved.

Finally, the LABEL function can be seen as a series of the following blocks:

- *non-minimality* check (lines 38–40),
- *reuse label* check (lines 41–43), and
- *compute label* operations (lines 44–49).

Note that this LABEL function of RBF-HS' is equal to the one used in Reiter's HS-Tree [20], except that the *duplicate* check is obsolete in RBF-HS'. The reason for this is that there cannot ever be any duplicate (i.e., set-equal) nodes in memory at the same time during the execution of RBF-HS. This holds because, for all potential duplicates n_i, n_j , we must have $|n_i| = |n_j|$, but equal-sized nodes must be siblings (depth-first tree exploration) which is why n_i and n_j must contain $|n_i| - 1$ equal elements (same path up to the parent of n_i, n_j) and one necessarily different element (label of edge pointing from parent to n_i and n_j , respectively).

Properties. RBF-HS is sound, complete and best-first, and allows to compute an arbitrary fixed number of diagnoses while requiring only linear memory:⁵

Theorem 1. *Let $dpi = \langle \mathcal{K}, \mathcal{B}, P, N \rangle$ be a DPI and let FINDMINCONFLICT be a sound and complete method for conflict computation, i.e., given dpi , it outputs a minimal conflict for dpi if a minimal conflict exists, and 'no conflict' otherwise. RBF-HS is sound, complete and best-first, i.e., it computes all and only minimal diagnoses for dpi in decending order of probability as per the cost-adjusted probability measure pr . Further, given a fixed number ld of diagnoses to be computed, RBF-HS requires space in $O(|\mathcal{K}|)$.*

3.1 RBF-HS Exemplification

The following example illustrates the workings of RBF-HS.

Example 5 (RBF-HS) Inputs. Consider a defective system described by $dpi := \langle \mathcal{K}, \mathcal{B}, P, N \rangle$, where $\mathcal{K} = \{ax_1, \dots, ax_7\}$ and no background knowledge or positive and negative measurements are given, i.e., $\mathcal{B}, P, N = \emptyset$. Let $\langle pr(ax_1), \dots, pr(ax_7) \rangle := \langle .26, .18, .21, .41, .18, .40, .18 \rangle$ (note: pr is already cost-adjusted, cf. Sec. 3). Further, let all minimal conflicts for dpi be $\langle ax_1, ax_2, ax_5 \rangle$,

⁵For more details and a proof of this theorem, see the extended version of this paper at <http://isbi.aau.at/ontodebug/publications>.




Figure 1: RBF-HS executed on example DPI (part I).

$\langle ax_2, ax_4, ax_6 \rangle$, $\langle ax_1, ax_3, ax_4 \rangle$, and $\langle ax_1, ax_5, ax_6, ax_7 \rangle$. Assume we want to use RBF-HS to find the $ld := 4$ most probable diagnoses for dpi . To this end, dpi , pr and ld are passed to RBF-HS (Alg. 1) as input arguments.

Illustration (Figures). The way of proceeding of RBF-HS is depicted by Figures 1 and 2, where the following notation is used. Axioms ax_i are simply referred to by i (in node and edge labels). Numbers \textcircled{n} indicate the chronological node labeling (expansion) order. Recall that nodes in Alg. 1 are sets of (integer) edge labels along tree branches. E.g., node $\textcircled{9}$ in Fig. 1 corresponds to the node $n = \{ax_2, ax_4\}$, i.e., to the assumption that components c_2, c_4 are at fault whereas all others are working properly. The probability $pr(n)$ (i.e., the original f -value) of a node n is shown by the black number from the interval $(0, 1)$ that labels the edge pointing to n , e.g., the cost of node $\textcircled{9}$ is 0.18. We tag minimal conflicts $\langle \dots \rangle$ that label internal nodes by C if they are freshly computed (*expensive*; FINDMINCONFLICT call, line 44), and by R if they result from a reuse of some already computed and stored (see list C in Alg. 1) minimal conflict (*cheap*; reuse




Figure 2: RBF-HS executed on example DPI (part II).

label check; lines 41–43). Leaf nodes are labeled as follows: “?” is used for open (i.e., generated, but not yet labeled) nodes; $\checkmark_{(\mathcal{D}_i)}$ for a node labeled *valid*, i.e., a minimal diagnosis named \mathcal{D}_i , that is not yet stored in \mathbf{D} ; $\times_{(Expl)}$ for a node labeled *closed*, i.e., one that constitutes a non-minimal diagnosis or a diagnosis that has already been found and stored in \mathbf{D} ; $Expl$ is an explanation for the non-minimality in the former, and for the redundancy of node in the latter case, i.e., $Expl$ names a minimal diagnosis in \mathbf{D} that is a proper subset of the node, or it names a diagnosis in \mathbf{D} which is equal to node, respectively. Whenever a new diagnosis is added to \mathbf{D} (line 16), this is displayed in the figures by a box that shows the current state of \mathbf{D} . For each expanded node, the value of the *bound* variable relevant to the subtree rooted at this node is denoted by a red-colored value above the node. By green color, we show the backed-up F -value returned in the course of each backtracking step (i.e., the best known

probability of any node in the respective subtree). Further, f -values that have been updated by backed-up F -values are signalized by green-colored edge labels, see, e.g., in Fig. 1, the left edge emanating from the root node of the tree has been reduced from 0.41 (f -value) to 0.09 (F -value) after the first backtrack. Finally, F -values of parents inherited by child nodes (line 23) are indicated by brown color, see the edge between node ⑭ and node ⑮ in Fig. 2.

Discussion and Remarks. Initially, RBF-HS starts with an empty root node, labels it with the minimal conflict $\langle 1, 2, 5 \rangle$ at step ①, generates the three corresponding child nodes $\{1\}$, $\{2\}$, $\{5\}$ shown by the edges originating from the root node, and recursively processes the best child node (left edge, f -value 0.41) at step ②. The *bound* for the subtree rooted at node ② corresponds to the best edge label (F -value) of any open node other than node ②, which is 0.25 in this case. In a similar manner, the next recursive step is taken in that the best child node of node ② with an F -value not less than $bound = 0.25$ is processed. This leads to the labeling of node $\{1, 4\}$ with F -value $0.28 \geq bound$ at step ③, which reveals the first (provenly most probable) diagnosis $\mathcal{D}_1 := [1, 4]$ with $pr(\mathcal{D}_1) = 0.28$, which is added to the solution list \mathbf{D} . Note that $-\infty$ is at the same time returned for node ③. After the next node has been processed and the second-most-probable minimal diagnosis $\mathcal{D}_2 := [1, 6]$ with $pr(\mathcal{D}_2) = 0.27$ has been detected, the by now best remaining child node of node ② has an F -value of 0.09 (leftmost node). This value, however, is lower than *bound*. Due to the best-first property of RBF-HS, this node is not explored right away because *bound* suggests that there are more promising unexplored nodes elsewhere in the tree which have to be checked first. To keep the memory requirements linear, the current subtree rooted at node ② is discarded before a new one is examined. Hence, the first backtrack is executed. This involves the storage of the best (currently known) F -value of any node in the subtree as the backed-up F -value of node ②. This newly “learned” F -value is signalized by the green number (0.09) that by now labels the left edge emanating from the root. Analogously, RBF-HS proceeds for the other nodes, whereas the used *bound* value is always the best value among the *bound* value of the parent and all sibling’s F -values. Please also observe the F -value inheritance that takes place when node $\{2, 4\}$ is generated for the third time (node ⑮, Fig. 2). The reason for this is that the original f -value of $\{2, 4\}$ is 0.18 (see top of Fig. 1), but the meanwhile “learned” F -value of its parent $\{2\}$ is 0.11 and thus smaller. This means that $\{2, 4\}$ must have already been explored and the *de-facto* probability of any diagnosis in the subtree rooted at $\{2, 4\}$ must be less than or equal to 0.11.

Output. RBF-HS immediately terminates as soon as the ld -th (in this case: fourth) minimal diagnosis \mathcal{D}_4 is located and added to \mathbf{D} . The list \mathbf{D} of minimal diagnoses arranged in descending order of probability pr is returned. \square

4 Evaluation

Dataset. As a test dataset for our experiments with RBF-HS we used twelve diagnosis problems from the knowledge-base debugging domain (Table 2) where RBF-HS’s features soundness, completeness and best-firstness are important requirements to diagnosis searches [11; 18; 26]. These problems were already analyzed in studies conducted by other works, e.g., [17; 26; 27], and represent particularly

Table 2: Dataset used in the experiments (sorted by 2nd column).

KB \mathcal{K}	$ \mathcal{K} $	expressivity ¹⁾	#D/min/max ²⁾
Koala (K)	42	$ALCON^{(D)}$	10/1/3
University (U)	50	$SOIN^{(D)}$	90/3/4
IT	140	$SROIQ$	1045/3/7
UNI	142	$SROIQ$	1296/5/6
Chemical (Ch)	144	$ALCHF^{(D)}$	6/1/3
MiniTambis (M)	173	$ALCN$	48/3/3
Transportation (T)	1300	$ALCH^{(D)}$	1782/6/9
Economy (E)	1781	$ALCH^{(D)}$	864/4/8
DBpedia (D)	7228	$ALCHF^{(D)}$	7/1/1
Opengalen (O)	9664	$ALCHF^{(D)}$	110/2/6
CigaretteSmokeExposure (Cig)	26548	$ST^{(D)}$	1566*/4/7*
Cton (C)	33203	SHF	15/1/5

1): Description Logic expressivity, cf. [25]; the higher the expressivity, the higher is the complexity of consistency checking (conflict computation).

2): #D/min/max denotes the number/the min. size/the max. size of minimal diagnoses for the DPI resulting from each input KB \mathcal{K} . If tagged with a *, a value signifies the number/size determined within 1200sec using HS-Tree.

challenging cases in terms of the complexity of consistency checking (e.g., a consistency check for $SROIQ$, cf. third column of Table 2, is 2-NEXPTIME-complete [28]). As Table 2 shows, the dataset also covers a spectrum of different problem sizes (number of axioms or components; column 2), logical expressivities (column 3), as well as diagnostic structures (number and size of minimal diagnoses; column 4). Note that every model-based diagnosis problem (according to Reiter’s original characterization [20]) can be represented as a knowledge-base debugging problem [19], which is why considering knowledge-base debugging problems is without loss of generality.

Experiment Settings. We evaluate RBF-HS in relation to Reiter’s HS-Tree [20], which is a state-of-the-art sound, complete and best-first diagnosis search that is as generally applicable as RBF-HS, due to its independence from the used theorem prover and from the logical language used to describe the diagnosed system.

In our experiments, we considered a multitude of different *diagnosis scenarios*. A diagnosis scenario is defined by the set of inputs given to Alg. 1, i.e., by a DPI dpi , a number ld of minimal diagnoses to be computed, as well as a (cost-adjusted) setting of the component fault probabilities pr . The DPIs for our tests were defined as $\langle \mathcal{K}, \emptyset, \emptyset, \emptyset \rangle$, one for each \mathcal{K} in Tab. 2. That is, the task was to find a minimal set of axioms (faulty components) responsible for the inconsistency of \mathcal{K} , without any background knowledge or measurements initially given (cf. Example 5). For the parameter ld we used the values $\{2, 6, 10, 20\}$. The fault probability $pr(ax)$ of each axiom (component) $ax \in \mathcal{K}$ was specified in a way the diagnosis search returns minimum-cardinality diagnoses first (cf. Sec. 3). As a logical theorem prover, we adopted Pellet [29].

To simulate as realistic as possible diagnosis circumstances, where the *actual diagnosis* (i.e., the *de-facto* faulty axioms) needs to be isolated from a set of initial minimal diagnoses, we ran five sequential diagnosis [10; 17] sessions for each diagnosis scenario defined above. At this, a different randomly chosen actual diagnosis was set as the target solution in each session. Note, running sequential diagnosis sessions instead of just applying a single diagnosis search execution to the DPIs listed in Table 2 has the additional advantage that multiple diagnosis searches, each for a different (updated) DPI, are executed during one sequential session and flow into the experiment results, which gives us a more




Figure 3: Experiment Results: x-axis shows KBs from Table 2 and parameter $ld \in \{2, 6, 10, 20\}$. y-axis pictures factor of memory saved and factor of more time needed by RBF-HS (vs. HS-Tree).

representative picture of the algorithms’ real performance.

A sequential diagnosis session can be conceived of having two alternating phases, that are iterated until a single diagnosis remains: diagnosis search, and measurement conduction. More precisely, the former involves the determination of ld minimal diagnoses \mathbf{D} for a given DPI, the latter the computation of an optimal system measurement (to rule out as many spurious diagnoses in \mathbf{D} as possible), as well as the incorporation of the new knowledge resulting from the measurement outcome into the DPI. Measurement computation is accomplished by means of a *measurement selection function* which gets a set of minimal diagnoses \mathbf{D} as input, and outputs one system measurement such that any measurement outcome eliminates at least one spurious diagnosis in \mathbf{D} . In our experiments, a measurement was defined as a true-false question to an oracle [5; 17; 18; 27], e.g., for a biological KB one such query could be $Q := \text{Bird} \sqsubseteq \text{FlyingAnimal}$ (“does every bird fly?”). Given a positive (negative) answer, Q is moved to the positive (negative) measurements of the DPI. The new DPI is then used in the next iteration of the sequential diagnosis session. That is, a new set of diagnoses \mathbf{D} is sought for this updated DPI, an optimal measurement is calculated for \mathbf{D} , and so forth. Once there is only a single minimal diagnosis for a current DPI, the session stops and outputs the remaining diagnosis. To determine measurement outcomes (i.e., to answer generated questions), we used the predefined actual diagnosis, i.e., each question was automatically answered in a way the actual diagnosis was not ruled out. As a measurement selection function we adopted the commonly used *entropy (ENT)* heuristic [10; 17], which selects a measurement with highest information gain.

To sum up: We ran five diagnosis sessions, each searching for a randomly specified minimal diagnosis, for each algorithm among RBF-HS and HS-Tree, for each DPI from Tab. 2, and for each number of diagnoses $ld \in \{2, 6, 10, 20\}$ to be computed (in each iteration of the session, i.e., at each call of a diagnosis search algorithm).

Experiment Results are shown by Figure 3, which compares the runtime and memory consumption we measured for RBF-HS and HS-Tree averaged over the five performed sessions (note the logarithmic scale). More specifically, the figure depicts the factor of less memory consumed by RBF-HS (blue bars), as well as the factor of more time needed by RBF-HS (orange bars), in relation to HS-Tree. That is, blue bars tending upwards (downwards) mean a better (worse) memory behavior of RBF-HS, whereas upwards (downwards) orange bars signify worse (better) runtime of RBF-HS. For instance, a blue bar of height 10 means that HS-Tree required 10 times as much memory as RBF-HS did in the same experiment; or a downwards orange bar representing the value 0.5 indicates that RBF-HS finished the diagnosis

search task in half of HS-Tree’s runtime. Regarding the absolute runtime and memory expenditure (not displayed in Figure 3) in the experiments, we measured a min/avg/max runtime of 0.04/24/744sec as well as a min/avg/max space consumption of 9/17.5K/1.3M tree nodes.

We make the following observations based on Figure 3:⁶

(1) *More space gained than extra time expended:* Whenever the diagnosis problem was non-trivial to solve, RBF-HS trades space favorably for time, i.e., the factor of space saved is higher than the factor of time overhead (blue bar is higher than orange one).

(2) *Substantial space savings:* Space savings of RBF-HS range from significant to tremendous, and often reach values larger than 10 and up to 50. In other words, the memory overhead of HS-Tree compared to RBF-HS for the same diagnostic task reached up to 4900 %.

(3) *Often also favorable runtime:* In 40 % of the cases RBF-HS exhibited even a lower or equal runtime compared with HS-Tree. In fact, the runtime savings achieved by RBF-HS reach values of up to more than 88 % (case D,20) while at the same time often saving more than 90 % of space. Note, also studies comparing classic (non-hitting-set) best-first searches have observed that linear-space approaches can outperform exponential-space ones in terms of runtime. One reason for this is that, at the processing of each node, the management (node insertion and removal) of an exponential-sized priority queue of open nodes requires time linear in the current tree depth [30]. Hence, when the queue management time of HS-Tree outweighs the time for redundant node regenerations expended by RBF-HS, then the latter will outperform the former.

(4) *Performance independent of number of computed diagnoses:* The relative performance of RBF-HS versus HS-Tree appears to be largely independent of the number ld of computed minimal diagnoses.

(5) *Performance dependent on diagnosis problem:* The gain of using RBF-HS instead of HS-Tree gets the larger, the harder the considered diagnosis problem is. This tendency can be clearly seen in Figure 3 where the diagnosis problems on the x-axis are sorted in ascending order of RBF-HS’s memory reduction achieved, for each value of ld . Note that roughly the same group of (more difficult / easy to solve) diagnosis problems ranks high / low for all values of ld .

5 Conclusions and Future Work

We introduced RBF-HS, a *general* (reasoner-independent and logics-independent) diagnosis (or: hitting set) search that computes minimal diagnoses (hitting sets) in a *sound* and *complete* way, and enumerates them in *best-first* order as prescribed by some preference function (e.g., minimum cardinality, maximal probability). In contrast to existing systems in model-based diagnosis, RBF-HS guarantees these three properties under *linear-space memory bounds*.

In experiments on a corpus of real-world diagnosis problems of various size, reasoning complexity, and diagnostic structure, we put RBF-HS to the test on minimum-cardinality diagnosis computation tasks. At this, we compa-

⁶Note, both RBF-HS and HS-Tree always return exactly the same diagnoses because they provenly have the same features (soundness, completeness, best-firstness). Hence, any savings observed do *not* arise at the cost of losing any theoretical guarantees.

red RBF-HS against HS-Tree, a state-of-the-art sound, complete and best-first hitting set algorithm which is equally general (i.e., reasoner-independent and logics-independent) as RBF-HS. The results testify that: (1) RBF-HS achieves significantly higher space savings than time losses in all non-trivial cases, and the performance gains tend to increase with increasing problem size and complexity; (2) in many cases, RBF-HS's improvements of memory costs are enormous, reaching savings of up to 98 %; (3) the memory advantages reached by RBF-HS mostly do *not* come at the cost of notable runtime increases; (4) in four out of ten cases, the runtime of RBF-HS was even lower than that of HS-Tree, and runtime savings reached values of up to more than 88 %.

Future work topics include (1) further evaluations of RBF-HS, e.g., when used to compute most probable diagnoses, in combination with other measurement selection heuristics [31; 32; 33; 34], or on diagnosis problems from other domains, such as spreadsheet [35] or software debugging [36], and (2) the integration of RBF-HS into our debugging tool *OntoDebug*⁷ [37].

Acknowledgments

This work was supported by the Austrian Science Fund (FWF), contract P-32445-N38.

References

- [1] Dietmar Jannach, Thomas Schmitz, and Kostyantyn Shchekotykhin. Parallel model-based diagnosis on multi-core computers. *JAIR*, 55:835–887, 2016.
- [2] Franz Wotawa. A variant of Reiter's hitting-set algorithm. *Inform. Process. Lett.*, 79(1):45–51, 2001.
- [3] Li Lin and Yunfei Jiang. The computation of hitting sets: Review and new algorithms. *Inform. Process. Lett.*, 86(4):177–184, 2003.
- [4] Russell Greiner, Barbara A Smith, and Ralph W Wilkerson. A correction to the algorithm in Reiter's theory of diagnosis. *Artif. Intell.*, 41(1):79–88, 1989.
- [5] Kostyantyn Shchekotykhin, Gerhard Friedrich, Patrick Rodler, and Philipp Fleiss. Sequential diagnosis of high cardinality faults in knowledge-bases by direct diagnosis generation. In *ECAI*, pp. 813–818, 2014.
- [6] Patrick Rodler and Manuel Herold. StaticHS: A variant of Reiter's hitting set tree for efficient sequential diagnosis. In *SoCS*, 2018.
- [7] Patrick Rodler. Reuse, Reduce and Recycle: Optimizing Reiter's HS-tree for Sequential Diagnosis. In *ECAI*, 2020.
- [8] Alexander Feldman and Arjan Van Gemund. A two-step hierarchical algorithm for model-based diagnosis. In *AAAI*, pp. 827–833, 2006.
- [9] Tom Bylander, Dean Allemang, Michael Tanner, and John Josephson. The computational complexity of abduction. *Artif. Intell.*, 49:25–60, 1991.
- [10] Johan de Kleer and Brian C Williams. Diagnosing multiple faults. *Artif. Intell.*, 32(1):97–130, 1987.
- [11] Christian Meilicke. *Alignment incoherence in ontology matching*. PhD thesis, Univ. Mannheim, 2011.
- [12] Markus Klein and Lars Nielsen. Evaluating some gain scheduling strategies in diagnosis of a tank system. *IFAC Proceedings Volumes*, 33(11):879–884, 2000.
- [13] Brian C Williams and Pandurang Nayak. A model-based approach to reactive self-configuring systems. In *AAAI*, pp. 971–978, 1996.
- [14] Brian C Williams and Robert J Ragno. Conflict-directed A* and its role in model-based embedded systems. *Discrete Appl. Math.*, 155(12):1562–1595, 2007.
- [15] Peter Zoetewij, Jurryt Pietersma, Rui Abreu, Alexander Feldman, and Arjan Van Gemund. Automated fault diagnosis in embedded systems. In *SSIRI*, pp. 103–110, 2008.
- [16] Richard E Korf. Linear-space best-first search: Summary of results. In *AAAI*, pp. 533–538, 1992.
- [17] Kostyantyn Shchekotykhin, Gerhard Friedrich, Philipp Fleiss, and Patrick Rodler. Interactive Ontology Debugging: Two Query Strategies for Efficient Fault Localization. *J. Web Semant.*, 12-13:88–103, 2012.
- [18] Patrick Rodler. *Interactive Debugging of Knowledge Bases*. PhD thesis, Univ. Klagenfurt. *CoRR*, abs/1605.05950, 2015.
- [19] Patrick Rodler and Konstantin Schekotihin. Reducing model-based diagnosis to knowledge base debugging. In *DX*, pp. 284–296, 2018.
- [20] Raymond Reiter. A Theory of Diagnosis from First Principles. *Artif. Intell.*, 32(1):57–95, 1987.
- [21] Alexander Felfernig, Gerhard Friedrich, Dietmar Jannach, and Markus Stumptner. Consistency-based diagnosis of configuration knowledge bases. *Artif. Intell.*, 152(2):213–234, 2004.
- [22] Johan de Kleer, Alan K Mackworth, and Raymond Reiter. Characterizing diagnoses and systems. *Artif. Intell.*, 56, 1992.
- [23] Ulrich Junker. QuickXPlain: Preferred Explanations and Relaxations for Over-Constrained Problems. In *AAAI*, pp. 167–172, 2004.
- [24] Patrick Rodler. Understanding the QuickXPlain Algorithm: Simple Explanation and Formal Proof. *CoRR*, abs/2001.01835, 2020.
- [25] Franz Baader, Diego Calvanese, Deborah L McGuinness, Daniele Nardi, and Peter F Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2007.
- [26] Aditya Kalyanpur. *Debugging and Repair of OWL Ontologies*. PhD thesis, Univ. of Maryland, 2006.
- [27] Patrick Rodler, Dietmar Jannach, Konstantin Schekotihin, and Philipp Fleiss. Are query-based ontology debuggers really helping knowledge engineers? *Knowl.-Based Syst.*, 179:92–107, 2019.
- [28] Bernardo Cuenca Grau, Ian Horrocks, Boris Motik, Bijan Parsia, Peter F Patel-Schneider, and Ulrike Sattler. OWL 2: The next step for OWL. *J. Web Semant.*, 6(4):309–322, 2008.
- [29] Evren Sirin, Bijan Parsia, Bernardo Cuenca Grau, Aditya Kalyanpur, and Yarden Katz. Pellet: A practical OWL-DL reasoner. *J. Web Semant.*, 5(2):51–53, 2007.
- [30] Weixiong Zhang and Richard E Korf. Performance of linear-space search algorithms. *Artif. Intell.*, 79(2):241–292, 1995.
- [31] Patrick Rodler, Kostyantyn Shchekotykhin, Philipp Fleiss, and Gerhard Friedrich. RIO: Minimizing User Interaction in Ontology Debugging. In *RR*, pp. 153–167, 2013.
- [32] Patrick Rodler and Wolfgang Schmid. On the impact and proper use of heuristics in test-driven ontology debugging. In *RuleML+RR*, pp. 164–184, 2018.
- [33] Patrick Rodler. On active learning strategies for sequential diagnosis. In *DX*, pp. 264–283, 2018.
- [34] Bernard ME Moret. Decision trees and diagrams. *CSUR*, 14(4):593–623, 1982.
- [35] Dietmar Jannach, Thomas Schmitz, Birgit Hofer, and Franz Wotawa. Avoiding, finding and fixing spreadsheet errors—a survey of automated approaches for spreadsheet QA. *J. Syst. Softw.*, 94:129–150, 2014.
- [36] Franz Wotawa. Fault localization based on dynamic slicing and hitting-set computation. In *ICSQ*, pp. 161–170, 2010.
- [37] Konstantin Schekotihin, Patrick Rodler, and Wolfgang Schmid. Ontodebug: Interactive ontology debugging plugin for Protégé. In *FoIKS*, pp. 340–359, 2018.

⁷See <http://isbi.aau.at/ontodebug>.