TEMPORAL PATTERNS AND MODAL STRUCTURE

Johan van Benthem

Institute for Logic, Language and Computation
University of Amsterdam

July 1998

Abstract

Temporal logic arose at the border of philosophy and linguistics. From the seventies onward, it became a major tool also in computer science and artificial intelligence, which have become the most powerful source of new logical developments since. We discuss some recent themes demonstrating new connections with modal logic. In the course of this, we point out some new types of open research questions.

to appear in
A. Montanari, A. Policriti & Y. Venema, eds.
special issue on Temporal Logic
Bulletin of the Interest Group for Pure and Applied Logics

1 Variety of Motivations

Temporal logic as an identifiable research area emerged in the 50s, largely due to the pioneering efforts of Arthur Prior. Born some decades later than modern modal logic, it turned out technically much like it. But unlike in the case of modal logic, from the start, different motivations fed into temporal logic, coming from philosophy (analysis of historical argumentations involving the structure of time), linguistics (description of tenses and other temporal expressions in natural languages), and to some extent also mathematics (theory of linear and branching orders, non-point-based geometries). This breadth reflects, of course, the ubiquity of time in virtually every intellectual endeavour. Van Benthem 1983 ("The Logic of Time") tries to systematize temporal logic across these various motivations, even taking up issues in relativistic physical space-time. In the 80s, temporal logic penetrated into further disciplines, such as cognitive psychology (cf. Jackson & Michon 1985). By far the most spectacular development, however, has taken place at the interface of logic with computer science and artificial intelligence, where temporal logic itself really changed its agenda and its depth of reach. Broad surveys across this more contemporary width are van Benthem 1983, 1989, 1995. This paper presents some recent trends along this spectrum, without scholarly completeness.

There may not even be one unified field of logical studies of time. After all, it is not quite clear that the different 'foster disciplines' are concerned with the same thing. Physicists deal primarily with real-world space-time, linguists and psychologists with 'representation time', and computer scientists with machine-induced 'process time'. At least, one has to be aware that these notions differ, and then, see if or how they relate. E.g., linguistic representation time has to 'fit' space-time to keep us attuned to our physical environment, and computational process-time needs to 'mesh' with space-time to ensure reliable on-line performance. Much of this conceptual interfacing is non-trivial, and remains to be done. Indeed, technically and sociologically, there are now different communities in temporal logic with different styles and agendas, such as those active in pure logic, computer science, or AI – while formal philosophers of time are yet another separate breed. Perhaps the highest status has been achieved by the CS-based variety of temporal logic developed by Amir Pnueli since the mid 70s, studying temporal specification of desired executions of complex distributed programs (cf. De Bakker, de Roever & Rozenberg, eds., 1989). This work was honoured by the 1996 Turing Award, the highest recognition in computer science. Even though this line still shows many traces of the spirit and tools of traditional temporal logic, it is definitely diverging. (E.g., the currently emerging work on specifying 'hybrid temporal systems' is mixing the analysis of 'logical time' with standard physical engineering techniques).

Another prominent line in computer science integrates concerns from temporal logic with modern process theories (dynamic logic, μ–calculus, process algebra), surveyed in Colin Stirling's contribution to this issue. (Cf. also Stirling 1992, Harel, Kozen & Tiuryn 1996.) Our presentation will not attempt to cover this broad panorama. Instead, we concentrate on a more traditional logical line in the field, emphasizing connections with some newer developments in modal logic. But on occasion, we find further evidence of the diversity of temporal structure. To mention one instance, in computing a fixed point for a monotone operator (the key to modal and general process theories), the natural approximation stages generate one more discrete 'cognitive evaluation time'!

To conclude, here is our view of temporal logic in a nutshell. This field sits at the interface of many uses of time, across different disciplines. Its business is to offer a conceptual framework in which one can model these views of time exactly, study their common features, but also, see how they might be integrated. As a spin-off, the various disciplines involved may also learn from each other's agendas. Our aim here is merely to make a few of all possible comparisons, and highlight some current developments. Moreover, we assume that the reader is familiar with the logical basics of the field.

2 Temporal Logic as Modal Logic

Prior's temporal logic may be viewed as a two-directional modal logic, with operators F ("at least once in the future") and P ("at least once in the past"). This makes temporal logic and modal logic very close mathematically. Traditionally, however, one stressed a philosophical difference in the 'direction of thought'. Given the relative concreteness of its subject matter, temporal logic constructs description languages for independently given temporal models, while modal logic is in the business of constructing models for given modal languages. (The syntax of modality, alas, is clearer than its semantics.) This particular distinction no longer works these days. Partly under the influence of temporal logic (Gabbay 1981), over the past decade modal logicians, too, have turned with zest to new language construction over Kripke models (or in CS speak: 'labeled transition systems'). Technical themes therefore have tended to converge between the fields. A typical example of this current trend is the theory of the 'process equivalence' of *bisimulation* in the analysis of expressive power for modal or temporal languages.

2.1 Simulations: From Matching Single States To Matching Tuples

Prior's original propositional F, P language is invariant for two-sided bisimulations, which are like standard modal bisimulations, but now with back-and-forth clauses for both successors and predecessors in the ordering. The resulting theory is largely the same as in the modal case. In particular, a first-order formula $\phi = \phi(t)$ in the language with binary $<$ (temporal precedence) and unary predicates P, Q, ... (one for each proposition letter p, q, ...) is P, F–definable iff it is invariant for two-sided bisimulation. (The purely modal proof of van Benthem 1976 goes through *in toto*.) This observation can be used e.g. for quick proofs of non-TL-definability by concrete counterexamples. To be sure, there are some minor technical differences here and there. For instance, the standard modal technique of 'tree unraveling' has to be modified to allow for branching patterns in two directions now, to make their behaviour more symmetric. Also, some very smooth Gentzen-style decidability reductions available for modal logic fail for the two-sided temporal language (Andréka, van Benthem & Németi 1998). Theoretically, Spaan 1993 treats Prior's logic as a *disjoint sum* of two separate monomodal logics with one 'connection axiom', making the two accessibility relations mutual converses. The latter linkage can affect complexity of satisfiability for temporal logics in general. (In particular, it may be higher than the maximum of the separate complexities for the two monomodal components.) On the other hand, temporal logic repays us by having a natural adjunction between the valid equivalence of (say) $P\phi \rightarrow \psi$ with $\phi \rightarrow G\psi$.

Early on, temporal logicians have started studying much stronger languages than Prior's original one. A pioneering example was the system with operators S, U, devised in Kamp 1966 to mirror the adverbs "since" and "until" of natural language. E.g., $U\phi\psi$ says about any point t that ϕ holds in some point t' later than t, while ψ holds in any point in between t and t'. ($S\phi\psi$ is the downward dual of this.) Kamp's Theorem says that, at least over Dedekind complete linear orders with unary point predicates, every first-order formula is definable in the S, U temporal language. This pilot result has spawned a whole series of expressive completeness theorems for temporal languages vis-à-vis first-order ones. These temporal languages have not yet been studied from a simulation perspective, and they pose some interesting questions. For instance, which stronger simulation is needed to characterise definability with S, U? (To see the difference with the above, there is an obvious temporal bisimulation in the earlier sense between a 2-cycle and a 1-cycle model, even though the former satisfies $UT\bot$ while the latter does not.) There is no evident Ehrenfeucht back-and-forth clause here, because of the double $\exists\forall$ quantifier combination in the truth condition for U. This particular simulation problem was solved recently in Kurtonina and de Rijke 1997. Their solution turns out to require a genuine extension of the game: temporal bisimulations now relate both states and *ordered pairs of states* across temporal models.

This move is really more general. As one increases expressive power of temporal languages, the natural bisimulations move away from the computer science paradigm of single state-links, and start relating pairs, triples, ... and in general, *finite tuples* of temporal points (cf. van Benthem 1991). Such simulations are a bit harder to visualise (van Benthem 1996 has concrete examples for processes). But in the limit, they move closer to 'potential isomorphisms', i.e., non-empty sets of finite partial isomorphisms satisfying the usual back-and-forth conditions which are natural model links for model theorists (Barwise 1975). The latter are also the natural tool for dealing with *many-dimensional modal logics*, interpreted at tuples of worlds up to a fixed finite length (cf. the monograph Marx & Venema 1996). Immerman & Kozen 1987 have deconstructed Kamp's Theorem in this finite-dimensional spirit. The key to expressive completeness for a finite temporal operator formalism turns out to be restriction to some *finite number of variables* used (free or bound) in the first-order language of its temporal model class. Standard modal logic requires only two variables, as may be seen by exercising a little care with the first-order 'standard translation' of its base language. And the same holds for Prior's original P, F language. Using more sophisticated Ehrenfeucht game-style arguments, Immerman & Kozen managed to show that three variables suffice for writing any first-order assertion (in a language with a binary ordering predicate plus unary properties of points) over linear orders, Dedekind complete or not.

The general model-theoretic notion which generalizes many bisimulations is that of a *potential k-isomorphism*. This is a non-empty family **PI** of partial isomorphisms, each of size at most k, which is closed under taking restrictions, and which satisfies the following back-and-forth property. If the size of a partial isomorphism F in **PI** is smaller than k , and we pick any object d in one model, then there is an object e in the other model such that $F \cup \{(d, e)\}$ is in **PI** . Van Benthem 1991 proves that any first-order formula is invariant for potential k–isomorphism iff it is definable by a first-order formula with at most k variables in all, free or bound. (Barwise & van Benthem 1997 give a new technique extending this type of result to infinitary languages.) On the other hand, we know from Gabbay 1981 that k–variable fragments of first-order logic have expressively equivalent k–dimensional temporal logics with some finite set of operators. By and large, standard temporal languages make do with three variables. That is, their compositional evaluation involves at most inspection of 3-point patterns of 'betweenness'. More complex 4-point patterns would be required, e.g., for defining suprema or infima in a branching temporal ordering. Conversely, semantic 'pattern invariance' may also help us in designing new languages, better suited for some temporal model class. For instance, instead of the somewhat traditional S, U language, the three-variable fragment of first-order logic *itself*, or its corresponding complete 3-dimensional modal operator language, is a good (and in general, a richer) alternative for the description of temporal structure. As a side benefit, its characterisation in terms of potential 3–isomorphisms does not need special $\exists\forall$ tricks.

Conclusion. In current research, there is a tight fit between either designing some language with a certain desired expressive power, or finding a simulation between models having just the right disregard for details of structure. Equivalence theorems may then state (there are various options) that a first-order formula is invariant for the simulation iff it is definable in the specified temporal language format. These are two natural ways of studying 'temporal structure', which are profitably developed in tandem.

2.2 Quantifier Bounds and Decidability

From expressive power, we turn to complexity. Much is known about decidability or more precise complexity of modal and temporal languages, over universal or restricted model classes (cf. Chagrov & Zakharyashev 1997, Wolter & Zakharyashev 1996). Here we only mention one recent general line of analysis, which covers both modal and temporal languages. Prior's F, P language is *decidable*, as may be shown by a simple filtration argument (Goldblatt 1987). A more sophisticated reason for this phenomenon is the following fact: the whole two-variable fragment of first-order logic is decidable (Mortimer 1975). Next, the three-variable fragment of first-order logic is undecidable, as it embeds all of relational set algebra. But then, what to do with the three-variable temporal formalisms, needed, e.g., to transcribe the typical operators S, U ?

Here one can use an alternative analysis, proposed in Andréka, van Benthem & Németi 1997, which focuses on bounded quantifier patterns $\exists \mathbf{y}\ (G(\mathbf{x}, \mathbf{y})\ \&\ \phi(\mathbf{x}, \mathbf{y}))$, where the atom $G(\mathbf{x}, \mathbf{y})$ serves as a so-called *guard*. The *Guarded Fragment* GF consists of all first-order formulas constructed from atoms with Booleans using this special schema of quantification. This fragment runs across all finite-variable levels. GF is decidable, indeed doubly-exponential time complete (Grädel 1997). Guarded analysis fits the basic temporal language just as well as the basic modal one. The reason is that, in the above guard atoms, we allow any order (or multiplicity) of the variables. Thus, there is no special preference for successors ('Rxy') over predecessors ('Ryx'). As it stands, however, this analysis does not apply to the richer temporal S, U language. For, the first-order transcription of $U\phi\psi$ is the provably non-guarded three-variable formula $\exists t'\ (t<t'\ \&\ \phi(t')\ \&\ \forall t''\ (t<t''<t' \rightarrow \psi(t'')))$. If we are to prove general decidability here, there must be some special feature which makes these formulas deviate from the 3-variable

fragment in general. The explanation turns out to lie in an extension of GF with *loose guards* (yielding 'pairwise guarded formulas' with conjunctions of atomic guards), in which every pair of variables from x, y must occur in some conjunctive guard atom. (Note that the above transcription of "until" is indeed pairwise guarded.) The Pairwise Guarded Fragment remains decidable, and therefore basic S, U logic is decidable via translation into the latter (van Benthem 1997A). This first-order translation may also be used for automated deduction in temporal logic, witness the resolution-style analysis of GF and its ilk in De Nivelle 1997. By now, there is a variety of guarded languages, providing further fine-structure for decidable temporal operators (van Benthem 1998A).

Conclusion. As with modal logics, decidability of basic temporal formalisms may be understood by reference to powerful decidable first-order fragments. These are of new kinds, different from the usual decidable 'fragments' studied in the literature.

Coda. This is not the whole story, however. Guarded syntax analysis does not explain the decidability of many *special temporal model classes*, such as transitive or linear orders. For some of these, defining conditions are indeed (pairwise) guarded – but more often, they are not. Of course, there need be no uniform explanation for all complexity features of temporal reasoning. But there may be one over the ubiquitous *transitive* models. Transitivity $\forall xyz\ ((Rxy \,\&\, Ryz) \rightarrow Rxz)$ is not pairwise guarded. Decidability of the two-variable first-order fragment does not apply either: transitivity essentially needs 3 variables. Then why is its basic temporal logic easily decidable? There are two possible routes here. One extends the syntactic scope of GF and its ilk, to find still broader decidability results. We doubt that this is feasible. Transitivity is dangerous, and can make first-order fragments undecidable (Börger, Grädel & Gurevich 1996). But there is another diagnosis. Recall that propositional dynamic logic PDL, and even the μ–calculus are decidable (Harel, Kozen & Tiuryn 1996). Now the basic modal logic K4 over transitive models (basic temporal logic is similar) is precisely the logic of any iteration modality $[a^*]$ over general models, without any special restrictions at all. This is a genuinely different route. The PDL language cannot define transitivity. Like the basic modal one, it is *invariant for bisimulation* (the infinitary conjunctions that define iteration do not affect this), while transitivity is not. Hence, general decidability results over transitive temporal models follow from the known decidability of PDL with a converse operator, or more generally, from that of the temporal μ–calculus. If we want to achieve further generality, though, we must enter the realm of current speculation.

Open Question   Find decidable fixed-point extensions of the Guarded Fragment.

2.3 Unraveling and Coexistence

The preceding items showed that extended modal logic and temporal logic can really be the same thing. But there are also natural settings in which we can *contrast* the two styles of thinking. For a concrete example, consider something as well-known as the *sequence unraveling* of any modal labeled transition system into a bisimilar tree. We can think of the model as some kind of machine, where modalities describe possible moves. By contrast, the unraveled tree is really a temporal model of 'possible histories'. This allows for comparisons between different languages over these bisimilar models: modal ML and temporal TL. We quote some relevant results from Andréka, van Benthem & Németi 1998, referring to both standard tree unraveling and 'ω–unraveling', a model construction procedure which simultaneously produces countably many copies of each daughter node (and its subtree):

$\mathbf{M} \equiv_{mod} \mathbf{N}$ iff $\text{tree}(\mathbf{M}) \equiv_{F,\,P} \text{tree}(\mathbf{N})$

$\mathbf{M} \equiv_{mod} \mathbf{N}$ iff $\omega\text{-tree}(\mathbf{M}) \equiv_{FOL} \omega\text{-tree}(\mathbf{M})$

Thus, modal equivalence of machines may be upgraded to first-order equivalence of their trees of histories. Here is an alternative analysis of the situation, not via linguistic equivalence, but in terms of bisimulation and isomorphism (where $\kappa > |\mathbf{M}|, |\mathbf{N}|$):

$\mathbf{M}$ <u>bisimilar to</u> $\mathbf{N}$ iff κ-tree($\mathbf{M}$) <u>isomorphic to</u> κ-tree($\mathbf{N}$)

A more elaborate treatment of this kind of result may be found in D'Agostino 1998, whose study of 'κ-branching bisimulations' generalizes the Barwise & Moss 1996 theory of circular sets. (The main operator of the matching languages counts alternative routes at a fork – which has more temporal than modal motivation anyway.)

But further forms of coexistence between the two perspectives may occur. The temporal style of thinking concentrates on *branches*, where histories take place. In particular, one may only be interested in those branches which represent 'desirable histories', satisfying special conditions of safety, liveness or fairness for the processes manifesting themselves (cf. Pnueli 1989). (The latter are unrealistic for human histories, alas.) Accordingly, one can restrict the branches in the unraveled model tree($\mathbf{M}$) to just the 'reasonable' runs in some sense. Also, one can perform a two-way unraveling of $\mathbf{M}$ to zigzagging time travels using both successors and predecessors, etcetera. All these temporal structures share one feature: they become richer than the original models. Most conspicuously, an unraveled tree($\mathbf{M}$) is a model for a richer temporal language, viz. a two-sorted *branching temporal logic*, interpreted over both points and branches. I know of no good comparisons between the austere modal logic of a labeled transition system and the richer branching temporal logic of its tree unraveling. But one would assume that the latter encodes all of the former, so that there should be effective translations between the two levels. (Incidentally, all information about **M** up to bisimulation is contained in its characteristic infinitary modal Scott-formula, constructed in Barwise & Moss 1996; cf. also van Benthem & Bergstra 1995. What about a similar characterisation for branching temporal structures?)

Conclusion. Standard modal languages over temporal models and branching temporal languages over their tree unravelings are natural companions. Intuitively, the former describe temporal 'potentialities', the latter temporal 'evolutions'. What are the systematic logical relations between these two views of temporal structure?

Coda. Sistla & Clarke 1985 even *define* truth of a linear-time temporal formula ϕ in a model **M**, s as truth of ϕ at s on at least one branch in the unraveling of **M**. (They then determine the complexity of this notion for various temporal languages.) This notion takes only a fragment of the full branching temporal logic over tree unravelings. Can one still reduce this special notion to truth of ordinary modal formulas in **M** ?

3 Temporal Geometry

Various systems extend the logical treatment of time to that of space. This is a natural move. In common sense reasoning, we often switch from temporal to spatial metaphors, and vice versa. In fact, some of the most beautiful axiomatisations for specific temporal models in the literature *combine* the two realms, such as the characterization of the complete modal logic of Minkowski space-time (Goldblatt 1980, Shehtman 1983). But this direction is not what we have in mind for this section. Our idea is rather that temporal models suggest new kinds of *non-spatial geometry*. Here is what we take to the essence of any 'geometrical' approach. One works with 'flat' *many-sorted pictures* of various geometrical objects, without higher set-theoretic constructions favouring one sort (say, primitive points) over another (say, lines construed as point sets). We review some geometrical objects that are naturally suggested by our preceding discussions.

3.1 Reification of Tuples in Many-Sorted Modal Logic

Usually, the only temporal objects considered are points. To be sure, the literature has seen a development of a competing (or co-existing) paradigm of *periods* or *intervals* (cf. van Benthem 1983), motivated largely by ontological preferences, plus the needs of linguistic semantics. But the latter move is orthogonal to what we have in mind here, viz. the recent phenomenon of *reification* of 'constructed' temporal objects such as pairs or tuples. One striking motivation for this move has been the desire to lower complexity of existing modal or temporal logics. Another has been a return from many-dimensional systems to standard one-dimensional (though many-sorted) modal languages.

A prime example of this trend is *Arrow Logic*, whose full version replaces Relational Set Algebra by a two-sorted modal logic. Its intended models have two kinds of primitive object: the original *states*, plus *arrows* standing for ordered pairs (transitions) with their relevant internal structure. Notably, arrows naturally 'compose' among themselves via ternary triangles, and they have points for their beginnings and ends. 'Propositions' will now stand for properties of states, and 'relations' for properties of arrows. A good introduction to the resulting research program is Venema 1996. Arrow Logic is close to more general algebraic logic, which mimicks the complete first-order language. Thus, one can reify both pairs and triples when modalizing the full three-variable first-order fragment, or more general tuples, as in Venema's Cylindric Modal Logic, or Vakarelov's brand of Arrow Logic which manipulates sequences. Tuples of arbitrary lengths then become modal or temporal objects per se, and our task is to analyse their basic mathematical structure - without assuming that they must be reducible to full Cartesian products of the underlying state set. Our general point is this. These moves are not a mere auxiliary analysis for standard first-order logic (though they do help us understand what is essential and what more accidental about the latter). They are a full-fledged, less complex, independent alternative, treating tuples as ontological first-class citizens. In particular, though first-order logic is *undecidable*, its reified versions usually are not. Thus, many-dimensional modal logic over set-theoretically composite tuples becomes many-sorted one-dimensional modal logic, say over 'paths'. Just as in real geometry, we do not derive the properties of primitive objects like lines from set-theoretic constructions out of points. We rather analyse them directly, using our spatial intuitions. The same is true for arrow logics and more general tuple logics. Nevertheless, it is fair to say that, beyond a few case studies, we do not yet understand the general complexity effects of this strategy. Our guiding idea is 'trading complexity for new objects'. Exactly what can be achieved in this fashion?

Summary. In addition to temporal points, point pairs may be viewed as independent temporal objects ('transitions'), while general tuples become primitive 'temporal paths'. This move lowers the complexity of standard many-dimensional temporal logics, and it invites us to describe more general decidable temporal core geometries extending the pair-based paradigm of Arrow Logic.

3.2 Reification of Paths and Branches in Temporal Logic

Another prime candidate for reification in temporal logic arises in so-called 'branching time' (cf. Section 2.3). In addition to points in time, one now has histories or branches. Usually, temporal logicians treat the latter as certain sets of points (maximal chains), and then have to cope with the complexities of the resulting second-order logic. But there is a way around this, by taking branches or more general paths seriously as temporal objects in their own right. Again, of course, this is exactly what happens in ordinary geometry. One concrete example of this move is the reification of computation paths, made in order to avoid infinitary notions, found in van Benthem, van Eyck & Stebletsova 1995 on 'stuttering bisimulations' (modelling, amongst others, learning by unwilling adolescents). Another case is the first-order treatment of branches in Zanardo 1990, who decreases dependence on set-theoretic existence axioms, and works with 'absolute' notions only. (The latter author is not to be confused with the Zanardo of the present volume, who goes to the opposite extreme, and ties branching temporal logic to some highly esotherical set theory.) For the sake of concreteness, here is an example.

Branching models are tuples $\mathbf{M} = (S, C, <, ON, R, V)$, with S a set of temporal *states*, C the *possible histories* of the system described by the model, $<$ *temporal precedence* among states, *incidence* $ON(s, \sigma)$ says that state s lies on history σ, and R is a modal *alternative relation* which says at each history and state which alternative histories are accessible. The assignment V evaluates proposition letters at states. Now, we can interpret the formulas of a branching modal-temporal language as follows, at pairs of a history σ and some state s on it:

- $\mathbf{M}, \sigma, s \models p$ iff $s \in V(p)$
- Boolean connectives have the usual truth conditions
- $\mathbf{M}, \sigma, s \models F\phi$ iff $\mathbf{M}, \sigma, s' \models \phi$ *for some* $s' > s$
- $\mathbf{M}, \sigma, s \models \Box\phi$ iff $\mathbf{M}, \sigma', s \models \phi$ *for some* σ' *with* $R(\sigma, s, \sigma')$

In this setting, temporal geometry may be brought out through correspondence with natural laws of a branching temporal language. First, as in standard temporal logic, pure F, P principles express conditions on single histories. In particular, the usual Prior axioms will make these linear orders. But now, a temporal-modal 'mixing law' like $(\Box Fq \wedge \Box Fr) \rightarrow \Box F(q \wedge \Box F\Box Pr)$ will express a geometrical principle *confluence* which may be computed by a modified correspondence substitution algorithm. It says, for the web of histories as seen from any state s lying on a branch h, that:

$$\forall s_1 \forall s_2 \forall h_1 \forall h_2 \ (\ (ONsh_1 \wedge ONsh_2 \wedge ONs_1h_1 \wedge ONs_2h_2 \wedge s<s_1 \wedge s<s_2)$$
$$\rightarrow \exists s' \ (s_1<s' \wedge s_2<s' \wedge \exists h_3 \exists h_4 \ (ONs_1h_3 \wedge ONs'h_3 \wedge ONs_2h_4 \wedge ONs'h_4) \ ))$$

Interesting geometrical properties may also be sought independently from the modal-temporal formalism. An example is 'fusion closure' (Stirling 1989): "for any state occurring in two histories, its past in the one and its future in the other may be glued together so as to form a new history". As it happens, this is expressed after all, by the modal-temporal formula $(\Box(Hp \wedge Gq) \wedge \Box(Hr \wedge Gs)) \rightarrow \Box(Hp \wedge Gs)$. But one can also develop state-branch theories directly in first-order logic. In particular, what are genuinely universal laws of History, and what are negotiable existence assumptions about the availability of branches that reflect second-order axioms of choice?

One can also look at the matter more generally, as advocated in van Benthem 1996A. Second-order logic or set theory hide complexity in an amorphous background, whereas an axiomatic first-order theory of branches or paths forces us to explicitly identify the basic structural principles that we want to work with. This amounts to doing our conceptual homework in temporal modeling more properly, instead of appealing to some grab bag. A further beneficial side-effect is that we can determine the price in logical complexity of mathematical existence assumptions. Once again we see a game of trading complexity for extra 'temporal objects': this time, making higher-order logics, if not decidable, then at least recursively axiomatizable.

3.3 Modal Geometry

Our analogy with geometry invites a logical comparison. In Hilbert's "Grundlagen der Geometrie" (1899) points and lines lived on a par. In Tarski's 'What is Elementary Geometry?' (1953) points have become the only primitive objects, while line-talk was reduced to statements involving pairs of objects. This was an excellent idea for algebraic and model-theoretic reasons, but it was not faithful to geometrical thinking. There are recent modal-style axiomatizations of traditional geometrical spaces: cf. Balbiani, Fariñas, Tinchev & Vakarelov 1997. But these accept classical geometry in toto, and do not innovate. (Also, these authors reduce everything to a one-sorted modal universe, incurring a lot of spurious complexity.) Earlier work by Goldblatt on modal topological structures may be relevant, too. But modal logics of space are not what we are advocating here. In our view, temporal geometry would include at least this:

Primitive objects:	*points, lines, paths*
Primitive relations:	*incidence, parallellism, orthogonality*, and others

The result is more complex than standard geometry, because paths are not straight lines. For instance, we need an account of polygons and other closed paths as well. Modal languages here would access these relations in the standard way, via modalities such as <ON> (true at some point on the current line, or some line through the current point), and SUM$\phi\psi$ (there exists some division of the current line into an initial ϕ–part and a final ψ–part). The division of labour in such a temporal geometry will be as in the earlier-mentioned Arrow Logic program: "decidable logics, undecidable mathematics". But the first priority seems to be plain structural analysis of the first-order base theory of points and paths on a par: i.e. just 'geometry'.

Conclusion. We have advocated development of independent temporal geometries, as mathematical theories of interest for their own sake. Concrete samples of this style of thinking are Rodenburg 1987 (branches in intuitionistic Beth models as histories in 'cognitive evaluation time'), van Benthem, van Eyck & Stebletsova 1995 (path bisimulation in computation time), and Pratt & Lemon 1997 (model-theoretic geometry of polygons in space – which also cover walks in temporal models).

4 Temporal Structures in Natural Language

Prior and his students had definite linguistic interests, which are reflected in the name 'tense logic' which was often used for the whole field. Through the 70s, there were lively discussions about temporal representation of natural language. These brought to light several inadequacies of basic temporal logic, and the need for richer formalisms (cf. Aqvist & Guenthner 1977, or the radical first-order alternative Needham 1975). But around 1980, Hans Kamp (...) put an end to the dominance of temporal logic in natural language semantics, changing the research agenda to temporal 'discourse representation theory' (Kamp 1979) which then became a general paradigm for linguistic interpretation (Kamp & Reyle 1993, Kamp & van Eyck 1997). Its fundamental idea is that temporal expressions serve as instructions for creating successive discourse representations, which can – but need not – be 'embedded' into standard models at some separate later stage. When viewed this way, in particular, verb tenses do not behave like the logical operators P, F at all. (The latter are closer to the temporal auxiliaries in perfect and future statements.) They rather serve as deictic expressions creating a sequence of temporal reference points, moving the representation forward. There has been little interaction between linguistic temporal semantics and temporal logic since. (Two interesting exceptions are Ter Meulen & Seligman 1994, Dünges 1998.) Nevertheless, important issues remain concerning natural language in a broader context of reasoning and action. How does logical time meet up with linguistic representation time, and also, with real physical time? Some of the sophisticated phenomena that arise in descriptive linguistics seem relevant and inspirational to temporal logic in the larger sense. For an up-to-date survey of these matters with a computational and cognitive slant, we refer to Mark Steedman's chapter on 'Temporality' in the "Handbook of Logic and Language" (J. van Benthem & A. ter Meulen, eds., 1997).

4.1 Temporal Constructions in Natural Language

One conspicuous feature of temporal expressions in natural language is their variety, including tenses, auxiliaries, or temporal adverbs. They do not live in one syntactic category of 'sentence operators'. Likewise, there seems to be a much greater ontological variety than what one finds in sparse logical models. Temporal discourse refers to, amongst others, points in time, extended intervals, but also events, or 'cases'. Thus, standard temporal logic would need to be extended to match this greater richness – in particular, with suitable generalized quantifiers for temporal adverbs such as "always", "often", or "usually". (One noteworthy exception to this formal neglect is Galton 1985, whose study of verb aspects took temporal logic much closer to real-life linguistic phenomena, and introduced a lot of interesting new logical operators in the process.) What makes these things rather difficult is a further complication. Temporal structure in natural language is usually intertwined with other phenomena (context, defaults, etc.). There are many subtleties in even such a relatively simple sentence as "Mary never finishes her soup". For instance, "always" does not mean at all times, and indeed, the exact assertion involves temporality intertwined with habits and generic expressions. Another example is the recent ABN AMRO poster where Dutch soccer legend Johan Cruyff says "I have to admit that there are some people who are better than me in some things". Clearly, the quantification here is over *kinds of* people, things, and occasions, rather than over individuals. Temporal logic has not yet begun to absorb such simple, yet linguistically well-known points about natural language. Nevertheless, occasional more sophisticated encounters occur. For instance, the challenging dissertation Leonoor Oversteegen 1989 combines so-called 'A-series' (P, F-based) and 'B-series' (precedence-event-based) structures from temporal logic to account for the dynamics of temporal discourse. Another fruitful interface is current dynamic semantics for natural language, discussed in the next section.

Conclusion. The variety of temporal expressions and temporal objects in natural language is much greater than what is covered in the usual temporal logics. Although this gap has been well-known for many years, hardly any active bridge building occurs.

4.2 Dynamics of Evaluation

Dynamic semantics says that semantic evaluation is a logical process over some kind of information states, and linguistic expressions are programs for evaluation or update, or yet other cognitive processes over these states. In the case of time these processes involve such typical phenomena as manipulating temporal perspective and exploiting persistence of information through time. A toy example of this kind of system can be extracted from Dekker 1993. We can interpret Prior's base language in the style of dynamic predicate logic (Groenendijk & Stokhof 1991). The semantic format is this:

$\mathbf{M}, t, t' \models \phi$ iff there exists a succesful evaluation of ϕ starting in t and ending in t'

Fix a basic Priorean temporal model $\mathbf{M} = (T, <, V)$. Each formula ϕ will now denote a binary transition relation $[[\phi]]$ on T, constructed via the following induction:

- *atomic propositions* function as instantaneous tests:
  $[[p]] \quad = \quad \{ (t, t) \mid t \in V(p) \}$

- *conjunction* becomes sequential composition of two successive tasks:
  $[[\phi\wedge\psi]] \quad = \quad \{ (t, t') \mid \textit{for some } t'', (t, t'') \in [[\phi]] \textit{ and } (t'', t') \in [[\psi]] \}$

- *futurity* involves making a step to the right and then starting again from there:
  $[[F\phi]] \quad = \quad \{ (t, t') \mid \textit{for some } t'', t<t'' \textit{ and } (t'', t') \in [[\phi]] \}$

  and the explanation for the *past* operator P is analogous toward the left

- one reasonable form of *negation* is a test for 'strong failure':
  $[[\neg\phi]] \quad = \quad \{ (t, t) \mid \textit{for no } t', (t, t') \in [[\phi]] \}$

For further discussion of this system, cf. van Benthem 1995. In particular, 'procedural' differences emerge between formulas which used to be equivalent in the basic logic. Thus, $Fp\wedge q$ is now read as an instruction to move first to some future point where p is the case, and then test whether q holds there. The net effect is a transition to some future point where both p and q hold. This outcome is similar to that of the formerly non-equivalent instruction $F(p\wedge q)$ moving us to some future point where successive tests for p and q are succesful. But it is quite dissimilar to the procedural effect of the permuted formula $q\wedge Fp$. These sequential effects mirror the sequential nature of temporal narrative. The resulting system turns the P, F language into a simple two-dimensional modal logic, to which earlier techniques apply. A general move from unary modal logic to binary systems can create undecidability (think of Relational Algebra), but this does not happen here. Dynamic tense logic remains decidable, something which can be shown by guarded first-order analysis. But then, one can take advantage of the richer structure of dynamic models, and introduce *new dynamic operators*, which may increase complexity again. (One could also use some form of Arrow Logic to perform this 'dynamification'.) The point of our presentation here is just to show how standard temporal languages acquire new functions under a different style of interpretation. We have demonstrated this for Prior's P, F language. But one can also dynamify the richer S, U language, which has a computational ring already, of imperative "while ... do ..." instructions. Indeed, judging from current discussions in the linguistic literature, interesting dynamifications may even be performed in a number of different ways.

Conclusion. Dynamic semantics provides a new view of – and use for – existing temporal formalisms. The broad technical properties of this move may be understood as a semantic shift from shorter to longer tuples in a many-dimensional modal setting. Viewed somewhat differently, this dynamic move turns languages making descriptive temporal statements into programming languages for verifying temporal statements.

4.3 Dynamics of Representation

Discourse Representation Theory takes an alternative view of the locus of dynamics. Discourse creates successive temporal representations, which are annotated syntactic structures that can be interpreted (later on) in standard models. A similar system are the 'dynamic aspect trees' of ter Meulen 1995. Van Benthem & ter Meulen 1997 discuss the resulting architectural issues more generally. They distinguish two broad strategies, dubbed *Dynamics of Interpretation* – turning texts into representation structures – and *Dynamics of Evaluation* – locating the temporal dynamics in the successive state-changes of semantic evaluation. The former is highlighted in a *Tree Calculus*, whose objects are trees with nodes that can be annotated for propositional information, while nodes are temporal locations that can stand in key relationships like precedence or inclusion. Dynamic processes include: informational update (learning more about some node), temporal update ('move rules' add nodes or temporal relations between nodes), and what may be called perspectival update: shifting the 'active/current node' (the center of attention). Like Kamp-style discourse representations, the resulting annotated trees can be 'embedded' into temporal models, as if they were (abstract) 'pictures' of certain parts of the latter. Such an embedding associates them with concrete logical statements. A simple system will demonstrate this way of thinking.

Consider finite trees with nodes and arrows that can be labeled for various relations. Intuitively, the nodes stand for temporal points or small intervals at which events can take place or states can hold. We assume that each tree has a designated initial point and one currently active node. The simplest case is the tree with one single node, which is both initial and current. Update conditions for such trees can be given in the format

UPD(ϕ) (T), where ϕ is some temporal statement, and T the current tree.

Atomic updates UPD(q) (T) are immediate: one writes q on the current node of T . Conjunctions trigger compositions of updates, as in the above dynamic semantics. Updates with a future statement UPD(Fϕ) (T) first add a new successor node (with a labeled <-link) to the current one, shift the 'current' marker to that new node, and then, on the newly obtained tree T' , proceed with UPD(ϕ) (T') . The instruction for a past statement Pϕ is similar. These successive instructions create trees that form small pictures of facts over time. Similar rules can be written for operators creating subinterval-nodes for other temporal operators. *Negations* $\neg\phi$ are treated somewhat differently. Intuitively, they do not enrich the temporal picture, but exclude certain continuations of it. (Thus, a negation is more like a general *constraint* or rule.) The update adds a new 'negation-link' from the current node to a 'forbidden tree' obtained by updating a single initial node tree with ϕ . This completes the tree construction. A proper semantic relationship between trees and temporal models is given by *succesful embeddings*, being maps from nodes to temporal points in the right temporal relations, and satisfying the right propositional information.

Remark. The recursive clause for negations says a forbidden tree cannot be embedded succesfully with its initial point equal to the current node. This is logical language. Natural language has more restricted negations in temporal discourse. It uses atomic polarities for verbs, plus 'general rules' for temporal adverbs ("always", "often", etc.).

An easy induction shows the following connection between the two dynamic systems:

Proposition For all Priorean F, P-formulas ϕ, and all temporal models **M**, the following two assertions are equivalent:

(i) **M**, t, t' $\models \phi$ holds, as defined in the preceding section,

(ii) the map $\{(s_{in}, t), (s_{out}, t')\}$ is a succesful embedding of the tree UPD (ϕ)('single node') into **M**.

More general equivalences between the two dynamic strategies can be obtained as well, since there is no constraint on the choice of states. This open end also recurs in general discussions of semantic compositionality: cf. Janssen 1997. (Cf. also Kaufmann 1998 on stack-based models as an intermediate between representation structures and models, whose top-level stack handling incorporates the 'right frontier constraint' found in several linguistic tree calculi. Further examples are in Kempson & Meyer Viol 1997.) In the current literature, one sees more baroque notions of state, which encode a lot of syntactic information. Nevertheless, the tree format has its attractions because of its more concrete computational character. Thus, it also suggests new dynamic operations without an immediate semantic counterpart, such as more complex shifts in the center of attention in temporal discourse (cf. the computational model of Kameyama 1996).

One issue of particular logical interest is the locus of *reasoning* in this dynamic set up. One can reason with trees as ordinary assertions about models via their embedding conditions. This activity may be called (upper case) **I**nference in the classical sense. But their is also (lower case) local **i**nference in a more dynamic sense, which interleaves with the very tree construction that determines what assertion we are making. (For instance, one sometimes has to infer in order to see which statement is being made.) Examples of the latter are so-called *spreading rules*, which transfer information from one location in the tree to another. E.g., a past statement $P\phi$ ("Mary cried") will be transferred to all rightward nodes representing later time points, and a statement $[\Downarrow]\phi$ (say a 'state assertion' about John's being tired) to all lower nodes representing subintervals. But in addition, there are more sophisticated backtracking *consistency constraints*, which determine whether some statement will be added at the current node, or whether it must be attached to some later one.

> "<u>Mary opened the door. She smiled.</u>" is interpreted as two events taking place at the same time. But "<u>Mary opened the door. She closed it.</u>" is interpreted as a sequence of events, because of the incompatibility between the two actions.

In general, attachment for a new event is calculated according to some algorithm involving closeness in the tree and consistency, which may have to operate recursively. Thus, for a realistic linguistic account of temporal interpretation and inference, we need a two-level logic, with fast inferential mechanisms at tree level spreading local information, and supplying tests for consistency-based interpretation algorithms, while more sophisticated but slower temporal inferences are postponed to the discourse level. There are many further issues here. These include placement of conclusions in trees according to 'focus', or deviant *ternary structural rules* for local and global inference. (Cf. van Benthem & ter Meulen 1997 for details.)

Conclusion. Temporal discourse in natural language mixes semantic interpretation with inference at various levels, thereby challenging the standard division of labour in temporal logic, whose syntax and semantics are traditional. This observation leads to many new research questions. A traditional issue would be to isolate low-complexity Horn Clause fragments, or Monotonicity fragments, for existing full temporal logics. (Cf. van Benthem 1986.) Another would be to set up new two-level calculi of inference, and more innovative temporal-logical architectures generally. Curiously, over the past two decades, the study of temporal reasoning and representation has been a major source of semantic innovation in natural language semantics, but not in logic itself!

5 Appendix: Further Interfaces

We conclude with a sketch of developments in other fields with some significance for the above discussion. No complete coverage of the vast literature is intended! Cf. the Temporal Logic chapters in the Handbooks Gabbay et al., eds., 1990/4, 1991/5 for that.

5.1 Computer Science

Temporal logic versus dynamic logic. There is a recurrent debate on 'temporal logic versus dynamic logic' as a paradigm for program semantics (Lamport, Pratt, Pnueli). Our plea has been for peaceful co-existence of abstract process time and the physical time of system evolution. (Compare our technical questions in Section 2.3 about the branching temporal logic of tree unravelings for state machines.)

Applied temporal logic. Standard temporal logic would be wise to keep up with the Pnueli tradition! The latter has developed a diverging agenda, highlighting model checking rather than consistency checking, which leads to different outcomes as to complexity, emphasizing fixed models rather than whole model classes, etcetera. Merges do occur, e.g., in practical computational systems mixing theorem proving with model checking (SRI, Stanford CS; cf. Sipma, Uribe & Manna 1996). Other interesting merges bring together temporal logics with rewrite systems (Denker 1998).

Metrical operators. Realistic practical modelling of processes and events will soon involve quantitative information about passage of time, as in $<k>\phi$: " ϕ will hold k time units from now", or "over some period of duration k ". This requires an additional calculus of quantities of some kind, tagged onto the temporal logic much as dynamic logic merges propositional structure with program structure (Montanari & de Rijke 1997). More generally, current 'hybrid systems' merge logical techniques with standard physical ones for the purpose of describing system evolutions in real time.

Granularity. One temporal situation may be described at different levels of temporal detail (seconds, days, centuries). Montanari 1996 has an elaborate many-level view of temporal structure, with matching logics and complexity results. (Cf. Lamport's 1985 many-level theory of temporal 'views' in an interval setting.) With such different levels, there is a special importance to logical operators effecting *connections* between them, which serve as shifting modalities that 'zoom in' and 'out' (Blackburn & de Rijke 1997).

5.2 Artificial Intelligence

Connections between spatial and temporal reasoning. The editorial to the IJCAI satellite volume Anger, Guesgen, Rodriguez & van Benthem 1997 has the main current issues.

Context logics. Buvac & McCarthy 1998 (and other papers in the research program of 'logical AI') discuss the structure of contexts, and shifts between them. This raises very much the same issues as emerged for temporal logics in the lively debate of the 70s about modal-style formalisms versus straight first-order translation in linguistic semantics. (For a discussion of analogies, cf. van Benthem 1998B.)

Belief revision. Process time and real time may interact. Katsuno & Mendelzon 1992 are famous for their distinction between temporal update for a changing world versus mere epistemic revision. (Ryan & Schobbens 1997 is an up-to-date modal analysis.) Temporal epistemic logic has also been proposed as an alternative theory of cognitive dynamics (Engelfriet 1996, Engelfriet & Treur 1996).

5.3 Physics and Philosophy of Science

Temporal logic in physical time. Temporal and modal structures in relativistic space-time have been investigated in Rakic 1997, which endows Minkowski models with additional common-sense structure. The result are models for richer temporal logics with primitive relations of present, future, and causal connection. These may be connected with issues in linguistic and logical semantics, witness Rakic' list of open questions. This kind of research is also a demonstration of the use for renewed contacts between logical semantics and work on space-time in the *philosophy of science*.

5.4 Cognitive Psychology

Temporal reasoning and representation have been long-standing concerns in cognitive psychology (cf. Jackson & Michon 1985). Several developments discussed in this paper seem relevant to human cognition – but, even more than with linguistics, this broader interface for temporal logic has been regrettably unlively so far.

6 Conclusions

We recapitulate the general themes and claims of this paper:

(1) *the diversity of uses of time and their underlying temporal patterns needs to be taken seriously*

(2) *temporal logic can provide a methodological unity across these, through (at least) techniques from general modal logic*

(3) *in particular, we need logical architectures for systematic connections between different temporal modellings*

(4) *to keep the field on a healthy diet of challenges, we need to absorb richer logical structures from natural language and cognition.*

With our rather non-standard survey we hope to have shown at least two things. Temporal logic is alive and well, and: it may still have most of its future ahead of it!

7 References

G. d'Agostino, 1998, *Modal Logic and Set Theory: Translation, Bisimulation, and Interpolation*, forthcoming dissertation, Institute for Logic, Language and Computation, University of Amsterdam.

A. Aliseda-Lera & R. van Glabbeek, eds., 1998, *Computing Natural Language*, CSLI Publications, Stanford.

H. Andréka, J. van Benthem & I. Németi, 1998, 'Modal Logics and Bounded Fragments of Predicate Logic', *Journal of Philosophical Logic* 27:3, 1998, 217–274.

F. Anger, H. Guesgen, R. Rodriguez & J. van Benthem, 1997, 'Space, Time and Computation: Trends and Problems', editorial, special issue on Spatial and Temporal Reasoning, *Intern. Journal of Applied Intelligence* Vol. 6:1, 5–9.

L. Åqvist & F. Guenthner, 1977, 'Fundamentals of a Theory of Verb Aspect and Events within the Setting of an Improved Tense Logic, In F. Guenthner & Ch. Rohrer, eds., *Sudies in Formal Semantics*, North-Holland, Amsterdam.

J. W. de Bakker, W.-P. de Roever & G. Rozenberg, eds., 1989, *Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency*, Springer, Berlin.

Ph. Balbiani, L. Fariñas del Cerro, D. Vakarelov, 1997, 'Modal Logics for Incidence Geometries', *Journal of Logic and Computation* 7:1, 59–78.

J. Barwise, 1975, *Admissible Sets and Structures*, Springer Verlag, Berlin.

J. Barwise & J. van Benthem, 1996, 'Interpolation, Preservation, and Pebble Games', Report ML–96–12, Institute for Logic, Language and Computation, Amsterdam. To appear in the *Journal of Symbolic Logic*.

J. Barwise & L. Moss, 1996, *Vicious Circles. On the mathematics of non-wellfounded phenomena*, CSLI Publications, Stanford.

R. Bäuerle, C. Schwarze & A. von Stechow, eds., 1979,*Semantics from Different Points of View*, Springer, Berlin.

J. van Benthem, 1976, *Modal Correspondence Theory*, dissertation, Mathematisch Instituut, University of Amsterdam, 148 pp.

J. van Benthem, 1983, *The Logic of Time*, Reidel, Dordrecht (second edition 1991).

J. van Benthem, 1986, 'The Ubiquity of Logic in Natural Language', in W. Leinfellner and F. Wuketits, eds., 1986, *The Tasks of Contemporary Philosophy*, Schriftenreihe der Wittgenstein Gesellschaft, Verlag Hölder-Pichler-Tempsky, Wien, 177–186.

J. van Benthem, 1989, 'Time, Logic and Computation', in J. W. de Bakker et al., eds., 1–49.

J. van Benthem, 1991, *Language in Action. Categories, Lambdas and Dynamic Logic*, North-Holland, Amsterdam.

J. van Benthem, 1995, 'Temporal Logic', in D. Gabbay et al., eds., *Handbook of Logic in Artificial Intelligence and Logic Programming*, Volume 4, 241–350.

J. van Benthem, 1996A, 'Complexity of Contents versus Complexity of Wrappings', in M. Marx, et al., eds., 203–219.

J. van Benthem, 1996B, *Exploring Logical Dynamics*, Studies in Logic, Language and Information, CSLI Publications, Stanford.

J. van Benthem 1997A, 'Dynamic Bits and Pieces', Report LP-97-01, Institute for Logic, Language and Computation, University of Amsterdam.

J. van Benthem 1997B, 'The Range of Modal Logic', Report ML–97–05, Institute for Logic, Language and Computation, University of Amsterdam. To appear in the *Journal of Applied Non-Classical Logics*, memorial issue for George Gargov.

J. van Benthem, 1998A, 'Process Operations in Extended Dynamic Logic', abstract for tutorial, *Proceedings LICS 98*, Indianapolis, IEEE Publications.

J. van Benthem, 1998B, 'Shifting Contexts and Changing Assertions', in A. Aliseda-Llera & R. van Glabbeek & D. Westerståhl, eds., 51–65.

J. van Benthem & J. Bergstra, 1995, 'Logic of Transition Systems', *Journal of Logic, Language and Information* 3:4, 247–283.

J. van Benthem, J. van Eyck & V. Stebletsova, 1994, 'Modal Logic, Transition Systems and Processes', *Journal of Logic and Computation* 4:5, 811-855.

J. van Benthem & A. ter Meulen, eds., 1997A, *Handbook of Logic and Language*, Elsevier, Amsterdam.

J. van Benthem & A. ter Meulen, 1997B, 'Two Strategies for Dynamic Interpretation', manuscript, ILLC, Amsterdam. Short version appeared in P. Dekker, M. Stokhof & Y. Venema, eds., *Proceedings 11th Amsterdam Colloquium*, ILLC, Amsterdam.

P. Blackburn & M. de Rijke, 1997, 'Zooming in, Zooming out', *Journal of Logic, Language and Information* 6, 5-31.

E. Börger, E. Grädel & Y. Gurevich, 1996, *The Classical Decision Problem*, Springer, Berlin.

S. Buvac & J. McCarthy, 1998, 'Formalizing Context (Expanded Notes)', in A. Aliseda & R. van Glabbeek, eds., 13–50.

A. Chagrov & M. Zakharyaschev, 1997, *Modal Logic*, Oxford University Press, Oxford.

P. Dekker, 1993, *Transsentential meditations. Ups and Downs in Dynamic Semantics*, Dissertation, Institute for Logic, Language and Computation, University of Amsterdam.

G. Denker, 1998, 'From Rewrite Theories to Temporal Logic Theories. A Distributed Temporal Logic Extension of Rewriting Logic', CS Lab, SRI International, Menlo Park.

P. Dünges, 1998, 'Eventualities in Time', Report Nr. 104, Computerlinguistik an der Universität des Saarlandes, Saarbrücken.

J. Engelfriet, 1996, 'Minimal Temporal Epistemic Logic', *Notre Dame Journal of Formal Logic* 37(2):233-259, 1996. Special Issue on Combining Logics. Edited by M. de Rijke and P. Blackburn.

J. Engelfriet & J. Treur, 1996, 'Specification of Nonmonotonic Reasoning', in D.M. Gabbay and H.J. Ohlbach (eds.), *Practical Reasoning*, Proceedings FAPR'96, Lecture Notes in Artificial Intelligence 1085, Springer, Berlin, 111–125.

D. Gabbay, 1981, 'Expressive Functional Completeness in Tense Logic', in U. Mönnich, ed., 91–117.

D. Gabbay, T. Maibaum & S. Abramsky, eds., 1990–1994, *Handbook of Logic and Computer Science*, Oxford University Press, Oxford.

D. Gabbay, C. Hoggar & J. Robinson, eds., *Handbook of Logic in Artificial Intelligence and Logic Programming*, Oxford University Press, Oxford, 1991–1995.

A. Galton, 1984,*The Logic of Aspect*, Clarendon Press, Oxford.

R. Goldblatt, 1980, 'Diodorean Modality in Minkowski Space-Time', *Studia Logica* 39, 219–236.

R. Goldblatt, 1987, *Logics of Time and Computation*, CSLI Lecture Notes, Stanford University, Chicago University Press.

V. Goranko, 1995, 'Hierarchies of Modal and Temporal Logics with Reference Pointers', department of mathematics, Rand Afrikaans University. To appear in *Journal of Logic, Language and Information* .

J. Groenendijk & M. Stokhof, 1991, 'Dynamic Predicate Logic', *Linguistics and Philosophy* 14, 39-100.

D. Harel, D. Kozen & J. Tiuryn, 1996, *Dynamic Logic*, book manuscript, Departments of Computer Science, Cornell University and Technion, Haifa.

D. Hilbert, 1899, *Grundlagen der Geometrie*, Teubner, Stuttgart.

J. Hintikka, ed., 1969, *The Philosophy of Mathematics*, Oxford University Press, Oxford.

N. Immerman & D. Kozen, 1987, 'Definability with Bounded Number of Bound Variables', *Proceedings IEEE 1987*, 236–244.

J. Jackson & J. Michon, eds., 1985, *Time, Mind and Behaviour*, Springer, Heidelberg.

Th. Janssen, 1997, 'Compositionality', in J. van Benthem & A. ter Meulen, eds., 417–473.

M. Kameyama, 1996, 'Indefeasible Semantics and Defeasible Pragmatics', in M. Kanazawa et al. , eds., 111–138.

H. Kamp, 1966,*Tense Logic and the Theory of Linear Order*, Dissertation, Department of Philosophy, University of California, Los Angeles.

H. Kamp, 1979, 'Instants, Events and Temporal Discourse', In R. Bäuerle et al., eds., 376–417.

H. Kamp & J. van Eyck, 1997, 'Representing Discourse in Context', in J. van Benthem & A. ter Meulen, eds., 179–237.

H. Kamp & U. Reyle, 1993, *From Discourse to Logic*, Kluwer, Dordrecht.

M. Kanazawa, Ch. Piñon & H. de Swart, eds., 1996, *Quantifiers, Deduction, and Context*, CSLI Publications, Center for the Study of Language and Information, Stanford.

H. Katsuno & A. Mendelzon, 1992, 'On the Difference between Updating a Knowledge Base and Revising it', in P. Gärdenfors, ed., *Belief Revision*, Cambridge University Press, 183–203.

S. Kaufmann, 1997, 'What's at Stack in Discourse', in p. Dekker et al., eds., *Proceedings of the 11th Amsterdam Colloquium*, ILLC, Amsterdam, 187-192.

R. Kempson & W. Meyer Viol, 1997, 'Parsing as Tree Construction', department of linguistics, School of African Studies, London.

N. Kurtonina & M. de Rijke, 1997, 'Bisimulations for Temporal Logic', *Journal of Logic, Language and Information* 6, 403-425.

L. Lamport, 1985, 'Interprocessor Communication. Final Report', SRI International & DEC Corporation, Menlo Park.

M. Marx & Y. Venema, 1996, *Multi-Dimensional Modal Logic*, Kluwer, Dordrecht.

M. Marx, M. Masuch & L. Pólos, eds., *Arrow Logic and Multimodal Logic,* Studies in Logic, Language and Information, CSLI Publications, Stanford.

A. ter Meulen, 1995, *Representing Time in Natural Language*, MIT Press, Cambridge (Mass.).

A. ter Meulen & J. Seligman, 1994, 'Dynamic Aspect Trees', in L. Polós et al., eds., *Applied Logic: How, what and why. Logical approaches to natural language.* Kluwer , Dordrecht, 287-320.

U. Mönnich, ed., 1981, *Aspects of Philosophical Logic*, Reidel, Dordrecht.

A. Montanari, 1996, *Metric and Layered Temporal Logic for Time Granularity,* Dissertation, Institute for Logic, Language and Computation, University of Amsterdam.

A. Montanari & M. de Rijke, 1997, 'Two-Sorted Metric Temporal Logic', *Theoretical Computer Science* 183, 187-214.

J. Needham, 1975, *Temporal Perspective*, Philosophical Studies, Department of Philosophy, University of Uppsala.

R. Nelken and N. Francez, 1996, 'Automatic Translation of Natural-Language System Specifications into Temporal Logic', in R. Alur and Th. Henzinger, eds., *Proceedings of the 8th International Conference on Computer Aided Verification,* Lecture Notes in Computer Science 1102, Springer Verlag, Berlin, 360–371.

H. de Nivelle, 1998, 'Resolution Decides the Guarded Fragment', Report CT-98-01, Institute for Logic, Language and Computation, University of Amsterdam. (New version: 'A Resolution Decision Procedure for the Guarded Fragment', in *Proceedings CADE-15*, pp. 191-204, 1998.)

L. Oversteegen, 1989,*Tracking Time*, Dissertation, Nederlands Instituut, Rijksuniversiteit Utrecht.

A. Pnueli, 1989, 'The Temporal Semantics of Concurrent Programs', *Theoretical Computer Science* 13, 45-60.

I. Pratt & O. Lemon, 1997, 'Ontologies for Plane, Polygonal Mereotopology', Department of Computer Science, University of Manchester, Report UMCS-97-1-1.

I. Pratt & D. Schoop, 1997, 'A Complete Axiom System for Polygonal Mereotopology of the Real Plane', Department of Computer Science, University of Manchester, Report UMCS-97-2-2.

N. Rakic, 1997, *Common Sense Time and Special Relativity*, Dissertation, Institute for Logic, Language and Computation, University of Amsterdam.

M. de Rijke, 1993, *Extending Modal Logic*, Dissertation, Institute for Logic, Language and Computation, University of Amsterdam.

P. Rodenburg, 1987, *Intuitionistic Correspondence Theory*, Dissertation, Mathematical Institute, University of Amsterdam.

M. Ryan & P-Y Schobbens, 1997, 'Counterfactuals and Updates as Inverse Modalities, *Journal of Logic, Language and Information* 6:2, 123–146.

V. Shehtman, 1983, 'Modal Logics of Domains on the Real Plane', *Studia Logica* 42, 63–80.

H. Sipma, T. Uribe & Z. Manna, 1996, 'Deductive Model Checking', *Proceedings 8th International Conference on Computer-Aided Verification*, LNCS 1102, Springer, Berlin, 209-219.

A. Sistla & E. Clarke, 1985, 'The Complexity of Propositional Linear Temporal Logics', *Journal of the ACM* 32:3, 733-749.

E. Spaan, 1993, *Complexity of Modal Logics*, dissertation, Institute for Logic, Language and Computation, University of Amsterdam.

C. Stirling, 1989, 'Expressibility and Definability in Branching and Linear Time Temporal Logics', Department of Computer Science, University of Edinburgh.

C. Stirling, 1992, 'Modal and Temporal Logics', in S. Abramsky, D. Gabbay & T. Maibaum, eds., 477-563.

A. Tarski, 1969, 'What is Elementary Geometry?' , in J. Hintikka, ed., 164–175.

D. Vakarelov, 1992, 'A Modal Theory of Arrows I', Report ML-92-04, Institute for Logic, Language and Computation, University of Amsterdam.

Y. Venema, 1995, 'Cylindric Modal Logic', *Journal of Symbolic Logic* 60, 591-623.

Y. Venema, 1996, 'A Crash Course in Arrow Logic', in M. Marx et al., eds., 3-34.

F. Wolter & M. Zakharyaschev, 1996, 'Advanced Modal Logic', School of Information Science, JAIST, Kanazawa. To appear in D. Gabbay, ed., *Handbook of Philosophical Logic*, second enlarged edition, Kluwer, Dordrecht.

A. Zanardo, 1990, 'Axiomatization of 'Peircean' Branching-Time Logic', *Studia Logica* 49, 183-195.