Model-based diagnosis in the real world: lessons learned and challenges remaining

Luca Console
Dip. Informatica, Universita di Torino
Corso Svizzera 185
1-10149. Torino (Italy)
email: lconsole@di.unito.it

Oskar Dressier
Occ'm Software GmbH,
D-82041 Deisenhofen (Germany)
email: dressler@occm.de
http://www.occm.de

Abstract

This paper discusses some trends in model-based diagnosis. We consider some recent applications and discuss why they were possible, the lessons we learned from them, the new impulse that they gave to research in the field and the new challenges that emerged from them.

1 Introduction

The aim of this paper is to discuss some recent stories concerning the application of model-based diagnosis to industrial domains and the lessons we (as individuals and as a community) learned from such stories.

Model-based Diagnosis (MBD in the following) is an approach to diagnosis that was proposed in the early 80's to overcome limitations of the traditional expert systems approach (see [Hamscher *et* ai., 1992] for a collection of papers and [Davis and Hamscher, 1988] for an introduction). The philosophy of the model-based approach can be sketched as follows. Diagnosis should be based on an objective model of the device (system) to be diagnosed. More specifically, different types of models can be considered: structural (concerning the physical or logical structure of a device), functional (describing the functions of a device), behavioral (describing how a device works, i.e., how its functions are achieved), teleological (describing the purposes of the use of a device), or a combination of them. Models should be reusable, in two ways. On the one hand, the same model of a device should be used for different problem solving tasks (such as diagnosis, simulation, reconfiguration, ...). On the other hand, models should be compositional: the model of a device should be usable in all the cases where the device is used as a component of a larger system.

Since the 80's a lot of work has been carried on and

- several approaches to modeling have been proposed and investigated;
- definitions of diagnosis and formal accounts of the task have been presented;
- algorithms for solving the problem and overcoming its complexity have been designed;
- applications have been developed.

MBD attracted researchers from many different fields of AI (and engineering) and played both the role of testbed of several approaches to KR and reasoning and the role of generator of new problems and approaches.

From an industrial point of view, diagnosis (or, more generally, guaranteeing the availability of systems) is a primary need. Indeed, several diagnostic systems have been implemented since the mid 70's and some of them were success stories for AI. More recently, however, some links became stronger and this led, at least in some areas, to important projects and applications. In particular, the two authors were involved in some of these projects.

In the paper we analyze the reasons for such important evolution pointing out why, in our view, MBD is ready for applications (Sect. 2), the applications that have been built and the lessons we learned from them (Sect. 3), the new research issues that arose from applications (Sect. 4), concluding with some considerations on the implications for future activity in the field (Sect. 5).

2 Why MBD is ready for application

In this section we discuss the reasons that in our view make MBD ready for real application. In order to do that, it is interesting to start by looking at some past seminal talks and papers that strongly influenced our approach to the field [Davis, 1982; 1983; 1984; Hart, 1982; Chandrasekaran and Milne, 1985] (and recent ones, such as [Hamscher, 1991; Struss, 1992b]). Those paper pointed out the potential advantages of the model-based approach, as well as opportunities and goals.

Did we achieve all of those goals and see all the opportunities? Certainly not. Some of them have been widely studied and we now have a clear understanding and available solutions (e.g., many of the problems concerning modeling, the definition of diagnosis and algorithms for computing diagnoses efficiently had major advances). On the other hand, there are many aspects (such as dealing with dynamics or integrating models) that are still open research problems.

However, what was achieved in these years is in some sense more than solving many specific problems. In our view we succeeded in establishing ourselves as a community, with its agreed language and basic definitions. The foundations of the field are well established and we share a clear specification of what the problem is and what its basic ingredients are. A good deal of theoretical work provided semantic foundations for our definitions and almost all of us have this semantics in mind when discussing about research topics or applications.

Let us analyze these foundations in more detail.

First of all, the diagnostic process relies on models which are in most cases component-oriented. A device is described in terms of its minimal replaceable or repairable components. For each type of component the model includes: a list of its variables (interface, internal or state variables, parameters), as well as a definition of its modes of behavior (including correct and fault modes). Then the behavior of the component is described via a set of relations (constraints). Both the nominal (correct) behavior and a set of known faulty behaviors can be described [de Kleer and Williams, 1989; Struss and Dressier, 1989], exploiting different modeling assumptions. For example, a model may be quantitative or qualitative (using different types of qualitative abstractions); dynamic or static (in many cases the dynamics of a system is abstracted for the sake of simplicity); crisp or probabilistic (where incomplete and/or imprecise information is represented using numeric information, as e.g., in probabilistic networks).

The models of the components constitute a library of basic units that can be used for building models of complex devices (which only need to specify the structure of the device in terms of its components). Such compositionality assumption is a fundamental one, allowing the re-use of models (and thus the possibility of building the diagnostic system for a new device without major efforts). Obviously, the assumption has some major consequences on modeling (e.g., the well-known "no structure in function" principle which requires that the model of a component does not depend on the contexts in which the component may be used in a complex device).

Although the component oriented approach is the most widely adopted, it is not the only one. There are cases, in fact, where it is not easy or interesting to isolate components or the behavior of a system is not easily obtained as the composition of the behavior of its components. This is the case, for example, when processes have to be modeled or when the structure of the system is complex, as, e.g., in physiological systems. Thus, there are approaches that are based on process-oriented models that describe the global behavior of a system, usually in terms of causal networks (see e.g., [Patil, 1981; Porcheron *et* al., 1994; Provan, 1998]).

The second foundation is the definition of diagnosis. There is a diagnostic problem whenever the observed behavior of a device is not in accordance with the expected behavior, i.e., the behavior that can be predicted from the model, assuming that all the components are in a normal mode (fault detection). Diagnosis corresponds to removing such an inconsistency and to explaining the faulty behavior. This amounts to finding which components of the system may be faulty (fault isolation) or, more precisely, finding an assignment of a mode to each component in such a way that the observed behavior is explained (fault identification). Different notions of explanation can be adopted, ranging from weaker ones based on consistency [Reiter, 1987; de Kleer *et* al., 1992] to stronger ones based on abduction [Poole, 1989]; see also the analysis in [Console and Torasso, 1991]. Besides the logical characterization, there are also set-theoretic [Peng and Reggia, 1991; Lucas, 1998] and probabilistic ones [Pearl, 1989].

Several other aspects of the diagnostic process have been analysed and formalized: e.g., diagnosis as an iterative process [Struss, 1992a] test/probe selection [de Kleer and Williams, 1987], repair [Priedrich *et al.,* 1994], integration with monitoring [Dvorak and Kuipers, 1989; Lackinger and Nejdl, 1991].

Characterizing diagnoses is a necessary and interesting problem but it is not sufficient. Several algorithms for computing diagnoses that are correct wrt the characterization have been designed. The efficiency issue has been tackled from different points of view. FVom the theoretical one there are analyses of the complexity of the problem and of the special cases in which it is tractable (e.g., see [Levesque, 1989; Bylander *et* a/., 1991; Eiter and Gottlob, 1995; FYiedrich and Nejdl, 1992]). FVom the practical point of view, several strategies for computing diagnoses efficiently have been devised. On the one hand, some approaches defined notions of "preferred diagnoses" (e.g., based on cardinality — single faults diagnoses are the preferred ones — or on other orderings or on fault probabilities) and algorithms for generating only those diagnoses [de Kleer, 1991; Raiman and de Kleer, 1995; Dressier and Struss, 1994]; other introduced focusing techniques for avoiding useless computations (e.g., [Console *et* a/., 1996]) or techniques for re-using inferences [Dressier and FVeitag, 1994]; others discussed how to exploit hierarchical models for focusing diagnosis [Genesereth, 1984; Mozetic, 1991]; yet others coupled MBD with other techniques such as case-based reasoning (for caching results, e.g. [Portinale and Torasso, 1996]), or machine learning. As regards the latter, there are approaches that use simple techniques to induce diagnostic rules from simulated cases (e.g., [Mozetic, 1990; Cascio *et at.,* 1999]), or approaches proposing a closer integration between MBD and learning (e.g., [Baroglio *et* a/., 1994]).

The research on MBD benefited from interactions with many other disciplines of AI and of computer science; to mention some: knowledge representation, non-monotonic reasoning, qualitative reasoning, constraint problem solving, probabilistic reasoning, machine learning, temporal reasoning, control theory, ... In other words, model-based diagnosis has been in the late 80's and 90's what heuristic diagnosis was ten years before: an area that originated new ideas that spread to other fields of AI (the use of the ATMS [de Kleer, 1986] and of abduction are two examples) and an area where new ideas have been tested in practice (e.g., qualitative or non-monotonic reasoning).

The fact that several aspects in the field are mature is demonstrated by the existence of papers that review the diagnostic, task systematically and at the knowledge level, mapping the techniques that can be used for solving diagnostic problems to the features of the application domains. Papers such as [Poole, 1989; Console and Torasso, 1991; ten Teije and van Harmelen, 1994; de Kleer *et* of., 1992; Cordier, 1998; Lucas, 1998; Brusoni *et al.,* 1998}) pointed out the different alternatives that can be considered for characterizing diagnoses, discussing their properties and comparing them. Moreover, some recent projects defined frameworks for selecting the most appropriate approach to diagnosis, given a specific application (e.g., the ARTIST [Leitch *et* ai., 1992; 1994] and PRIDE [Chantler *et al.,* 1998] projects).

Finally, but not less important, MBD has been a meeting point for theoretical and application work. This is not to say that this is not true of other areas of AI, but what happened in MBD is that many people have been active on both sides. This led to interesting research projects and applications. This is a very important topic and we shall devote the next section to it.

In conclusion, let us return to the initial question "why is MBD ready for applications?". A tentative answer is: "Because we have reasonable foundations for the field and understanding of how a diagnortic problem can be tackled". This means that, given a problem, we can envisage how it could be tackled, which techniques could be adopted and which aspects of the problem could be critical or beyond available techniques (and thus would need either to be abstracted or solved in some approximate way). And we have the instruments to actually build systems. Thus, although there are many open problems for research, we can tackle applications without cheating and claiming that everything is solved, but also without having to consider each new application starting from scratch and as a sort of bet.

3 What we did learn from applications

After the discussion in section 2, it is interesting to analyse some of the applications that have been built recently and the lessons we learned from them, focusing on a specific field, automotive systems, on which we gained experience over the last years.

Model-based diagnosis has been applied to many different fields. Historically, one can recognize two mainstreams. On the one hand, there is the approach that focused on the diagnosis of technical systems (mainly circuits in the early years), based on a structural decomposition of the device to be diagnosed and on a description of the (correct or faulty) behavior of each type of component. On the other hand, it is important to mention the important works on medical diagnosis. However, the latter had no strong impact on applications (for several reasons whose analysis is outside the scope of this paper), but had important impacts on the genesis of MBD (consider experiences such as ABEL [Patil, 1981] and CASNET [Weiss et al., 1978]), although the community recognizes itself mostly in the first mainstream.

As mentioned above, electronic devices were the main application used in the early work to experiment ideas and techniques (see most of the papers in [Hamscher *et a*/., 1992]). The diagnosis of copiers led the Xerox Pare group to interesting research on efficient algorithms and focusing strategies. As an example of industrial application, it is worth mentioning the system developed by Hewlett Packard [Allred *et* al., 1991]; CDIAG by Dassault Electronique [Loiez and Taillibert, 1997] is a commercial system for the diagnosis of different types of circuits (including analog one; this shows that it is interesting to have working applications also for domains in which there are still interesting open problems for research).

Although electronic devices were one of the main field for experimenting ideas, most of the applications that are currently on the field come from other areas. For example, the experience in TIGER (an EU founded project) originated a commercial system for the diagnosis of gas turbines that is used in several locations Trave Massuyes and Milne, 1997]. Interesting work was performed on power generation and distribution systems: two applications that are being moved to the field are the DIAPO system by Electricite de Rrance (for nuclear power plants (Porcheron *et* al., 1994]) and a system by Cise for power distribution lines [Lamperti and Pogliano, 1997].

Aerospace is another important area of application (on which many details are not disclosed). Among different applications, it is interesting to mention the role that model-based reasoning and diagnosis are playing in the NASA REMOTE AGENTS project for autonomy in space [Doyle, 1997]; particularly interesting is the LIVINGSTONE system [Williams and Nayak, 1996] for the diagnosis and reconfiguration of space autonomous agents. Other interesting aerospace applications have been developed by Rockwell (see [Provan, 1998J for a recent example); Aerospatiale is experimenting the CDIAG system mentioned above in an automatic test equipment for analog and hybrid boards; similarly Boeing developed interesting applications.

Let us now turn to discussing the automotive domain. There are several reasons why diagnosis (and MBD in particular) became more and more important in this field. First of all, the increasing complexity of the cars (especially from the electronic point of view) called for more sophisticated diagnostic techniques (both on-board the car and in the workshop). Second, legislation required the presence of diagnostic systems in the Electronic Control Unit (ECU) of the car. Third, competition between manufactures led them to investigate new features for attracting customers and for augmenting their satisfaction. Thus, the interest of car manufacturers is growing and they are looking at model-based reasoning as one way of managing the increasing complexity of cars and the high maintenance costs and unnecessary downtime deriving from such a complexity.

The application to the automotive field is very promising from a commercial point of view, but also challenging from the technical point of view, under several viewpoints. First of all, the systems to be diagnosed are-complex, but many of them can be managed with state-of-the-art technologies. Second, different physical domains are involved, ranging from mechanical, to hydraulic, electric and electronic systems. Thus the field is an ideal one for applying and experimenting modeling and diagnostic techniques. Third, most of the critical devices to be diagnosed (especially on-board) are dynamic feedback system with an active control (performed in most of the cases by the software running on the ECUs); diagnosing these system is a very challenging problem and stimulated a lot of work from the research perspective. The application on-board the car imposed several new problems and constraints, mainly due to the resources that are available on-board and to the real-time constraints.

In summary, the application to such a domain led MBD researcher to confront themselves with many interesting problems and, as we shall discuss later, this had a very positive feedback on research. This was due to the possibility of working on real systems, together with the end users. Several results came out of this co-operation. First of all, there is a number of systems, of various nature, that are deployed and used in practice. Among them, it is worth mentioning IDEA, developed by Centro Ricerche Fiat and currently used in more than 1500 workshops around Italy to diagnose some 80 different types of subsystems within a car [Cascio and Sanseverino, 1997]. Another interesting system is MDS by Daimler-Chrysler, which is going to become a commercial product. AUTOSTEVE is a system for the automatic generation of the FMEA (Failure Modes Effect Analysis) in case of electrical circuits [Price, 1998] used by Ford and Jaguar. An approach to diagnose Chevrolet engines is discussed in [Mosterman *et al.,* 1998]. RAZ'R by Occ'm is a modeling and diagnosis development and runtime environment; it is used by Bosch and Volvo within the VMBD project (see below).

Last but not least it is fundamental to mention the VMBD project (funded by the European Union), which involved a number of partners including car manufactures (Daimler-Benz, Fiat, Volvo), suppliers (Bosch, Dassault Electronique, GenRad, Magneti Marelli and Occ'm) and Universities (Aberystwyth, Paris XIII and Turin), and whose goal was to make a step forward in the use of model-based diagnosis both on-board the car and in the workshop. The project was successful from several points of view. As an immediate effect, it led to the development of three prototype applications, based on state of the art techniques, and demonstrated on real cars: the Common Rail fuel delivery system (on a Lancia car), the DTI fuel delivery system and the automatic transmission (on a Volvo car). The models we adopted are mostly qualitative ones (some simple quantitative models had to be used for measurement interpretation and for fault detection). For the two fuel delivery systems we adopted models based on qualitative deviations [Malik and Struss, 1996]. The components are modeled in terms of qualitative differential equations that include appropriate parameters whose values allow for the representation of fault modes of the components. The equations for deviations (in which the variables represent deviations of quantities from expected values) are obtained via algebraic transformations. This form of modeling proved to be very useful since in most cases it is relevant to reason in terms of variables deviating from their expected values, rather than in term of absolute values.

For the off-board system we adopted standard MBD; for the on-board one we experimented the direct use of MBD (for the DTI using Raz'r) and a precompilation approach: in the Common Rail diagnostic trees are synthesized starting from the solutions to a set of simulated cases computed using the model-based approach. This allowed us to produce diagnostic systems that can be easily implemented on current ECUs (see [Cascio *et al.,* 1999] for details).

Probably the most important (even if not immediate) effect of the project is that it contributed to pushing MBD; manufactures and suppliers had the opportunity of having a close look at the technology, using it on their systems and viewing the advantages (and problems) that it can lead to; as a result, some of them decided to invest on MBD. The project was a unique opportunity for us as researchers, not only to experience our ideas and techniques, but also to face new problems and generate a huge number of new research issues. Finally, the project stimulated the interest of other manufactures and suppliers and will probably originate a number of new projects, in several different directions.

In summary, we have that although automotive applications are complex and there are still open research problems, the systems that have been implemented in the last years witness the fact that applications solving real problems can be built. The same story could repeat in other application areas, provided that similar conditions are met. This would require the involvement of end-users in application projects, in co-operation with research centers and suppliers of solutions. In order for this to happen, as naive as it may sound, it is important that researchers work actively in trying to involve companies in their local environment and start projects on real applications. In this sense, in the European context, the role of the funding from the European Union has been fundamental in the automotive domain and could play the same role also in other domains.

From a technical point of view, it is important that the experiences gained in the past projects are re-used. In fact, they can tell a lot on which problems can be solved and which are currently beyond the state of the art and about which techniques are suitable for a given problem. In particular, it would be very important to have frameworks for analysing problems and tools for supporting the construction of model-based systems, without having to re-implement most of the software every time. Some advances in this direction are currently being made but many efforts are still needed. In particular, it would be interesting and important to design open architectures for diagnostic systems and to make available software modules (e.g., as COM or CORBA components) that can be combined to create MBD applications.

4 Applications stimulated new research

In the previous sections we stressed the idea that it is important to work on applications even if not all the research problems are solved. In this section we want to stress another issue: building applications is not the end of research. Actually, working on applications is a big advantage for research, rising new relevant research issues. This indeed has been the case for model-based diagnosis, in our experience. Thus, it is fundamental that researchers work on both theory and applications; keeping the two aspects too distant is a big mistake, especially for researchers who risk to spend a lot of time and efforts on artificial problems, each one attracting the interest only of a closed community.

Thus, let us analyse in more detail some of the research issues that arose from the recent applications of model-based diagnosis. As in previous sections, we shall base the discussion on our experience in the automotive domain (and obviously the discussion will be partial).

Different modeling techniques have been used in the applications discussed in the previous section. Each model relies on some assumptions, being an abstraction of the actual system. Although the problem of choosing the appropriate models, given the requirements of the diagnostic task, has been pointed out already in [Davis, 1983], there is still a lot to be done for clarifying the issue. In particular, it would be interesting to have a general characterization and criteria that could help in the choice of the model that is most suitable for a given problem. This does not mean, however, that new modeling techniques are not needed. Several applications, for example, pointed out that although qualitative models are very useful, some form of quantitative reasoning would be important, especially in the early phases of the diagnostic process (fault deletion in particular) and for discrimination purposes. Thus, research on the integration of qualitative and quantitative reasoning (or on semi-quantitative reasoning) is very important.

An ideal output of research on modeling should lead to a sort of "intelligent model library", i.e., a repository of models of components that could be used in different applications. Reusability is in fact an important issue which still needs a lot of research. In fact, although in principle models of components are reusable, in practice this is a difficult problem. In order to improve reusability one should have different models of the same component, at different levels of granularity and based on different modeling assumptions. The repository should be intelligent and should provide guidelines to select the most appropriate model for a given application.

One specific research issue on which a lot of attention concentrated since the early age is the problem of dealing with dynamics [Hamscher and Davis, 1984]. This is in fact a critical aspect to be taken into account in most applications. Several approaches to deal with dynamic (and time-varying) behavior have been proposed (see the discussion in [Brusoni *et al.*, 1998]) and there is a very active debate (in which the experiments on applications play a fundamental role). It is worth noting that some of the proposals introduced in the MBD community techniques that are used in other fields; two examples are the use of Bond graphs in [Mosterman and Biswas, 1997] and of techniques coming from system engineering, as e.g., in [Chantler *et al.*, 1996]. Closer connections with these (and other) fields is a topic for future research.

The work on automotive applications pointed out that the diagnostic process, especially when performed on-board the car, is quite complex. The need of integrating diagnosis (in its restricted sense of interpretation or explanation of observations) and planning for scheduling (and performing) tests or actions was recognized already several years ago (e.g., [Friedrich *et* al., 1994; Struss, 1992a]). However, in many cases the problem is more articulated and difficult than anticipated and there are aspects and opportunities that have not been considered yet. For example, in several systems it is interesting and important to perform active tests when the operating conditions allow them (e.g., when a car is running at a constant engine speed and the pedals are in specific positions). This requires the ability of scheduling such tests (to avoid performing them in critical situations), performing some action and interpreting its effects. There are many opportunities for research on this problem, linking work on planning and theories of action and change and on perception (e.g. vision or sensing) more closely to the diagnostic task.

The future of diagnosis in automotive domains is the application on-board the car. In VMBD we experimented the use of MBD for on board diagnosis but in order to sit on current technology ECUs (which have a few kilobytes of memory), we also considered the use of MBD to generate decision trees to be used on-board. MBD proved to be very flexible for such a task. However, the possibility of actually using model-based diagnosis on-board has to be investigated further. This, in fact, would be more flexible, opening new opportunities (and a closer integration with the off-board system to which the on-board one could pass useful information). Would it be conceivable to define a run-time model-based system, with reduced functionality but that could be implemented in next generation ECUs? Will the idea of having an on-board car-PC be pursued by car manufactures and suppliers? This would be an fundamental opportunity for the growth of our field.

Still considering on-board applications, it would be interesting to investigate tele-diagnostics, having a satellite connection between a car and an assistance center (for performing diagnosis or actions on the car or suggesting the driver what to do or how to reach the closest workshop, which, at the same time could be alerted). Such a centralized model would be important also for simplifying the maintenance of models which could be done only on the server. An important role could be played also by approaches to preventive diagnosis, which needs further studies. For example, one could recognize a non-optimal behavior of some component (possibly exploiting wear models), and tune the system to avoid such behaviors or warn the user or just keep track of the problem until the next maintenance.

However, the most important research (and practical) issue risen by research on automotive domains does not concern a specific step within the diagnostic process, but rather the role of diagnosis along the design manufacturing and service chain. What emerged clearly is that the current situation in which diagnosis is not integrated along this chain is not working. Building models for diagnosis *a-poateriori,* when devices are manufactured is not a good choice. It makes diagnosis more difficult (especially as regards modeling) and does not allow to exploit all the potential advantages of the model-based approach. Indeed, model-based diagnosis should be integrated along the process; in such a way, the diagnosabiiity of the system being designed could be tested and design choices (e.g., selecting which sensors to use and where) could be influenced by these tests. This requires the availability of a library of models of basic components, (library that could be extended when new components are in turn designed) and the integration of model-based diagnosis into design tools. In this way, the models for diagnosis could be an automatic result of the design process, as well as other aspects such as, e.g., the FMEA of the device to be used in documentation manuals. This integration would provide significant benefits also for design since it would avoid the redesign of systems which do not meet diagnosability requirements. Some ideas in the direction of integrating diagnosis and design were presented in the past (e.g., the work on design for testability [Williams and Parker, 1983] and on sensor placement [Scarl, 1994]). However, this is a topic requiring many efforts, possibly involving people working on design and companies that produce design tools.

The list of new topics could be longer, but providing it is not the goal of this paper. Our goal is to stress that many of these problems would have not been recognized without working on real with end users. This work provided a new important impulse to research, impulse that would have been impossible without the applications.

5 Some implications for the future

In the previous sections we discussed the important role that applications had in the recent advances in MBD. They allowed researchers to experiment and validate their techniques; this work pointed out a number of new important research issues. Thus, applications allowed the field to grow with new impulses, opening new research opportunities and avoiding to be confined on sterile, purely academic issues.

Working on real applications requires the meeting of people from the research part and users with application problems. The recent successes in the field of MBD point out that this requires efforts from both sides. First of all, there must be an effort to create the opportunity for the meeting; after that there must be an effort in understanding each other. A number of conditions facilitated these aspects in the case of MBD. First of all, we stress once again the importance of presenting ourselves to companies as a community with a strong common background. This is not to say that we agree on everything, but this means that the probability of having completely different answers on basic questions and problems when discussing with different researchers is low. This avoid confusion into potential "customers". Second, we are a community which already involves people from the application side, working in corporate research centers or in end-user companies. The participation to the annual "Workshop on Principles of Diagnosis" is not only academic. This was the result of big efforts; we always tried to involve application people, maybe organizing special events. This interest is now a witness for attracting other people from the application side.

Finally, on the European side an important role has been played by funding from the European Union. A part from many specific protects, an important initiative is the MONET [Monet, J network of excellence on model-based systems and qualitative reasoning. MONET involves some 70 nodes, including Universities, research centers and companies and its role is the diffusion and promotion of model-based techniques. The network is very active in this task, especially towards companies, creating thus the background for facilitating the meeting between research and application and thus the conditions for new projects. Obviously, for that to happen, it is important that all researchers in the field look around for finding potential applications and that companies invest on such applications. Past experiences show that there is a chance that the investments do not get lost.

Prom the research point of view, the analysis in the paper points out that there are many open issues and opportunities for co-operation with many other fields of AI, exactly as it was in the past.

6 Acknowledgements

The authors are indebted to the partners in the VMBD project (BE 95/2128). Special thanks to D. Theseider Dupre, P. Struss, P. Torasso, for their comments.

References

[Allred *et* al., 1991] D. Allred, Y. Lichtenstein, C. Preist, M. Bennet, and A. Gupta. AGATHA: an integrated expert system to test and diagnose personal computers boards. In R. Smith and A. Carlisle Scott, editors, *Innovative Applications of AI* 5, pages 87-104. AAAI Press, 1991.

[Baroglio *et* al., 1994] C. Baroglio, M. Botta, and L. Saitta. WHY: a system that learns using causal models and examples. In R. Michalski and G. Tecuci, editors, *Machine Learning, an Artificial Intelligence Approach (vol IV).* Morgan Kaufman, 1994.

[Brusoni *et al.,* 1998] 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.

[Bylander *et al.,* 19911 T. Bylander, D. Allemang, M. Tanner, and J. Josephson. The computational complexity of abduction. *Artificial Intelligence,* 49(1-3) :25-60,1991.

[Cascio and Sanseverino, 1997] F. Cascio and M. Sanseverino. IDEA (Integrated Diagnostic Expert Assistant) model-based diagnosis in the the car repair centers. *IEEE Expert, 12(6),* 1997.

[Cascio *et al.,* 1999] 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), 1999.

[Chandrasekaran and Milne, 1985] B. Chandrasekaran and R. Milne. Special section on reasoning about structure, behavior and function. In *Sigart Newsletter,* volume 93, pages 4-55. 1985.

[Chantler *et al.,* 1996] M.J. Chantler, S. Daus, T. Vikatos, and G.M. Coghill. The use of quantitative dynamic models and dependency recording for diagnosis. In Proc*. 7th Int. Work, on Principles of Diagnosis,* pages 59-68, 1996.

[Chantler *et ai,* 1998] M.J. Chantler, G.M. Coghill, Q Shen, and R.R. Leitch. Selecting tools and techniques for model-based diagnosis. *Artificial Intelligence in Engineering,* 12:81-98,1998.

[Console and Torasso, 1991] L. Console and P. Torasso. A spectrum of logical definitions of model-based diagnosis. *Computational Intelligence,* 7(3): 133-141,1991.

[Console *et ai,* 1996] L. Console, L. Portinale, and D. Theseider Dupre. Using compiled knowledge to guide and focus abductive diagnosis. *IEEE Transactions on Knowledge and Data Engineering,* 8(5):690-706, 1996.

[Cordier, 1998] M.-O. Cordier. When abductive diagnosis fails to explain too precise observations: an extended spectrum of definitions of MBD based on abstracting observations. In *Proc. 9th Int. Work, on Principles of Diagnosis,* pages 24-31, 1998.

[Davis and Hamscher, 1988] R. Davis and W. Hamscher. Model-based reasoning: Troubleshooting. In H.E. Shrobe, editor, *Exploring Artificial Intelligence,* pages 297-346. Morgan Kaufman, 1988.

[Davis, 1982] R. Davis. Expert systems: Where are we? And where do we go from here? *AI Magazine,* 3(2):3-22,1982.

[Davis, 1983] R. Davis. Diagnosis via causal reasoning: paths of interaction and the locality principle. In *Proc AAAI88,* pages 88-94, 1983.

[Davis, 1984] R. Davis. Diagnostic reasoning based on structure and behavior. *Artificial Intelligence,* 24(1-3):347~410,1984.

[de Kleer and Williams, 1987] J. de Kleer and B.C. Williams. Diagnosing multiple faults. *Artificial Intelligence,* 32(1):97-130,1987.

[de Kleer and Williams, 1989] J. de Kleer and B.C. Williams. Diagnosis with behavioral modes. In *Proc. 11th IJCAI,* pages 1324-1330, Detroit, 1989.

[de Kleer *et al.,* 1992] J. de Kleer, A. Mackworth, and R. Reiter. Characterizing diagnoses and systems. *Artificial Intelligence,* 56(2-3):197-222,1992.

[de Kleer, 1986] J. de Kleer. An assumption-based TMS. *Artificial Intelligence,* 28(2):127-162,1986.

[de Kleer, 1991] J. de Kleer. Focusing on probable diagnoses. In *Proc. AAAI 91,* pages 842-848,1991.

[Doyle, 1997] R. Doyle. The emergence of spacecraft autonomy. In *Proc. AAAI 97,* pages 756-761,1997.

[Dressier and Freitag, 1994] 0. Dressier and H. Freitag. Prediction sharing among time and contexts. In *Proc. AAAI 94,* pages 1136-1141,1994.

[Dressier and Struss, 1994] O. Dressier and P. Struss. Model-based diagnosis with the default-based diagnosis engine: Effective control strategies the work in practice. In *Proc. 11th ECAI,* pages 677-681,1994.

[Dvorak and Kuipers, 1989] D. Dvorak and B. Kuipers. Model-based monitoring of dynamic systems. In *Proc. 11th IJCAI,* pages 1238-1243, Detroit, 1989.

[Eiter and Gottlob, 1995] T. Eiter and G. Gottlob. The complexity of logic-based abduction. *Journal of the ACM,* 42(l):3-42,1995.

[Priedrich and Nejdl, 1992] G. Friedrich and W. Nejdl. MOMO: Model-based diagnosis for everybody. In W. Hamscher, L. Console, and J. de Kleer, editors, *Readings in Model-based Diagnosis,* pages 143-149. Morgan Kaufman, 1992.

[Priedrich *et* al., 1994] G. FViedrich, G. Gottlob, and W. Nejdl. Formalizing the repair process - extended report. *Annals of Mathematics and Artificial Intelligence,* II(l-4):187-202,1994.

[Genesereth, 1984] M.R. Genesereth. The use of design descriptions in automated diagnosis. *Artificial Intelligence,* 24(l-3):411-436,1984.

[Hamscher and Davis, 1984] W. Hamscher and R. Davis. Diagnosing circuit with state: an inherently underconstrained problem. In *Proc. AAAI 84,* pages 142-147,1984.

[Hamscher *et al.,* 1992] W. Hamscher, L. Console, and J. de Kleer, editors. *Readings in Model-Based Diagnosis.* Morgan Kaufmann, 1992.

[Hamscher, 1991] W. Hamscher. Principles of diagnosis: Current trends and report on the first international workshop. *AI Magazine,* 12(4):15-23,1991.

[Hart, 1982] P.E. Hart. Directions for AI in the eighties. *Sigart Newsletter,* 79:11-16,1982.

[Lackinger and Nejdl, 1991] F. Lackinger and W. Nejdl. Integrating model-based monitoring and diagnosis of complex dynamic systems. In *Proc. 12th IJCAI,* pages 1123-1128, Sydney, 1991.

[Lamperti and Pogliano, 1997] G. Lamperti and P. Pogliano. Event-based reasoning for short circuit diagnosis in power transmission networks. In *Proc. IJCAI 97,* pages 446-451* 1997.

[Leitch *et* dL, 1992] R Leitch, H. Freitag, Q. Shen, P. Struss, and G. Ibrnielli. Artist: A methodological approach to specifying model-based diagnostic systems. In G. Guida and A. Stefanini, editors, Jn*dustrial Applications of Knowledge-Based Diagnosis,* pages 183-208/Elsevier, 1992.

[Leitch *et* of., 1994) R. Leitch, M. Chantler, Q. Shen, and G. Coghill. A preliminary specification methodology for model-based diagnosis. *Annals of Mathematics and Artificial Intelligence,* 11(1-4):11~32,1994.

[Levesque, 1989] EL J. Levesque. A knowledge-level account of abduction. In *Proc. 11th IJCAI,* pages 1061-1067, Detroit, 1989.

[Loiez and T&llibert, 1997] E. Loiez and P. Taillibert. Polynomial temporal band sequences for analog diagnosis. In *Proc. IJCA1* 97, pages 474-479,1997.

[Lucas, 1998] P. Lucas. Analysis of notions of diagnosis. *Artificial Intelligence,* 105(1-2) :295~343,1998.

[Malik and Struss, 1996] A. Malik and P. Struss. Diagnosis of dynamic systems does not necessarily require simulation. In *Proc. 7th Int. Work, on Principles of Diagnosis,* pages 147-156,1996.

[Monet,] Monet. *The European network of excellence in model-based systems and qualitative reasoning.* http://monet.aber.ac.uk.

[Mosterman and Biswas, 1997] P.J. Mosterman and G. Biswas. Monitoring, prediction and fault isolation in dynamic physical systems. In *Proc. AAAI 97,* pages 100-105,1997.

[Mosterman *et* o/., 1998] P.J. Mosterman, G. Biswas, and E. Manders. A comprehensive framework for model based diagnosis. In *Proc. 9th Int. Work, on Principles of Diagnosis,* pages 86-93, 1998.

[Mozetic, 1990] I. Mozetic. Diagnostic efficiency of deep and surface knowledge in KARDIO. *Artificial Intelligence in Medicine,* 2(2):67~83, 1990.

[Mozetic, 1991] I. Mozetic. Hierarchical model-based diagnosis. *Int. J. of Man-Machine Studies,* 35(3):329-362,1991.

[Patil, 1981] R. Patil. Causal representation of patient illness for electrolyte and acid-base diagnosis. Technical Report LCS-267, MIT, Cambridge, MA, 1981.

[Pearl, 1989] J. Pearl. *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufmann, 1989.

[Peng and Reggia, 1991] Y. Peng and J. Reggia. i4Muc*tive inference models for diagnostic problem solving.* Springer-Verlag, 1991.

[Poole, 1989] D. Poole. Normality and faults in logic-based diagnosis. In *Proc. 11th IJCAI,* pages 1304-1310, Detroit, 1989.

[Porcheron *et al.,* 1994] M. Pordberon, B. Rieard, J.-L. Busquet, and P. Parent. DIAPO: a case study in applying advanced AI techniques to the diagnosis of a complex system. In *Proc. ttth ECAI,* pages 43-47,

[Portinale and Torasso, 1996] L; Portinale and P. Torasso. On the usefulness of reusing diagnostic solutions. In W. Wahlster,editor, *Proc. ECAI 96,* pages 137-141. J. Wiley, 1996.

[Price, 1998] C. Price. Function-directed electrical design analysis. *Artificial Intelligence in Engineering,* 12:445-456,1998.

[Provan, 1998] G. Provan. Modeling and diagnosing rocket engines: a multimodeling approach. In *Proc. 9th Int. Work, on Principles of Diagnosis,* pages 102-109,1998.

[Raiman and de Kleer, 1995] O. Raiman and J. de Kleer. Trading off the cost of inference vs. probing in diagnosis. In *Proc. 14th IJCAI,* pages 1736-1741, Sydney, 1995.

[Reiter, 1987] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence,* 32(1):57~96, 1987.

[Scarl, 1994] E. Scarl. Sensor placement for diagnosability. *Annals of Mathematics and Artificial Intelligence,* II(I-4):493-510,1994.

[Struss and Dressier, 1989] P. Struss and O. Dressier. Physical negation - integrating fault models into the general diagnostic engine. In *Proc. 11th IJCAI,* pages 1318-1323, Detroit, 1989.

[Struss, 1992a] P. Struss. Diagnosis as a process. In W Hamscher, L. Console, and J. de Kleer, editors, *Readings in Model-based Diagnosis,* pages 408-418. Morgan Kaufman, 1992.

[Struss, 1992b] P. Struss. Knowledge-based diagnosis: an important challenge and touchstone for AI. In *Proc. 10th ECAI,* pages 863-874, Vienna, 1992.

[ten Teije and van Harmelen, 1994] A. ten Teije and F. van Harmelen. An extended spectrum of logical definitions of diagnosis. In *Proc. 5th Int. Work, on Principles of Diagnosis,* pages 334-342, 1994.

[Trave Massuyes and Milne, 1997] L. Trave Massuyes and R. Milne. Gas turbines condition monitoring using qualitative model-based diagnosis. *IEEE Expert,* 12(3):21-31,1997.

[Weiss *et al,* 1978] S. Weiss, C. Kulikowski, S. Amarel, and A. Safir. A model based method for computer-aided medical decision making. *Artificial Intelligence,* 11(1-2):145-172,1978.

[Williams and Nayak, 1996] B. Williams and P. Nayak. A model-based approach to reactive self-configuring systems. In *Proc. AAAI 96,* pages 971-978,1996.

[Williams and Parker, 1983] T.W. Williams and K.P. Parker. Design for testability: A survey. *Proceedings of the IEEE,* 71(1):98-112,1983.