ECAI 2014
T. Schaub et al. (Eds.)

doi:10.3233/978-1-61499-419-0-1091

Exploiting the Semantic Web for Systems Diagnosis

Anika Schumann and **Freddy Lécué** and **Joern Ploennigs**[1]

Abstract. Diagnosis is the task of explaining abnormal behaviors of systems like telecommunication, transportation or energy systems. Given a sequence of observations the problem is to determine, online, all faults that are in line with these observations. Many approaches tackle this problem but they either require domain expertise or a formal description of how observations and faults are connected. This limits their scope to the diagnosis of well-understood faults. We address the problem of diagnosing faults that may occur for the first time and present a new diagnosis approach that integrates techniques for analyzing semantic descriptions of observations and faults.

1 Introduction

Many approaches have been developed for automatically diagnosing problems in complex systems like fault-tree based systems, expert systems, and model-based approaches [1, 4, 5, 6]. However, they all assume a closed world scenario where the set of possible causes that could explain the effects is well defined and where cause-effect relationships can (at least with unlimited computational resources) be established.

However, in practice one may also want to diagnose other causes. Consider, for example, the telecommunication domain to which model-based diagnosis approaches have been successfully applied [3]. They can be used to identify which routing note might be blocked but they cannot help explain why the note is blocked. The root cause here might be the large mobile phone usage of people in a soccer stadium. Automatically identifying such root causes requires extracting semantic information from social web sites or event planers.

We tackle the problem of diagnosing such root causes based on possibly new observable events for applications where the semantic descriptions of known and new events are known, but where the cause-effect relationships of new fault events are unknown. For example, we may know that *soccer game at olympia stadium at 4:00pm* is a new event that is happening today for the first time but we may not know whether this event is part of a diagnosis result, i.e. whether it can actually *cause blockage of routing note* N_1 *at 4:00pm*. What we can resort to are the cause-effect relationships of known events like *open air concert at 8:30pm* and *high utilization of routing note* N_1 *at 8:30pm* and semantic techniques for evaluating semantic descriptions of events.

An illustration of our approach is shown in Figure 1. It consists of three steps: (i) the semantic matching of observations, (ii) the computation of diagnosis candidates using existing methods, and (iii) the semantic matching of faults. The integration of existing diagnosis methods into our approach ensures that we can solve all diagnosis problems that existing methods can solve. For our example the only problem that an existing approach can solve is the diagnosis of observations S', since its diagnosis result, $F1$, is the only one out of the currently possible ones for which the cause-effect relationships are explicitly known and captured.

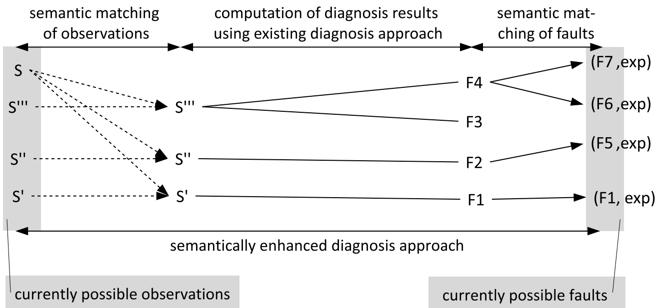

Figure 1. Semantically enhanced Diagnosis for the currently possible observable sequences S, S', S'', and S''', and the currently possible fault sets $F1$, $F5$, $F6$, and $F7$.

In contrast, our approach can retrieve diagnosis results for all the currently possible observations shown in Figure 1. For example, if S''' is observed then we explain its occurrence by faults $F6$ and $F7$.

2 Semantic Matching for Systems Diagnosis

We consider that all *fault* events (faults for short) and *observable* events (observations for short) are semantically represented and organized respectively in ontologies $\Sigma_f^{\mathcal{T}}$ and $\Sigma_o^{\mathcal{T}}$, both part of a more general terminology $\mathcal{T}$ in a Description Logic $\mathcal{L}$. The ontologies $\Sigma_f^{\mathcal{T}}$ (resp. $\Sigma_o^{\mathcal{T}}$) contain the descriptions for faults $\Sigma_{f'}^{\mathcal{T}}$ (resp. observations $\Sigma_{o'}^{\mathcal{T}}$) for which the cause-effect relationships are known and for those that are currently possible, i.e. $F^{\mathcal{T}} \subseteq \Sigma_f^{\mathcal{T}}$ (resp. $O^{\mathcal{T}} \subseteq \Sigma_o^{\mathcal{T}}$). This allows us to explore these semantic representations and to search for approximations of observable o' and fault f' events that most closely match known events o and f.

In a first step we define function $\tilde{g}$ that returns all events that are semantically similar. $\tilde{g}$ is computed with respect to (1), where o' is a matched observation of a new event o with respect to $\tilde{g}$. Any observation o' in $\Sigma_{o'}^{\mathcal{T}}$ which intersects the description of the new event o is returned as a potential approximation, which ensures that o' and o have descriptions in common.

$$\tilde{g}(o) \doteq \{o' \in \Sigma_{o'}^{\mathcal{T}} \mid \Sigma_o^{\mathcal{T}} \not\models o' \sqcap o \sqsubseteq \bot\} \quad (1)$$

For example, the observation $[\{\exists util.Low \sqcap \exists time.\{\text{4:00pm}\} \sqcap \{N_1\}\}]$ and new ones like $[\{\exists util.High \sqcap \exists time.\{\text{8:30pm}\} \sqcap \{N_1\}\}]$ are semantically different in utilisation and time, but identical in location, respectively.

When matching sequences of observations S and S' we need to ensure that its observations match pairwise and thus that they are

[1] IBM Research, Ireland, email: (firstname.lastname)@ie.ibm.com

solutions to g' as defined below:

$$g'(S = [o_1, \ldots, o_k]) \doteq \{(S' = [o'_1, \ldots, o'_k]) | S' \text{ is an observable sequence in system } G \text{ and } \forall o'_i \in \{o'_1, \ldots, o'_k\}\ o'_i \in \tilde{g}(o_i)\}) \quad (2)$$

Furthermore we seek to retrieve only those sequences of observations as possible matches for S that are most similar to S using concept abduction [2]. Concept abduction captures what is overspecified by new event o (resp. f') with respect to existing observation o in $\Sigma_o^{\mathcal{T}}$ (resp. $\Sigma_f^{\mathcal{T}}$).

Definition 1 ***(Concept Abduction Problem -** CAP)* *Let $\mathcal{L}$ be a DL; $\Sigma_o^{\mathcal{T}}$ be a set of axioms in $\mathcal{L}$; o, o' be two satisfiable descriptions in $\Sigma_o^{\mathcal{T}}$. A CAP i.e., $o \backslash o'$ consists of finding a concept $X \in \mathcal{L}$ such that (i) $\Sigma_o^{\mathcal{T}} \models o' \sqcap X \sqsubseteq o$ and (ii) $\Sigma_o^{\mathcal{T}} \not\models o' \sqcap X \equiv \perp$.*

Illustrated for the context of one observation o, Definition 1 can be extended for the context of sequences of observations i.e., $S = [o_1, \ldots, o_k]$ and $S' = [o'_1, \ldots, o'_k]$. The dimension of S and S' needs to be similar and a CAP problem related to S and S' is defined as a sequence of k *CAP* problems noted $S \setminus S'$: *CAP*(o_1, o'_1) i.e. $o_1 \setminus o'_1, \ldots,$ *CAP*(o_k, o'_k) i.e. $o_k \setminus o'_k$.

Consider, for example, that we observe a high utilisation of a routing node at 8:30pm, i.e. $S = [\{\exists util.High \sqcap \exists time.\{\text{8:30pm}\} \sqcap \{N_1\}\}]$. Suppose that such a high utilisation has never before been observed at that time and node. What has been observed and what is captured in the diagnoser are the sequences:

$S' = [\{\exists util.Low \sqcap \exists time.\{\text{8:30pm}\} \sqcap \{N_1\}\}]$,
$S'' = [\{\exists util.Low \sqcap \exists time.\{\text{4:00pm}\} \sqcap \{N_1\}\}]$.

By computing the concept abduction we get:

$S \setminus S' = [\{\exists util.High\}]$,
$S \setminus S'' = [\{\exists util.High \sqcap \exists time.\{\text{4:00pm}\}\}]$.

Thus we retrieve that S' is most similar to S since its differences to S is only a subset of that of S'', i.e. $S \setminus S' \sqsubset S \setminus S''$. Considering all most similar observable sequences as relevant for the diagnosis is in line with the classical diagnosis approach that seeks to compute all possible diagnosis results. Omitting those sequences that are clearly less similar than others, like S'', is also in line with classical diagnosis as the more similar sequences more accurately represent the actual observations. Hence, the faults causing these observations, will represent also more accurately those that cause the current observations. Diagnosis-relevant observable sequences for S are thus defined as follows:

Definition 2 ***Relevant observable sequences** ($g(S)$)*
Let $\mathcal{L}$ be a DL; $\Sigma_o^{\mathcal{T}}$ be a set of axioms in $\mathcal{L}$; $S = [o_1, \ldots o_k]$ a sequence of events in $\Sigma_o^{\mathcal{T}}$ and g' a function as defined in (2). The set of relevant observable sequences that could possibly help to explain S is defined as:

$$\begin{aligned} g(S = [o_1, \ldots o_k]) \doteq \{(S' = [o'_1, \ldots o'_k], S \setminus S') | \\ (S') \in g'(S) \textit{ and } \nexists S'' \in g'(S) \textit{ s.t.} \\ \textit{either } S \setminus S'' \sqsubset S \setminus S' \\ \textit{or } S \setminus S'' = S \setminus S' \textit{ and } S'' \setminus S \sqsubset S' \setminus S\} \end{aligned} \quad (3)$$

Note that function g does not only return the most similar observable sequences but also the properties of S which these sequences do not satisfy. These semantic differences are taken into account when matching faults. For example, let an observable event o be approximated by an event o' that occurred later and let f' be the fault that explains o'. Then we seek to approximate f' by an event f that occurred earlier since o also occurred before o'. The matching function for faults is defined as follows:

Definition 3 ***Relevant fault Sets** ($h(F', S \setminus S')$)*
Let $\mathcal{L}$ be a DL; $\Sigma_f^{\mathcal{T}}$ be a set of axioms in $\mathcal{L}$; $\hat{F} = \{f_1, \ldots f_k\}$ a set of faults in $\Sigma_{f'}^{\mathcal{T}}$ and $S \setminus S'$ a semantic description in $\Sigma_o^{\mathcal{T}}$. The set of relevant fault sets is defined as:

$$\begin{aligned} h(F', S \setminus S') \doteq \{(\hat{F} = \{f_1, \ldots f_k\}, F' \setminus \hat{F}, S \setminus S') | \\ \hat{F} \subseteq F^{\mathcal{T}} \textit{ and } \Sigma_f^{\mathcal{T}} \not\models \hat{F} \sqcap \tilde{F}' \sqsubseteq \perp \textit{ and } \nexists F'' \subseteq F^{\mathcal{T}} \textit{ s.t.} \\ \textit{either } \tilde{F}' \setminus F'' \sqsubset \tilde{F}' \setminus \hat{F} \\ \textit{or } \tilde{F}' \setminus F'' = \tilde{F}' \setminus \hat{F} \textit{ and } F'' \setminus \tilde{F}' \sqsubset \hat{F} \setminus \tilde{F}'\} \\ \textit{where } \tilde{F}' \textit{ is a description that has all properties of } F' \\ \textit{that are not covered in } S \setminus S' \textit{ and all properties of} \\ S \setminus S' \textit{ that are defined in } F^{\mathcal{T}}. \end{aligned} \quad (4)$$

With the ability to retrieve the closest matches for observable sequences and fault sets we can compute the set of the semantically most accurate explanations that are consistent with the currently observed event sequence S as follows:

$$\hat{D}(S) = h(D(S'), S \setminus S') \quad (5)$$

where $S' \in g(S)$ and $D(S')$ is the function that returns the diagnosis result for S' based on an existing diagnosis approach. Thus (5) is correct given that $D(S')$ ensures that the cause-effect relationships between faults and observations are indeed correctly established and that functions g and h consider indeed all events that have some semantic similarity with the ones to be matched (see (1) and second line of (4)). Only those events are omitted that match the original events strictly less accurately than others (see lines 3 and 4 of (3) and (4)).

3 Conclusions

We presented a novel approach, integrating semantic and diagnostic reasoning techniques, for diagnosing events that might have never happened before. In the context of pure AI diagnosis our approach is the first one that can diagnose new events based on historic diagnosis information and the difference between current and historic observations. This was achieved by integrating existing semantic matching functions. Future work includes extending our framework to computing a ranking of diagnosis candidates by evaluating similarities between descriptions of observations and faults.

REFERENCES

[1] Johan de Kleer, Alan K. Mackworth, and Raymond Reiter, 'Characterizing diagnoses and systems', *Artificial Intelligence*, **56**(2–3), 197 – 222, (1992).

[2] Tommaso Di Noia, Eugenio Di Sciascio, Francesco M. Donini, and Marina Mongiello, 'Abductive matchmaking using DLs', in *IJCAI*, pp. 337–342, (2003).

[3] Y. Pencolé, M. O. Cordier, and L. Rozé, 'A decentralized model-based diagnostic tool for complex systems', *International Journal on Artificial Intelligence Tools*, **11**(3), 327–346, (September 2002).

[4] A. Poucet, S. Contini, K. E. Petersen, and N. K. Vestergaard, 'An expert system approach to systems safety and reliability analysis', in *Fault Detection and Reliability: Knowledge Based & Other Approaches*, eds., M. G. Singh, K. S. Hindi, G. Schmidt, and S. Tzafestas. Pergamon Press, (1987).

[5] R. Reiter, 'A theory of diagnosis from first principles', *Artificial Intelligence*, **32**(1), 57–95, (1987).

[6] N. Viswanadham and T. L. Johnson, 'Fault detection and diagnosis of automated manufacturing systems', in *Proc. 27th IEEE Conferencce on Decision and Control*, pp. 2301–2306, (1988).