

Lecture Notes in Artificial Intelligence

1003

Subseries of Lecture Notes in Computer Science

Edited by J. G. Carbonell and J. Siekmann

Lecture Notes in Computer Science

Edited by G. Goos, J. Hartmanis and J. van Leeuwen

Springer

Berlin

Heidelberg

New York

Barcelona

Budapest

Hong Kong

London

Milan

Paris

Santa Clara

Singapore

Tokyo

P. Pandurang Nayak

Automated Modeling of Physical Systems



Springer

Series Editors

Jaime G. Carbonell
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3891, USA

Jörg Siekmann

University of Saarland
German Research Center for Artificial Intelligence (DFKI)
Stuhlsatzenhausweg 3, D-66123 Saarbrücken, Germany

Author

P. Pandurang Nayak
NASA Ames Research Center, Recom Technologies
Moffett Field, CA 94035, USA

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Nayak, P. Pandurang:

Automated modelling of physical systems / P. Pandurang
Nayak. - Berlin ; Heidelberg ; New York ; Barcelona ;
Budapest ; Hong Kong ; London ; Milan ; Paris ; Santa Clara ;
Singapore ; Tokyo : Springer, 1995

(Lecture notes in computer science ; 1003)

Zugl.: Diss.

ISBN 3-540-60641-6

NE: GT

CR Subject Classification (1991): I.2, I.6, F.1, J.2, C.5.4

ISBN 3-540-60641-6 Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1995
Printed in Germany

Typesetting: Camera ready by author

SPIN 10487123 06/3142 - 5 4 3 2 1 0 Printed on acid-free paper

To Mala and my parents

Foreword

It is a privilege for me to write the foreword to this important, powerful, and elegant piece of Artificial Intelligence research. In both substance and method, it is the quintessence of what today's AI research should be like.

The prelude to doing an important work is to seek out an important problem. Nayak's problem area is important for two reasons.

- For most tasks whose solution requires intelligence, the AI will have to be able to reason about the world of physical objects, physical systems, and engineered artifacts. The path from the earliest AI programs (e.g. the Logic Theorist and the Geometry Theorem Proving Program) to the expert systems of the 1980s and 1990s has not included much work on physics and engineering problem solving. The exception of course is the thinly populated field of qualitative physics research, of which Nayak's work is an example.
- In expert systems ("where the rubber meets the road" for most of applied AI), the models of domains are very extensive and detailed. Indeed, it is the hallmark of expert systems that their knowledge bases are relatively large and their reasoning methods simple (usually just Aristotelian and/or Bayesian logic). A narrow pinnacle of knowledge (literally a pinnacle of excellence) allows the expert system to perform at high levels of competence for a narrow class of problems. But from pinnacles it is possible to fall precipitously, for example when the problem presented lies outside of the narrow band of knowledge. What is needed is a soft landing on a much more gentle slope of generalized knowledge. Such knowledge contains the generalized models and equation systems that we learn in courses on physics and engineering. Of course, this "generalized" knowledge is also somewhat domain-specific, but its breadth is much greater than the specialists' knowledge used in expert systems. The increased scope and power of the "broad" knowledge does not come for free. To use it, the reasoning processes can no longer be simple – more powerful methods are needed. To develop the necessary representation and reasoning methods is the goal of qualitative physics research, and the contribution of Nayak's research.

Here are some specific points to keep in mind as you read this monograph.

- Carefully chosen problem representations are crucial to effective problem-solving. Amarel, as early as the 1960s, discussed this issue extensively in his landmark paper “On representations of problems of reasoning about actions” [Amarel, 1968], where he provides various representations for the Missionaries and Cannibals problem. The choice of problem representations is equally crucial in engineering problem-solving, where one has to carefully balance the veracity of a model against its complexity. An overly detailed model can be too complex, while simple models may miss important problem features. Hence, in modeling physical systems, it is important to identify just the relevant physical phenomena, and to identify just the appropriate level of detail to model each phenomenon. Nayak’s research elegantly treats and combines both the theoretical and practical sides of this problem.
- On the theoretical side, the research develops a precise formalization of the problem of selecting adequate models of physical systems. It casts the problem as a search problem, and provides clear definitions for the search space and the goal criterion. Such a precise formalization is essential for applying computational tools to the problem. It then uses the formalization to analyze the difficulty of finding adequate models, and for understanding the different reasons for its intractability. Finally, it uses the analysis of the reasons for intractability to identify special cases of the problem that can be solved very efficiently. These special cases stem from the identification and use of a special class of approximations called causal approximations. Importantly, the research shows that causal approximations are very common in modeling the physical world, making the special cases broadly applicable and useful.
- On the practical side, Nayak’s work analyzes the representation and reasoning issues that arise in building practical systems for constructing adequate device models. It develops a class-level description for device models that facilitates knowledge base construction and supports focused generation of device models. It develops a novel order of magnitude reasoning algorithm for reasoning about a device’s behavior, and incorporates it into the efficient model selection algorithm developed above. The resulting algorithm has been implemented in Common Lisp.
- Nayak tests and validates his methods and algorithms by running his implementation on a large knowledge-base of device models. The range of devices is not only interesting but significant, and the testing is a necessary activity to ground the research in reality and to add credibility to Nayak’s approach.

In recommending Nayak’s work to the Association for Computing Machinery for its dissertation award, I remember using a word that I rarely use: “brilliant.” In reviewing the work again as a prelude to writing this foreword,

I am once again excited by the work and am reminded of why I used the strong B word. Proceed with haste to the first page so that you can find out for yourself. You will be richly rewarded.

October 1995

*Edward A. Feigenbaum
Professor of Computer Science, Stanford University
and Chief Scientist, United States Air Force*

Preface

Effective reasoning about complex physical systems requires the use of models that are adequate for the task. Constructing such adequate models is often difficult. In this dissertation, we address this difficulty by developing efficient techniques for automatically selecting adequate models of physical systems. We focus on the important task of generating parsimonious causal explanations for phenomena of interest. Formally, we propose answers to the following: (a) what is a model and what is the space of possible models; (b) what is an adequate model; and (c) how do we find adequate models.

We define a model as a set of *model fragments*, where a model fragment is a set of independent equations that partially describes some physical phenomenon. The space of possible models is defined implicitly by the set of applicable model fragments: different subsets of this set correspond to different models. An adequate model is defined as a simplest model that can explain the phenomenon of interest, and that satisfies any domain-independent and domain-dependent constraints on the structure and behavior of the physical system.

We show that, in general, finding an adequate model is intractable (NP-hard). We address this intractability, by introducing a set of restrictions, and use these restrictions to develop an efficient algorithm for finding adequate models. The most important restriction is that all the approximation relations between model fragments are required to be *causal approximations*. In practice this is not a serious restriction because most commonly used approximations are causal approximations.

We also develop a novel *order of magnitude* reasoning technique, which strikes a balance between purely qualitative and purely quantitative methods. The order of magnitude of a parameter is defined on a logarithmic scale, and a set of rules propagate orders of magnitudes through equations. A novel feature of these rules is that they effectively handle non-linear simultaneous equations, using linear programming in conjunction with backtracking.

The techniques described in this dissertation have been implemented and have been tested on a variety of electromechanical devices. These tests provide empirical evidence for the theoretical claims of the dissertation.

Acknowledgements In preparing this dissertation I have been greatly influenced by a number of different people, and have benefited enormously from

my discussions with them. I would like to thank my advisor, Ed Feigenbaum, for his guidance, encouragement, and support. He gave me the freedom to pursue my research interests, while always reminding me of the importance of real problems. Leo Joskowicz participated actively in my research and helped me clarify the ideas. Sanjaya Addanki, Richard Fikes, Rich Keller, and Jean-Claude Latombe provided valuable feedback on various aspects of this research. Surajit Chaudhuri and Alon Levy shared the ups and downs of graduate student life with me and helped to make research genuinely fun. Members of the How Things Work project, Bob Englemore, Tom Gruber, Yumi Iwasaki, and Jim Rice, provided a challenging and educational environment for research.

Grace Smith, Michelle Perrie, and Margaret Timothy provided superb administrative support. IBM supported me financially during the course of this research through an IBM Graduate Technical Fellowship. Additional support for this research was provided by the Defense Advanced Research Projects Agency under NASA Grant NAG 2-581 (under ARPA order number 6822), by NASA under NASA Grant NCC 2-537, and by IBM under agreement number 14780042.

On a more personal note, my friends from IIT and Stanford made my years at Stanford enjoyable and memorable. My family in the US helped alleviate any home-sickness I might have felt. My parents and my sister supported me with their love and encouragement throughout my life. My parents taught me that in any endeavor, it is the effort that is important, not the result. It is because of them that I am what I am. And last, but by no means least, Mala's constant friendship, support, and love has made this research possible. This dissertation is as much a result of her efforts, as it is a result of mine.

P. Pandurang Nayak
October 1995

Table of Contents

1. Introduction	1
1.1 Models and tasks	1
1.2 Problem statement	4
1.3 Proposed solution: An overview	4
1.3.1 What is a model and what is the space of possible models?	4
1.3.2 What is an adequate model?	5
1.3.3 How do we find adequate models?	8
1.4 Contributions	10
1.5 Readers guide	11
2. Models and model fragments	13
2.1 Models of the behavior of physical systems	13
2.1.1 Parameters	13
2.1.2 Equations	14
2.2 Multiple models	15
2.3 Model fragments	16
2.3.1 What is a model fragment?	16
2.3.2 Advantages of model fragments	17
2.3.3 Composing model fragments	18
2.3.4 Relations between model fragments	19
2.4 Space of possible models	22
2.4.1 Device structure	22
2.4.2 Structural abstractions	23
2.4.3 Possible models of a component	24
2.5 Model fragment classes	24
2.5.1 What are component and model fragment classes?	25
2.5.2 Typographic conventions	25
2.5.3 Defining component and model fragment classes	26
2.5.4 Difference between component and model fragment classes	31
2.6 Summary	31

3. Adequate models	33
3.1 Tasks and models	33
3.2 Causal explanations	34
3.2.1 Importance of causal explanations	34
3.2.2 Types of causal explanations	34
3.3 Causal ordering	35
3.3.1 Loops in the causal ordering	36
3.3.2 Equations	36
3.3.3 Computing the causal ordering	37
3.4 Consistency and completeness of models	46
3.4.1 Model consistency	46
3.4.2 Model completeness	47
3.5 Representing the phenomenon of interest	48
3.6 Constraints from the structural context	50
3.6.1 Structural context	50
3.6.2 Constraints	52
3.7 Constraints from the behavioral context	54
3.7.1 Behavioral context	54
3.7.2 Constraints	54
3.7.3 Thresholds in behavioral constraints	56
3.8 Simplicity of models	57
3.9 Summary	58
4. Complexity of model selection	61
4.1 Formalizing the problem	61
4.1.1 Formalizing the input	61
4.1.2 Problem statement	67
4.2 Complexity analysis	68
4.2.1 Problem size	69
4.2.2 Preliminaries	69
4.2.3 The SELECT MODEL FRAGMENTS problem	70
4.2.4 The SELECT ASSUMPTION CLASSES problem	75
4.2.5 The SATISFY CONSTRAINTS problem	76
4.2.6 The intractability of finding causal models	80
4.3 Other complexity results	81
4.3.1 Fixed causal orientations	81
4.3.2 Finding coherent models	82
4.4 Summary	85
5. Causal approximations	87
5.1 Upward failure property	88
5.1.1 Efficient model selection algorithm	89
5.1.2 Discussion	93
5.2 Preliminary restrictions	93
5.3 Causal Approximations	95

5.3.1	Definitions	95
5.3.2	Causal approximations and the upward failure property	97
5.3.3	Causal approximations are common	101
5.4	Selecting assumption classes	101
5.4.1	Causal approximations are not enough	102
5.4.2	Parameter ownership	103
5.4.3	Ownership constraints	103
5.4.4	Monotonicity of causal relations	104
5.4.5	Discussion	106
5.5	Individually approximating model fragments	106
5.6	Expressivity of constraints	108
5.7	Efficiently simplifying a coherent model	109
5.7.1	Simplifying a model by approximating	110
5.7.2	Simplifying a model by dropping model fragments	113
5.8	Summary	117
6.	Differential equations	119
6.1	Canonical form	120
6.2	Approximating differential equations	121
6.3	Exogenizing differential equations	123
6.4	Equilibrating differential equations	123
6.5	Monotonicity of causal relations	125
6.6	Efficiently equilibrating differential equations	135
6.6.1	Equilibrating differential equations can be hard	136
6.6.2	Self-regulating parameter	137
6.6.3	Discussion	142
6.7	Summary	143
7.	Order of magnitude reasoning	145
7.1	Motivating example	146
7.2	Order of magnitude reasoning in NAPIER	147
7.2.1	Inference rules in NAPIER	148
7.2.2	Set of simultaneous equations	149
7.2.3	Backtracking algorithm	150
7.2.4	Example	151
7.3	Order of magnitude reasoning is intractable	153
7.4	Approximation algorithms in NAPIER	157
7.4.1	Ordering the equations	157
7.4.2	Loss of accuracy	158
7.5	Error estimation	160
7.5.1	Estimating the error of heuristic rules	160
7.5.2	Alternate order of magnitude rules	162
7.5.3	Discussion	164
7.6	Related work	165
7.7	Summary	166

8. Model selection program and results	169
8.1 Finding an initial causal model	170
8.1.1 Component interaction heuristic	170
8.1.2 Finding an initial causal model	173
8.2 Simplifying the model	181
8.3 Implementation and results	181
8.3.1 Overview of the knowledge base	181
8.3.2 Overview of the devices	182
8.3.3 Results	185
8.4 Summary	187
9. Related work	189
9.1 Automated modeling of physical systems	189
9.1.1 Compositional modeling	189
9.1.2 Graphs of models	191
9.1.3 Fitting approximations	192
9.1.4 Critical abstractions	192
9.1.5 Model-based diagnosis with multiple models	193
9.1.6 Reasoning about model accuracy	193
9.1.7 Microscopic ontologies	194
9.2 Logical approaches	195
10. Conclusions	197
10.1 Summary and contributions	197
10.2 Future work	198
A. Examples of causal approximations	201
B. Example devices	207
B.1 Bimetallic strip temperature gauge	207
B.2 Bimetallic strip thermostat	208
B.3 Flexible link temperature gauge	209
B.4 Galvanometer temperature gauge	210
B.5 Electric bell	211
B.6 Magnetic sizing device	212
B.7 Carbon pile regulator	213
B.8 Electromagnetic relay thermostat	214
B.9 Tachometer	215
B.10 Car distributor system	216
C. Composable operators	217
C.1 Influences	217
C.2 Kirchhoff's laws	218
C.3 Other methods	220

D. Matchings in bipartite graphs	223
References	227

List of Figures

1.1	A temperature gauge	2
1.2	The possible models of a wire.	3
1.3	A possible model of the temperature gauge	5
1.4	Causal ordering of the parameters	6
1.5	Algorithm for finding a minimal causal model.	9
2.1	A temperature gauge	15
2.2	A set of equations describing the temperature gauge	16
2.3	Model fragment describing a wire as a resistor.	17
2.4	Model fragment describing a wire as an ideal conductor.	17
2.5	Model fragment describing the temperature dependence of the wire's length.	17
2.6	Model fragments describing electrical conduction in a wire.	20
2.7	Approximation relation between the electrical conduction model fragments.	21
2.8	Model fragments describing a wire's resistance.	22
2.9	Structural description of the temperature gauge.	23
2.10	The Resistor model fragment class.	26
2.11	The Two-terminal-electrical-component model fragment class.	28
3.1	A set of equations with two onto causal mappings	39
3.2	Graph representing a set of equations	42
3.3	Causal ordering algorithm	43
3.4	A possible model of the temperature gauge	44
3.5	Bipartite graph representing the equations in Figure 3.4.	45
3.6	The direct causal dependencies generated by the causal mapping corresponding to the complete matching shown in Figure 3.5.	45
3.7	Algorithm for checking whether a model can explain the expected behavior.	49
4.1	Mapping of parameters to equations	74
5.1	Function <i>find-minimal-causal-model</i>	90
5.2	Model fragments describing electrical conduction in wire-1	96
5.3	Model fragments describing a wire's resistance.	96

XX List of Figures

5.4	Causal orderings generated from M_1 and M_2	102
5.5	Causal ordering generated from M'_1	102
5.6	Model fragments in the Electrical-conductor(wire-1) assumption class	103
5.7	The function <i>simp-by-approximating</i>	111
5.8	Simplifying a model by dropping model fragments	116
6.1	Partial causal mappings H and H'	125
6.2	A motivating example	126
6.3	Example of the graph G'	130
6.4	The resulting onto causal mapping H''	133
6.5	Assumption classes and model fragments	136
6.6	Model fragment describing the velocity of a falling raindrop	139
6.7	Model fragment describing the raindrop's terminal velocity	139
6.8	Assumption classes, and their model fragments, describing the temperature of objects 1 and 2, and the heat path between them ..	143
6.9	Causal ordering before and after equilibration	143
7.1	Ionization reactions that occur on dissolving AH in water	146
7.2	Rules for order of magnitude reasoning	148
7.3	A backtrack tree	150
7.4	Alternate rules for order of magnitude reasoning	163
8.1	Algorithm for finding an initial causal model	172
8.2	Components and their initial models	173
8.3	Component models after the expected behavior parameters have been included	176
8.4	Component models after the heuristic coherence constraints have been repeatedly satisfied	178
8.5	The initial causal model	180
8.6	The two stages of simplifying the initial causal model	182
B.1	Bimetallic strip temperature gauge	207
B.2	Bimetallic strip thermostat	208
B.3	Flexible link temperature gauge	209
B.4	Galvanometer temperature gauge	210
B.5	Electric bell	211
B.6	Magnetic sizing device	212
B.7	Carbon pile regulator	213
B.8	Electromagnetic relay thermostat	214
B.9	Tachometer	215
B.10	Car distributor system	216
D.1	Matchings in a bipartite graph	224
D.2	Algorithm for finding a maximum matching	225

List of Tables

1.1 Examples of causal approximations 9

2.1 Composable operators..... 19

7.1 NAPIER’s run times on an Explorer II, with and without causal ordering. 157

7.2 Maximum value of Δ for each example. 159

7.3 Properties of the causal ordering graph..... 160

8.1 Number of components and operating regions used in each device. 183

8.2 Summary of experimental results 185

8.3 Number of model fragments in the most accurate model, the initial causal model, and the minimal causal model constructed by the program. 186

8.4 Number of model fragments that were dropped and approximated in simplifying the initial causal model. 187