

## Preface

This book provides a definition and study of a knowledge representation and reasoning formalism stemming from conceptual graphs, while focusing on the computational properties of this formalism.

Knowledge can be symbolically represented in many ways. The knowledge representation and reasoning formalism presented here is a graph formalism – knowledge is represented by *labeled graphs*, in the graph theory sense, and reasoning mechanisms are based on *graph operations*, with graph homomorphism at the core.

This formalism can thus be considered as related to semantic networks. Since their conception, semantic networks have faded out several times, but have always returned to the limelight. They faded mainly due to a lack of formal semantics and the limited reasoning tools proposed. They have, however, always rebounded because labeled graphs, schemas and drawings provide an intuitive and easily understandable support to represent knowledge.

This formalism has the visual qualities of any graphic model, and it is *logically founded*. This is a key feature because logics has been the foundation for knowledge representation and reasoning for millennia. The authors also focus substantially on *computational facets* of the presented formalism as they are interested in knowledge representation and reasoning formalisms upon which knowledge-based systems can be built to solve real problems. Since object structures are graphs, naturally graph homomorphism is the key underlying notion and, from a computational viewpoint, this moors calculus to combinatorics and to computer science domains in which the algorithmic qualities of graphs have long been studied, as in databases and constraint networks.

The main notions are intuitively introduced from a knowledge representation viewpoint, but the authors have also tried to give precise definitions of these notions along with complete proofs of the theorems (readers only require a few simple mathematical notions, which are summarized in the appendix).

This book does not present a methodology for knowledge representation nor describe actual applications (except for one chapter) or software tools. It presents theoretical bases that merge numerous previous studies carried out in the conceptual

graph domain since Sowa's seminal book and links basic reasoning issues to other fundamental problems in computer science.

In a nutshell, the authors have attempted to answer the following question: “*How far is it possible to go in knowledge representation and reasoning by representing knowledge with graphs (in the graph theory sense) and reasoning with graph operations?*”

## Organization

The book is divided into 3 parts, 13 chapters and 1 appendix.

The first part is devoted to the kernel of the formalism. In Chap. 2, Basic conceptual Graphs (BGs) are defined. BGs are structured by a subsumption relation defined by the homomorphism notion between BGs, as well as by elementary specialization and generalization operations. In Chap. 3, Simple conceptual Graphs (SGs) are introduced. SGs can be sketchily defined as BGs augmented with equality. Chapter 4 is devoted to set and FOL semantics for SGs. The soundness and completeness of homomorphism with respect to entailment relation and FOL deduction is proven. Chapter 5 relates BG homomorphism with homomorphisms of other structures (e.g., hypergraphs, relational structures, conjunctive database queries) and solving a constraint network.

The second part develops computational aspects of basic conceptual graphs. Chapters 6, and 7 are devoted to algorithms for BG homomorphism. Chapter 8 presents other specialization and generalization operations that are not covered in the first part, e.g., least generalization, maximal join and other extended joins.

The third part pools the kernel extensions. All of these extensions are provided with logically sound and complete reasoning mechanisms based on graph homomorphism. Chapter 9 focuses on nested conceptual graphs, which are hierarchically structured graphs. In Chap. 10, the important rule notion, which allows representation of implicit knowledge, for instance, is studied. The introduction of rules gives the powerfulness of a computability model to the formalism. Positive and negative constraints are considered in Chap. 11, and the BG family of models, combining facts, rules and constraints, is presented. Chapter 12 is devoted to negation in conceptual graphs. The last chapter presents semantic annotations—a currently favored and potentially fruitful application.

Finally, the Appendix summarizes the basic mathematical notions used in the book.

Each chapter begins with an overview of its content. It ends with bibliographical notes devoted mainly to studies in the conceptual graph domain. Concerning the references, the authors chose to cite studies from other domains throughout the text and provide the conceptual graph references in the bibliographical notes.

## Implementation

This book is about a KR formalism that can serve as a theoretical foundation for knowledge-based systems. A distinction is made between a KR formalism (e.g., a fragment of first order logic) and a KR programming language (e.g., PROLOG). Even when a KR programming language is based on a KR formalism, the KR language presents variations to the KR formalism (limitations, e.g., Horn clauses or extensions, second order features, etc.), and it has a concrete syntax, contains specific programming constructs (e.g., the cut in PROLOG) and is equipped with software tools. Between a KR formalism and a programming language implementing this formalism, there may be KR tools and platforms. This is the current situation of the graph-based KR formalism presented in the book. CoGITaNT [cog01a] contains a library of C++ classes that implement most of the notions and reasoning mechanisms presented here. CoGITaNT also contains a graphical interface and a client-server architecture. CoGITaNT thus allows a programmer to build knowledge-based systems grounded on the formalism presented in this book. Let us also mention COGUI [cog01b], a graphical user interface dedicated to the construction of a knowledge base and which integrates CoGITaNT.

## Audience

The book is intended for computer science students, engineers and researchers. Some of the materials presented in this book have been used for several years at different academic levels, ranging from AI courses for graduate students to professional and research master's level courses. The mathematical prerequisites are minimal, and the Appendix outlines the mathematical notions used in the book.

Here are some suggestions for reading paths of this book depending on readers' specific interests.

Chapter 1 (introduction), Chap. 2 (basic conceptual graphs), Chap. 3 (simple conceptual graphs), the two first sections of Chap. 4 (set and first order logic semantics) are the core chapters and should be considered as the basis.

For programming and algorithmic purposes the following materials can be added to the base: Chap. 6 (basic algorithms for BG homomorphism), the last section of Chap. 5 (relationships with constraint programming techniques), Chap. 7 (techniques for trees and other tractable cases), Chap. 8 (algorithms for other specialization/generalization operations, especially maximal joins), Chap. 10 (rule processing), and Chap. 12 (algorithms for processing BGs with atomic negation).

For modeling purposes the following parts can be added to the base: Chap. 9 (nested conceptual graphs), Chap. 10 (definition and use of rules), Chap. 11 (definition and use of constraints and their combination with rules), Chap. 12 (definition and use of atomic negation) and Chap. 13 (semantic annotations).

For more theory-oriented readers, expressivity, decidability and complexity results, as well as stating equivalence with problems of other domains, are presented throughout the book, except in the last chapter.

## Acknowledgments

We are indebted to John Sowa, who is the founding father of conceptual graphs, and the present book would not have existed without his pioneering work [Sow84].

We began working on conceptual graphs in 1992, and over the years many of our colleagues and students have contributed, in one way or another, to this work.

We would like to thank:

Friends and colleagues who have supported our approach over the years: Franz Baader, Marie-Christine Rousset, Pierre Marquis, Fritz Lehmann, Daniel Kayser, Michel Habib, Guy Chaty;

Colleagues or students with whom we have directly collaborated on topics covered in this book: Jean-François Baget, Boris Carbonneill, Olivier Cogis, David Genest, Olivier Guinaldo, Ollivier Haemmerlé, Gwen Kerdiles, Michel Leclère, Anne Preller, Eric Salvat, Geneviève Simonet, Rallou Thomopoulos;

Colleagues from the conceptual graph community with whom we have had fruitful discussions: Tru Cao, Fritjof Dau, Gerard Ellis, Jean Fargues, Bikash Gosh, Fritz Lehmann, Robert Levinson, Dickson Lukose, Guy Mineau, John Sowa, Mark Willems, Vilas Wuwongse.

Our approach has benefited from numerous collaborations on applied research projects, and we would like to thank: Christian Baylon, Alain Bonnet, Bernard Botella, Olivier Corby, Patrick Courounet, Rose Dieng, Steffen Lalande, Philippe Martin, Patrick Taillibert.

We acknowledge Jean-François Baget, David Genest, Alain Gutierrez, Michel Leclère, Khalil Ben Mohamed, Nicolas Moreau, Fatiha Saïs, Eric Salvat, who reviewed parts of this book, with special thanks to Geneviève Simonet (some of her comments could be the seeds for a new book ...).

We would also like to thank David Manley for rectifying our Franglais and Beverley Ford and her team at Springer London for their support.

Obviously, all remaining errors are our fault alone. We welcome corrections and suggestions on all aspects of the book; please send these to us at:

Michel.Chein@lirmm.fr and Marie-Laure.Mugnier@lirmm.fr. A web site companion to the book can also be queried at:

<http://www.lirmm.fr/gbkrbook>

La Boissière,

*Michel Chein and Marie-Laure Mugnier*  
March 2008