



## AUTOMATED DIAGNOSIS

LUCA CONSOLE • PIETRO TORASSO

Dipartimento di Informatica - Università degli Studi di Torino

## 1 Introduction

Diagnosing a disease in a patient or singling out a fault in a complex device has often been considered a hard task requiring a significant amount of competence and experience. So, it is not surprising that diagnosis has been one of the main tasks (probably the most relevant) investigated when the first knowledge-based systems (KBS) have been developed in the late seventies. KBS supporting medical diagnosis have played a major role not only to make popular the adoption of production rules but also to develop the notion itself of frames. If we look carefully at these diagnostic systems we realize that in most cases the knowledge is represented as patterns of the form

*Symptoms*  $\Rightarrow$  *Diagnostic hypothesis*

In many cases the relations are much more complex involving a number of intermediate conclusions and in some cases also a taxonomy of diagnostic hypotheses. An in-depth analysis of the structure of these diagnostic systems showed that diagnosis has been reduced to classification, that is the goal of the diagnostic system is to identify the diagnostic hypothesis which is better supported by the data available for the specific diagnostic case. In particular, [13] showed that many of the early diagnostic systems follow a pattern called *Heuristic classification*.

While this approach has resulted very successful in some domains, there are a number of open questions concerning the ability of handling multiple faults, of dealing with symptom masking (two faults may compensate each other so that symptoms are cancelled), of taking into consideration the evolution of the faults.

In the second half of eighties a number of pioneers started to propose a quite different approach to diagnosis: instead of exploiting knowledge coming from human experts, they proposed to use a model of the system to be diagnosed. Since then the notion of Model-Based Diagnosis (MBD) has been introduced (see [27] for an overview of the work done in the first years and the special issue of AI Mag. for recent developments [1]). While the notions

of model can be quite diverse (see following paragraph), in most cases the model was intended to be a qualitative model of the behavior of the system, relating variables corresponding to the status of components of the systems, their internal states, inputs and outputs. Given such models, predictions of the system behavior can be generated. Thus, the diagnostic task consists in evaluating whether there is some discrepancy between the actual values of the observations and the values of the observations predicted by the model when the inputs are known and the health status of the system is OK. If such a discrepancy is singled out, then the diagnostic process has to single out which parts (or components) of the system are faulty and which is (are) the specific fault(s).

The goal of the diagnostic process can then be articulated more precisely: we can distinguish among *fault detection* (automatic detection of deviations from the nominal behavior), *fault localization* (singling out one or more components where the fault is localized) and *fault identification* (singling out the specific fault).

Let us introduce the notions of model and diagnosis by means of a simple example: an AND gate in a digital circuit. The basic knowledge about its behavior can be represented by means of a set of clauses (see Figure 1). The first clause can be read as: for any  $X$  that is an AND gate and that is not faulty ( $OK(X)$ ) and whose two inputs have value 0, then the output of  $X$  has value 0. The following three clauses complete the description of the behavior of any AND gate when the gate behaves nominally. The nominal behavior is essential for fault detection. Let us assume that in a specific AND gate  $A_1$  the two inputs have value 0 and 1 respectively (i.e.,  $In_1(A_1, 0)$  and  $In_2(A_1, 1)$ ). By using the model of the generic gate AND we can infer that expected value of  $A_1$  is 0 (i.e.,  $Out(A_1, 0)$ ). If we have observed (measured) that  $Out(A_1, 1)$  there is discrepancy so a fault has been detected. If one has at disposal just the model of the nominal behavior, the only conclusion is *not*  $OK(A_1)$ , that is the  $A_1$  does not behave properly. However, if one has at disposal the fault model of a generic



$$\begin{aligned}
 &AndG(X) \wedge OK(X) \wedge In_1(X, 0) \wedge In_2(X, 0) \Rightarrow Out(X, 0) \\
 &AndG(X) \wedge OK(X) \wedge In_1(X, 1) \wedge In_2(X, 0) \Rightarrow Out(X, 0) \\
 &AndG(X) \wedge OK(X) \wedge In_1(X, 0) \wedge In_2(X, 1) \Rightarrow Out(X, 0) \\
 &AndG(X) \wedge OK(X) \wedge In_1(X, 1) \wedge In_2(X, 1) \Rightarrow Out(X, 1) \\
 &AndG(X) \wedge Sa0(X) \Rightarrow Out(X, 0) \\
 &AndG(X) \wedge Sa1(X) \Rightarrow Out(X, 1)
 \end{aligned}$$

Figure 1: the nominal and faulty model of an AND gate

AND gate, also the step of fault identification is possible. In particular, Figure 1 reports the faulty behavior when the AND gate is stuck at 0 (represented by the predicate  $Sa0(X)$ ) or is stuck at 1 ( $Sa1(X)$ ). Coming back to the case of  $A_1$ , it is easy to see that the detected discrepancy between the predicted value and the observed one can be now explained and it is possible to infer  $Sa1(A1)$ , that is we have been able to identify the fault. This trivial example shows that there is a relation between the amount of knowledge available for modeling the system to be diagnosed and the kind of results that MBD can provide.

The domain knowledge about a generic AND gate can be considered as a part of a model library relative to combinatorial digital circuits. If such a library contains models for gates of type OR, NOT, XOR, etc., it is easy to see that a model of a specific combinatorial circuit can be built by exploiting the model library and the description of the circuit in terms of components and connections among components.

If we have to diagnose a circuit, it is easy to imagine that the detection of one or more discrepancies may be explained in many different ways, that is the occurrence of a fault in different logical gates may explain the observations. In some complex case, no single fault (that is the fault of a single logical gate) can explain the observations, so the diagnosis has to involve a number of concomitant faults. In this case it is no more possible to conceive diagnosis as a form of classification. On the contrary, diagnosis has to be seen as a synthetic task where the solution has to be assembled by combining the nominal and faulty behavior of the components in order to explain the observations.

## 2 Formal issues

One of the achievements of MBD is the formalization of the diagnostic process; precise semantics can be given to terms such as "explain observations". The pioneering work in [45] and in [23] started an analysis of the notion of diagnosis. In the following we report a formalization that is quite general and is valid for static systems (see next section for the temporal dimension). Let us start from the definition of the notion of system description.

**Definition 1** A *System Description* (SD) is a pair:

- $\mathcal{SV}$  is the set of variables partitioned in *INPUTS* (system inputs), *COMPS* (components of the system), *INTVARS* (non observable internal variables and

*OBS* (observable variables).  $DOM(v)$  denotes the finite domain of variable  $v \in \mathcal{SV}$ ; in particular, for each  $C \in COMPS$ ,  $DOM(C)$  consists in the list of *behavioral modes* for  $C$  (an *ok* mode and a set of fault modes)

- *DT* (Domain Theory) is a set of logical formulas defined over  $\mathcal{SV}$  representing the *instantaneous* behavior of the system (under normal and abnormal conditions).

**Definition 2** A *Diagnostic Problem* is a tuple  $DP = (SD, \mathcal{X}, \mathcal{Y})$  where:

- $SD$  is a System Description
- $\mathcal{X}$  is an instantiation of the *INPUTS* variables, that is the actual value of the inputs
- $\mathcal{Y}$  is an instantiation of the *OBS* variables, that is the actual value of the observable internal variables.

We are now in the position of defining the notion of diagnosis. Actually, the formal framework allows to distinguish between different characterizations of diagnosis.

**Definition 3** Given a diagnostic problem  $DP = (SD, \mathcal{X}, \mathcal{Y})$ , let  $\mathcal{D}$  be an instantiation of *COMPS* (that is  $\mathcal{D}$  specifies a behavioral mode for each component  $C \in COMPS$ ).  $\mathcal{D}$  is a consistent diagnosis for  $DP$  iff  $DT \cup \mathcal{X} \cup \mathcal{Y} \cup \mathcal{D} \not\vdash \perp$

This definition requires that the behavioral modes of the components of the system do not cause any contradiction with the actual observations  $\mathcal{Y}$  when the inputs are  $\mathcal{X}$ .

A stronger notion of explanation imposes the additional requirement that the diagnosis  $\mathcal{D}$  entails the observations (abductive diagnosis).

**Definition 4** Given a diagnostic problem  $DP = (SD, \mathcal{X}, \mathcal{Y})$ , let  $\mathcal{D}$  be an instantiation of *COMPS*.  $\mathcal{D}$  is an abductive diagnosis for  $DP$  iff  $DT \cup \mathcal{X} \cup \mathcal{Y} \cup \mathcal{D} \not\vdash \perp$  and  $\forall m(x) \in \mathcal{Y}; DT \cup \mathcal{X} \cup \mathcal{D} \vdash m(x)$

Abductive diagnosis is stronger than consistency-based diagnosis: all abductive diagnoses are consistency-based diagnoses, but the vice-versa is not true. [20] shows that the consistent characterization and the abductive one of diagnosis are just two extremes of a spectrum where various form of integration are possible: the most suitable characterization of the diagnosis depends on the amount of knowledge available for the system to be diagnosed<sup>1</sup>.

Other formal characterizations of diagnosis have been proposed: [32] discusses a set theoretic characterizations while [26] a probabilistic one.

## 3 Modelling systems

The formal characterization discussed above is a knowledge level reference for diagnostic problem solving. In order to move into the real world, other problems have to

<sup>1</sup>The model of the nominal behavior is sufficient for consistency-based diagnosis while an abductive characterization requires the model of both the nominal and faulty behavior of the system



be solved such as defining ontologies for modeling real systems, defining computational approaches to diagnostic problem solving, integrating diagnosis in the overall operational life of systems and in the process of designing systems. In this section we focus on the first topic; the others will be tackled in the following sections.

Model-based diagnosis relies on a *model* of the system to be diagnosed. A model is an abstraction of the system and may or may not capture some of its features. Deciding the features to be captured is a critical issue, depending on the goals of the diagnostic process and on the types of repair/recovery actions that can be performed after fault detection/identification/isolation [21]. Let us consider a simple example concerning the electronic/hydraulic system controlling the landing system of an aircraft. If the diagnostic process only aims at detecting anomalies to activate a backup system, then we only need a coarse model allowing us to detect that the gear is not opening. On the other hand, if we aim at compensating fault, e.g., by changing activation of pumps or opening/closing valves to maintain functionality, then we need a more detailed model allowing us to detect the specific part (e.g., pump or pipe) which is failing.

Considering physical systems, modeling the structure of the system, i.e., the set of its components and their interaction may be sufficient for some tasks. Notice, however, that also selecting the granularity of components may lead to alternative choices (e.g., consider again the hydraulic system: a pump may be regarded as a basic component or may be in turn decomposed into sub-components). Usually such a choice is guided by the repair strategies that identify the minimal replaceable/repairable components. Adopting a structural model only is usually not enough and models of the system's function or of the system behavior have to be considered. A basic principle adopted by most approaches to MBD is decomposition; the description of behavior/function is given for each component type independently of the others and of the systems in which components instances will be used. This principle, called *no function in structure* is a very important one as it allows having re-usable models and to build *model libraries*, that is chunks of models that can be reused in whatever system that has instances of the same component type.

Models may describe the nominal behaviour of component types (e.g., in a correct pipe flows at the two extremes are the same and pressure is constant) or the behavior in the presence of faults (e.g., in a leaking pipe there is a pressure drop and flows are not balanced).

Let us now discuss how to describe the function/behavior of a system. Also in this case different assumptions can be made, leading to different types of models. In the following we briefly sketch, in a partial way, some of the basic choices and modeling ontologies.

First of all, in almost all the cases the models are qualitative, that is the variables describing the components features (e.g., input, output, state, etc.) assume values in finite

sets corresponding to qualitative abstractions of their actual domains [48]. As an extreme example, we may model the pressure or flow in an hydraulic model as only absent ( $0$ ) or present ( $1$ ). Such distinctions, in fact, are usually sufficient to perform diagnosis where one does not need to simulate precisely the system behavior, but only to detect if the behavior differs from the expected one. In many cases rather than representing absolute values the qualitative domains represent orders of magnitude or deviations with respect to the expected behavior. For example, one may simply represent that in a normal pipe pressure and flow do not deviate while in a leaking pipe the output flow and the pressure are less than they should be (on the other hand, in a clogged pipe the output flow is less than it should be and the pressure is more than it should be).

Qualitative forms of simulation have been defined and theoretical and computational approaches developed to represent notions such as causation, functional dependencies (see [55]). Causal models, in particular, have been adopted in many systems to represent influences (this is common especially in many work on physiological systems and in medical applications). Moreover, also hybrid models mixing semi-quantitative (or even quantitative) and qualitative models have been used when more precise models are needed (e.g., in the detection phase where one needs to determine whether the value read on a given sensor is as expected or is deviated).

When modeling the behavior of a system (or component) different assumptions can be made about its evolution across time. The simplest assumption is to abstract time by building static models in which all phenomena are modeled atemporally and diagnosis is performed on a snapshot of sensor readings. This may be sufficient in some cases, especially when time is not necessary to discriminate among alternative behavior or when it is not necessary to interpret measurements gathered across time. In other cases these assumptions cannot hold and the temporal behavior needs to be modeled. There is a large variety of temporal constraints: for example one may add them on the system behaviour (e.g., delay between input and output) or may model the time-varying behaviour of systems (e.g., the fact that component may be in different modes across time, moving from correct to abnormal and vice-versa, as in transient faults) or may include in the model a notion of state and then represent the dynamic behavior of systems (e.g., the case when output across time depend on the input and state of a component).

It is worth noting that new formal characterizations which extend those in the previous section have been proposed in order to deal with temporal dynamic systems. [6] proposes an extension of the classical logical formalization; many other formalizations are based on the adoption of some discrete event approach, such as Petri net, automata or process algebras (see e.g., [47, 39]).

Italian researchers brought important contributions in this field. At University of Torino (UniTo) different kinds



of modeling have been investigated: causal models [19], component oriented models based on deviations (with Centro Ricerche Fiat [8]), process algebras [15]). Causal modeling including time [4] and hybrid modeling have been studied in Pavia, with special focus on medical applications [3] and material sciences [28]. Approaches based on discrete event systems have been proposed in Brescia for the atemporal [2] and temporal case [31], with application to large dynamic systems.

#### 4 Efficient diagnostic problem solving

The problem of computing model based diagnoses is an hard task. This kind of analysis was another important aspect of the formalization of the diagnostic task. In fact, the solution space to search is quite large:  $2^N$  where  $N$  is the number of components of the system to be diagnosed and we restrict the task to fault localization (i.e., each component has just two behavioral modes: OK and Abnormal). If we want to deal with fault identification, the complexity grows since the solution space becomes  $m^N$ , where  $m$  is the average number of behavioral modes of each component. It is not surprising that in general the complexity of solving a diagnostic problem has been proved to be an NP-hard problem for both the consistent and abductive characterizations [7]. A further source of complexity derives from the very nature of diagnostic problem solving: a diagnostic problem cannot be considered solved when a diagnosis has been determined, but just when all the diagnoses have been computed, so any recovery /repair /decision process has got the complete picture of what has gone wrong. In general, the number of diagnoses can be exponential in the number of components and even if we introduce some preference criterion among diagnoses (such as dealing just with the minimum cardinality faults) the number of preferred diagnoses can be still exponential. So, in principle we have to deal with intractability both in time and space.

In recent years a number of approaches have been investigated in order to practically alleviate the computational complexity of computing the diagnoses. One of the main approaches is the development of strategies computing preferred diagnoses only. As said above there are many different ways for defining preference among diagnosis but most of them are based on the intuitive idea that one diagnosis  $D$  is preferred over diagnosis  $D'$  if  $D$  involves less faults than  $D'$  (or a subset of the faults involved in  $D'$ ). This criterion has also a probabilistic motivation if we assume independence among faults (an assumption often done). For the task of diagnosis in most cases it is not strictly necessary to have at disposal a precise measure of the fault probability as also very rough estimates are sufficient [22]. Algorithms which focus just on leading diagnoses may get good performance but in some cases may have to depend on (costly) backtracking in order not to lose solutions.

As stated above time complexity is just part of the problem. Also space can be a problem when we have to rep-

resent the entire set of possible diagnoses. One way to deal with the problem is to devise some methods for representing in a compact way a large number of alternative diagnoses. An interesting approach was proposed by [24]. More recently the notion of scenario was introduced in [52]. The adoption of symbolic methods (in particular the use of Ordered Binary Decision Diagrams) has also been investigated. OBDDs can be exploited not only to represent in a compact way the set of diagnoses, but also to encode the domain theory (the behavior of the components under nominal and faulty conditions) as well as the possible evolutions of the faults over time.

A quite different approach for dealing with the complexity of diagnosis involves the use of multiple models (e.g., [11] and [51] for Italian contributions).. An approach that is been investigated is the exploitation of structural abstraction for hierarchical diagnosis. The basic idea is to have a number of descriptions of the same system at different levels of detail: in particular, in structural abstraction a macro-component aggregates a number of components for providing a more abstract representation of the system. This structural abstraction can be used for focusing diagnosis: fault localization is done first at the most abstract level and the results obtained in one level are used to focus diagnosis at the more detailed levels. In this way, in most cases, there is a significant reduction of the overall computational cost. Starting from the seminal work in [37], the approach has been widely adopted and refined.

A significant Italian contribution to the field of hierarchical diagnosis is the one in [10]; they can derive automatically a simpler version of a diagnostic problem by exploiting the amount of available observations for the problem at hand. The problem of automatically learning a set of abstractions (instead of exploiting the ones provided by a human expert) has been addressed by [46] and [53] for static systems. Compilation strategies for performing efficient problem solving is another solution which has been investigated in the atemporal [17] and temporal case [18].

Multiple models do not mean necessarily that all the models share the same formalisms. There is a long tradition in diagnosis to adopt and combine different representation formalisms. Among the integration of different reasoning paradigms for diagnostic problem solving it is worth mentioning the one between Case-Based Reasoning (CBR) and Model-Based Reasoning [34, 42]. CBR is based on the idea that new problems can be solved by looking at a set of problems already solved, determining the ones that are most similar to the one at hand and adapting the solution of the retrieved cases to be a solution for the new problem. Models at different levels of abstraction have been studied by CISE [25] with application to power generation and distribution networks; similarly multimodeling with focus on fuzzy models has been studied at Politecnico di Milano [5].



## 5 Off-line VS On-line and beyond diagnosis

Diagnosis is a task which should be integrated in the overall lifetime of a system. Usually diagnosis is a mean rather than a goal: the actual goal is preserving system functionality (this may have different meanings, such as availability of the system, reliability of its performance, preserving correct behaviour, avoiding critical situations . . .). This means that diagnosis has to be part of the monitoring / interpretation / repair / recovery / reconfiguration process.

Two very different situations have to be considered: off-line and on-line diagnosis. In the former case, faults have occurred, measurements have been logged or may be taken on the system and the ultimate goal is to restore system functionality by repairing it. A typical example is when a car is brought to a workshop and diagnosis is a mean to decide how to repair it (by replacing components or repairing them or reconfiguring the system).

In the latter case, the aim is at monitoring the system during its operation and trying to detect faults as soon as possible and then performing some recovery action to preserve the system functionality as much as possible, preventing possibly dangerous or critical situations. Using again the car example, a typical case is the diagnostic function added to the electronic control unit which detects anomalies while the car is operating and decides recovery or compensation actions (e.g., limiting performance in case of failures to the braking system). Most of the devices and systems nowadays have sophisticated electronic control and need control/diagnostic software (the evolutions of cars is paradigmatic but other domains such as the aerospace one are even more critical). The two situations are not exclusive and one may have both on-line and off line diagnostic systems; in such a case the former should be able to feed information to the latter.

Off-line and on-line diagnosis impose different requirements to the designer of the diagnostic system. First of all the goal is different and can be stated as Repair vs recovery. The two tasks usually need different types of models or models at different level of granularity. For example, in the on-line case one may even be not interested in localizing the fault as far as the recovery action to be performed is anyway the same. In the on-line case, moreover, abstracting from temporal aspects and dynamic behavior is usually impossible. Moreover, the integration with activities such as tracking/monitoring and planning for recovery is very important. Second, several constraints are imposed by the need to be on-line: for example, real-time response may be needed or at least response before degeneration of the system to a critical and dangerous state.

Off-line vs on-line diagnosis has been studied by Centro Ricerche Fiat (CRF) and UniTo. At CRF a system for off-line diagnosis has been developed in the late 80s (IDEA [9]). The problems arising in embedding diagnostic software in ECU (Electronic Control Units) on-board vehicle has been studied in the European project VMBD

which proposed a methodology to compile fault trees from a model-based system [8, 16]. Similarly compilation of OBDDs has been studied by UniTo for on-line monitoring of a team of robots [35]. Compilation to improve efficiency has also been studied at University of Brescia [30]. CISE studied on-line approaches to monitoring and diagnosis of power generation systems [25].

## 6 Putting diagnosis in the context

Developing the diagnostic agent for a system is part of the overall system life-cycle. In most technical domains many activities are performed during the design of a system, mainly in connection with the definition of its control subsystem. In the aerospace or automotive industry, for example, models of the system are usually adopted to simulate their behavior and the behavior of the control software (the models are quantitative and tools such as Matlab/Simulink are employed).

Unfortunately in current practice not much attention is paid to aspects related to the design of the diagnostic system or even to the analysis that the system is actually diagnosable (e.g., that it contains an appropriate number and placement of sensors that provide data for interpreting and diagnosing the system behaviour). FMECA (Failure Mode Effect and Criticality Analysis) is usually performed after the layout of the system and control strategies have been defined<sup>2</sup>. The definition of diagnostic software is usually performed taking the results of FMECA as the main input.

The need of integrating diagnostic issues in the design process has been recently advocated at industrial and research level [14]. Approaches to integrated tasks such as diagnosability analysis, sensor placement, automated FMECA analysis and support for the design of the diagnostic agent have been proposed. They integrate quantitative models used for control design with qualitative ones and adopt integrate different methodologies (e.g., qualitative simulation, analysis of discrete event systems, model checking, to mention only some of them).

CRF in cooperation with UniTo studied the integration of diagnosis into the design process in the IDD project [40]; Alenia and UniTo developed automatic generation of FMECA from models [41], while researchers at IRST exploited model-checking for verifying diagnosability [12].

## 7 Applications

Model-based diagnosis has been applied to many different areas such as, for example, automotive, aerospace, power generation and transmission systems, telecommunications, medicine, ... Surveying the applications areas is beyond

<sup>2</sup>FMECA is the process of determining, according to a standardized process, the effects of each single fault to the components of a system and the sensors for detecting such faults. The process is performed manually by expert engineers and produces tables which are part of the system design documentation.



the scope of this paper; the interested reader is referred to [36] for more details. Methodological and application perspectives met several times in the model-based community (especially the European one) thanks to the Monet Network of excellence on Model-Based Systems<sup>3</sup>. This integration between theory and application has been very fruitful as it allowed researchers to face realistic application problems and allowed companies to get aware of the technology and to transfer it into internal projects and applications. Thus, ideas from model-based systems can be found today in many applications, especially in automotive, aerospace and telecommunication domains. Some paradigmatic case studies:

*Tiger*, developed by Intelligent Applications [54, 50] is a system for diagnosing Gas Turbines installed in several locations around the world and is probably one of the most important AI success stories of the last decades. NASA applied model-based techniques for monitoring, diagnosis and reconfiguration during the experiment run on board of the spacecraft Deep Space 1 [38]. *IDEA* developed by Centro Ricerche Fiat (CRF) was installed and used in more than 1500 Fiat Diagnostic Centers. *Autosteve* [43] developed by University of Wales (and then by First Earth, bought by Mentor Graphics; another success story for AI) is successfully used in the Ford group to support FMECA generation.

Other relevant industrial applications of diagnosis in Italy have been developed by Alenia (diagnostic systems to support the staff maintaining military aircrafts) Telecom Italia (diagnosis of telecom networks) and CISE- ENEL (diagnosis of power generation process).

While MBD has been mainly applied to technical domains, other approaches to diagnosis have been widely used in the medical domain: particularly relevant is the adoption of Bayesian Networks for dealing with uncertainty [33] and formalisms (such as Neural networks) where knowledge can be learned by data [29]. However, in recent years a lot of attention has been given to tools supporting Clinical Guidelines where medical diagnosis is seen just as a step in the complex activity of patient management and clinical decision making (for Italian contributions to the field see, e.g., [44] and [49]).

## 8 Conclusions

Diagnosis has always been a very active area of research inside AI and, in particular, an area where different methodologies met and where research and applications integrated several times. The Monet Network funded by the EU had in the last decade a very important role to create a strong European research area and community. The International Workshop on Principle of Diagnosis (DX) presents every year the advances in research and the applications being developed.

<sup>3</sup><http://monet.aber.ac.uk>

## REFERENCES

- [1] Special issue on model-based systems. *AI Magazine*, 24(3), 2003.
- [2] P. Baroni, G. Lamperti, P. Pogliano, and M. Zanella. Diagnosis of large active systems. *Artificial Intelligence*, 110(1):135–183, 1999.
- [3] R. Bellazzi, R. Guglielmann, L. Ironi, and M. Stefanelli. Qualitative models and fuzzy systems: an integrated approach for learning from data. *Artificial Intelligence in Medicine*, 1:5–28, 1998.
- [4] C. Berzuini. Modeling temporal processes via belief networks and petri nets, with application to expert systems. *Annals of Math. and Artif. Int.*, pages 39–64, 1990.
- [5] A. Bonarini and P. Sassaroli. Opportunistic multimodel diagnosis with imperfect model. *Information Sciences*, 103:161–185, 1997.
- [6] V. Brusoni, L. Console, P. Terenziani, and D. Theseider Dupre. A spectrum of definitions for temporal model-based diagnosis. *Artificial Intelligence*, 102(1):39–79, 1998.
- [7] T. Bylander, D. Allemang, M. Tanner, and J. Josephson. The computational complexity of abduction. *Artificial Intelligence*, 49(1-3):25–60, 1991.
- [8] F. Cascio, L. Console, M. Guagliumi, M. Osella, A. Panati, S. Sottano, and D. Theseider Dupre. Generating on-board diagnostics of dynamic automotive systems based on qualitative deviations. *AI Communications*, 12(1):33–44, 1999.
- [9] F. Cascio and M. Sanseverino. IDEA (Integrated Diagnostic Expert Assistant) model-based diagnosis in the car repair centers. *IEEE Expert*, 12(6), 1997.
- [10] L. Chittaro and R. Ranon. Hierarchical model-based diagnosis based on structural abstraction. *Artificial Intelligence*, 155(1-2):147–182, 2004.
- [11] L. Chittaro, C. Tasso, and E. Toppano. Putting functional knowledge on former ground. *Applied Artificial Intelligence*, 8(2):239–258, 1994.
- [12] A. Cimatti, C. Pecheur, and R. Cavada. Formal verification of diagnosability via symbolic model checking. In *Proc. IJCAI 2003*, pages 363–369, 2003.
- [13] W.J. Clancey. Heuristic classification. *Artificial Intelligence*, 25(3):289–350, 1985.
- [14] L. Console and O. Dressler. Model-based diagnosis in the real world: lessons learned and challenges remaining. In *Proc. 16th IJCAI*, pages 1393–1400, 1999.
- [15] L. Console, C. Picardi, and M. Ribaud. Process algebras for system diagnosis. *Artificial Intelligence*, 142(1):19–51, 2002.
- [16] L. Console, C. Picardi, and D. Theseider Dupre. Temporal decision trees: Bringing model-based diagnosis on board. *Journal of Art. Int. Research*, 19(1):469–512, 2003.
- [17] L. Console, L. Portinale, and D. Theseider Dupre. Using compiled knowledge to guide and focus abductive diagnosis. *IEEE TKDE*, 8(5):690–706, 1996.
- [18] L. Console, P. Terenziani, and D. Theseider Dupre. Local reasoning and knowledge compilation for efficient temporal abduction. *IEEE TKDE*, 14(6):1230–1248, 2002.



- [19] L. Console, D. Theseider Dupre, and P. Torasso. A theory of diagnosis for incomplete causal models. In *Proc. 11th IJCAI*, pages 1311–1317, 1989.
- [20] L. Console and P. Torasso. A spectrum of logical definitions of model-based diagnosis. *Computational Intelligence*, 7(3):133–141, 1991.
- [21] R. Davis. Diagnostic reasoning based on structure and behavior. *Artificial Intelligence*, 24(1-3):347–410, 1984.
- [22] J. de Kleer. Using crude probability estimates to guide diagnosis. *Artificial Intelligence*, 45(3):381–391, 1990.
- [23] J. de Kleer and B.C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97–130, 1987.
- [24] G. Friedrich. Theory diagnoses: A concise characterization of faulty systems. In *Proc. 13th IJCAI*, pages 1466–1471, 1993.
- [25] M. Gallanti, L. Gilardoni, G. Guida, A. Stefanini, and L. Tomada. Integrating shallow and deep knowledge in the design of an on-line process monitoring system. *Intr. J. of Man-Machine Studies*, 5:641–664, 1987.
- [26] H. Geffner and J. Pearl. Distributed diagnosis of systems with multiple faults. In *Proc Third IEEE Conf. on AI Application*, pages 156–162, 1987.
- [27] W. Hamscher, L. Console, and J. de Kleer, editors. *Readings in Model-based Diagnosis*. Morgan Kaufmann, 1992.
- [28] L. Ironi and S. Tentoni. A model-based approach to the assessment of physicochemical properties of drug delivery materials. *Computers and Chemical Engineering*, 27(6):803–812, 2003.
- [29] I. Kononenko. Machine learning for medical diagnosis: history, state of the art and perspective. *Artificial Intelligence in Medicine*, 23(1):89–109, 2001.
- [30] G. Lamperti and M.Zanella. Diagnosis of discrete-event systems by separation of concerns, knowledge compilation, and reuse. In *Proc. ECAI 04*, pages 838–842, 2004.
- [31] G. Lamperti and M.Zanella. Diagnosis of discrete-event systems from uncertain temporal observations. *Artificial Intelligence*, 137(1-2):91–163, 2005.
- [32] P. Lucas. Analysis of notions of diagnosis. *Artificial Intelligence*, 105(1-2):295–343, 1998.
- [33] P. Lucas, L. van der Gaag, and A. Abu-Hanna. Bayesian networks in biomedicine and health-care. *Artificial Intelligence in Medicine*, 30(3):201–214, 2004.
- [34] C. Marling, M. Sqalli, E. Rissland, H. Munoz-Avila, and D. Aha. Case-based reasoning integrations. *AI Magazine*, 23:69–86, 2001.
- [35] R. Micalizio, G. Torta, and P. Torasso. On-line monitoring and diagnosis of multi-agent systems: A model based approach. In *Proc. ECAI 04*, pages 308–313, 2004.
- [36] R. Milne, editor. *Applications Compendium*. Monet, <http://monet.aber.ac.uk>, 2000.
- [37] I. Mozetic. Hierarchical model-based diagnosis. *Int. J. of Man-Machine Studies*, 35(3):329–362, 1991.
- [38] N. Muscettola, P. Nayak, B. Pell, and B.C. Williams. Remote agent: To boldly go where no ai system has gone before. *Artificial Intelligence*, 103(1-2):5–47, 1998.
- [39] Y. Pencole and M.-O. Cordier. A formal framework for the decentralised diagnosis of large scale discrete event systems and its application to telecommunication networks. *Artificial Intelligence*, 164(1-2), 2005.
- [40] C. Picardi, R. Bray, F. Cascio, L. Console, P. Dague, O. Dressler, D. Millet, B. Rehfus, P. Struss, and C. Vallee. IDD: Integrating Diagnosis in the Design of automotive systems. In *Proc. ECAI 2002*, pages 628–632, 2002.
- [41] C. Picardi, L. Console, F. Berger, J. Breeman, T. Kanakis, J. Moelands, E. Arbaretier, S. Collas, N. De Domenico, E. Girardelli, O. Dressler, P. Struss, and B. Zilbermen. AUTAS: a tool for supporting FMECA generation in aeronautic systems. In *Proc. ECAI 2004*, 2004.
- [42] L. Portinale, D. Magro, and P. Torasso. Multi-modal diagnosis combining case-based and model-based reasoning: a formal and experimental analysis. *Artificial Intelligence*, 158:109–153, 2004.
- [43] C. J. Price. AutoSteve: Automated electrical design analysis. In *Proc. ECAI 2000*, pages 721–725, 2000.
- [44] S. Quaglini, M. Stefanelli, G. Lanzola, V. Caporusso, and S. Panzarasa. Flexible guideline-based patient careflow systems. *Artificial Intelligence in Medicine*, 22(1):65–80, 2001.
- [45] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–96, 1987.
- [46] M. Sachenbacher and P. Struss. Task-dependent qualitative domain abstraction. *Artificial Intelligence*, 162:121–143, 2005.
- [47] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Diagnosability of discrete event systems. *IEEE TAC*, 40:1555–1575, 1995.
- [48] Peter Struss and Chris Price. Model-based systems in the automotive industry. *AI Magazine*, 24(4):17–34, 2004.
- [49] P. Terenziani, G. Molino, and M. Torchio. A modular approach for representing and executing clinical guidelines. *Artificial Intelligence in Medicine*, 23(3):249–276, 2001.
- [50] Tiger. *Real time situation assessment of dynamic systems*. <http://www.intapp.co.uk>.
- [51] P. Torasso. Multiple representations and multi-modal reasoning in medical diagnostic systems. *Artificial Intelligence in Medicine*, 23:46–69, 2001.
- [52] P. Torasso and G. Torta. Compact diagnoses representation in diagnostic problem solving. *Computational Intelligence*, 21:27–68, 2005.
- [53] G. Torta and P. Torasso. Automatic abstraction in component-based diagnosis driven by system observability. In *Proc. IJCAI*, pages 394–402, 2003.
- [54] L. Trave-Massuyes, T. Escobet, and R. Milne. Model-based diagnosability and sensor placement. application to a frame 6 gast turbine subsystem. In *Proc. IJCAI 2001*, pages 551–556, 2001.
- [55] L. Trave-Massuyes, L. Ironi, and P. Dague. Mathematical foundations of qualitative reasoning. *AI Magazine*, 24(4):91–106, 2004.