


ERIK T. MUELLER

COMMONSENSE REASONING

An Event Calculus Based Approach



MK
MORGAN KAUFMANN

SECOND
EDITION

Commonsense Reasoning

An Event Calculus Based Approach

Second Edition

Erik T. Mueller

IBM Watson Group and IBM Research



AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

Morgan Kaufmann is an imprint of Elsevier



Acquiring Editor: Steve Elliot
Editorial Project Manager: Kaitlin Herbert
Project Manager: Punithavathy Govindaradjane
Designer: Matthew Limbert

Morgan Kaufmann is an imprint of Elsevier
225 Wyman Street, Waltham, MA, 02451, USA

Copyright © 2015, 2006 Elsevier Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without permission in writing from the publisher. Details on how to seek permission, further information about the Publisher's permissions policies and our arrangements with organizations such as the Copyright Clearance Center and the Copyright Licensing Agency, can be found at our website: www.elsevier.com/permissions.

This book and the individual contributions contained in it are protected under copyright by the Publisher (other than as may be noted herein).

Notices

Knowledge and best practice in this field are constantly changing. As new research and experience broaden our understanding, changes in research methods, professional practices, or medical treatment may become necessary.

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information, methods, compounds, or experiments described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors, assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions, or ideas contained in the material herein.

Library of Congress Cataloging-in-Publication Data

Mueller, Erik T.

Commonsense reasoning : an event calculus based approach / Erik T. Mueller. – Second edition.
pages cm

Includes bibliographical references and index.

ISBN 978-0-12-801416-5 (paperback)

1. Commonsense reasoning–Automation. 2. Artificial intelligence–Mathematics. 3. Logic, Symbolic and mathematical–Data processing. I. Title.

Q338.85.M84 2015

153.4'3–dc23

2014029954

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library.

ISBN: 978-0-12-801416-5

For information on all Morgan Kaufmann publications
visit our web site at www.mkp.com



Working together
to grow libraries in
developing countries

www.elsevier.com • www.bookaid.org

Commonsense Reasoning

For Susanna and Matthew

Praise for Commonsense Reasoning

A comprehensive exposition of reasoning about actions and change using the circumscription-based Event Calculus. The book has an excellent up-to-date bibliography on actions and change.

Chitta Baral, Arizona State University

Central to the idea of Artificial Intelligence is getting computers to understand simple facts about people and everyday life-what we call Common Sense. Amid the technical discussions about inference algorithms and knowledge representation, a larger question arises: What have we actually learned in the past 30 years about how to put Commonsense knowledge in computers? Look no further than Erik Mueller's Commonsense Reasoning for a deep and insightful survey of the state of the art in this topic. Some say that Commonsense defies logic; here Mueller shows that logic, at least, can put up a good fight.

Henry Lieberman, MIT Media Laboratory

Erik Mueller has given the most thorough treatment of commonsense knowledge and reasoning yet to appear.

John McCarthy, Stanford University

The strength of this book is that it uses a uniform representation formalism, the event calculus, to solve a variety of commonsense reasoning problems. Researchers will find the book an inspiring tool which provides many ideas for applications of action formalisms. Thanks to both the exemplary presentation style and numerous examples, the book is also well-suited for teachers and students alike.

Michael Thielscher, Dresden University of Technology

Developing systems that can perform actions and deal with change is a major challenge in intelligent system design, because it requires the construction of sophisticated models for knowledge representation and reasoning. This book provides important ideas and methods which can be used to model commonsense reasoning about events in complex and dynamic environments. The content is well thought out, and difficult topics are presented in highly accessible ways. The author tells a compelling story that highlights the utility of event calculus for applications that require commonsense models of action and change.

Mary-Anne Williams, University of Technology, Sydney, and Knowledge Representation and Reasoning Inc.

Foreword to the First Edition

An eminent professor of logic once said to me, “Why do you bother devising all those little predicate calculus theories? We already know that first-order predicate calculus is capable of expressing almost anything, so what is the point?” This question typifies the attitude of a certain breed of logician, for whom the quintessence of intellectual endeavour is the study of the metalevel properties of various formalisms—their expressive power, their computational limitations, and the relationships between one formalism and another. Without doubt such work is, from an academic point of view, noble and worthwhile. So I did wonder, for several minutes, whether the eminent logician was perhaps right. But then it occurred to me that no one ever says, “Why do you bother giving your undergraduates all those little programming exercises? We already know that Turing machines can compute almost anything, so what is the point?”

The point, of course, is that the gap between believing something to be possible and knowing how to achieve it is very wide indeed. There is an art to programming, and learning how to do it well takes many years. If the eminent logician had not retired to his office before allowing me a return blow, I would have replied that what goes for programming also goes for logic: There is an art to the use of logic for knowledge representation, and learning it requires much practise. So it is surprising that more books aimed at teaching this art do not exist. Fortunately, we can now add to this small corpus the admirable volume by Erik Mueller that you are hopefully about to read. But this book is more than just a guide to building complex representations of knowledge in logic because its target is an area that might be thought of as the nemesis of artificial intelligence, namely common sense.

One of the starkest lessons of AI research in the twentieth century was that it is those aspects of human behavior that we most take for granted that are the hardest to emulate on a computer. A two-year-old child who finds a chocolate bar hidden in his mother’s bag is performing a feat of common sense that our most sophisticated AI systems would be utterly incapable of matching. It is true that we now have programs that can defeat chess grandmasters—but only at chess. To a cognitive scientist, the most remarkable thing about a chess grandmaster is that, having played a great game of chess, she can then go and make a cup of tea. The very same biological apparatus that has mastered chess had long beforehand mastered the everyday physical world of solids, liquids, gravity, surfaces, and shapes, not to mention the everyday social world of interaction with her peers.

What is the right way to approach the daunting problem of endowing computers and robots with common sense? There’s no shortage of opinions on this question among AI researchers and cognitive scientists. Perhaps we should be inspired by biology. Perhaps we should imitate evolution. Perhaps we should eschew nature and instead embrace mathematical formalisms such as logic. If we were empirical scientists, there would be a right and a wrong answer, whether or not we yet knew

which was which. But insofar as we are engineers, there can be many right answers. With a mere half century of thinking behind us—a very short time in the history of ideas—we still do not know how far the symbolic approach exemplified by this book can take us towards human-level artificial intelligence. But we do know that the symbolic approach makes for elegant designs with provable properties in a wide range of application areas where systems with a little extra intelligence have the edge. So Erik Mueller's book is necessary and timely, and I hope it gains the wide readership it deserves.

Murray Shanahan
Imperial College London
July 2005

Preface

Commonsense reasoning is the sort of reasoning we all perform about the everyday world. We can predict that, if a person enters a kitchen, then afterward the person will be in the kitchen. Or that, if someone who is holding a newspaper walks into a kitchen, then the newspaper will be in the kitchen. Because we make inferences such as these so easily, we might get the impression that commonsense reasoning is a simple thing. But it is very complex.

Reasoning about the world requires a large amount of knowledge about the world and the ability to use that knowledge. We know that a person cannot be in two places at once, that a person can move from one location to another by walking, and that an object moves along with a person holding it. We have knowledge about objects, events, space, time, and mental states and can use that knowledge to make predictions, explain what we observe, and plan what to do.

This book addresses the following question: How do we automate commonsense reasoning? In the last few decades, much progress has been made on this question by artificial intelligence researchers. This book provides a detailed account of this progress and a guide to automating commonsense reasoning using logic. We concentrate on one formalism, the event calculus, that incorporates many of the discoveries of the field. Although the event calculus is defined by a handful of classical predicate logic axioms, it enables reasoning about a wide range of commonsense phenomena.

WHY COMMONSENSE REASONING?

Why study commonsense reasoning? The first reason for studying commonsense reasoning is practical. Automated commonsense reasoning has many applications ranging from intelligent user interfaces and natural language processing to robotics and vision. Commonsense reasoning can be used to make computers more human-aware, easier to use, and more flexible.

The second reason for studying commonsense reasoning is scientific. Commonsense reasoning is a core capability of intelligence that supports many other high-level capabilities. The ability to understand what is happening in a story, for example, crucially involves commonsense reasoning. By studying commonsense reasoning we can gain a greater understanding of what intelligence is.

APPROACH

The approach to commonsense reasoning taken in this book is not shared by all researchers. My approach can be characterized by the following assumptions.

I assume, along with most cognitive scientists, that commonsense reasoning involves the use of representations and computational processes that operate on those representations.

I assume, along with researchers in symbolic artificial intelligence, that these representations are symbolic.

I assume, along with researchers in logic-based artificial intelligence, that commonsense knowledge is best represented declaratively rather than procedurally.

I use the declarative language of classical many-sorted predicate logic with equality.

I do not claim that the methods for commonsense reasoning presented in this book are the methods used by humans. This book presents one way of automating commonsense reasoning. How humans perform commonsense reasoning is an interesting topic, but it is not the topic of this book. There is evidence both for and against the use of logic in human reasoning.

INTENDED AUDIENCE

This book is intended for use by researchers and students in the areas of computer science, artificial intelligence, mathematics, and philosophy. It is also intended for use by software designers wishing to incorporate commonsense reasoning into their applications. The book can be used as a graduate-level textbook for courses on commonsense reasoning and reasoning about action and change, as well as a reference work for researchers working in these areas. It will be of interest to those using logic as their primary technique, as well as those using other techniques. This book can also be used as a supplementary graduate-level or advanced undergraduate textbook for courses on knowledge representation and artificial intelligence. I assume the reader has some familiarity with predicate logic, although reviews of predicate logic are provided in Chapter 2 and Appendix A.

ROADMAP

This book consists of 20 chapters and six appendices. The chapters are organized into eight parts.

Part I describes the foundations of the event calculus.

Part II deals with various commonsense phenomena. Chapter 3 discusses the effects of events. Chapter 4 discusses the triggering of events by conditions. Chapter 5 discusses the commonsense law of inertia. Chapter 6 discusses the indirect effects of events. Chapter 7 discusses continuous change. Chapter 8 discusses concurrent events. Chapter 9 discusses nondeterministic effects of events.

Part III deals with important commonsense domains. Chapter 10 presents axiomatizations of relational and metric space, and discusses reasoning about object identity, space, and time. Chapter 11 presents axiomatizations of the mental states of agents, including beliefs, goals, plans, and emotions.

Part IV discusses default reasoning.

Part V deals with programs and applications. Chapter 13 discusses the Discrete Event Calculus Reasoner program used to solve event calculus reasoning problems, Chapter 14 discusses several real-world applications, and Chapter 15 discusses the use of answer set programming for commonsense reasoning.

Part VI reviews logical and nonlogical methods for commonsense reasoning and discusses their relationship to the event calculus. Chapter 16 reviews logical methods, Chapter 17 reviews nonlogical methods, and Chapter 18 discusses the use of unstructured information for commonsense reasoning.

Part VII discusses the acquisition of commonsense knowledge.

Part VIII presents my conclusions.

MATERIAL COVERED

The skills that make up human commonsense reasoning are complex, and the body of research related to it is large. Because no book can realistically cover every aspect of commonsense reasoning, a choice had to be made about what this book would cover. The coverage of this book was determined by the following considerations.

Most instances of commonsense reasoning involve action and change because action and change are pervasive aspects of the world. It is crucial for any method for commonsense reasoning to deal with action and change. Therefore, a large part of this book is devoted to this topic. In addition to reasoning about action and change, or the domain of time, this book covers two other significant domains of commonsense reasoning: space and mental states, including emotions, goals, plans, and beliefs. This book also covers default reasoning and reasoning about object identity.

Over the last few decades, researchers have developed a number of logics for commonsense reasoning. It would take much time and space to cover all of these in detail. Hence, this book concentrates on one logic, the event calculus, which incorporates many of the features of the other logics. The reader who understands the event calculus will be well equipped to understand the other logics. They are closely related to the event calculus, and some are provably equivalent to the event calculus—see Appendices D and E. Chapter 16 compares the event calculus with other logics, and detailed Bibliographic notes throughout the book discuss research performed using other logics.

Several types of commonsense reasoning are not covered by this book. Commonsense reasoning using probability theory is not covered because this is not a well-developed area. But this book does cover nondeterminism, and some work on the use of probability theory for commonsense reasoning is reviewed (in Section 17.3). Although this book covers most features of the event calculus, it does not cover

continuous change described using differential equations; this book does, however, cover continuous change described by closed-form expressions.

SUPPLEMENTAL MATERIALS

WEB SITE AND REASONING PROGRAMS

The book web site at decreasoner.sourceforge.net contains additional material related to this book. This includes links to tools that can be downloaded, such as the Discrete Event Calculus Reasoner program discussed in Chapter 13 and answer set programming tools discussed in Chapter 15.

EXERCISES AND SOLUTIONS

Exercises are provided at the end of Chapters 2 through 19. Solutions to selected exercises are provided in Appendix F. Solutions to further exercises are available online to instructors who have adopted this text. Register at www.textbooks.elsevier.com for access.

TEXT AND FIGURE ACKNOWLEDGMENTS

Portions of the book, *The Snowman*, by Raymond Briggs, courtesy of Random House.

Portions of the book, *Solving the Frame Problem*, by Murray Shanahan, courtesy of MIT Press.

Portions of the article “A Logic of Actions” by Patrick J. Hayes, in *Machine Intelligence*, Vol. 6, courtesy of Edinburgh University Press.

Portions of the article, “Event Calculus Reasoning Through Satisfiability” by Erik T. Mueller, in *Journal of Logic and Computation*, courtesy of *Journal of Logic and Computation* and Oxford University Press.

Portions of the article, “Story understanding through multi-representation model construction” by Erik T. Mueller, in *Text Meaning: Proceedings of the HLT-NAACL 2003 Workshop*, courtesy of Association for Computational Linguistics.

Portions of the article, “Discrete event calculus with branching time” by Erik T. Mueller, in *Logical Formalizations of Commonsense Reasoning: Papers from the AAAI Spring Symposium*, courtesy of the Association for the Advancement of Artificial Intelligence.

Portions of the article, “Event calculus and temporal action logics compared” by Erik T. Mueller, in *Artificial Intelligence*, courtesy of Elsevier.

Portions of the discussion, “[Protocol of online discussion about the article A logical account of the common sense informatic situation for a mobile robot]” by Murray Shanahan, courtesy of “Electronic News Journal on Reasoning about Actions and Change.”

Figure, “Shanahan’s Circuit,” reprinted from *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, by Murray Shanahan, “The ramification problem in the event calculus,” p. 145, 1999, with permission from Elsevier.

Figure, “Thielscher’s Circuit,” reprinted from “Ramification and Causality,” by Michael Thielscher, p. 19, 1996, with permission from Berkeley, CA: International Computer Science Institute.

About the Author

Erik T. Mueller is a Research Staff Member at IBM Watson Group and IBM Research. He is a member of the IBM team that developed Watson,TM a natural language question answering system that won a two-game Jeopardy! match against two Jeopardy! grand champions. His previous books include *Daydreaming in Humans and Machines* and *Natural Language Processing with ThoughtTreasure*. He received an S.B. in Computer Science and Engineering from the Massachusetts Institute of Technology and an M.S. and Ph.D. in Computer Science from the University of California, Los Angeles. He is currently developing WatsonPaths, a reasoning system based on Watson, and Watson for Healthcare, which will help healthcare professionals diagnose, treat, and manage patients.

Acknowledgments to the First Edition

It is my great pleasure to acknowledge the people who have contributed to this book, although space constraints do not permit me to refer to everyone by name. I am particularly indebted to the inventors of the event calculus. This book owes its existence to the pathbreaking work of Robert A. Kowalski, Marek J. Sergot, Murray Shanahan, Rob Miller, and other researchers in the field of reasoning about action and change.

I am privileged to work in the academically stimulating environment of the IBM T. J. Watson Research Center. Special thanks are due to Doug Riecken, the creator of the Commonsense Computing department in which I work. I am grateful for the support of Guruduth S. Banavar, Arthur C. Ciccolo, Thomas A. Cofino, Paul Horn, Charles W. Lickel, Alfred Spector, John Turek, Jürg von Känel, and Ellen J. Yoffa. Over the years, I have learned a great deal about formal knowledge representation from Leora Morgenstern. My research benefited immensely from many conversations with Tessa A. Lau, Daniel Oblinger, and David Canfield Smith. Other colleagues who helped in a number of ways were Rangachari Anand, Karen Appleby, Richard J. Cardone, David Ferrucci, Leiguang Gong, Robert Hoch, Xuan Liu, Mark Podlaseck, Moninder Singh, Vugranam Sreedhar, Malgorzata Stys-Budzikowska, and Christopher Welty.

I feel very fortunate to receive continuing inspiration from the faculty, staff, and students of the MIT Media Laboratory. I particularly value the support and friendship of Marvin Minsky and Push Singh, who made it possible for me to begin to work seriously in the area of commonsense reasoning by inviting me to spend 6 months as a research scientist at the Media Lab. Both Marvin and Push have helped improve my thinking about this area in innumerable ways. Other people from the Media Lab who have helped me in my work are Barbara A. Barry, Walter Bender, Nicholas L. Cassimatis, Timothy Chklovski, Glorianna Davenport, Robert W. Lawler, Henry Lieberman, Xinyu Hugo Liu, Betty Lou McClanahan, Nicholas Negroponte, Deb Roy, Warren Sack, Oliver G. Selfridge, and Ted Selker.

Other researchers have been very generous in assisting me. Conversations with Benjamin Kuipers helped me to define a research strategy for commonsense reasoning. The excellent paper of Murray Shanahan and Mark Witkowski on event calculus planning through satisfiability sparked my interest in this technique. My work has vastly improved as a result of the helpful comments and criticism I received from those who served as anonymous referees of my papers on the event calculus. I am grateful to John McCarthy for his ongoing support. I thank Geoff Sutcliffe for his collaboration on the use of automated theorem proving for event calculus reasoning. Discussions with Srini Narayanan have also been helpful.

I also acknowledge the reviewers, Chitta Baral, Arizona State University; Patrick Hayes, Florida Institute for Human and Machine Cognition; Michael Thielscher,

Technische Universität Dresden; Mary-Anne Williams, University of Technology, Sydney; and an anonymous reviewer, who took the time to read the entire manuscript and provide me with feedback. Vladimir Lifschitz's helpful comments significantly improved the section on the $C+$ action language.

I have had the good fortune to work with an extremely talented group of people at Morgan Kaufmann and Elsevier. Special thanks go to Denise E. M. Penrose, publisher, whose strong support for this project has never faltered. The editorial and production process was a smooth and efficient experience thanks to the excellent work of Valerie Witte, editorial assistant, and of Dawnmarie E. Simpson, project manager. I also thank Julie F. Nemer for copyediting the book, and Dartmouth Publishing for producing many of the figures.

I am grateful for the support and friendship of Alexandra Guest, Anne-Marie Hainault, Karen M. Kensek, Lo-Ann Lai, Christopher A. Perez, Joseph E. Pingree, Lorenzo A. Sadun, and Uri Zernik. My thanks also go to Eyal Amir, Patrick Doherty, Jeffrey Epstein, Dov M. Gabbay, Douglas B. Lenat, Jane Spurr, and Jeff Wolfers, who assisted this project in many ways.

My greatest debt is to my family. I especially thank my parents, Diana E. Mueller and Robert E. Mueller, for their love and enthusiasm for my work. I am grateful for the love and moral support of my sister, Rachel A. G. Mueller-Lust, and brother-in-law, Andrew D. Mueller-Lust. I thank my in-laws, David Hackett Fischer and Judith H. Fischer, who read the manuscript and provided helpful feedback. My thanks also go to Ilus Lobl, John H. Fischer, Victor Lobl, Anne W. Fischer, Frederick C. Turner, and Althea W. Turner. Sophie and Nicole Fischer provided welcome distractions. Words cannot express my love and gratitude to my wife, Susanna F. Fischer, and our son, Matthew E. Mueller, for the happiness we share.

Erik T. Mueller

July 2005

Acknowledgments to the Second Edition

I would like to thank all those at Elsevier who made the second edition happen. Special appreciation goes to Steve Elliot, publisher, for suggesting a second edition and for his editorial help. Kaitlin Herbert, editorial project manager, and Punitha Govindaradjane, project manager, made the publishing process enjoyable.

I would like to thank the reviewers of parts of this book. Joohyung Lee and Ravi Palla provided useful comments on the chapter about commonsense reasoning using answer set programming. Catherine Havasi and Rob Speer gave helpful suggestions on the chapters about commonsense reasoning using unstructured information and acquisition of commonsense knowledge.

I am grateful to Michael Dyer, Douglas Metzler, Rob Miller, Murray Shanahan, and Stuart C. Shapiro for their helpful input on the second edition. I thank Henry Lieberman for fruitful discussions about AnalogySpace.

This work would not be possible without the kind support of Jennifer Chu-Carroll, Eric W. Brown, and the entire Watson Group at IBM.

I am so happy to be in the delightful company of my wife Susanna and our son Matthew.

Erik T. Mueller
September 2014