CORA Cork Open Research Archive
Cartlann Taighde Oscailte Chorcaí

Title	A general characterization of Model-Based Diagnosis
Authors	Provan, Gregory
Publication date	2016-01
Original Citation	Provan, G. (2016) 'A general characterization of Model-Based Diagnosis' (from proceedings of the Twenty-second European Conference on Artificial Intelligence (ECAI 2016),The Hague, The Netherlands, 29th August - 2nd September), Frontiers in Artificial Intelligence and Applications, 285, pp. 1565-1566. doi:10.3233/978-1-61499-672-9-1565
Type of publication	Conference item
Link to publisher's version	http://www.ecai2016.org/ - 10.3233/978-1-61499-672-9-1565
Rights	© 2016, the Authors and IOS Press. This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0) - https://creativecommons.org/licenses/by-nc/4.0/
Download date	2026-11-01 21:49:15
Item downloaded from	https://hdl.handle.net/10468/3770

UCC
University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

ECAI 2016
G.A. Kaminka et al. (Eds.)

© 2016 The Authors and IOS Press.
This article is published online with Open Access by IOS Press and distributed under the terms of the Creative Commons Attribution Non-Commercial License 4.0 (CC BY-NC 4.0).
doi:10.3233/978-1-61499-672-9-1565

A General Characterization of Model-Based Diagnosis

Gregory Provan[1]

Abstract. The Model-Based Diagnosis (MBD) framework developed by Reiter has been a strong theoretical foundation for MBD, yet is limited to models that are described in terms of logical sentences. We propose a more general framework that covers a wide range of modelling languages, ranging from AI-based languages (e.g., logic and Bayesian networks) to FDI-based languages (e.g., linear Gaussian models). We show that a graph-theoretic basis for decomposable system models can be augmented with several languages and corresponding inference algorithms based on valuation algebras.

1 A GENERAL MBD FRAMEWORK

We propose a framework for extending the Reiter MBD approach [7] by integrating several MBD and FDI approaches within a decomposable graphical framework in which the modeling language and inference are specified by a valuation algebra [6].

More formally, we define an MBD framework using a tuple $(\mathcal{G}, \mathcal{T}, \Gamma)$, where $\mathcal{G}$ is a factor graph [5], $\mathcal{T}$ is the diagnosis task, and Γ is a valuation algebra [6]. The *factor graph* $\mathcal{G}$ specifies a system topology in terms of a decomposable relation Ψ, defined over a set $\boldsymbol{V}$ of variables, such that $\Psi = \otimes_i \psi_i(\boldsymbol{V}_i)$, where ψ_i is a sub-relation, $\otimes$ is the composition operator and $\boldsymbol{V}_i \subseteq \boldsymbol{V}$. This decomposition can be mapped to a graph, e.g., a DAG for a Bayesian network (BN) or an undirected graph for a Markov network. The *diagnosis task* is given by the tuple $\mathcal{T} = (\mathcal{D}, \boldsymbol{y}, \mathcal{R})$, where $\mathcal{D}$ is the task specification; $\boldsymbol{y}$ is the required system measurement; and residual $\mathcal{R}(\Psi, \boldsymbol{y})$ indicates a discrepancy between observation $\boldsymbol{y}$ and model-derived prediction $\hat{\boldsymbol{y}}$ using some distance measure. The valuation algebra Γ specifies (1) the underlying language for the diagnosis system, and (2) the inference necessary to compute the diagnosis for the task $\mathcal{T}$, e.g., multiple-fault subset-minimal diagnosis or Most-Probable Explanation.

This decomposable representation can encode a wide range of diagnosis models, including propositional logic models, FDI dynamical systems models, as well as stochastic models (Bayesian networks, HMMs, and linear Gaussian models). Previous work, e.g., [4] has shown AI-based approaches to diagnosis [7, 2] can be described by valuation algebras. Here, we extend this to include FDI approaches based on ordinary differential equations (ODEs), and show the importance of system structure and diagnosis task in specifying the full diagnosis representation.

This framework has several outcomes. First, it enables a clear separation between models and inference (although the two are linked). Specifically, the model structure encoded as a factor graph that governs inference complexity. For example, tree-structured factor graphs are all poly-time computable. Second, the factor graph encoding of models clarifies the structural difference between AI-based approaches and FDI-based approaches.

[1] University College Cork, Cork, Ireland, email: g.provan@cs.ucc.ie. Supported by Science Foundation Ireland (SFI) Grant SFI/12/RC/2289.

2 REPRESENTING MULTIPLE MODEL TYPES

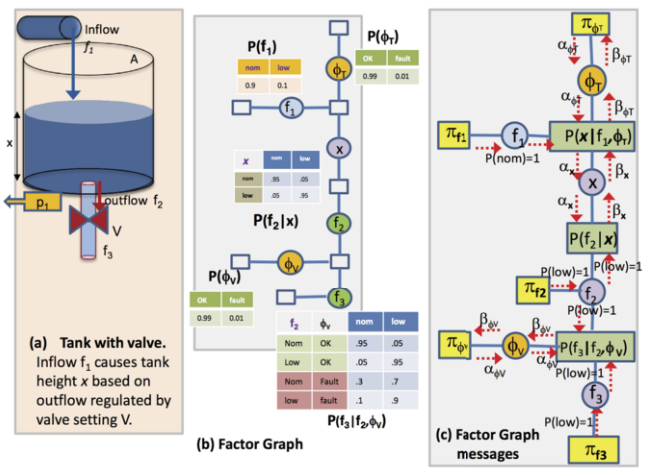


Figure 1. Bayesian network for controlled tank example.

We can represent a system (or plant) model Ψ using a factor graph, which represents the physical connectivity of Ψ in terms of a structured decomposition of Ψ into sub-relations. Consider Figure 1(a), which shows an example of a tank with a valve, where we control the level x in the tank by controlling the inflow f_1 and the valve state V. There are two possible failures in the system: (1) the tank can leak, with failure mode ϕ_T, and (2) the valve can malfunction, with failure mode ϕ_V. This example has variables $\{f_1, f_2, f_3, x, \phi_T, \phi_V\}$ and three relations: $\psi_1(f_1, \phi_T, x)$, $\psi_2(x, f_2)$, and $\psi_3(f_2, f_3, \phi_V)$. ψ_1 represents how the tank's fluid level x is governed by inflow f_1 and fault (leak) ϕ_T, ψ_2 represents how outflow is governed by fluid height x, and ψ_3 represents the valve's impact on flows f_2, f_3.

Given such a decomposition, we can represent the modelling language as a valuation over $\psi_1 \otimes \psi_2 \otimes \psi_3$. For example, if we choose a probabilistic algebra then we obtain a diagnostic BN model, for which Figure 1(b) shows the structure and valuation $P(\boldsymbol{V}) = P(f_1)P(\phi_T)P(\phi_V)P(f_1|\phi_t, x)P(x|f_2)P(f_3|f_2, \phi_V)$. $P(x|f_1, \phi_T)$ defines the conditional dependence of tank level x on the inflow f_1 and the tank fault-state ϕ_T, and (2) $P(f_3|f_2, \phi_V)$ the conditional dependence of flow from this system, f_3, on the tank outflow f_2 and the valve fault-state ϕ_V.

	Approach	Language	Structure	Task	Semi-Ring	RA	FI	Complexity
Atemporal	Reiter	Prop. Logic	DAG	Dx	$\langle\{0,1\},\wedge,\vee\rangle$	$\Psi^{\downarrow\emptyset}$	$\Psi^{\downarrow\phi}$	NP-hard
	ATMS	M-V Logic	DAG	all Dx	$\langle 2^{\mathcal{P}},\cap,\cup\rangle$	$\Psi^{\downarrow\emptyset}$	$\Psi^{\downarrow\phi}$	NP-hard
	Qualitative	Q-Relation	arbitrary	Dx	$\langle 2^{\mathcal{P}},\oslash,\odot\rangle$	$\Psi^{\downarrow\emptyset}$	$\Psi^{\downarrow\phi}$	NP-hard
	Constraint Network	Constraint	arbitrary	all Dx	$\langle 2^{\mathcal{P}},\times,+\rangle$	$\Psi^{\downarrow\emptyset}$	$\Psi^{\downarrow\phi}$	NP-hard
	BN-Posterior	PGM	DAG	$P(\boldsymbol{\phi}\|\boldsymbol{y})$	$\langle[0,1],\times,+\rangle$	$\Psi^{\downarrow\emptyset}$	$\Psi^{\downarrow\phi}$	NP-hard
	BN-MPE	PGM	DAG	MPE	$\langle[0,1],max,\times\rangle$	$\Psi^{\downarrow\emptyset}$	$\Psi^{\downarrow\phi}$	NP-hard
Temporal	DES	M-V Logic	arbitrary	all Dx	$\langle 2^{\mathcal{P}},\cap,\cup\rangle$	$\Psi^{\downarrow\emptyset}$	$\Psi^{\downarrow\phi}$	NP-hard
	DBN	PGM	DAG	$P(\boldsymbol{\phi})$	$\langle[0,1],\times,+\rangle$	$\Psi^{\downarrow\emptyset}$	$\Psi^{\downarrow\phi}$	NP-hard
	HMM	PGM	tree	$P(\boldsymbol{\phi})$	$\langle[0,1],\times,+\rangle$	$\Psi^{\downarrow\mathcal{R}}$	$\Psi^{\downarrow\phi}$	$O(n)$
	FDI	ODE	bipartite	Dx	$\langle[0,1],\times,+\rangle$	$\Psi^{\downarrow\mathcal{R}}$	$\Psi^{\downarrow\phi}$	NP-hard
	Kalman filter	PGM	tree	MPE	$\langle[0,1],\times,+\rangle$	$\Psi^{\downarrow\mathcal{R}}$	$\Psi^{\downarrow\phi}$	$O(n^3)$
	Particle filter	PGM	arbitrary	MPE	$\langle[0,1],\times,+\rangle$	$\Psi^{\downarrow\mathcal{R}}$	$\Psi^{\downarrow\phi}$	NP-hard

Table 1. Classification of Diagnosis Approaches using Generalized Approach. RA and FI denote Residual Analysis and Fault Isolation, respectively. Shaded rows denote stochastic methods, and unshaded rows denote deterministic methods. BN denotes Bayesian-network. For the language, Prop and M-V denote propositional and multivalued respectively; Q-Relation denotes Qualitative Relation; PGM refers to probabilistic graphical model, and ODE to Ordinary Differential Equation. For the task, Dx corresponds to a single diagnosis, $P(\boldsymbol{\phi})$ to the probability distribution over $\boldsymbol{\phi}$, and MPE to Most-Probable Explanation.

We perform inference in Ψ by message-passing and valuation updating [6]. Figure 1(c) shows how we can compute a diagnosis (i.e., evaluate $\phi_T,\ \phi_V$) by passing messages among the nodes starting with the control and observation settings $\mathcal{S}$. If all assignments in $\mathcal{S}$ are nominal (*nom*), the "diagnosis" is $P(\phi_T = fault) = .004$ and $P(\phi_V = fault) = .004$. If scenario $\mathcal{S}$ has control assignment f_1=*nom*, and observation of $f_2,\ f_3$ both as *low*, we obtain a diagnosis given by $P(\phi_T = fault) = .067$ and $P(\phi_V = fault) = .009$, i.e., a faulty tank is the most likely diagnosis.

We can represent several different tank models by keeping the same decomposable structure and changing the valuation. For example, we obtain a qualitative model by replacing (1) the conditional probability tables with qualitative relations, and (2) passing qualitative messages (e.g., {+,-,0} rather than discrete-valued probabilities), and using qualitative inference rather than Bayesian updating.

Table 1 summarizes the properties of several models characterized by our framework, defining the language, model structure, the underlying semi-ring, and the inference complexity. The language and task can be characterized by the valuation semi-ring (Z, O_1, O_2), which consists of a set Z and two operations (O_1, O_2) [3]. The last column of Table 1 shows the inference complexity for which the primary determinant is the topology [1]: any non-tree topology is likely to be NP-hard for computing a task requiring at least one multiple-fault diagnosis, whereas tree topologies are poly-time solvable for the majority of languages and tasks.

A **valuation** is a measure over the possible values of a set $\boldsymbol{V}$ of variables [3]. Each valuation ψ refers to a finite set of variables $d(\psi) \subseteq \boldsymbol{V}$, called its domain. Given the power set $\mathcal{P}$ of $\boldsymbol{V}$ and a set ψ of valuations with their domains in $\mathcal{P}$, we can define 3 key operations: (1) *Labeling*: $\psi \mapsto d(\psi)$, which returns the domain of each valuation; (2) *Combination*: $(\psi_1, \psi_2) \mapsto \psi_1 \otimes \psi_2$, which specifies functional composition, i.e., the aggregation of data from multiple functions; (3) *Projection*: $(\psi, V) \mapsto \psi^{\downarrow V}$ for $V \subseteq d(\psi)$, which specifies the computation of a query (set of variables) of interest.

Given an observation $\boldsymbol{y}$, we specify diagnosis within a valuation algebra as a two-step process of: (1) residual analysis (RA); and (2) fault isolation (FI).

Residual analysis: This inference depends on the type of residual. AI logic-based approaches compute RA using a consistency check, denoted $\Psi^{\downarrow\emptyset}$. FDI continuous-valued systems compute RA as $\mathcal{R} = |\hat{\boldsymbol{y}} - \boldsymbol{y}|$, where $\hat{\boldsymbol{y}}$ is the model's prediction. Residual-specific FG structure may be necessary to enable us to compute $\Psi^{\downarrow\mathcal{R}}$.

Fault isolation: Isolating a diagnosis is equivalent to projecting the marginal over the fault-mode variables $\boldsymbol{\phi}$, denoted $\Psi^{\downarrow\phi} = (\psi_1 \otimes \cdots \otimes \psi_n)^{\downarrow\boldsymbol{\Phi}}$. Diagnostic inference requires all 3 valuation operations, in particular combination and projection. The task also may change the FG structure and operations required. For example, different operations are required for computing a posterior distribution $P(\boldsymbol{\phi}|\boldsymbol{y})$ as opposed to the Most Probable Explanation (MPE).

Given an observation $\boldsymbol{y}$ and prediction $\hat{\boldsymbol{y}}$, the typical objective of a diagnosis process is to identify the system fault-state that minimises the residual vector: $\phi^* = \operatorname{argmin}_{\phi\in\Phi} \mathcal{R}(\Psi, \boldsymbol{y})$. The full paper generalizes the inference metric to define our diagnosis task as jointly minimizing the accuracy (based on $\mathcal{R}$) and the inference complexity.

3 CONCLUSION

This article has presented a general framework for MBD that integrates several approaches developed within different communities, most notably the AI and FDI communities. By characterizing MBD using the triple $(\mathcal{G}, \mathcal{T}, \Gamma)$ we show structural similarities in MBD techniques using the underlying graph $\mathcal{G}$. The valuation algebra Γ enables us to demonstrate the operations and message-passing techniques underlying the MBD approaches. As a consequence, we are able to identify similarities among MBD approaches, thereby paving the way for a more holistic approach to MBD and potential cross-pollination of MBD inference techniques.

REFERENCES

[1] Adnan Darwiche, 'Utilizing device behavior in structure-based diagnosis', in *IJCAI*, pp. 1096–1101. Citeseer, (1999).

[2] Johan De Kleer, 'An assumption-based TMS', *Artificial intelligence*, **28**(2), 127–162, (1986).

[3] J. Kohlas, *Information Algebras: Generic Structures for Inference*, Springer, Heidelberg, 2003.

[4] Jürg Kohlas, Rolf Haenni, and Serafín Moral, 'Propositional information systems', *Journal of Logic and Computation*, **9**(5), 651–681, (1999).

[5] Hans-Andrea Loeliger, 'An introduction to factor graphs', *Signal Processing Magazine, IEEE*, **21**(1), 28–41, (2004).

[6] Marc Pouly and Jürg Kohlas, *Generic Inference: A Unifying Theory for Automated Reasoning*, John Wiley & Sons, 2012.

[7] Raymond Reiter, 'A theory of diagnosis from first principles', *Artificial intelligence*, **32**(1), 57–95, (1987).