

PREFACE

MODAL LOGIC

This Handbook documents the current state of modal logic, a lively area of logical research which was born in philosophy, but which has since made its way into mathematics, linguistics, computer science, AI, and even economic game theory. As with other thriving scientific endeavours, it is not easy, and perhaps not even fruitful, to give an official definition of the subject. From the earliest days of modern modal logic, about a century ago now, there were many different interpretations, formalisms, and applications, and new developments have only added to this diversity. On the other hand, modal logic is also a remarkably coherent field in many ways, and its practitioners have no difficulty recognising research — and colleagues! — as being ‘modal’ in spirit. As editors, we see two broad perspectives that help give rise to this coherence. Writing with very broad strokes of the pen, we might say that the following two conceptions of the field have been particularly influential among modal logicians:

- *Modal logics as formalisations of modalities.* Many natural notions in language and science have a ‘modal’ character, in that they talk about possibility and necessity in some space of relevant situations. This was true for the original philosophical study of metaphysical modality, but it is equally true of modal logics of time, space, obligation, conditionality, knowledge, computation, and action, which have permeated other fields. Under this view, modal logics model natural reasoning concerning ubiquitous notions, and in doing so, they expand the descriptive scope of ‘standard’ logic. Technically, then, modal logics are obtained from standard logical systems (like classical propositional or predicate logic, intuitionistic logic, and so on) by adding new, *non-truth-functional operators* (that is, *modalities*). The non-truth-functional nature of the operators, reflecting the larger space of relevant situations, typically leads to systems richer than the underlying logic.
- *Modal logics as fragments of standard logics.* But surprisingly, another viewpoint on modal logic has become equally prominent among its practitioners. Under this view, modal logics inherit their semantics from the standard semantics of classical first or even higher-order predicate logic, but they restrict expressive power by using *operators instead of explicit quantification*. The term “fragment” carries no negative connotation of poverty here: curbing expressive power leads to systems with logical properties rather different from those of standard logic; the decidability of many modal logics is a striking example. The mathematical study of modal logics in this vein has brought to light a delicate balance between the expressive power and computational complexity of logical systems in general. That is, from this perspective modal logic is better viewed as a methodology for tapping into a core theme in standard logic. It is this second, more technical perspective, which provides much of the mathematical coherence of the field today.

The first perspective emphasises the *descriptive range* of modal logic as the study of key concepts and the reasoning patterns they give rise to. The second perspective emphasises the methodological aspect of *fine-structure*: modal languages bring to light the inner structure of classical systems. But these views are not in conflict. As the Handbook makes abundantly clear, most active research directions in modal logic take something from both. Indeed, the most widespread semantics of modal logic in terms of relational models (used in virtually every chapter of the Handbook) provides a setting in which these perspectives coexist fruitfully. Moreover, the two views help us better understand historical contributions made in the field. For example, the modal system **S4** was introduced in the 1920s as an analysis of the concept of necessary implication. But the limited expressive power of the formalism as an account of implicative structure turned out to be the key to **S4**'s wide range of other applications, and its attractive mathematical behaviour. In short, both perspectives on modal logic are widely applicable, and both have proved historically robust. Let's take a closer look at them, and see how they are related.

Modal logic as the study of old and new modalities Modal logic, conceived of as the *formal* study of modalities was invented in philosophy almost a century ago — though the *informal* study of modalities can be traced back much earlier: through the work of the medieval logicians, and back to the ancient Greeks. The first modal operators were introduced in order to solve the paradoxes of material implication and to obtain logics of necessity and possibility; the key figure here is C. I. Lewis, who published his pioneering work in 1918. Putting his idea in modern notation, we take some logical formula φ , and by prefixing it with a \Box or a \Diamond symbol we obtain the expressions $\Box\varphi$ (“the proposition φ is necessary”) and $\Diamond\varphi$ (“the proposition φ is possible”). That is, the box and diamond notation enables us to assert fundamentally new modes of truth concerning the information expressed by φ , namely that it is necessary or possible.

In 1933, Kurt Gödel, driven by concerns in the foundations of mathematics, used modal operators to formalise the notion of mathematical provability. In particular, his work enabled intuitionistic logic to be reduced to classical logic extended with a provability operator, and the resulting logic turned out to be Lewis's system **S4**. A striking result indeed, but the general point is this: once again, modalities are being used to express fundamentally new modes of truth concerning a piece of information. In particular, now $\Box\varphi$ means that φ is provable, and $\Diamond\varphi$ means that it is consistent.

These early examples of applying modalities to logical formulas to make assertions concerning a novel mode of truth are only the tip of the iceberg. In the decades following the work of Lewis and Gödel, many modal operators were introduced and investigated, all dealing with truth in some space of possible situations. Tense logic (or temporal logic) arises with the addition of modalities like “eventually” or “earlier”. Deontic logic adds modalities “it is permitted, or obligatory, that”. In epistemic logic we make use of modalities like “it is known that”, either for single agents or for groups. And conditional logics analyse further species of conditional reasoning far beyond Lewis's original account. This way of thinking about modal logic and modalities underlies the work of some of the field's most prominent pioneers, including G. H. von Wright, Arthur Prior, Jaakko Hintikka, Hans Kamp, and David Lewis.

Then the torch passed to other disciplines. In particular, temporal, dynamic and epistemic logics found their way into computer science, AI, and economic game theory. Temporal logics, of both branching and linear time, are now used in industry for automated verification of hardware and software. Epistemic, temporal and conditional operators are the main ingredient of knowledge-based programming. And modal logics of active agents with knowledge, beliefs, and desires form a theoretical backbone of modern accounts of intelligent distributed computing.

Pioneers of modal methods in computer science include Edmund Clarke, Joe Halpern, Zohar Manna, Robin Milner, Rohit Parikh, Amir Pnueli, Vaughan Pratt, and many others — including quite a few of the authors and commentators in this Handbook. But again, diversity reigns, and creation of new modal formalisms for novel reasoning purposes continues unabated.

Summing up: it is entirely reasonable to say that modal logics are formalisms used to represent and reason about the plethora of modal notions that underlie, among other things, distributed computations and intelligent actions, and their corresponding modes of truth. They achieve this by making use of new operators, called modalities, whose truth-conditions involve access to some larger space of relevant situations, such as worlds, times, theories, or computational states. If you view modal logic in this descriptive way, you will be in excellent company.

Modal logic and the fine-structure of classical logics The invention of graph-based relational semantics (by Jaakko Hintikka, Stig Kanger, and Saul Kripke) in the late 1950s and early 1960s showed that standard modal logics could be regarded as fragments of first or second-order predicate logics. The underlying idea is straightforward. Suppose we read $\Box\varphi$ as “necessarily φ ” and $\Diamond\varphi$ as “possibly φ ”. Drawing on an idea that dates back to the work of Leibniz, we could view “necessarily φ ” as a claim that φ is true in all possible worlds, and “possibly φ ” as a claim that φ is true in some possible world. Thus, modal operators perform quantification without making use of explicit variables and binding.

This idea, when expressed mathematically, has turned out to be the most significant milestone in the history of modal logic. For present purposes, the crucial idea in the above is just this. Because necessity means truth at *all* worlds, \Box becomes linked to the universal quantifier \forall , and because possibility means truth at *some* world, \Diamond becomes linked to the existential quantifier \exists . That is, necessity and possibility have been analysed in terms of classical quantification. The idea that modal operators are essentially concealed forms of classical quantification is fully general: for example, we can think of “eventually φ ” as meaning “there is *some* future time at which φ holds”, and we can think of “after performing a certain action, φ ” as meaning “at *every* state which is accessible by performing a certain kind of action, φ holds”.

But this is not all there is to the analogy. Viewed in this way, modal logics might just be different notation for classical ones! The creative difference is that the quantification in modal languages tends to be bounded in some way to ‘relevant’ or ‘accessible’ situations lying beyond the current one. In other words, we are working on structured universes of worlds, computational states, or what have you — and access is *mediated*. Together with the quantifier analogy, this bounded access explains two things. First, a number of properties of modal logics follow at once from those of their classical quantificational counterparts. Second, as the fragments of classical logic that modal operators correspond to typically have *less* expressive power than full first-order predicate logic, this results in many *new* properties. For example, the semantic invariances between models appropriate for modal expressive power are not those of classical logic, but rather turn out to be various forms of *bisimulation*, which preserve local properties of worlds and their transition patterns. Moreover, the study of fragments of classical systems and translations from modal to predicate logic (and sometimes back) has yielded satisfying explanations of why so many modal logics are decidable (unlike classical predicate logic), and why their computational complexity is often relatively low — and in a more activist way, modal analysis has led to the discovery of many new decidable fragments of classical logics. Indeed, the systematic design of modal languages stronger than the formalisms bequeathed to us by the founding fathers has been a major theme in recent research.

First-order languages have been the most traditional companions of modal ones, but the same

points apply to modal fragments of higher-order languages, and — a significant development in recent years — to modal fragments of classical languages with fixed-point operators expressing iterative or recursive structures in action, computation, and knowledge representation.

Summing up: it is also entirely reasonable to say that modal logics are fragments of classical logics, which somehow strike an optimal balance between expressive power and computational simplicity. And if you view modal logic in this way (that is, as a laboratory for fine-structure) you will also be in excellent company.

Nowadays, few modal logicians would feel compelled to choose between the two perspectives just outlined. Their respective virtues are clear, and most researchers have assimilated both. Given the current explosion of practical applications and fundamental ideas in the field (amply documented in this Handbook), dogmatism concerning the nature of modal logic is becoming increasingly unsustainable. Let us emphasise this point a bit more.

Contemporary modal logic Modal logic today is a vast family of studies of modal notions, with the original philosophical and mathematical motivations still alive, but with an increasing symbiosis with other fields, and in particular, with computer science. Indeed, its interface with computer science (and more generally, informatics) is extremely broad, ranging from hardware and software verification, to ontologies in medical and bio-informatics, and the analysis of query languages for XML documents. Moreover, it also takes in commonsense reasoning in AI, covering issues ranging from representing and reasoning about space and time to modelling complex interactive multi-agent information systems. Nowadays, modal structures seem to occur everywhere, just as they did in the creative explosion of modality in the philosophical logic of the 1950s and 1960s. This independent (re-)discovery of modal operators in different settings is one of the strongest arguments for the stability and naturalness of the modal stance. Here are three striking illustrations of contemporary rediscoveries.

The first example is *description logic*, a branch of knowledge representation and reasoning in AI; nowadays it supplies many of the formalisms used to fix terminologies in medical and bio-informatics, and has been proposed as the language for annotating web pages to develop a semantic web. Since the late 1970s, the description logic community has articulated its fundamental research goal with great clarity: to obtain fragments of predicate logics which are computationally well-behaved but still have the expressive power required in knowledge representation applications. Intriguingly, around 1990 it became apparent that many of the logics obtained by pursuing this goal were in fact modal logics in a different notational guise. Bounded quantification turns out to be as fundamental for description logics as it is for standard modal logics. In this case, of course, the accessible objects are not worlds or situations but individual objects from some application domain (for example, the biological function of a DNA sequence). But many of the underlying ideas are the same, and this observation has opened the doors to joint work with the modal logic community, with benefits to both fields.

Another area where modal structure is currently surfacing is in the abstract study of *processes* in the emerging field of *coalgebra*. Though modally inspired process theories like dynamic logic, temporal logic, and process algebra have a long history in computer science, coalgebra adds a new twist. Starting from the work by Peter Aczel, Jon Barwise, and others on generalised set theory, coalgebra has now become a theory of finite and infinite processes, with deep connections with universal algebra and other parts of mathematics and theoretical computer science. The crucial feature here is that processes need not be bottom-up inductive, but can instead be top-down co-inductive streams of events. One gets to know a process through observation of events, chewing off the head of the event stream. The surprising discovery has been that such processes

and their observational analysis again shows clear modal patterns, leading to rapidly developing interfaces between coalgebra, modal logic, and universal algebra.

Finally, a third independent rediscovery of modal notions occurred in economic game theory. In the 1970s, Robert Aumann and others introduced formal models of *interactive knowledge* of agents in order to account for the reasoning underpinning the Nash equilibrium solutions that would be found by rational players of a game. Disregarding some differences in notation and style, the resulting formalisms turned out to be epistemic logics from the philosophical tradition, with operators for various forms of collective knowledge of groups. Over the past three decades, logical analysis of games has become another flourishing interface, with studies of beliefs and preferences in modal languages, and the development of dynamic logics of actions that can change modal attitudes as a game proceeds. Moreover, this modal study of games has now largely merged with that of computational processes, as games are naturally viewed as goal-driven multi-agent forms of computation. Traditional game theory was the mathematics of equilibrium, using methods from analysis and dynamical systems. The modal stance that is now emerging provides a natural level of fine-structure to go with this.

Despite this diversity of modal structures, there are also strong unifying tendencies, especially in the *mathematical metatheory* of the field, which got into its stride in the 1970s with work by Wim Blok, Kit Fine, Dov Gabbay, Rob Goldblatt, Larisa Maksimova, Steve Thomason, and others. Model theory of bisimulation and related frame constructions is one important strand here; among other things, it yielded broad definability techniques for matching modal languages with classical ones, and enabled the interpolation properties of modal languages to be charted. A second major strand is algebraic semantics and the duality between modal algebras and relational structures, which built on seminal work by Bjarni Jónsson and Alfred Tarski from the 1950s; since the rediscovery of their work in the 1970s, as universal algebra has grown in sophistication, so have its ties with modal logic. Another unifying force was the development of genuinely ‘modal’ techniques with wide applicability for proving completeness, axiomatizability and decidability results. And since the 1980s, further unifying mathematical themes have emerged: these include the study of the computational complexity of reasoning and its relation to succinctness; the exploration of the relation between various forms of tree automata, logical games, and the expressivity of modal languages; and the study of logic combinations and related model constructions, which has led to a theory, still under active development, of various types of products of modal logics.

Another unifying tendency is the undeniable fact that the members of this growing modal family keep influencing each other. The word “applications” has a uni-directional ring to it, but it is a fact of life that every road can be walked both ways. For example, the action-oriented modal perspectives that are so prominent at the computer science interface have now crossed back into philosophy, giving rise to theories of information update and belief revision, which describe how agents come to acquire knowledge or change problematic beliefs. Moreover, setting up such systems also brings in conditional logics from the philosophical tradition; these are now seen as underlying belief revision and non-monotonic reasoning in deep and surprising ways.

This interdisciplinary and interactive setting is the stage where the drama of contemporary modal logic is played out. Modal structures are being studied in a growing number of areas, and often they seem to arise almost like naturally occurring phenomena; no premeditation by the modal logician is required. And at the same time, in response to this dramatic expansion, modal logicians have had to adopt a far wider range of technical ideas and tools than ever before, tools that lie beyond the placid waters of traditional textbook introductions to the field. It was against this exciting and challenging background that this Handbook was conceived.

THE HANDBOOK OF MODAL LOGIC

This Handbook presents a detailed overview of the main lines of research in contemporary modal logic. The editors have tried to present a fair picture of the modern scene, and one that (to the extent possible in a one-volume handbook) reflects the scene in its entirety. Moreover, the selection of authors has been made with a view toward representing the most active and creative research communities worldwide

The tricky question is: what is the best compromise between “detailed” and “overview”? We felt the field would be best served by a single volume handbook; that is, we opted for judicious selection and bounded access. Not that it would have been difficult to design a multi-volume handbook. On the contrary, the most frequent request we had from our authors was for more generous pages limits; the pull towards detail is strongly felt in a field such as modal logic, and quite rightly so. Many of the most treasured results and insights of the field are the results of years of painstaking work. In a sense, every student of the subject has to retrace these intellectual journeys; short cuts aren’t possible.

But the evident need for an accessible overview of the whole, leaving deeper access to valleys and caves to a second stage, suggests that the one volume choice was correct. One of the points that emerged most strongly during the Handbook’s preparation was just how unified modal logic still is. To be sure, some of its branches are now highly technical, whereas other branches are better thought of as conceptual investigations which use the language of modal logic as an aid to precision. Furthermore, some work emphasises generality and takes mathematical criteria as its primary guide, whereas other areas may be highly specific in their focus and take their cue from applications. But in spite of such differences, the field remains stubbornly coherent, and surprisingly comprehensible. It is not an exaggeration to claim that most researchers in modal logic have at least a nodding acquaintance with the majority of the topics discussed in this Handbook, feel that this sort of extensive knowledge is useful, and would like their students to have a map of the terrain at least as wide. This Handbook is an attempt to provide such a map. To put it another way, it tries to gather together the background assumptions, the working knowledge, the mathematical techniques, and the general world view that add up to that somewhat elusive entity “contemporary modal logic”, and to bring it together in a digestible form. We hope it will provide exactly the sort of snapshot of the field that will serve as intellectual nourishment for the next generation of researchers in, and users of, modal logic.

We made no serious effort to impose notational or other kinds of uniformity on the authors. This was partly for pragmatic reasons: it was always evident that the attempt to impose a standard notation would please nobody — and who is to say that linguistic diversity is worse than linguistic, or even cultural, uniformity? But there are deeper reasons for our hands-off stance. Research in contemporary modal logic takes place in a shifting environment that jostles the borders of many fields. One modal logician’s interests may lead to the frontiers of linguistics or automated reasoning, another to the foundations of computation and games, and another towards purely mathematical issues concerning topological spaces. Moreover, these interests keep converging, and diverging, in new and often unpredictable ways; we have given some telling examples already. Had this Handbook appeared five years earlier, the borders between various fields would have been drawn somewhat differently, and it is entirely possible that in another five years many will have to be drawn yet again. In the face of such flux, imposing uniformity would be an artificial exercise. Modal logic is a sea, and the point of this Handbook is to help the reader learn to swim in it; pretending it is a swimming pool distorts the truth, and what’s worse, is unhelpful.

USING THE HANDBOOK

We suspect that most of our readers will be accustomed to navigating their way through weighty research tomes, and will require little in the way of advice. In particular, readers who already know something about modal logic should simply consult the Table of Contents and start where it looks most interesting; while there are cross-references between the chapters, each is, to a great extent, self contained, so this is a viable strategy. Moreover, the Handbook can be used as a reference to the field. In particular, the index gives detailed entries for most common logics, notions, and results.

But we have also attempted to make the Handbook accessible to less experienced readers. Now, we should say right away that by “less experienced” we mean less experienced in *modal* logic. This is a technical volume, and readers without technical background are going to find it difficult. Thus our less experienced reader is someone who already has some understanding of what modern logic is about, and why and where it is useful, and who wants to find out something about what modern modal logic has to offer.

The Handbook is structured to provide an answer to such readers. The book has 21 chapters, and is divided into four parts: Basic Theory, Advanced Theory, Variations and Extensions, and Applications. Although independent, together they tell a story, the story of contemporary modal logic. This story starts with the basic tools and techniques, takes the reader to the outer reaches of the underlying mathematical theory, surveys the key points where the modal approach is being adapted and extended, and finishes by examining the various applications which modal logic serves and from which it draws inspiration. Let’s take a closer look at how all this unfolds over the course of the Handbook.

Part 1. Basic Theory The chapters in Part 1 lay the foundation for later ones. Together they present an overview of the most fundamental themes, techniques and results in contemporary modal logic. The growing impact of computer science is clearly reflected in the choice of topics: two chapters are wholly devoted to complexity, decision methods, and implementation.

Chapter 1. Modal Logic: A Semantic Perspective. Patrick Blackburn and Johan van Benthem.

This chapter discusses the semantic ideas underlying modern modal logic, and in particular, Kripke semantics — or relational semantics, as it now (more informatively and fairly) tends to be called. It introduces the basic model theoretic constructions in a modern way, explores links between modal logic and classical (predicate) logic, both on models and on frames, and examines the extent to which the key semantic ideas transfer to richer modal logics and languages while maintaining a relatively low computational complexity. It also introduces some alternative viewpoints: algebraic semantics, neighbourhood semantics, and topological semantics.

Chapter 2. Modal Proof Theory. Melvin Fitting.

Modal proof theory is the study of syntactic calculi, defined in terms of symbol manipulation, for performing modal logical reasoning. How can such systems be designed? Is there an interesting range of design choices? And how can the syntactic ideas underlying proof calculi be linked with the semantic ideas introduced in Chapter 1? This chapter answers these questions by introducing a wide range of proof styles, and discussing modal completeness theory, the fundamental bridge between proof-theoretical and semantic investigations.

Chapter 3. Complexity of Modal Logic. Maarten Marx.

The basic modal language, when interpreted over relational models, can be regarded as a decidable fragment of classical logic. But this observation immediately leads to a host of further questions. Given that it is decidable, how difficult is it to compute with? That is, what is the computational complexity of determining validity, or of performing more modest tasks like model checking? And what are the parameters that affect modal complexity results, and what happens when we play with their settings? This chapter, an introduction to the computational complexity of modal logic, provides some fundamental answers.

Chapter 4. Computational Modal Logic. Ian Horrocks, Ullrich Hustadt, Ulrike Sattler, and Renate Schmidt.

Although Chapter 2 introduced modal proof theory, and Chapter 3 studied the computational complexity of modal logic, only with this chapter do we reach the heartland of computational modal logic: how to build modal inference systems that are efficient in practice. Although it surveys a number of topics, this chapter concentrates on two fundamental issues: how resolution and tableaux methods can be adapted to modal logic, and how these methods are related.

Part 2. Advanced theory The chapters in Part 2 provide a deep and wide ranging theoretical analysis of modal logic that is broad enough to apply to many application areas. In some cases they provide deeper perspectives on topics already introduced in Part 1, but often they introduce ideas barely hinted at in earlier chapters. Taken together, they present the central core of contemporary insight into the mathematical structure of modal logic.

Chapter 5. Model Theory of Modal Logic. Valentin Goranko and Martin Otto.

At the heart of relational semantics is the idea of interpreting modal languages over relational structures by viewing them as fragments of first-order predicate logic or some stronger formalism. This perspective is not only intuitively attractive, it also makes available to modal logic the results and tools developed in such areas as classical model theory and finite model theory. This chapter shows, in great detail, how such tools can be put to work to gain a deep mathematical understanding of modal model theory, and what makes it *sui generis*.

Chapter 6. Algebras and Coalgebras. Yde Venema.

This chapter develops in detail the algebraic semantics of modal logic and introduces an alternative coalgebraic approach. Algebraic semantics, which has thrived as a research area since the early 1970s, is important because it makes it possible to apply general techniques from universal algebra to the study of modal logic. The approach has given rise to some of the most penetrating analyses of the mathematics of modality. The more recent coalgebraic approach, which also links up with category theory, is valuable because it offers a uniform mathematical setting in which to analyse dynamic systems in terms of modal logic.

Chapter 7. Modal Decision Problems. Frank Wolter and Michael Zakharyashev.

Modal logic is decidable — or at least it is when interpreted on the class of all models. But change the interpreting class of models, and you change the logical validities, and decidability is typically lost when the structural conditions come too close to ‘danger zones’ such as tiling patterns, arithmetic, or other structures allowing for Turing machine computation. This chapter is a detailed examination of how such properties as decidability, the finite model property, and finite axiomatisability are distributed across the lattice of normal modal logics. The emphasis is on providing general results, and drawing attention to important open questions.

Chapter 8. Modal Consequence Relations. Marcus Kracht.

The notion of consequence, which tells us when a conclusion follows from given premises, is a fundamental logical concept, and in the setting of modal logic it can be defined in a number of different ways. This chapter surveys some of the most important ideas, covering in detail such topics as local versus global consequence, reducing multimodal consequence to monomodal consequence, interpolation theorems, and the admissibility of rules. As with Chapter 7, the emphasis is on providing general results which apply across a wide range of logics.

Part 3. Variations and Extensions The main focus of the chapters in Parts 1 and 2 was on relatively simple propositional modal systems based on (collections of) \diamond and \square modalities. Such systems are historically central but they don't exhaust the kinds of logic that now go under the name "modal logic". The chapters in Part 3 introduce some of the extensions and variations of the basic modal technology that the reader is likely to encounter.

Chapter 9. First-order Modal Logic. Torben Braüner and Silvio Ghilardi.

First-order modal logics are modal logics in which the underlying propositional logic has been replaced by a first-order predicate logic. These are one of the oldest forms of modal logic, and arguably the most philosophically important. They also pose some of the most difficult mathematical challenges. This chapter first surveys basic first-order modal logics, and then examines recent attempts to find a general mathematical setting in which to analyse them.

Chapter 10. Higher-order Modal Logic. Reinhard Muskens.

The basic ideas of modal logic have also been extended to higher-order settings, and indeed, extended in a number of different ways. This chapter motivates such extensions, some of them from linguistic semantics in the tradition of Richard Montague, examines some of the more historically influential ones, indicates some of the difficulties that can arise in the transition to higher-order logic, and finally shows how these difficulties can be overcome.

Chapter 11. Temporal Logic. Ian Hodkinson and Mark Reynolds.

Temporal logic is one of the classic branches of modal logic and is currently one of the most active. It has been remarkably fruitful in the issues it has raised (what kinds of temporal structure should we work with?), the results it has given rise to (it is the source of some of the most interesting expressivity results in modal logic), and as an applied tool (contemporary model checking technology is based on temporal logic). This chapter will introduce the reader to the key issues of this important and diverse area.

Chapter 12. Modal μ -Calculus. Julian Bradfield and Colin Stirling.

In the late 1960s, pioneers in reasoning about programs adopted some key ideas of modal logic. They repaid the debt handsomely. Among other things, they developed dynamic logic (used in several chapters of this Handbook), and the modal μ -calculus, one of the most interesting modal formalisms to have emerged in the last two decades. This provides second-order expressive power sufficient to generalise the most common temporal logics, but is still decidable and has the finite model property. It raises many intriguing issues about the interface between modal logic, complexity theory, and automata theory.

Chapter 13. Description Logic. Franz Baader and Carsten Lutz.

Modal logic is sometimes thought of as an intrinsically intensional logic, suitable only for applications such as reasoning about necessity, possibility, and knowledge. But description logics (which developed from pioneering work in the AI community) are undeniably modal logics and, as the description logic community has shown in impressive detail, are extremely well suited for

reasoning about ordinary individuals and the relations between them. This chapter is a detailed introduction to one of modal logic's closest neighbours.

Chapter 14. Hybrid Logics. Carlos Areces and Balder ten Cate.

Standard modal logics use modalities for talking about the relations in relational structures, but don't contain mechanisms for talking about particular worlds. Hybrid logic arises when mechanisms for naming and asserting identity of worlds are added; to give an analogy, they are to standard modal systems what first-order languages with equality are to equality-free languages. This chapter surveys the proof theory, expressivity, and complexity of a number of the better known hybrid logics, thereby giving a snapshot of the logical territory lying between the basic modal languages and their classical companions.

Chapter 15. Combining Modal Logics. Agi Kurucz.

The idea of combining modal logics (for example, a modal logic of time with a modal logic of knowledge) is natural for many applications. But how can modal logics be combined, and what happens when you combine them? This chapter surveys two key combination methods (fusions and products) in detail, shows how various properties do (or do not) transfer from the individual logics to the combination, and briefly examines a number of other combination methods. The properties of combined logics turn out to depend in subtle ways on those of their components plus the particular method of combination.

Part 4. Applications Historically, modal logic has been profoundly influenced by its applications, which have been extremely diverse in nature. The chapters in Part 4 survey the key application domains, thereby showing where modal logic comes from, where it has visited along the way, and also indicating areas to which it is likely to return.

Chapter 16. Modal Logic in Mathematics. Sergei Artemov.

Mathematics is one of modal logic's oldest application areas. In particular, the pioneering work of Gödel in the 1930s showed that modal logic offered an important perspective on the notion of mathematical provability, and (more recently) modal logics of proof have been developed. But modal logic also gives rise to natural logics of space and dynamic systems, and even turns out to be a tool with applications in set theory. This chapter surveys these themes. In doing so, it emphasises an intriguing duality in interpreting the modal box: either as the universal quantifier "in all worlds", or as the existential "there exists a proof". Just when and how such accounts converge is a deep metamathematical issue.

Chapter 17. Automata-Theoretic Techniques for Temporal Reasoning. Moshe Y. Vardi.

Many modal and temporal logics can be viewed as fragments of monadic second-order logic over trees in a suitable signature, so there is a clear theoretical link (via Rabin's celebrated decidability theorem) between modal logic and automata theory. But this link turns out to have practical repercussions for computational applications. In particular, by viewing temporal formulas as giving rise to what are known as "alternating automata", we gain a theoretically transparent but also practical perspective on both validity and model checking, one of the most significant applications of contemporary modal logic.

Chapter 18. Intelligent Agents and Common-Sense Reasoning. John-Jules Meyer and Frank Veltman.

Modal logics have been used in AI in a number of different ways. This chapter discusses two of its more important roles there. The first is as a logic of agents, and here the chapter takes the reader from basic epistemic and deontic logic to multi-agent logics of beliefs, desires, and

intentions. The second is as a model of common sense reasoning, and here the chapter covers modal treatments of counterfactual conditionals and non-monotonic reasoning in a variety of guises, including default reasoning.

Chapter 19. Applications of Modal Logic in Linguistics. Lawrence S. Moss and Hans Jörg Tiede.

Modal logic is best known in linguistics for the light it throws on semantics; indeed Richard Montague's use of higher-order modal logic for this purposes is widely considered to be the starting point of modern natural language semantics. Recently, however, modal logic has also been used to analyse syntactic structure, and interesting links with formal language theory have thereby emerged. This chapter discusses both topics, providing a sophisticated view of modern interfaces between logic and natural language.

Chapter 20. Modal Logic for Games and Information. Wiebe van der Hoek and Marc Pauly.

Game-theoretic ideas have long played an influential role in analysing various branches of logic, but the focus of this chapter is on using modal logics to describe and reason about games. After introducing the basic ideas of game theory, it systematically investigates how modal logic can be used to do this. Three main topics are discussed: modeling imperfect information and multi-agent information update via dynamic epistemic logics; reasoning about game structure through operations for combining games; and logics of collective action and the power of coalitions of agents over time.

Chapter 21. Modal Logic and Philosophy. Sten Lindström and Krister Segerberg.

Modal logic was born in philosophy, and though it has since travelled widely, it still retains important links with the discipline. This chapter first discusses the historical heartland of philosophical modal logic: namely, the scope and limitations of modal logic as an account of necessity and possibility. It then examines two more recent topics: modal logic and the logic of belief change, and modal logic as a logic of action.

Together, these chapters present a broad picture of modal logic today. Of course, some choices had to be made, and some bias may remain. In particular, the emphasis throughout has been on relational graph-style models. The editors fully acknowledge that there are other important traditions, such as modal logics based on non-classical logics, proof-theoretic semantics, as well as the more general neighbourhood semantics. Some of these have a historical pedigree reaching back to the 1930s, and they are still very much alive. These approaches do occur at various places in the book, but we have not made them fundamental to the Handbook's architecture. We think this is a fair reflection of the bulk of current research, but times may change.

Some readers may also think that this Handbook has a bias toward propositional model systems, leaving predicate-logical versions underrepresented. This may be true to some extent — however, the chapters on modal predicate logic, modal logic in philosophy, and also temporal logic give a clear account not only of the basic theory of modal predicate logic, but also of recent developments of its mathematics and its relevance to computer science. But there is also a more defiant stance. Predicate logic itself can be profitably viewed as a modal logic of variable assignment and assignment change. And in that light, modal predicate logic is not some privileged enrichment of modal logic. It is really about combining a propositional modal logic of worlds and one of variable assignment. And the reader can learn a lot about both, and about combination methods, from the pages of this Handbook!

Once these chapters are seen together, many questions may at once be formulated concerning further relationships between them. The editors considered adding some remarks on this topic,

as diversity tempered by the resulting need for system comparison are a major driving force for innovation in the field. But in the end, we have decided to let the chapters, and their juxtaposition, speak for themselves.

Finally, we would like to remark that modal logic is as much a community of living people as a family of systems. The results and insights in this Handbook exist only because of a long line of distinguished researchers who have shaped modal logic and its interfaces with computer science and other fields. The Hall of Fame of our field certainly extends beyond the grand old names of the classical period; just read this Handbook, and you will come to know the grand new names by their fruits.

COMMENTATORS

Following an idea of Dov Gabbay's used in many publications, we made use of the designated commentator system in this Handbook. That is, in addition to editorial feedback, we attempted to find, for each chapter, a reader who could provide the kind of feedback that would inspire the authors during the writing process. In some cases we chose commentators with special expertise in the topic of the chapter. In other cases, we felt that the comments offered by someone working in a somewhat different area might be more appropriate and helpful. Moreover, whenever possible, we chose authors of other chapters as commentators, as we felt this would improve the Handbook's coherence.

We are extremely grateful to everyone who agreed to undertake this task; in many cases the input of a commentator acted as precisely the catalyst needed to help the full potential of a chapter to emerge. Our commentators were (Chapter 19 had no commentator):

<i>Chapter 1: Aleksander Chagrov.</i>	<i>Chapter 2: Heinrich Wansing.</i>
<i>Chapter 3: Martin Otto.</i>	<i>Chapter 4: Carlos Areces.</i>
<i>Chapter 5: Maarten de Rijke.</i>	<i>Chapter 6: Aleksander Kurz.</i>
<i>Chapter 7: Rosalie Iemhoff.</i>	<i>Chapter 8: Ian Hodkinson.</i>
<i>Chapter 9: Melvin Fitting.</i>	<i>Chapter 10: Grigori Mints.</i>
<i>Chapter 11: Valentin Goranko.</i>	<i>Chapter 12: Yde Venema.</i>
<i>Chapter 13: Silvio Ghilardi.</i>	<i>Chapter 14: Ulrike Sattler.</i>
<i>Chapter 15: Carsten Lutz.</i>	<i>Chapter 16: Rineke Verbrugge.</i>
<i>Chapter 17: Julian Bradfield.</i>	<i>Chapter 18: Reinhard Muskens.</i>
<i>Chapter 20: Giacomo Bonanno.</i>	<i>Chapter 21: Vincent Hendricks.</i>

FURTHER INFORMATION

We have set up a home page for the Handbook at:

<http://www.csc.liv.ac.uk/~frank/MLHandbook>

We will make available there any corrections that may need to be made, and news concerning future Handbook-related developments. We welcome feedback from our readers.

It would not be useful to attempt to list all the workshops, conferences, and journals where work on modal logic is published; given what we have said about its wide range of applications and techniques, it will come as no surprise that such work may be made public in a wide variety of forums. However it is worthwhile taking this opportunity to mention two workshops specifically

devoted to modal logic. The first is *Advances in Modal Logic* (AiML), modal logic's main event, which is held every two years. You can find out more about this event, and the associated book series at:

<http://www.aiml.net>

Here we'll simply say that AiML attempts to bring together scholars working in all areas of modal logic and its applications. The second workshop is *Methods for Modalities* (M4M), see

<http://m4m.loria.fr>

M4M is also held every two years. It is more practically oriented than AiML, focusing on the development of computational tools and results for modal logic.

ACKNOWLEDGEMENTS

Editing a handbook is a lengthy task, and sometimes, very modally, it seems as if all the things that could go wrong take the opportunity to actually do so. Fortunately, we had the benefit of a great deal of support, which generally enabled us to dissolve problems before they became too pressing. Indeed, in retrospect, it's clear we had a relatively easy time of it, and we would like to thank all the people who made this possible.

First, and foremost, we would like to thank the authors of the Handbook's chapters. They showed enormous enthusiasm for the project, many of them going out of their way to attend a meeting to set the Handbook in motion (in March 2004, Amsterdam; the meeting was kindly sponsored by NWO, KNAW and other Dutch scientific organisations). We received a great deal of valuable feedback from them, and they showed good-humoured patience when faced by style file changes, demands for quick attention to page proofs, and requests for additions or alterations to their chapters. It is a privilege to work with such a lively and cooperative community. The resulting Handbook is a testimony to their efforts. Our deepest thanks to them all.

We were also extremely lucky with the support we received on the production side. First, our thanks to Arjen Sevenster, of Elsevier, for his enthusiastic support in making this Handbook happen right from the start, including judicious alternations of leaving us be and prodding, and to Andy Deelen, our contact at Elsevier, for her prompt responses to our questions. We are also extremely grateful to Carlos Areces, Christian Günsel, Eric Kow, Oliver Kutz and Dirk Walther for the technical help they gave us with the website, the style files, and the indexing. Special thanks are due to Balder ten Cate for organising the meeting in Amsterdam, and to Guido Governatori, who wrote the \LaTeX style file used to produce the final version of the Handbook, and who responded quickly and generously to our urgent requests.

And, finally, we come to Jane Spurr. Without her, it is hard to imagine how we could have put the Handbook together. It was Jane who grafted the 21 chapters into one big \LaTeX file overcoming many evil obstacles, who dealt with final corrections, and who found ways of making our numerous formatting requests work. Thanks to her unflagging good humour and patience, the final production stage became a positive pleasure.

Patrick Blackburn
Johan van Benthem
Frank Wolter