



Mémoire pour l'obtention de l'Habilitation à Diriger des Recherches
Spécialité Informatique

Representing Content Semantics, Ontology, and their Interplay

Laure Vieu

Soutenue le 10 Décembre 2009 devant la commission d'examen :

Nicholas Asher	DR, IRIT-CNRS	Directeur de Recherche
Claudette Cayrol	Professeur, Université Paul Sabatier	Examineur
Francis Corblin	Professeur, Université Paris 4 Sorbonne	Examineur
Jérôme Euzenat	DR, INRIA Grenoble Rhône-Alpes	Rapporteur
Antony Galton	Reader, University of Exeter	Rapporteur
Nicola Guarino	Senior Researcher, ISTC-CNR	Examineur
Alex Lascarides	Reader, University of Edinburgh	Rapporteur

Institut de Recherche en Informatique de Toulouse (UMR5505)
Université Paul Sabatier (Toulouse III) – 118, route de Narbonne – F-31062 Toulouse cedex 09

Contents

Foreword	3
1 Representing content in logical frameworks	7
1.1 Linguistic content: from words to discourse	8
1.1.1 Lexical semantics	9
1.1.2 Discourse semantics	12
1.2 Knowledge content	13
1.2.1 Temporal, spatial, and qualitative reasoning	14
1.2.2 Ontologies	15
1.3 Methodological issues	17
2 Ontology-driven lexical semantics: The semantics of spatial expressions	21
2.1 Relational and functional approach	22
2.2 Functional features and pragmatic effects	24
2.3 Relations and entities	26
2.4 Empirical investigations	28
3 A compositional semantics interlude: Locating adverbials	31
3.1 Scope	31
3.2 Syntactic position in the sentence	33
3.3 Relational semantics and referents	35
4 Discourse semantics	37
4.1 Fundamentals of SDRT	38
4.2 Subordinating and coordinating relations	42
4.3 The rhetorical nature of discourse relations	43
4.4 Discourse topics and forward-looking segments	47
4.5 Study of some Discourse Relations	49
4.5.1 Narration	49
4.5.2 Background	50
4.5.3 Entity-Elaboration	52
4.6 The spatio-temporal structure of discourse	53
4.6.1 Discourse effects of locating adverbials in IP-adjunct position . . .	54
4.6.2 Spatiotemporal interpretation of locating adverbials	58
4.7 Towards the next SDRT: a more empirical approach	59
5 Ontology of concrete entities	63
5.1 Concrete entities in a foundational ontology	64
5.1.1 Properties and individuals	64
5.1.2 Space and time	65
5.1.3 Categories and number	69

5.1.4	Towards a comprehensive foundational ontology	70
5.2	Region-based representations of space	71
5.2.1	Mereotopology	72
5.2.2	Beyond mereotopology: geometry and time	76
5.3	Parthood relations	79
6	Ontology of social reality	83
6.1	Roles and concepts	84
6.2	Artefacts	86
6.3	Action	88
7	Perspectives: The mental and social content of discourse	91
7.1	Enriching SDRT with an explicit account of communicative actions	91
7.1.1	Empirical evidence	92
7.1.2	Theoretical account	95
7.2	Contribution of discourse semantics to knowledge representation	97
	References	101

Merci

à Nicholas Asher, directeur de recherche officiel, et à Nicola Guarino, co-directeur officieux, pour m'avoir guidée en tant d'occasions et dans tant de domaines, y compris dans ce travail de synthèse ;

aux rapporteurs Jérôme Euzenat, Antony Galton et Alex Lascarides pour m'avoir fait l'honneur de relire ce manuscrit avec attention ;

aux examinateurs Claudette Cayrol et Francis Corblin pour avoir accepté d'évaluer mon travail ;

à mes amis de la première heure et compagnons de tout ce parcours, Michel Aurnague et Myriam Bras, pour perpétuer ce miracle qu'est la transformation du travail en bonheur de la découverte et de la création collectives ;

à tous mes autres amis, collaborateurs et collègues qui ont contribué à enrichir ma recherche et mes points de vue ;

à mes parents qui sont à l'origine de mon chemin, bien plus qu'ils n'osent le penser.

Et pour tant tant de choses, de la philosophie aux manicaretti, merci à Claudio.
Les montagnes se rapprochent.

Foreword

I discovered research with a background in mathematics and computer science, but with a curiosity for other fields, especially in humanities, that made me chose a topic related to natural language proposed by Mario Borillo at IRIT, back in 1987. This is when I discovered linguistics and semantics, as well as the whole of cognitive sciences. Right from the start I have been in direct contact with the linguists of ERSS, the laboratory directed at that time by Andrée Borillo, and with which I got involved in several research projects. Mario communicated me the sense of interdisciplinary studies, insisting though on rigour and precision brought by formal tools, resonating with my own inclinations.

I started the first period of my research work with a “DEA” (master) and a PhD on the semantics of spatial expressions, more particularly the lexical semantics of spatial prepositions. My propensity to focus on fundamental issues in the understanding of the observed phenomena, and my predilection for formal theories to analyse such issues, oriented my research at the interface between linguistic and logic-based knowledge representation rather than towards the development of new software in computational linguistics. The search for an adequate representational framework for the semantics of spatial expressions naturally brought me to the emerging field of qualitative spatial reasoning and its subfield called *mereotopology*.

In this first period, I met by best colleagues and friends Myriam Bras and Michel Aurnague. Together we learned the game of collaborative research, and together we kept discovering new linguistic domains. A large part of my work has been done with either one, or both. In particular, the foundations of my approach to spatial semantics and spatial representation have been developed in close collaboration with Michel Aurnague. Even so much later, each of us three specializing in a different area, it still is a great pleasure to tackle new issues and develop new solutions together. Thanks to Andrée and Mario Borillo, we were put in contact with leading scientists of many fields and disciplines. The “Time, Space and Movement” workshop series we organized together from 1989 to 1995 were extremely stimulating events which gave me the opportunity to meet many of my future collaborators and even colleagues. This is so that following Myriam Bras, I met Nicholas Asher and discovered discourse phenomena beyond the mere intersentential anaphora. My understanding of formal semantics dramatically increased then. With Nicholas Asher, I also enhanced my knowledge in philosophy and logic, as he helped me develop my first representation theorem in mereotopology. Around 1993, I also met Nicola Guarino who introduced me to the realm of formal ontology in metaphysics and its offsprings as foundational ontologies for knowledge representation. I had already come into contact with ontological issues, but mostly through the writing of formal semanticists and computer scientists. In fact, I worked for long on mereotopologies without even knowing about the fundamental book of Peter Simons [Sim87]. It is with Nicola Guarino and his group that I refined my approach to knowledge representation and could start contributing to formal ontology myself beyond mereotopologies.

After my PhD, my interest in botany- and agronomy-related fields drove me to jump into another world. I joined the artificial intelligence laboratory of the National Insti-

tute of Agronomical Research (INRA), where I attempted to apply qualitative spatial reasoning methods and mereotopology-based representations to the modelling of erosion processes during a short period (1991-1993). I will not present further this work here, since it didn't really have an impact on my later work.¹

I joined CNRS in 1993, and, although I kept exploiting and refining my initial work, a new period for me began. Together with Michel Aurnague, we set up several interdisciplinary projects with psycholinguists, in particular to test hypotheses at the semantics-pragmatics interface in lexical semantics. In this period too, with Myriam Bras and Anne Le Draoulec, I turned to yet another empirical method, as we began to extract examples from corpora, thus gaining so much more insights on the phenomena under scrutiny. Next, I started working on discourse with Myriam Bras, Nicholas Asher and Michel Aurnague, at first just for exploring the contribution of the lexical semantics of space to discourse structure, switching from spatial prepositions to motion verbs and especially locating adverbials. I quickly got caught by the diversity and complexity of the discourse dimension to which a large part of my work is still devoted today. Finally, the strengthening of the ontological bases of my work lead to a further systematization of the mereotopology domain together with Claudio Masolo.

During this period I had my first experiences in PhD supervision with Jean-Christophe Aurisset (1993, pragmatic factors in spatial lexical semantics [Aur93]) and Philippe Muller (1995, continuity in region-based space-time [Mul95]), both directed by Mario Borillo. I also partially contributed to supervising the PhDs of Claudio Masolo (2000, methodologies for comparing axiomatic theories [Mas00]), Maureen Donnelly (2001, mereogeometries [Don01]), and more recently, Laurent Prévot (2004, acknowledgement in route description dialogues [Pré04]).

The last period of my work started after 2000 with two important changes. On the one hand, I began addressing foundational issues in the discourse theory developed by Nicholas Asher and Alex Lascarides, SDRT, thus contributing to the present progress in the theory itself. This work has been in part triggered by ten years of refinement and improvement of the early account of locating adverbials in discourse. On the other hand, I decided to take a step beyond the spatio-temporal and material domain in formal ontology, and address social reality. For this, I obtained a “mise à disposition” to Nicola Guarino’s Laboratory of Applied Ontology of the ISTC-CNR in Trento, Italy, where I am since 2003. Beyond my personal work and collaborations, my two labs, IRIT and ISTC, had much to share. I thus created in 2004 the Interdisciplinary Laboratory on Interacting Knowledge Systems (ILIKS, <http://www.loa-cnr.it/iliks>), a “virtual lab”, to join them as well as departments of the University of Trento. The topic of interaction was chosen as a paradigm unifying many hot topics in AI, several of which I have contributed to: obviously discourse and dialogue, but also roles within organisations, social concepts, and action. It is on this last topic, action, that I directed my last two PhDs with the University of Trento, that of Robert Trypuz [Try07] and that of Nicolas Troquard [Tro07],² both obtained in 2007.

¹The relevant papers are: [MCV92, BMCVW92, Vie93, VMC94].

²Nicolas Troquard’s PhD has been done in co-tutorship with Andreas Herzig for Université Paul

This document aims to present the points of view I have adopted on the different topics addressed in my work and to situate them within the literature. It also aims at clarifying the relationships between these different topics and points of view, explaining the development of my scientific journey since my PhD, and hopefully showing the overall coherence of my work. The purpose is therefore not so much to present the details of the technical issues addressed and the results obtained. The reader is invited to refer to the publications for these.

I will first situate my work at the interface between formal semantics in linguistics and formal ontology in knowledge representation (Chapter 1). I will then present my approach and contribution to lexical semantics (Chapter 2), and discourse semantics (Chapter 4), with an incursion in the compositional semantics of locating adverbials in between (Chapter 3). My work in formal ontology is split into the ontology of concrete entities (Chapter 5) with in particular my contribution to the ontology of space, and the ontology of social reality (Chapter 6). I will conclude by presenting the programme for my future research (Chapter 7), in which I will address the mutual enrichment between the ontology of social reality and discourse semantics.

Sabatier, but I directed Nicolas on my own in behalf of the University of Trento.

1 Representing content in logical frameworks

The work reported in this document is situated within two areas in two disciplines, namely, formal semantics in linguistics and formal ontology for knowledge representation in computer science. These two areas are closely related: ontology is about *what there is* while semantics is about *referring to what there is*. In fact, what characterizes the work I have been carrying out for over 15 years in both disciplines is the interest towards what I will call “*content*”. Be it content of linguistic messages or content of knowledge, I have been seeking to represent *what* objects are described using *what* properties and relations, and *what* facts about them are considered true, obviously at a suitable level of generality and disregarding contingent facts.

For reasons of rigour and clarity, and for leaving open the possibilities of reasoning with such representations and exploiting them in computational applications and yet staying neutral in this respect (to be unduly short), I have chosen to represent content in formal frameworks, specifically logical frameworks.

But why using the term “content”, and not more simply the widespread terms of “meaning” or “semantics”? It turns out that what is usually called meaning is a different matter. Formal linguistic studies of meaning, or formal semantics, are mainly concerned with how meaning emerges from linguistic messages, that is, how different sub-expression and message *forms* recursively combine to convey different meanings, in truth-conditional terms. In other words, the study of meaning and semantics is largely dedicated to the study of the relationship between the form of messages and truth, the so-called syntax-semantics interface. Similarly, most of knowledge representation is dedicated to the study of different sorts or forms of knowledge, like necessary, uncertain or typical knowledge, and the formalisms adequate to represent and reason over these forms.

By using the expression “representing content”, I want to focus on what actually *fills* such forms —individual terms, predicates or even propositions— and denotes objects in a world and relations between them. This situates this work as complementing, and obviously depending on, core studies in formal semantics and knowledge representation.

We will see examples illustrating why I distinguish form and content in semantics in the next two parts of this chapter, first in linguistics and then in knowledge representation. But to already shed some light on this distinction it is worth to recall Husserl’s characterization of formal logic as the study of the (non-contingent) “interconnections of truths” and formal ontology as the study of the (domain-independent) “interconnections of things”, formal logic and formal ontology being distinct but mutually inseparable [Hus70]. In contemporary first-order logical terms, it corresponds to the distinction between the semantics of logical vocabulary (connectives, operators, quantifiers) and that of (the universally relevant part of) the non-logical vocabulary (predicates, constants). In other words, it is the distinction between a logic that characterizes the meaning of the logical vocabulary through either model-theoretic semantics or a proof theory, and a theory within a given logic that characterizes the meaning of (part of) the non-logical

vocabulary through axioms.³ In my view, much emphasis has been given to the first type of concern and relatively less to the second, both in knowledge representation and in formal semantics. And it is in both these fields that the work presented here tries to address this second concern, to which I globally refer under the expression “representing content”. So doing, I slightly extend the scope of Husserl’s formal ontology, that focused on those general notions that apply across all domains, such as *parthood* or *dependence*, to also cover some “core” domains, i.e., specific domains that are of general relevance, like e.g., time, space or material entities. Studying core domains is of particular importance to help characterizing the part of the non-logical vocabulary denoting the main categories of referents constituting the domain of quantification, i.e. focusing on *what* is the object of our messages and knowledge, on top of what are the relations that structure this domain.

1.1 Linguistic content: from words to discourse

The focus on the content of linguistic messages obviously implies selecting semantics as primary interest within the classical series ‘phonology, morphology, syntax, semantics, pragmatics’ defining the field of linguistics.

Yet, as just evoked, formal semantics, the branch of formal linguistics dealing with semantics initiated by Montague [Mon74], is essentially concerned with the syntax-semantics interface and with those linguistic ingredients whose contribution to meaning can be characterized in terms of operations on truth like connectives, operators and quantifiers. Formal semantics has analyzed major semantic phenomena so as to uncover what are the correct logical forms modelling the meaning of sentences and to what extent these forms can be constructed in a compositional way following the syntactic structure.⁴ So, for instance, formal semantics is concerned with how determiners involve quantification and with what scope, to what extent negation and conjunctions match logical connectives, how tense denote temporal modalities, and which entailment relation corresponds to presupposition⁵. But in the logical forms thus built, only the logical structure or skeleton matters and the meaning of the non-logical vocabulary is irrelevant. Sentences like *a cat eats a rat* and *a mat facilitates a nap*, having the same syntactic structure, will be analyzed into the formula $\exists x \exists y (P(x) \wedge Q(y) \wedge R(x, y))$ ⁶

³Depending on the terminology, the term “axiom” may be reserved for characterizing the logical vocabulary in some proof theory, and the term “meaning postulate” for characterizing the non-logical vocabulary [Car56]. I adopt here a wide use of the terms “axiom” and “axiomatic theory”.

⁴In Montague’s view, the construction of the logical formulas is an unnecessary step, and serves the purpose only to reuse the standard model-theoretic interpretation of classical logic. Discourse semantics, especially Discourse Representation Theory [KR93], as we will see later, has shown that an intermediate representation between language and models actually is necessary. And as advocated in [BB05], the logical representation of meaning is a must even at the sentence level if we are to do *computational* semantics and perform inferences from texts.

⁵In agreement with the contemporary view [PP02], I assume a large view on formal semantics, covering its tight relationship with formal pragmatics and therefore linguistic phenomena like presupposition.

⁶Or $\exists x \exists y \exists e (P(x) \wedge Q(y) \wedge R(e, x, y))$ in a Davidsonian approach [Dav67]. I disregard here tense and the possible generic interpretation that would involve other quantifiers.

where P , Q and R are either cat' , rat' and eat' , or mat' , nap' and $facilitate'$. The burden of distinguishing between the two sentences is therefore left to the interpretation function for appropriately mapping the predicates P , Q , and R to different sets. Formal semantics constrains this function on constants and predicates only as far as their logical type is concerned.⁷ Yet general enough differences of a semantic nature appear that go well beyond such logical types. Taking such differences into account could explain, for instance, why *a rat eats a cat* and *a cat eats a mat* are acceptable though unusual, while *a mat eats a rat* is not, and why the latter still makes more sense than *a nap facilitates a mat* which is outright impossible even in some fantasy world.

This schematic view on formal semantics is clearly an over-simplification. The programme initiated by Montague in [Mon73] did include a series of meaning postulates to adequately constrain the interpretation function,⁸ as [Dow79] well pointed out. However, this part of the programme has received much less attention in the discipline that developed under the name “formal semantics” than it deserved. Let us see anyway what has been done in this direction and better situate my own work along the way.

1.1.1 Lexical semantics

Filling in the logical forms, that is, paying attention to predicates and examining what constraints rule their meaning, implies first of all focusing on “content words” (as opposed to “function words”) that is, those words that make up the lexicon.

Lexical semantics as such has until quite recently received rather little attention in formal linguistics. Lexical semantics is largely developed in descriptive linguistics (e.g., [Cru86]), cognitive linguistics (e.g., [Tal83], [Jac83]⁹) and computational linguistics (e.g., [Pus91]), all branches of linguistics in which logic plays a small if no role at all. However, on the one hand, formal semantics requires lexical semantics to fill in the logical forms in non-arbitrary ways. More importantly still, lexical semantics is also needed to fully analyze some phenomena that traditionally fall within the range of formal semantics. A few classical examples: the aspectual classes of verbs are involved in analyzing the semantics of tense [Dow79]; genitive constructions implicitly refer to a range of relations which requires careful lexical and ontological studies to uncover [Bar95]; and presuppositions are often triggered by lexemes like the verb *to stop* or the noun *bachelor*. This simply confirms that form and content interact to produce the meaning of a sentence. On the other hand, I hold that the logical approach, with the systematic use of inferences to uncover meaning regularities, is a powerful tool to do in-depth lexical semantics. Formal semantics and lexical semantics therefore shouldn't ignore each other.

⁷This logical type is often expressed in recursive functional types constructed out of two basic types, e for “entity” (the type of a constant, e.g., a proper name) and t for “truth-value” (the type of a sentence). This way, logical types characterize not only in which set the denotation of a predicate is included, but how terms compose to yield complex expressions.

⁸Some postulates were intended to grasp some ontological differences between common nouns like *horse* and *temperature*. Others aimed at expressing the relationships between predicates of a same domain like *seek'* and *find'*.

⁹Jackendoff doesn't posit his own work within cognitive semantics but within “conceptual semantics”. From the viewpoint of formal semantics, the difference is slight.

Lexical semantics and formal semantics. In fact, these two subdisciplines have not completely ignored each other, and, as one can expect, the quick picture just sketched disregarded notable exceptions, i.e., formal semantics and lexical semantics do overlap. In particular, the work of Partee & Borschev on the semantics of genitives and relational nouns explicitly aims at integrating lexical and formal semantics [PB98]. However, this remained a quite isolated initiative; the most part of the overlap I’m referring to is of a less intended kind.

On the one hand, important work in lexical semantics, on selectional restrictions, semantic roles, and admissible syntactic alternations, is concerned with regularities in how lexemes denoting a predicator and its arguments may or may not compose into a sentence (see, e.g., [Lev93]). The main focus in this area is on exploring to which extent the meaning of a verb determines its syntactic behaviour, so the impact on formal semantics is indirect.

On the other hand, significant studies in formal semantics are dedicated to semantic phenomena that involve lexical distinctions. Works on determiners and number, notably those on the mass/count distinction and those on plurals, collections and generics, showed the necessity to distinguish several classes of predicates denoted by nouns and verbs (see e.g., [Pel79, Lin83, Car77]). In addition, they showed the need to analyse the nature of the domain of reference through structuring relations like parthood and constitution thus clearly entering the field of formal ontology. Similarly, the study of *aktionsarten*, the aspectual classes of verbs mentioned above, has been much developed by formal semanticists (e.g., [Dow79]). Yet, it focuses on revealing which combinations of verbs, tense and adverbial modification are possible, correlating this linguistic behaviour with the way eventualities¹⁰ referred to unfold in time, something characterized by the admissibility or not of some kinds of inferences.

It is this latter kind of semantic study, exploring a given content domain by analyzing both the reference to that domain and the inferential behaviour of some lexical entries, that most inspired my own work in lexical semantics. I indeed approached lexical semantics exploiting two fundamentals of formal semantics: its referential character¹¹ and its logical account of linguistic interpretation. I have also tried to stick to compositionality

¹⁰Following [Bac86b] and many semantic works since then, I use the term “eventuality” to refer to temporal entities that are perfective or imperfective occurrences of all sorts, e.g., (proper) events, processes and states. In formal ontology, “event” often takes on a generic sense synonymous with “eventuality”. In this work, both terms will appear, making clear, I hope, what is meant exactly in all uses of the ambiguous “event”.

¹¹The principle by which linguistic expressions denote individuals, relations and propositions in some world external to language itself —something not taken for granted in some areas of lexical semantics especially in cognitive semantics and in the French school of terminology studies. Structuralism-inspired linguists maintain that the meaning of a word lies in its relationships with the rest of the lexicon. Cognitive linguists adopt a specific view on the referential character of language, in which the worlds referred to are mental representations of reality, part of which cannot be accounted for by standard mathematical structures used in model theory. However, to account for the very possibility of linguistic communication and other collective actions, some common representation of the world needs to be assumed, a common world that cannot be confined into the subjectivity of anyone. For this purpose, the mature and precise model-theoretic approach goes a long way, especially if we select cognitively adequate structures.

as far as possible when I studied phrases composed of several words.

The picture of the relationships between formal and lexical semantics would not be complete without mentioning the relatively recent work around the “Generative Lexicon” [Pus91, Pus95] which aims at describing the lexicon so as to account for regularities such as polysemy which is involved in coercion, a phenomenon affecting compositionality [Pus98]. The semantic components of a lexeme in the generative lexicon are represented through a “qualia structure”, a set of four features inspired by Aristotle. Those features appear to have an ontological flavour but are chosen a priori, not on the basis of ontological studies, and even though they are useful in many cases, they are not adequate for all categories (e.g., the “telic” role makes sense for artefacts, actions, and little more) and obviously they are quite restrictive so many semantic aspects are missing. Somewhat departing from this proposal, Asher has proposed to model polysemy and coercion phenomena using basic ontological categories (or types) and a “dot” operator to build types of complex entities (i.e., combining entities of two different types): in [Ash07, Ash10], he develops a logic of types and “dot types” to account for polysemy in composition. Although not aiming at a global analysis of the lexicon, my more focused work contributed to this effort by showing on detailed studies of some lexical domains what ontological categories and relations are adequate.

The need for ontology. I believe useful at this point to stress the fact that semantic studies concerned with content that adopt a referential approach need to assume *an ontology* of the content domain under consideration, in addition to assuming a definite logical framework. By “choosing an ontology” I mean postulating the existence of a specific domain of reference, a domain of entities possibly constituted by different categories, and formally characterizing that domain through an axiomatics based on a number of relations on this domain. Formal ontology methods enable the development of this axiomatics, i.e., Montague and Dowty’s meaning postulates alluded to above, in a principled way and not on an ad-hoc basis. This requirement is well described in Bach’s “natural language metaphysics” programme [Bac86b], which encompasses some studies on plurals, mass terms and generics, and of several studies in the semantics of time that built on Davidson’s defense for events [Dav67]. With the exception of those early studies in formal semantics that were well anchored in the philosophy of language, selecting an ontology in current formal semantics studies is more often than not done with little attention to the fact that the ontology adopted can be more or less appropriate, i.e., adequate to how language describes that domain. For instance, if we take the temporal domain, which is involved when dealing with semantics of tense and temporal adverbials, time can be assumed to be a set of instants, the standard choice, a set of intervals or only an abstract construct characterized implicitly by a set of eventualities. Temporal relations on such different domains are then either just precedence or a family of relations (e.g., Allen’s relations for intervals [All84]). Finally, time can be assumed to be linear or branching, dense or discrete, etc. In brief, the situation hasn’t changed since Bach’s paper. One sees that even for a rather simple and well-explored domain like time there is no single obvious choice, and in fact, all possible choices are present in

the literature. I hold that it is important that ontological choices be made explicit and not taken for granted, and moreover, that such choices be motivated as far as possible on the basis of linguistic evidence, thus contributing to Bach’s programme.

This more than 20-years old programme unfortunately hasn’t developed that much as a general project of uncovering what ontological assumptions are presupposed by the semantics of natural languages, in other words, the ontological commitment of natural languages. Some work has been done for some domains, but no-one has proposed a global picture, stressing the importance of the fact that the alternatives faced when studying a single domain impact those of another domain.¹² Linguistic evidence in one domain doesn’t straightforwardly point to a single option as regards “the” ontological assumption presupposed by language, which in fact is underspecified. Choices have to be made. However, choices made for each separate domain are not independent from the others; all together they should provide a coherent, complete, picture of the world. My own work tried to complement the quite established lexically and ontologically-aware formal semantic studies mentioned above with the explicit goal to contribute to building such a complete picture. This picture is of course just one theory, one of the plausible renderings of natural language’s ontological commitment, without any claim that this is “the” one and only ontology.

My position fortunately doesn’t stand completely alone. On the one hand, some authors situate explicitly their work at the boundaries between formal semantics and formal ontology, although still focusing on specific areas and developing only partial systems (e.g., [Mol97]). On the other hand, in computational semantics other authors are stressing the need for well-founded ontological bases and do look for comprehensive schemes (e.g., [CR06] that adopts the foundational ontology DOLCE, [BHR95, BTF07] and [NR04]). Even developers of general lexical resources like WordNet [Fel98] are now aware of the importance of improving the ontological coherence of such resources [GGMO03, HCG+10].

1.1.2 Discourse semantics

Since [Kam81] and [Hei83a], twentyfive years of dynamic semantics have proved at length that meaning issues are not confined to the sentence, that discourse, a non-arbitrary sequence of sentences, is the relevant unit for many semantic phenomena, and that the formal semantics approach could be extended beyond Montague’s Grammar. There again, major efforts have been dedicated to the problem of constructing the right log-

¹²Borrowing an example from next section illustrates this claim. In DRT, the domain of quantification is separated in four sorts: times, events, states, and all the rest, in some sense subdividing the logical type e of entities into four subtypes. The motivation for distinguishing events and states builds on Davidson’s work and aspectual evidence (especially in French), and the restriction of the attention to temporal entities is probably due to the origins of DRT dedicated to the temporal structure of discourse, i.e., the temporal relations between discourse referents. But in a coherent larger picture, it would be necessary to distinguish further sorts among “the rest”, in a principled way. For instance, in DRT, events are assumed to have a time as location, but they are also located spatially (and I have contributed to show that space is also relevant to the discourse structure, see Section 4.6); the distinction between events and states has strong links with the count/mass distinction among material entities, etc.

ical form of a discourse: quantification scope is the core issue raised by intersentential anaphora and the famous ‘donkey sentences’ addressed both in Discourse Representation Theory (DRT) [KR93] and in Dynamic Predicate Logic [GS91].

Nevertheless, new content also arises at the discourse level, often in an implicit way, that is, without any word serving as a direct marker of such content. I believe this was first shown in [Hin81, KR83] regarding the anaphoricity of tenses, yielding temporal relations such as precedence or overlap between the eventualities of different sentences, something formally accounted for in DRT. This approach has been systematized in Segmented Discourse Representation Theory (SDRT) [Ash93, LA93, AL03], by borrowing the notion of *discourse relations* from discourse theories not rooted in formal semantics such as Rhetorical Structure Theory (RST) [MT88]. In a coherent discourse, a variety of discourse relations link its discourse segments (from basic clauses to complex stretches of text). Those relations characterize the discourse structure and may have, e.g., a temporal, causal, mereological or argumentative flavour. They model the semantic content specific to discourse, which concerns both “what is described” in the message and “the reasons why” the message is expressed, thereby extending the focus of attention to phenomena occurring at the semantics-pragmatics interface.

My work in discourse semantics has naturally focused on discourse relations, their nature, their contribution to meaning, and how they are linguistically expressed. In addition to grammatical features like tense, discourse relations are triggered by a variety of lexical cues: quite obviously adverbs and adverbials to which the category of discourse markers or connectives like *because* or *but* belong, but also adverbials that are not connectives (e.g., locating adverbials), verbs and nominals. This last point brought me to study the interaction between some families of lexical elements and discourse structure.

As established quite early in SDRT, beyond linguistic features proper, discourse structure is affected by “commonsense” (or “world” or “encyclopedic”) knowledge, especially typical knowledge like “when something is pushed, normally, it moves”, which is involved in the temporal reversal of events in the second variant of the following example.

- (1) a. Paul fell. Max helped him up.
 - b. Paul fell. Max pushed him.
- (from [AL03])

Therefore, when focusing on semantic content, in addition to specifying the ontology of the domain referred to, one should characterize commonly shared knowledge relative to that domain. Both these features call for knowledge representation.

1.2 Knowledge content

The area of artificial intelligence called knowledge representation has, just like formal semantics did, focused more on knowledge form rather than knowledge content. The bulk of the work in this field is in fact dedicated to specific kinds of knowledge and reasoning. It largely focuses on the epistemic status of knowledge, as with beliefs, or with uncertain,

vague, and typical knowledge.¹³ Such kinds of knowledge are best characterized by how one reasons about them, and accordingly the field has extensively studied dedicated modes of reasoning (e.g., reasoning about belief change, reasoning under uncertainty, reasoning with typical knowledge admitting of exceptions) and developed corresponding families of non-classical logics. For instance, commonsense knowledge exploited in SDRT as evoked above mainly comes under the form of typical knowledge and accordingly this theory makes use of a non-monotonic logic. Most of these reasoning formalisms aim at solving epistemic issues, not ontological issues regarding what I call content; they are neutral on *what* one believes, *what* is uncertain, or *what* is typical, and they even do not say much regarding *what* is a belief.¹⁴

1.2.1 Temporal, spatial, and qualitative reasoning

There are of course several important exceptions to this picture. In particular, the representation of time, situations and “actions”¹⁵, space, and what is called “qualitative physics” [Bob84, WdK90], “commonsense knowledge” [Dav90, Mor06] or even “naive physics” [Hay78, Hay85b] have historically constituted important areas of the field. Theories of time, space, causation, quantities, etc. have been proposed thus characterizing corresponding ontologies of these domains. In many occasions though, the focus has rather been put on reasoning about how time, space or action affect *other* knowledge. For instance, when reasoning about persistence and change in what is often called temporal reasoning, time is rather seen as a modality (i.e. when some proposition or “fluent” is true) and not as the primary object of study (see, e.g., [HM87, Rei01] as opposed to, e.g., [AH85]; both concerns are addressed in e.g., [McD82, All84]). It has though been shown that the nature of time does matter for reasoning correctly about change [Gal90]. So here again, in some occasions, the underlying ontologies have either been taken for granted or chosen in order to give the specific type of reasoning focused on good properties rather than for their own sake.

These domains have an indubitable general relevance, and this is one of the reasons why I have contributed to the development of the area called “qualitative spatial representation and reasoning”. It can also be noted that domains like time and space correspond to well-studied mathematical structures which make the development of corresponding theories or logics simpler than for a priori less explored (but not less relevant) domains like, say, material objects or social entities. But even in such a favorable con-

¹³On the respective roles of epistemology and ontology in knowledge representation and reasoning, see [Gua95].

¹⁴This is not to say that all belief logics only aim at solving epistemic issues. There are multimodal logics which relate different operators and do express some content knowledge about what these operators represent. This is the case, e.g., for logics relating beliefs of collectives and beliefs of individuals, and logics relating together beliefs, desires and intentions, which do say something about those concepts. And I surely do not imply here that, more generally, modal logics do not represent content; see Section 1.3 regarding the use of modal logics to represent domains like time or action.

¹⁵What are called actions in frameworks such as situation-calculus are only events, or in fact event types. There is no agent and no intention explicitly involved. Such theories usually focus on causation between events and states called effects, when some states called precondition hold or do not hold.

text, one soon realizes how complex and varied theories can be. To avoid developing just one option chosen more or less arbitrarily or by sticking only to formal properties such as completeness and categoricity, one has two strategies.

The first is to use empirical evidence from the concrete problems to be solved using knowledge representation. Here, when the aim is to describe the domain of reference associated to some semantic phenomena, I will use linguistic and more generally cognitive evidence.¹⁶ Such evidence provides a measure of adequacy by examining how straightforward the interpretation function is for terms denoting individuals and predicates denoting relations. Therefore, the two parts of my work are closely related: on the one hand doing semantics requires knowledge representation, especially formal ontology; on the other hand, developing adequate ontological theories requires empirical evidence, and I especially rely on linguistic observations. In fact, several of my papers are a combination of semantics and ontology.

The second, much more ambitious, strategy is to explore whole families of theories, something which I have tried to do in some occasions. Focusing on the contrasting properties that generate alternatives and analysing as much as possible the theoretical consequences of selecting one or the other help drawing a map of that family so as to orient the choice of an adequate theory in presence of some specific constraints. When a domain has been modelled through mathematical structures, formal translations and mappings further strengthen such maps (see [Mas00], a PhD that I contributed to supervise). For the many domains yet unexplored with algebraic tools, the variety of existing philosophical analyses provides invaluable departure points for trying to develop families of theories. A further advantage of paying attention to philosophy when doing knowledge representation is that it is only in metaphysics and formal ontology that one can find an attention to the coherence of comprehensive views of reality and to constraints relating choices in separate domains.

1.2.2 Ontologies

I obviously cannot ignore here the trend in knowledge representation specifically dedicated to representing content through the development of the so-called “ontologies”, which constitutes the second notable exception to the brief description of the field given above. This new branch of the field has grown in an exponential way in the last fifteen years in particular to answer the urgent interoperability requirements to develop the semantic web vision. In this context, ontologies are computational artefacts aimed at making formally explicit the intended meaning of a certain (technical or general) vocabulary [Gru93]. Unfortunately, the theories that can be associated to these artefacts may or may not deserve to be called ontologies as I used the term above, for they can be anything from a simple taxonomy of concepts (unary predicates related by implication, i.e. class inclusion) to a full-fledged axiomatic theory characterizing the structure of a domain and its relations [GOS09]. In fact, what we very often get is a “lightweight ontology” that is just a taxonomy, which provides a very limited characterization of a

¹⁶See also Section 1.3 below and especially footnote 19 regarding the use of linguistic evidence.

domain with limited inferential power.

The efforts to develop ontologies in a principled way are relatively low in contrast with efforts dedicated to the development of languages to serve as a standard, such as description logics (DLs) and especially OWL. DLs are fragments of first-order logic that have good computational properties and are particularly adequate to efficiently support taxonomic reasoning [BHS09, BCM⁺03]. So again, and this might come as a surprise in a field calling itself ‘ontology engineering’, the focus has actually borne more on form than on content proper. No question that this work was necessary, but it appears that the community tends to rely on the assumption that it is easy to fill in those formal languages with knowledge content. Either any domain expert would be able to formally describe his domain and create an ontology, suitably guided by methodologies and editing tools, or ontologies would be automatically extracted from data such as technical documents. Indeed, ontology learning is another main activity in that area [AGDS08, CMSV09]. It has though be shown that a number of principles need to be respected to avoid inconsistency and other unwanted consequences [Gua98, WG01]. Errors are easy to make because such principles call for non-trivial formal ontology notions and methods that certainly do not belong to common sense and are difficult to embed in extraction and learning methods. In fact, it is far from being not enough to be an expert of a domain to be able to build a theory of that domain. As a comparison, certainly not all competent speakers of a given language are able to write a grammar of that language. To further illustrate this point, let me just note that one far too often finds ontologies that confuse the basic taxonomic “is-a” relation (class inclusion) and the “instance-of” relation (membership). Proper guidelines to enforce ontological coherence are still limited to the OntoClean method [GW04] which is anyway not trivially implemented.

Of course, there have been several tentatives to propose standard ontologies, for instance Cyc/OpenCyc [LG90] (see also <http://www.opencyc.org/>) and SUMO [NP01] (see also <http://suo.ieee.org/SU0/SUMO/index.html>). These are much more than simple taxonomies and aim at covering large chunks of knowledge from very general concepts, i.e., those constituting the so-called “top-level” or “upper-level”, the part of an ontology which depends on formal ontological distinctions rather than directly on empirical domain data (see Chapter 5), to specific domains. Cyc in particular aims at being an encyclopedic knowledge base of all facts needed for commonsense reasoning so covers detailed domains that range from bicycles to shooting stars to Peter Pan. Unfortunately, these standards are not really convincing as far as well-foundedness and coherence is concerned, and in applications they have not achieved a breakthrough. In addition, however carefully designed it could be, any standard is unlikely to be satisfyingly adequate to all purposes, and indeed, a part of the computational ontology community is now explicitly rejecting them (see e.g., [VC98]). Instead of a unique standard, one may rely on semi-automatic or automatic alignment procedures to reach interoperability between systems based on different ontologies [ES07, Noy09]. This approach is the only one applicable in the case of legacy information system reuse. However, the ontologies to be aligned usually being just taxonomies, perhaps extended with parthood relations, the structural information is clearly insufficient to grasp the intended meaning of each concept and

establish correct mappings.¹⁷ Some techniques then rely on the lexical meaning encapsulated in the concept labels and make use of lexical resources to resolve ambiguities, but this is highly prone to induce errors, not just because word-sense disambiguation doesn't reach perfection especially when there is no linguistic context to exploit, but also because those lexical resources themselves are not ontologically irreproachable.

There is however another approach yet to reach interoperability. The WonderWeb project [MBG⁺03] aimed at building a library of “foundational” ontologies, quite similar to the strategy of developing the map of a family of theories that I described in the previous section. The vision of this project was not to provide a single standard ontology, but a variety of alternative carefully designed formal ontologies, already aligned by manually establishing their theoretical relationships. The advantages and inconveniences of each would be explained, so as to make the choice of an adequate ontology easier. The focus in this project was on foundational ontologies, that is, fully developed axiomatic theories whose domain is supposed to cover the whole of reality, subdivided into general categories, and characterized through general or formal relations (formal in Husserl's sense, i.e., relations applicable across all sub-domains) such as parthood, dependence and constitution. The project contains three such foundational ontologies, one of which, DOLCE [MBG⁺03, BM09] (see also <http://www.1oa-cnr.it/DOLCE.html>), has been largely used in a variety of applications (see e.g., [GGMO03, TLG⁺06, OAH⁺07]). I have situated some of my work in this framework, with the objective to contribute to the next, extended, release of DOLCE (see Section 5.1.4).

1.3 Methodological issues

A number of principles are used in the different disciplines my work is based on. It goes without saying that first comes a principle of consistency that cannot tolerate deviations.¹⁸ Next is adequacy, a must in linguistics and cognitive science: the representations should account as much as possible for the observed behaviour of language or cognition, and in particular should aim at modelling neither less nor more linguistic expressions than those attested. Such account must be done in an as transparent and direct as possible way. This constraint leads in particular to the compositionality principle in formal semantics, avoiding arbitrary intermediate paraphrasing of a sentence to obtain its logical representation. In formal ontology, adequacy is also the aim, with empirical bases varying across authors from personal intuition or (their version of) common sense, requirements from other areas of philosophy, linguistic evidence,¹⁹ to

¹⁷This is of course a controversial statement. But even though structuralist linguists working in the terminology field hold that the meaning of a term lies in its relations with the other terms, these relations are much richer because they are based on equivalence of terms in a paradigmatic analysis, that is, they exploit the complexity of linguistic structures.

¹⁸I obviously ignore here the epistemological tasks of fusing different, incoherent, sources of knowledge, and of maintaining the consistency of a collection of contingent facts evolving through time.

¹⁹Some philosophers hold that the discipline ontology, as the objective study of what exists in reality, has nothing to do with cognition nor language (see, e.g., [Var02, Var07] for the pitfalls of using linguistic evidence in ontology, and [Hof07] for a fresh look on this issue grounded in the notion of semantic information partition (topic-focus)). Although this is certainly not the place to argue about metaphysics

accepted facts in experimental sciences. The transparency constraint is also adopted by some authors for whom the description of reality should be made at a mesoscopic or middle level, the level corresponding to human perception and cognition of reality, as opposed to a description in terms of, e.g., physical subatomic particles. Adequacy—to language and cognition—also is my main principle, which imposes the search for sufficiently expressive representation frameworks.

Clearly, there are a number of other principles that are followed, and in some occasions, these may clash with adequacy. We have just seen the transparency constraint, with compositionality that finds its limits at the semantics-pragmatics interface. Two other principles have great importance: parsimony, especially in formal ontology, and computational efficiency in knowledge representation.

Parsimony is the principle that a theory should not be more complex than strictly necessary, especially in terms of number of entities or number of entity categories that are assumed to exist, cf. the well-known Ockham’s razor. However, the tension between adequacy and parsimony is resolved in the formal ontology literature in a large variety of ways from extreme unifiers or reductionists to extreme multipliers (cf. [PV00] for a review of positions in the domain of events). I generally adopt a middle position [VBM08], that is, I retain that parsimony is a valid principle but that adequacy, along with transparency, has in general priority over it.

Computational efficiency is a prominent principle in computer science, for good reasons. In logical approaches to knowledge representation, this implies a minimum of using decidable logics. Unfortunately, the clash with expressivity is here a major one, as first-order logic is not decidable while it is a minimum for an adequate level of description of content both in formal semantics and formal ontology. Indeed, the formulas obtained by a standard formal semantic analysis of natural language sentences are formulas of an at least first-order logic.²⁰ Focusing on the content of those formulas and on the formal ontologies modelling the corresponding domains requires the ability to predicate and quantify over these domains (cf. Quine’s motto “to be is to be the value of some variable” [Qui48]). This means developing first-order axiomatic theories in classical predicate logic—the choice I have made in all my work—, rather than, for instance, developing dedicated modal propositional logics. I thus favour expressivity, deliberately ignoring efficiency. My position is that with an adequate first-order theory in hand, one can in a second step extract decidable fragments, for instance in some description logic. Developing in a second step specific fragments supporting efficient reasoning for specific applications presents the advantage to have a clear knowledge of the simplifications that are thus made and of the limitations of these fragments, something which is not guaranteed when efficiency is put forward and expressivity is constrained from the start.

Insisting on the difference between first-order theories and (modal) logics might seem

and its methods—and I am surely not qualified to do it—I wish to take a pragmatic stance on this issue, observing that consensus is very rarely achieved in the community thus showing that personal biases exist, and that it is not clear how far one can free oneself from perceptual, linguistic, cultural and other cognitive influences in philosophical analyses.

²⁰In some occasions, first-order modal logic, and in others even higher-order logic, e.g., for some generalized quantifiers like *most*.

artificial. It has been shown, for the temporal and spatial domains in particular, that some first-order theories are equivalent to some modal propositional logics, in the sense that they share the same mathematical structures for models. The advantage of developing modal propositional logics is that, already having a restricted expressivity, it is easier to extract a decidable fragment as a basis of possibly efficient reasoning algorithms. On the other hand, the advantage of developing first-order theories, even when full first-order is not required, is that the domain is directly reflected in the language, so transparency can be preserved more easily [Qui48]. In addition, first-order logic –or *classical* logic– has become a standard in some areas of knowledge representation and so is to be preferred also for reasons of favouring communication, comparison, and integration. To further enhance communication and comparison, where it made sense, I have tried to demonstrate representation theorems linking the theories developed to well-known mathematical structures (see sections 5.2 and 6.3).

Even though I chose to favor expressivity and renounced to pay attention to efficiency, I tried to stick to first order (sometimes modal first-order logics that are reducible to first-order logic), with the aim to leave open the possibility to extract decidable fragments. Higher-order logic is actually needed in some areas of formal semantics, but for a computational semantics objective, choosing first-order logic as a general framework has clear advantages without being unreasonably restrictive, as convincingly advocated in [BB03, BB05].²¹ Perhaps some specific limited second-order logic would on the whole be more adequate, but what would be such a framework is not obvious and would require specific studies. Clearly, the decidable monadic second order logic is not adequate, as one simply cannot renounce to binary relations both in formal semantics and in formal ontology. The first-order interpretation of second-order logic via the “general models” [vBD83] could well be the answer. However, with the aim to make future implementation of fragments and integration possible, I have chosen to already apply at the language level the reification strategy used in the general models when I had to quantify over concepts (see Section 6.1).

Regarding the use of the lambda-calculus in compositional semantics, which arguably is not first-order, I adopt the utilitarian viewpoint defended in [BB05]. Lambda-calculus is for me a simple tool or “glue language” for building the first order formulas corresponding to elementary discourse representation structures, and from this point of view can be seen as a simple extension of first-order logic. So, whatever the semantics provides for terms as, e.g., intensional verbs, I view as a separate question from what there is.

²¹The domain itself might require an higher-order framework, when e.g., analysing mathematical texts dealing with cardinality. My position is that concentrating on general language and the everyday world, that is, letting these very special cases aside, doesn’t make the whole approach useless.

2 Ontology-driven lexical semantics: The semantics of spatial expressions

The work I have done in lexical semantics is characterized by its initial focus on spatial expressions, principally locative expressions involving spatial prepositions. This focus extended across the years in two directions: towards the temporal dimension with temporal and spatio-temporal locating adverbials (see Chapter 3) and towards the material and functional dimension with parthood expressions (see Section 5.3 in Chapter 5). As we will see, these three areas are tightly related; they all concern basic relations between concrete entities, i.e., entities that are located in space-time, as opposed to abstract ones.

The specificities of this work are first of all methodological. As explained above, I stick to a formal semantic, thus logical, approach using observations of the inferential behaviour of the lexemes under examination as significant empirical data, with an attention to both the ontological status and the cognitive relevance of the entities and predicates involved in the representations. Similar mixtures of interests have appeared repeatedly in the literature, probably first within the Hamburg-based LILOG project which dedicated a whole sub-project to spatial expressions and spatial knowledge [HR91]. However, I believe that my constant attention to ontological issues is rather singular. For instance, considering recent works that definitely posit themselves within formal semantics, ontology concerns are absent from Zwarts’s work on spatial expressions and that of Pratt and Francez on temporal prepositions [Zwa97, PF01].

Locative sentences describing static spatial relations are usually expressed in French and in English through the simple pattern “NP_{located-object} être (to be) P NP_{reference-object}” (e.g., *le chat est sur le tapis*, *the cat is on the mat*) in which the spatial relation is denoted by a spatial preposition among which *à* (*at*), *sur* (*on*), *dans* (*in*), *derrière* (*behind*), and complex ones like *à droite de* (*to the right of*) and *au milieu de* (*in the middle of*). In such patterns, following part of the tradition, we will say that the subject NP denotes the “located object”, and the NP complement to the preposition, the “reference object”.²² Spatial prepositions have therefore received considerable attention in studies on the semantics of spatial expressions, and constitute, together with motion verbs, the bulk of the studies in lexical semantics regarding space.²³

After the study of the preposition *dans* (*in*) during my PhD, together with Michel Aurnague, who had studied *sur* (*on*) at the same time, we generalized our work in [AV93a, AV93b, AV95b, AVB97], gathering the principles that guided it and extending it, largely building on the work of Michel Aurnague, to other prepositions (the projective prepositions, e.g., *in front of*, *below*, *to the right of*, etc.) [Aur95b] and other spatial expressions (the so-called “internal localization nouns”, e.g., *top*, *front*, *middle*, *corner*) [Aur96].

I will in the rest of this chapter focus on how our approach is situated in the literature,

²²Trajector/landmark and figure/ground are also found.

²³One could question the fact that studying prepositions belongs to lexical semantics, because prepositions are a closed-class so in principle not part of the lexicon. However, it can easily be argued that spatial prepositions carry at least as much meaning as a syntactic function.

motivating its main features. First, it belongs to the trend of functional approaches, but rather uncommonly, it adopts a relational view on spatial expressions. Second, it distinguishes the role of pragmatics among what is generally considered as functional features. This helps identifying precisely what relations constitute the semantics of spatial expressions; relations and the nature of the entities that are their arguments is the topic of the third section. Because of our ontological orientations and the reference approach founding formal semantics, we paid a special attention to such entities, avoiding the usual geometrical idealization.

Finally, I will address the work, still largely in collaboration with Michel Aurnague, dedicated to confront and enrich our analyses with empirical data gathered through psycholinguistic experiments that we contributed to set up.

2.1 Relational and functional approach

Spatial expressions have been largely studied in cognitive semantics and in psycholinguistic studies, for they are the ideal locus to examine how the cognitive faculties of perception and motion control are related to language. Even though situating our work in a logical framework with the aim to complement formal semantics does not fit the cognitive semantics programme, most relevant analyses we built on belonged to that area. Among these studies (so disregarding its logical features), our work distinguished itself by adopting a *relational* and *functional* approach to spatial prepositions, a position induced by the search for ontological and semantic precision.

The functional approach was emerging at the time, in contrast with the geometric approach traditionally adopted but still widely used today from cognitive semantics to computational semantics. In the geometric approach the located and reference objects are idealized as abstract geometric entities (points, lines and 2D or 3D regions with simple shapes) situated one with respect to the other through geometric relations denoted by the prepositions. For instance, *in* is generally assumed to denote inclusion and *on* contact (omitting here subtleties due to dimension) [Lee69, Her86, Tal83, Zwa97]. It shouldn't be surprising that even in cognitive semantics the formal tool of geometry has played such a prominent role. In this trend, geometric schemata or "image-like" representations have been proposed as a major modality of mental representations, that corresponding to perception and visual memory and reasoning.

This classical geometric approach is in addition made "referential" in [Ben75], [Bie88] and [Jac83, Chapter 9], the latter being very influential in current lexical semantics as well as in computational linguistics [CGN89] and in formal semantics [Zwa97]. In the referential approach, the whole preposition phrase is assumed to denote a third referent,

usually some spatial region²⁴ in which the located object is situated.²⁵ The preposition itself then, instead of a relation, denotes a function that applies to the denotation of its complement NP, the reference object, to yield that spatial region. In other words, there is a unique relation of spatial location, somehow corresponding to an inclusion, and the difference in meaning between prepositions is grasped through different geometric functions determining different regions with respect to the reference object: an inside region for *in*, a below region for *below*, etc.²⁶

The relational approach, in which each preposition denotes directly a specific spatial relation whose arguments are the located and the reference objects, has several advantages over the referential one, in addition to being obviously more parsimonious. First, not all geometric relations are equivalent to a mere inclusion in some spatial region. This is the case of the relation of contact involved in the semantics of *on* or *against*. Since large objects can be *on the table*, the region corresponding to this PP is a region above the table thick enough to include many smaller objects not in contact with the table (say, a flying fly) and thus not *on the table*. And clearly not all parts of the object on the table are on the table, while those parts necessarily are in the “on-the-table-region” if the whole object is on the table. Second, PP modifiers denoting a distance relation, as in *the cat is five centimeters behind the chair*, can hardly be taken into account in a compositional way with the referential approach. The modifier cannot denote a function transforming the region into another, for the reference object is no longer accessible to compute the correct distance [Zwa97]. In addition, just like contact, distance relations cannot be correctly handled by an inclusion relation in a region if the located objects

²⁴These spatial regions are sometimes conceived of as abstract regions in an absolute “container space”. However, they are entities in some sense dependent on the reference objects: they are descriptively dependent from it and they necessarily follow its movements. One can even take them to be existentially dependent on concrete objects, thus rejecting the very idea of a container space along with Leibniz. In my work (see Section 5.1) I have called entities like *the inside of / the space in the cup* or *the space underneath the table* “space portions” while I used the term “region” more generically to refer to the spatial extension of concrete entities (in the early papers, I have used quite awkwardly the terms “individual” or “spatio-temporal referent” instead of “region”). In this section though, I stick to the use of “region” instead of “space portion”, as usually done in the literature on spatial prepositions.

²⁵Zwarts [Zwa97] actually adopts a purely geometric approach at first sight very similar to the classical referential approach except that instead of a region, the PP denotes a set of vectors. Vectors can be used as a geometric representation of a relation between two points, especially adequate for projective and distance relations.

²⁶Most formal works based on this approach use in addition a basic “Place” function associating to any concrete entity the spatial region it occupies, and then use an inclusion relation between the spatial region occupied by the located object and the spatial region denoted by the PP anchored on the reference object. This apparently requires some absolute container space of which both spatial regions are part. However, spatial expressions denote relative positions (the place of the wine stays included in the region denoted by the static PP *in the bottle* wherever the bottle happens to be) and indeed no absolute space imposes itself as cognitively relevant in all situations. The only cognitively adequate solution along these lines is to use a container space anchored to the reference object, which means that the place function applied to the located object is a different function for each reference object, and so is of little use. In my work, I have often used a functional notation of “ST” or “STref” (spatio-temporal referent) quite similar to a Place function at first sight, but it was just a means to denote the equivalence class of those entities that are spatio-temporally colocated, as a statue and the piece of matter it is made of.

are extended: when the cat is five cms from the chair, surely not all parts of the cat are at five cms of the chair, and yet they are included along with the whole cat in the same “five-cms-from-the-chair-region”.²⁷ With the relational approach, distance modifiers simply add through conjunction a further distance relation between the located object and the reference one. Finally, the referential approach is by construction limited to an analysis in geometric terms, while the relational approach can easily be extended beyond geometry. And many authors, building on Herskovits’s seminal descriptive work [Her86], have shown that, in addition to or even prior to geometry, functional features are central in the semantics of spatial expressions [Van86, Cov92, CRCL99, GFC99]. For instance, *in* involves the relation of containment and *on* that of support (more on this below).

The alleged benefits of the referential approach are of two kinds. First, it would give a unified account of all prepositions. But the apparent unification is at the cost of considering that some prepositions have a null semantics: *at* and *in* wouldn’t contribute anything when the reference object is a location (and even when it is a material object in some accounts that consider the inside as part of the object). Second, it would reflect in the composition the syntactical asymmetry between located and reference objects [Jac96]. This asymmetry is supposed to appear semantically through the well-known figure/ground phenomena that show that there are no exact inverses among spatial relations: *the cat is on the mat* is not equivalent to *the mat is under the cat*, nor *the bicycle is near the house* to *the house is near the bicycle*, to take classical examples [Tal83]. But what is at stake here is not really a semantic feature of spatial expressions. It is a quite general pragmatic principle acting over the informational partition into focus and background. It is more informative and thus more relevant to use locative sentences to situate mobile objects with respect to fixed ones, whose location is more likely to be known, rather than the opposite.

2.2 Functional features and pragmatic effects

In fact, a further characteristics of our approach has been to isolate pragmatic effects, distinguishing them as far as possible from semantic features especially from functional ones, yielding what we called a “three-level approach”: geometric, functional and pragmatic.

In the classical approach, both functional and pragmatic features were considered as sources of deviation from the “core semantics” which is purely geometric. In some sense, any non-geometric feature of spatial expressions is considered as a pragmatic or conventional phenomenon modifying the standard meaning [Her86, Jac83, Zwa97]. This confusion can be explained on the one hand by the bias towards image-like representations in cognitive semantics and on the other hand by the prominent role of functional features in the application of pragmatic rules as will be seen below.

²⁷Zwarts’s proposal actually suffers from a similar problem as his semantics involves a universal quantification on all the points inside the region corresponding to the located object. For him, all points in the cat’s region have to be at five cms of the chair.

However, functional properties certainly are involved in the core semantics in the sense that there are configurations for which no communication context makes these functional features irrelevant, and this is so even when there is no deviation from the canonic geometric situation. The involvement of functional features in the semantics may happen in two ways. In the first, a functional relation directly holds between the two related objects, for instance, support for *on* is required when the located object is not above the reference object (*a cabinet on the wall* can in no context describe a cabinet standing against the wall). The second operates by restricting or contributing to determine the geometric relation between the objects.²⁸ For instance, *in* applies when a located object is simply situated in the inside of a reference object (even when there is no functional relation of containment involved, as with a fly flying in a bird cage), but the inside is a region determined by the containment capacity of the reference object, if any (in the previous example, the possibility to contain birds); the purely geometric convex-hull function often used is not enough to characterize the inside of a containing object (in the example above, the bird cage may have a very contorted shape) [Her86]. Projective prepositions (e.g. *in front of*) provide another example. They are notoriously ambiguous between an intrinsic use and a deictic one (among other contextual uses). In the intrinsic use, the semantics of these prepositions is based on an intrinsic “frame of reference” anchored to the reference object; some functional characteristics of the reference object determine the projective axes forming this intrinsic frame of reference (e.g., the usual direction of motion determines the frontal axis of a car) —see Section 2.4 below.

Several other deviations from a purely geometric account are instead primarily the effect of general pragmatic principles, as we have seen for the figure/ground effect. I have thus examined how Grice’s maxim of quantity [Gri75] applies in this domain, observing tendencies to choose a maximal located object and a minimal reference object (e.g., *the engine is in the garage* is not appropriate when *the whole car is in the garage*). Another group of pragmatic effects is based on the typicality or canonicity of a situation, itself largely influenced by the function of artefacts that are containers, support providers, etc. This is the case of the “fixation” principle by which the typical use of artefacts restricts their inside or their supporting surface (e.g., although a sticker *on a piece of furniture* can be attached anywhere on it, a sticker *on the table* is primarily assumed to be on the main supporting surface of the table). Similarly, canonicity is involved in differentiating the appropriateness of distinct locative sentences: there is a strong preference for the description that better matches the canonical situation. The canonicity of a situation is in particular higher when the functional relation associated to the preposition (e.g., containment for *in* and obstruction to perception for *under* [Van86]) is present; thus, the fact that there is no containment when an apple is lying on the table under an upside-down bowl rules out the use of *the apple is in the bowl* to the benefit of *the apple is under the bowl*. Similarly, the presence of a functional interaction between the located and the reference objects favors the intrinsic use of projective preposition *in front of* over the

²⁸This second type of functional effects could be accommodated in the referential approach by contributing to improve the function denoted by the preposition.

deictic use favored in its absence [CRR96]. In some occasions though, the canonicity of a geometric configuration may overtake the role of function: e.g., an object held by a string above the table so that it just touches the table can in some contexts be said to be *on the table* even though the table does not support it (although with less confidence than if the object simply lies on the table [GFC99]). The influence of these general pragmatic principles on reasoning patterns involving spatial prepositions has been assessed and confirmed in Jean-Christophe Aurisset’s dissertation, which I supervised [Aur93].

2.3 Relations and entities

Expressing the semantics of spatial prepositions with a formal semantics objective requires characterizing the semantics of the primitive relations involved in the representation. Our three-level approach is based on geometric and functional primitive relations, the latter actually encompassing all non-geometrical relations between entities, which obviously range far beyond the notion of “function” of artefacts or biological organisms (for instance, support is a physical notion related to gravity). Analyzing such relations implies first of all to understand what are their relata, i.e., what sort of entities they have as arguments.

As far as geometric relations are concerned, let’s first observe that influential works in cognitive semantics making use of semantic primitives rely on vague notions like “PLACE” or “PATH” without further characterization, and this of course is insufficient for our purposes [Jac83, Tal83]. Among those that precisely define the entities used as relata, most works reuse standard geometry assuming the relata are elements in some formal space (usually \mathbb{R}^3) with a more or less idealized shape. Some proposals, often motivated by cognitive idealizations, assume relata are points, lines and surfaces,²⁹ parallelepipeds, or complex volumes constructed out of cylinders [MN78]. Other proposals stick instead to the physical reality in which material entities come in all shapes and are spatially extended, and so take as relata regular sets of \mathbb{R}^3 . Our cognitive-ontological approach led to yet another proposal, at the same time compatible with the spatial extension imposed by reality and with the fact that abstract constructs such as sets in \mathbb{R}^3 are cognitively improbable: we take as relata the concrete entities themselves, without assuming they are constituted of an infinity of 0-dimensional points. If we restrict our attention to geometric relations and properties holding at a given time, this strategy amounts to use “point-less” theories of space in which the primitive entities are extended regions. This approach, appropriate not only for language but for many applications based on commonsense or qualitative knowledge, was emerging in the early 90s and led to a whole new area at the intersection between artificial intelligence, formal ontology and logic, called mereotopology and mereogeometry. This area and my

²⁹The cognitive idealizations used in many cognitive linguistics proposals embody granularity-sensitive perspectives which can easily be challenged within language itself: in a same sentence, a same referent can switch from the alleged point-like view imposed by the preposition *at* to the extended view imposed by the preposition *in* (and the parthood relation involved in the bridging anaphor), as with *I am at home, in the living room*. Our view is that possible effects due to idealizations are best taken into account as pragmatic effects rather than semantic ones.

contributions to it will be described in Section 5.2.

Regarding functional features, it is quite surprising to see that the literature that has discussed extensively the categorization of spatial concepts denoted by spatial prepositions and verbs has not focused much on the categories of spatial entities involved in spatial expressions. This can perhaps be explained by the fact that these prepositions do not appear to have selectional restrictions, that is, they seem to admit relations of many sorts. However, observing admissible inferential patterns involving spatial expressions leads to draw ontological distinctions about the entities related beyond their spatial dimension. For instance, I analyzed possible and impossible transitive patterns involving the preposition *dans* (*in*). This preposition has widely been assumed in the literature to be transitive, and it must indeed be so for those adopting a purely geometric approach in which *in* denotes inclusion. Nevertheless, transitivity is impaired in several cases: in the trivial case in which the inclusion is only partial:

- (2) a. *il y a une abeille dans la rose* (there is a bee in the rose)
- b. *la rose est dans le vase* (the rose is in the vase), but
- c. *l'abeille n'est pas dans le vase* (the bee is not in the vase)

but also in cases in which entities of different categories determine different sorts of “insides”:

- (3) a. *Léa est dans l'île de Lampedusa* (Lea is the island of Lampedusa)
- b. *Lampedusa est dans la mer Méditerranée* (Lampedusa is in the Mediterranean sea),³⁰ but
- c. *Léa n'est pas dans la mer* (Lea is not in the sea), at least not in a strict sense.

Such distinctions were very generic and actually not focused on the notion of function of entities itself, appropriate only for a limited class of entities (especially artefacts), but on the behaviour of the so-called functional relations that actually concern the structural and physical properties of the (concrete) entities related. They are twofold. First there is the need to categorize concrete entities into basic categories or types, thus distinguishing material objects from immaterial ones (also called space portions) and locations, which behave in different ways as far as containment is concerned. Michel Aurnague also showed that such distinctions are syntactically marked in Basque and are necessary to account for the internal localization nouns in French and other languages [Aur95a, Aur96, Aur04].³¹ The categories of this linguistically adequate ontology is discussed in Section 5.1. Second, structural distinctions, that is, those regarding the internal arrangement of parts, are relevant because on the one hand the specifically structured collective entities again behave differently with respect to functional relations, and on the other hand cases of parthood are often described through spatial prepositions (e.g.,

³⁰A use of *dans* between locations that doesn't involve parthood as is usually the case (e.g. with *Lampedusa is in Italy*), but an enclave relation.

³¹Locations are also distinguished from objects, e.g. containers, by the fact that locations but not objects can combine with the preposition *à* / *at* [Van88, Aur95c].

the brain is in the head, Lea is in the crowd, the switch is under the appliance). Studies of the parthood relations and my work regarding them are the subject of Section 5.3 in the Chapter on the ontology of concrete entities, but let's note here that my colleague Michel Aurnague did a linguistic descriptive analysis of the parthood expressions in French (relational nouns like *partie, part, portion, membre* etc.) and Basque (genitive cases *-ren* and *-ko*) that corroborates this proposal [Aur98, Aur02, Aur04]. One will see in that Chapter 5 that actually these two sorts of distinctions, categorial and structural, are very much related: the structural properties of entities contribute to characterize the categories themselves, in line with the hypothesis founding Husserl's programme stating that formal relations such as (general) parthood provide the basis for determining the fundamental categories (see also [CV99]).

Finally, other formal ontological relations beyond parthood were required to model correctly the semantics of spatial expressions. This is especially the case of the relation of dependence that attaches space portions to the material entity that determines it (e.g., the inside of a container depends on it). Taking into account this dependence relation immediately explains the known transitivity puzzle regarding holes and the preposition *in*: *there is a hole in the sheet, the sheet is in the drawer, but there is no hole in the drawer*.

To fully analyze apparently simple lexemes like spatial prepositions, we thus see that an in-depth observation of their behaviour is required along with a wide and general ontological basis.

2.4 Empirical investigations

Cognitive semantics has inspired and sometimes worked closely with psycholinguistics to put their hypotheses to the test. The semantics and pragmatics of spatial expressions easily lend themselves to tests based on the matching between a linguistic expression and an image representing a spatial scene. In particular, several psycholinguistic studies have tested and confirmed the effect of functional relations in the applicability of prepositions, from the early nineties on, i.e., after the influential linguistic work of Herkovits and Vandeloise [Cov92, GFC99]. As far as the projective prepositions —pervaded by the ambiguity of their intrinsic, deictic and contextual uses— are concerned, additional more specific results can be found. It has been shown that an active functional interaction between the located and reference objects favors the use of the intrinsic use of prepositions like *in front of* [CRR96]. [CPSR01] examines the importance of geometric and functional features for the appropriateness of projective prepositions along the vertical dimension, much less subject to intrinsic effects because of the pervading effect of gravity; both are shown to be relevant and functional effects appear even when the situation is geometrically canonical.

I contributed to this trend within two successive French national projects (1997-1998 et 2000-2001) in collaboration with linguists and psycholinguists. Although the collaboration between cognitive linguists and psycholinguists is quite common, the confrontation of a precise logical analysis of spatial expressions to empirical data is not often found.

We turned our attention towards the relative importance of geometric (e.g. shape) and functional features in the determination of the reference frames themselves, for which it is crucial to decide on what is *the front* of an entity. The question had only been partially addressed at the time [CRI93, Tve96]. Our work [ACV⁺07] built on Michel Aurnague’s analysis of internal localization nouns (ILNs) [Aur96], that is, the semi-open class of lexemes enabling the identification of a systematic arrangement of oriented parts within almost any concrete object: the relational nouns *front*, *back*, *top*, *bottom*, *right side*, *left half*, *middle*, *corner*, *edge*, *tip*, etc. While such an internal structuring system is largely determined by the shape of the object for its division into relevant halves, portions of surfaces, etc., the functional features and the canonical situations in which the object is involved usually determine how to label those halves, surfaces, etc. into *front/back*, *top/bottom*, etc.

We ran a series of three experiments with two types of tasks: a pointing task on a tactile screen to indicate the part of the object depicted on the screen corresponding to the ILN presented, and a yes/no judgement task to test the appropriateness of the ILN proposed to signify the highlighted part of the object presented on the screen. The ILNs under test were the pair *avant* and *devant* (*front (part or half)* and *front surface*). The results proved, as expected, that most fronts are identified by two classes of functional features, typical interaction (static, e.g., the facade of a house, the screen of a TV set) and typical direction of motion (dynamic, e.g., the forward part of a bicycle or an animal). The first experiment also showed an unexpected phenomenon that was confirmed by the following ones: control objects such as parallelepipeds with no apparent function are attributed fronts following two strategies: “interaction” or deictic attribution (selecting the side facing the subject, expected) and “motion” attribution (selecting one of the two smaller lateral sides in case of an elongated parallelepiped, unexpected). These two strategies thus appear to replicate the two main classes of functional frontal attribution. They act as surrogates because the experiments proved that both types of real functional effects, static and dynamic, outweigh the geometrical and presentation effects involved in either substitute strategy. Finally, we proved that ILNs differ in their sensitivity to the static and dynamic types of functional features on stimuli involving objects presenting both—but diverging—static and dynamic functional features. While both *avant* (front) and *devant* (front surface) prefer a front induced by motion rather than interaction, *devant* is more sensitive than *avant* to static interaction. Such a result is in line with the conclusions of an experiment that proves that a pair of semantically similar English prepositions *above* and *over* differ in their sensitivity, this time to functional features (favored by *over*) and geometrical ones (favored by *above*) [CPSR01].

The last pointing experiment (whose results are not yet fully exploited) aimed at proving that ILNs do not simply denote parts of objects, but present some of the characteristic features of locations, as their linguistic behaviour shows (combination with preposition *at*, among others) [Aur96]. Like locations (see Section 5.1), the entities denoted by ILNs are parts of material objects that are fixed in the framework constituted by the object, and have an adjacent space portion associated to them: *the fly is at the corner of the mat* ambiguously describes either a situation in which the fly is on the

mat's corner or is next to it. The experiment therefore compared the sets of pointing touches for ILNs and component nouns denoting spatially coinciding object parts, and ILNs contrasted with components in yielding point clouds extending beyond the object's boundaries.

The project in which this work was carried out generally aimed at analysing the role of the categories of spatial entities in processing spatial expressions, a role that had been little investigated [AHV05, AHV07b]. The multidisciplinary team aimed at confronting the several viewpoints of linguistics, psycholinguistics, formal semantics and formal ontology to gain new insights; part of the results were presented during a final conference which gave rise to an edited book [AHV07a]. A central issue was to understand the significance beyond psycholinguistics of the debate between universalism (cognitive categories and concepts are innate and universal) and linguistic relativism (categories and concepts depend on one's native language — the Sapir-Whorf hypothesis) [GL96]. Obviously, this issue affects formal ontology studies only for those that take a cognitively-oriented stance. This famous debate is currently being reconsidered at the light of new findings regarding infant's cognitive capacities and language development and an intermediate position is establishing itself as consensual. This moderate position roughly posits that universal basic categories and concepts constitute a core knowledge that is either innate or emerge very early on the basis of prelinguistic elementary perceptual and motor capacities. Language intervenes in the next developmental phase of construction of more elaborate categories and concepts, selecting relevant ones only [Spe03] (see also [Aur04] for a discussion and [Jac96] for a quite similar position, not based on experimental data). The large diversity of concept systems that can be observed would then rest on a common basis. Formal, foundational, ontology addresses precisely categories and relations at a general and basic level. Paying attention to the ontological foundations of lexical semantics, even observing a few languages only, might give results of some general value. Indeed, the ontological system that I proposed originally in my PhD for dealing with a few prepositions of French only, and then refined in collaboration with Michel Aurnague (see Section 5.1), proved adequate to handle spatial concepts as expressed in Basque, a totally unrelated language [Aur04]. This is not to say that we should necessarily aim at a single universal cognitive ontology (most probably data about core knowledge or core knowledge itself are too underspecified to decide on fine-grained axiomatic options), but rather to note that working with empirical linguistic data doesn't condemn one to the idiosyncrasies of a specific language if one does not stay on the surface of the lexicon. And that it is worth also paying a close attention to empirical data on infant's conceptual system that are being gathered, now freely from long-dominant Piagetian theoretical positions. However, the results available do not yet clearly point at the categories that would constitute an adequate foundational or top-level ontology; for instance, the five core knowledge systems currently identified [SK07] do not suffice to explain the categorization capabilities observed in pre-linguistic children [Pau02a].

3 A compositional semantics interlude: Locating adverbials

The observed similarities of the linguistic expressions used for spatial and temporal location processes have repeatedly induced an hypothesis of a metaphorical transfer of spatial concepts to the temporal domain. Many prepositions, for instance, have both a spatial and a temporal meaning (e.g., *à midi/Toulouse* (*at noon/Toulouse*), *dans deux minutes/kilomètres* (*in two minutes/kilometers*)). Without taking issue on the metaphor hypothesis, let's observe that the class of "locating adverbials", i.e. those adverbials that have the capacity to locate events and more generally eventualities³² either in time or in space, is very homogeneous. Temporal and spatial³³ locating adverbials have a similar internal structure, from which their semantics can be calculated with global compositional rules, and they play a common locating role in the sentence.

In collaboration with Myriam Bras, Michel Aurnague and Nicholas Asher, I studied these locating adverbials from their syntax to their role in discourse through their compositional semantics. A preliminary series of studies [AAB⁺94, AAB⁺95, AABV95] provided a first semantic analysis and showed for the first time the importance of their role in establishing what we called the spatio-temporal structure of discourse. A subsequent effort aimed at improving in all directions our original analyses, starting with establishing a compositional semantics base [ABVA01] that is the object of this chapter, followed by a revision of our discursive analysis which will be discussed in Section 4.6 of the next chapter.

Here, I first clarify what we mean by locating adverbials and what is the scope of this study. I then examine important semantic issues related to the syntactic position of these adverbials in the sentence. I conclude by a discussion of the relational approach followed, a non-standard choice.

3.1 Scope

The study at the syntax-semantics interface presented in [ABVA01] built on our previous work. That is, the work on spatial prepositions presented above and, more importantly, the work of Myriam Bras on the semantics of temporal locating adverbials, which are complex expressions such as *le lundi suivant* (*the following monday*), *depuis deux jours* (*for two days*) or *quelque temps après* (*some time later*) that locate eventualities [BM93, Bra08].

Moving from the study of spatial prepositions to that of locating adverbials implies that I extended my domain of study simultaneously in several directions: i) from the spatial domain to the temporal one,³⁴ ii) from the simple use of prepositions to complex

³²Following part of the literature I will use the term eventualities to denote temporal entities such as events and states, the two categories of perdurants that are used in DRT and SDRT (see next chapter).

³³In the literature, one often finds the term "locative adverbial" instead of "spatial adverbial". I refrain from using "locative adverbial" because I want to focus on those adverbials, either temporal or spatial, that are "locating" (a term used in e.g., [KR93]), and the two terms may be confusing.

³⁴It will be seen in Section 4.6.2 that the shift covers also the spatio-temporal domain.

adverbial phrases, possibly including distance modifiers and possibly recursively built, such as in *trois mètres à gauche de la mairie* (*three meters to the left of the town hall*), *deux jours après la réunion d’avant les vacances* (*two days after the meeting before the holidays*) and *jusqu’après Noël* (*until after Christmas*), and iii) from locating spatial entities in sentences of the form “NP_{located-object} to be P NP_{reference-object}” as seen above to locating eventualities through prepositional phrases used as VP or sentence modifiers.

No previous work had attempted at a common compositional account of temporal and spatial locating adverbials of this degree of complexity.³⁵ The semantics of temporal adverbials nevertheless is an area quite well explored but, apart from [BM93], the subtleties of the phenomena involved in the quite basic mechanism of location have been rather overlooked in recent work aimed at describing the compositional semantics of adverbials. The most detailed work of this kind is dedicated to combined effects of tense, aktionsart and (essentially duration) adverbials [Hit93] (after [Dow79]), to adverbials that in fact are subordinate clauses like *before John came* [Joh94],³⁶ and to temporal quantification phenomena in prepositional phrases like *during every meeting* or in habitual clauses like *when I take a shower* [dS91, Joh94, PF01]. Only the latter examine compositional semantics issues internal to the adverbial itself, but regarding quantification and not so much locating features, especially not anaphoricity. An exception in this picture is [Mói01], which partly builds on our own previous work. As far as spatial adverbials are concerned, much less proposals are available. [Mai95] focuses on the different roles locating spatial adverbials may have in the sentence (see just below), but considers adverbials with a simple internal structure.

We focused on possibly complex adverbials that locate unique eventualities in a specific way. This means that we considered adverbials of location and not those of measure (duration or spatial extent), without quantification, and that do not modify plural events. The adverbials we studied have the general syntactic structure: NP_{spec} P NP_{comp}.³⁷ The NPs can be of two types: of type “loc” if denoting a temporal or a spatial location, and of type “ext” if denoting a temporal or spatial extent, i.e., a measure. In the pattern above, the NP_{spec} always is of type ext, and the NP_{comp} most often is of type loc (when it denotes the reference object) but type ext is also found: compare *depuis midi* (*since noon*) and *depuis deux minutes* (*since two minutes ago*). Accordingly, several compositional patterns combine the semantics for the P and the different NPs into a homogeneous semantics for the adverbial. The locating process often requires the resolution of an indexical or anaphor in the NP_{comp} as in \emptyset *lundi* (*on/∅ Monday*), \emptyset *le lendemain* (*the next day*), *après ∅* (*after*) or *deux km plus loin ∅*

³⁵Since then, Francez and Steedman [FS06] have proposed an extension of the study of temporal locating adverbials in [PF01] to cover both types of locating adverbials, called in that paper “contextual”, but they still focus on quantification issues.

³⁶In a discourse approach like ours, adverbial clauses are considered as particular discourse segments.

³⁷Even though *plus tard* (*later*) and *plus loin* (*further*) are not prepositions but adverbs, they enter the same pattern with an implicit NP_{comp}, and contribute in the same way to the semantics, so we took them into consideration. The NP_{spec} is optional, but the NP_{comp} can have a null realisation, and even the preposition itself can be empty: *quelques mètres à droite de la maison* (*a few meters to the right of the house*) is prototypical, but we also have \emptyset *lundi* (*on/∅ Monday*) and *depuis ∅* (*since then*).

(*2 km further*), so we did pay an important attention to anaphora and deixis phenomena, and more generally to discursive effects. The discursive dimension is absent from the literature mentioned above, but is taken in consideration in [Hin86] and of course in [KR93]. Unlike the DRT treatment of temporal locating adverbials in [KR93], our study uses standard compositional semantics at the sentence level, to be in a second step situated in the framework of SDRT, as will be seen in Section 4.6.1.

3.2 Syntactic position in the sentence

Our proposal also differed from [Hin86, Dow82, KR93] and from most traditional works on temporal or spatial adverbials which disregarded the effects of possible different syntactic positions of the PP in the sentence. Assuming that PPs usually are VP modifiers and have a VP-adjunct position, syntacticians have shown that preposed PP are not VP-adjuncts that have moved, but IP-adjuncts, that is, sentence modifiers (see, e.g., [Rei81]).³⁸ Along with [Mai95, Hit93, Hit97, dS99, Fre03, Ver04]³⁹ we observed variations in meaning induced by these two possible syntactic positions, VP adjunct and IP adjunct, that needed to be accounted for.

While in VP-adjunct position a locating adverbial plays its standard role of locating the eventuality described by the VP, in IP-adjunct position, it modifies a clause and so says when or where the proposition denoted by the clause is true (see, e.g., [Mai95]). The induced differences are visible in some cases in the semantics at the sentence level.

First of all, and this is a generic phenomenon, adverbials in IP-adjunct position do not fall under the scope of negation, while those in VP-adjunct position may do so [dS99]:

- (4) a. *Marie ne mangea pas à 8h (mais à 10h)* (Marie did not eat at 8 (but at 10))
 b. *A 8h, Marie ne mangea pas (mais à 10h *Ø/si)* (At 8, Marie did not eat (but at 10 *Ø/she did))

The adverbial position also has effects on the semantics of tense, especially the past perfect which introduces both an event and its resulting state. In the following examples, the VP-adjunct locates the eating event, while the IP-adjunct preferably specifies when the state of having already eaten is going on:⁴⁰

- (5) a. *Marie avait mangé à 8h* (Marie had eaten at 8)
 b. *A 8h, Marie avait mangé* (At 8, Marie had eaten)

³⁸Preposed adverbials are sentence modifiers, while postposed adverbials are ambiguous between VP modifiers (the usual case) and sentence modifiers [Joh94, Mai95]. The view adopted here assumes adverbials are adjuncts, but the syntactic function of adverbials is not consensual, and the debate goes on [AER04]. Note that spatial PPs may also appear as verb complements, as in *Léa est allée jusqu'au Pôle Nord* (*Lea went up to the North Pole*) or *Paul put the suitcase in the bedroom*; this latter case falls out of the scope of the work discussed here.

³⁹The suggestion that these different syntactic positions have different semantic effects actually appears much earlier [Kun75, Cha84], even in formal semantics [Par84].

⁴⁰Contrarily to what is sometimes said, the reading in which the IP-adjunct specifies when the eating event has occurred is also available in some contexts with a series of past perfect clauses, as in *A 8h, Marie avait mangé. A 10h, elle était allée au lit.*

In addition, for [Hit97] some temporal adverbials in IP-adjunct position take on a *specific* meaning, that is, they refer to an anchored time, while this meaning may be left non-specified in VP-adjunct position. Duration temporal adverbials, for instance *for three years*, may acquire a locating flavour: in, e.g., *For three years Mary will live in Amsterdam* the adverbial means “for three years from now” [Dow79]. According to her, locating adverbials would also behave that way: the anaphoric adverbial *at noon* in VP-adjunct position as in *Mary will be in her office at noon* may refer to the noon of some unspecified day, while it is necessarily specified, i.e., the anaphor is resolved, when used in IP-adjunct position as in *At noon, Mary will be in her office*. This phenomenon is not confirmed in French, anaphoric adverbials always call for anaphora resolution. The unspecific reading is unavailable in the French equivalent *Marie sera à son bureau à midi*, perhaps because the explicitly unspecific *un midi* (*at/∅ some noon*) exists in French. In addition, the unspecific reading of *un jour* (*one day*) persists in IP-adjunct position and is even the only one available. In any case, the behaviour of anaphoric adverbials is better analyzed observing the larger discourse context, and this is precisely what we will do regarding locating adverbials in IP-adjunct position in Section 4.6.1.

In fact, the most conspicuous effects of the IP-adjunct position for adverbials occur at discourse level, which is to be expected given their role external to the clause. The preposed IP-adjunct position gives locating adverbials the role of “setting the scene” for a whole discourse segment, thus transforming them in what is called “frame adverbials” (it seems [Kun75, Cha84] were the first to write about this role and to give them this name; see also [Cha97, Cha03, Dic01]). For instance, in (6a), but not in (6b),⁴¹ the three events in the two sentences are clearly located in “that summer”.

- (6) a. *Cet été-là, François épousa Adèle. Jean-Louis partit pour le Brésil et Paul s’acheta une maison à la campagne.* (That summer, François married Adèle. Jean-Louis left for Brazil and Paul bought a house in the countryside.)
 b. *François épousa Adèle cet été-là. Jean-Louis partit pour le Brésil et Paul s’acheta une maison à la campagne.* (François married Adèle in the summer of that year. Jean-Louis left for Brazil and Paul bought a house in the countryside.)

As will be seen in Section 4.6 where we will account for this framing role, the semantic contribution we propose for locating adverbials in their standard position of VP-adjunct is kept identical when they are in IP-adjunct position. That is, we hold along with [Joh94, Hit97] that to maintain compositionality it is unnecessary to assume adverbials are ambiguous. The general form of the semantics of locating adverbials is:

$$(7) \quad \lambda P \lambda x (P(x) \wedge Q(x))$$

where, when used as VP-adjunct, P will grasp the verb predicate, x will be substituted with the eventuality, and $Q(x)$ is the contribution of the adverbial itself. Since we

⁴¹Actually, (6b) seems incoherent, as one wonders what is the link between the event of the first sentence and those in the second. More on this in Section 4.6.1.

adopt a relational approach as argued just below, Q is a predicate of an eventuality which actually is a relation between the eventuality and some reference object (possibly underspecified if the adverbial is anaphoric).

3.3 Relational semantics and referents

All the proposals in the literature on temporal locating adverbials assume a semantics in which an adverbial characterizes an interval of time in which the eventuality is located. The location relation is not always temporal inclusion; to account for aspectual phenomena, the inclusion is sometimes reversed or only an overlap is used [KR93].⁴² The adverbial denotes a temporal interval either directly, as with the null preposition on calendar times e.g. *lundi (on Monday)*, or through a temporal function applied to the reference object, i.e., some temporal referent (an eventuality or a temporal interval), as with *before the explosion* or *after this summer*. Similarly, for spatial locating adverbials, [Mai95] follows the standard conceptual approach to spatial prepositions: the location of a located object, in this case an eventuality, in a region denoted by the adverbial and determined by a spatial function applied to the reference object.

To sum up, the whole literature adopts the referential approach described in Section 2.1. The resulting semantic accounts of locating adverbials are certainly not parsimonious, that is, they assume more referents than necessary. Moreover, they present the same defects as observed for the referential approach to spatial prepositions: they have difficulties in dealing correctly with distance and contact (in time, contact is immediate precedence or succession) because the reference object is no longer available to establish the precise relation denoted by the preposition with the located eventuality. This problem is perhaps less conspicuous for temporal adverbials than for spatial ones. It nevertheless appears for instance in expressions of immediate succession like in *Last month, Lea read the Iliad. She devoured the Odyssey immediately after that* for which it is hardly possible to identify an interval that forces the second reading event to be immediately after the first. It also appears when dealing with “spanning” adverbials built on prepositions like *from*, *to*, *since*, or *until*, as in *From 2 to 3 Brutus watched Caesar*. For instance, [Par90] assumes that the eventualities, events or states, are included in the time interval, thus not accounting for the fact that states fill up the interval. On the other hand, [KR93] assumes that differently from events, states include the time interval, thus ensuring the whole coverage of the interval but inadequately allowing the state to start before the interval begins. [PF97] resolves the problem by quantifying on instants in the interval and asserting that the state holds at each of these instants and only these; this is a costly solution in terms of ontological commitment, as it relies on the existence of instants in addition to intervals. As far as distance modifiers of *before* or *later* are concerned, in principle the distance between the located eventuality and the reference entity cannot be grasped through a mere inclusion, but the situation is less

⁴²Some approaches that are clearly (neo-)Davidsonian like [Par90] do not locate directly the eventualities with respect to these time intervals. Just as some use a “Place” function to indirectly locate the reference object with spatial prepositions, they assume another function that gives the time of any eventuality, and this time is put in relation with the time interval denoted by the adverbial.

clear than in the spatial case. In fact, most distance temporal adverbials based on clock or calendar units convey an implicit granularity. The unit used also vaguely constrains the duration of the event: an event occurring *two days later* is assumed to last no more than a day, otherwise we would rather locate the *start* of the event.⁴³ A corollary is that an eventuality located anywhere in a whole unit interval can be said to be at a distance multiple of that unit: an event located *two days later* can usually have occurred at any hour of the day. These two facts would tend to favour the referential approach, but note that the use of modifiers like *exactly* can enforce a precise measure impossible to ensure with the referential approach.⁴⁴

In addition to providing a more correct semantics, the relational approach avoids assuming the systematic reference to extended space regions and time intervals (or even 0-dimensional space points and instants) whose existence might be questionable. In the relational approach, a direct location is made through temporal or spatial relations between an eventuality (the located object) and a reference object which is either another eventuality, a physical object, a temporal location (also called time), or a spatial location (also called location). Just as spatial locations like *Toulouse* or *my garden* are not regions in some absolute container space, a temporal location or time like *this morning* or *the year 2009* is not just a temporal interval in some absolute container time. A clock or calendar time has non purely temporal properties, something that could be described as “functional” following the terminology used in spatial semantics. The relational approach therefore allows for a more parsimonious ontology that doesn’t assume the existence of absolute space and time, although it is of course compatible with such a richer view. For a discussion on the ontological choices regarding the nature of space and time, see Section 5.1.2.

⁴³While *Lea read the Odyssey two months later* is fine, we would say *Lea started reading the Odyssey two minutes later*. Compare with *The cat is miaowing five cms behind Lea*.

⁴⁴Another argument in favour of a referential approach is presented in [Mói01], namely the fact that adverbials can be constructed recursively. Our own treatment of recursive adverbials like *jusqu’après les vacances (until after the holidays)* involves the coercion of *après les vacances* into a time situated *après les vacances*. This time is shorter than the whole interval denoted by *après les vacances* in the referential approach, and in fact yields a more appropriate reading.

4 Discourse semantics

Addressing the discursive dimension of language when coming from lexical semantics, even formal-semantic oriented and with developments into compositional semantics, may appear as a non obvious move. Such a jump began as a simple contribution to extending to the spatio-temporal dimension the work of Myriam Bras and Nicholas Asher on the interrelationships between discourse structure (the set of rhetorical relations between discourse segments) and temporal structure (the set of temporal relations between discourse referents) in French texts [BA94]. The latter built on the first study, contemporary with the birth of SDRT [Ash93], to overcome the limitations of the DRT analysis of tenses,⁴⁵ demonstrating the necessity and the feasibility of a formal analysis in terms of discourse relations [LA93].

The relevance of the spatial dimension for discourse structure certainly is less obvious than the temporal one which is motivated by the established hypothesis of the anaphoricity of tenses [Rei47]. Yet, we showed in [AAB⁺94, AAB⁺95] that discourse relations such as *Narration* and *Background* which are primarily temporal also involve spatial consequences, and that the role of spatial adverbials in discourse is very similar to those of temporal adverbials. So not only the lexical semantics of spatial adverbials had an impact at discourse level that needed to be analysed, but my knowledge in spatial representation contributed to reconsider the semantics of some discourse relations.

My interest in this new form of linguistic “content” embodied by discourse relations rapidly grew into an interest in the theory SDRT itself. This theory whose development was just beginning was on the one hand anchored in formal semantics (just as DRT on which it was built), and clearly assumed on the other hand that all sort of content, based on grammatical or lexical elements or related to conversational conventions and commonsense knowledge, was relevant to discourse semantics, thus perfectly matching my own inclinations. Contributing to make such an ambitious framework grow and clarify some of its foundational aspects was in itself exciting.

In this chapter, after briefly presenting and situating SDRT in the literature, I will first examine several theoretical aspects to which I have contributed: in Section 4.2, the definition of the subordinating or coordinating nature of discourse relations which determines their contribution to the hierarchical structure of discourse; in the next section, the consequences of the rhetorical nature of discourse relations; in Section 4.4, the role of discourse topics in SDRT and the introduction of a new use for them; finally, the improvement of the modelling of some standard discourse relations and the introduction of a new one. Then, in Section 4.6, I will review the work analyzing the discursive role of adverbials and more generally what we have called the spatio-temporal structure of discourse. Although the latter study started and motivated many of the theoretical developments, the presentation of such developments is a prerequisite to properly describe the current analysis of adverbials in discourse. I will conclude in Section 4.7 by describing current studies that have the potential to make SDRT take a significant new turn and pointing at questions to be answered.

⁴⁵Such limitations were already described in DRT work [KR83, KR93].

4.1 Fundamentals of SDRT

Segmented Discourse Representation Theory (SDRT) [Ash93, AL03] belongs to the dynamic semantics trend stemming from formal semantics and initiated in the early 80s, of which Discourse Representation Theory (DRT) [Kam81, KR93] is a major representative. Dynamic semantics showed that the discourse representations that stand in between natural language texts or dialogues and their interpretations in model theory are not just a technical support that could be dispensed off, like logical forms are for Montague Grammar. Building discourse representations constitute an indispensable step of interpretation if we are to account for a number of phenomena like intersentential anaphora, clearly at the discursive level, as well as quantifier scope issues within the so called “donkey-sentences”, of a more classical flavour in formal semantics. From this new viewpoint on interpretation, meaning is no longer primarily expressed in terms of truth-values, but in terms of *context-change potentials* [Hei83b, Chi94]. The “context” in this view is the representation of the previous discourse, a proposition in dynamic logic. More illuminating perhaps, and in line with the classical DRS “box” representations, is to say that contexts are constituted by a set of existentially quantified *discourse referents* (or a partial assignment function) and a logical formula establishing conditions on those referents (or a set of possible worlds).

I obviously cannot explain here the details of DRT, nor attempt at describing the semantics of DRSs and how they are built. I refer the reader to the main reference on DRT [KR93]. However, as grounds for comparison, I just recall what would be the “box” representation of a simple DRS for a simple text:

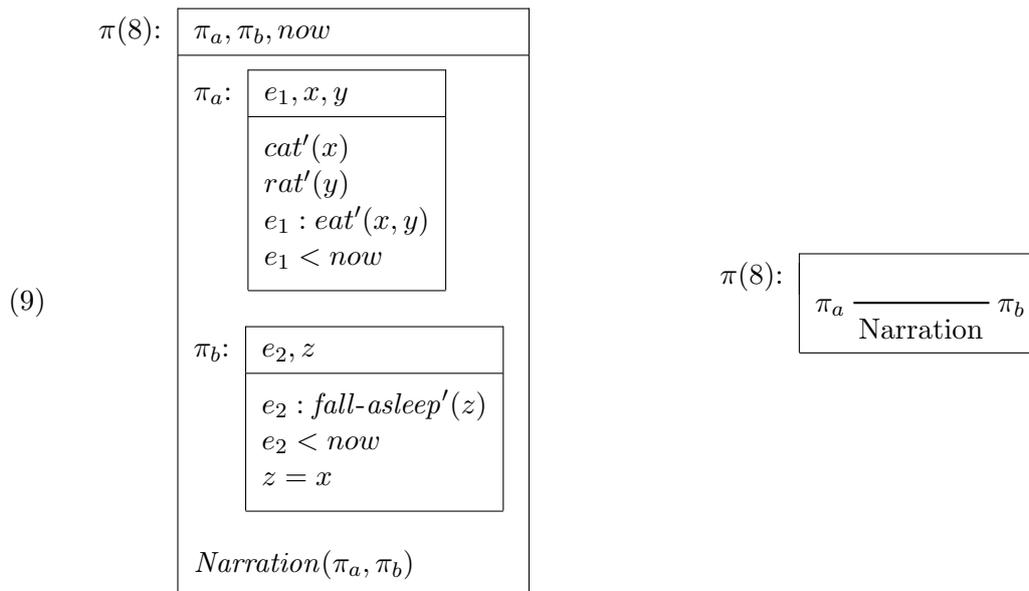
(8) A cat ate a rat (a). It fell asleep (b).

e_1, e_2, x, y, z, now
$cat'(x)$
$rat'(y)$
$e_1 : eat'(x, y)$
$e_1 < now$
$e_2 : fall-asleep'(z)$
$e_2 < now$
$z = x$
$e_1 < e_2$

Largely building on DRT, SDRT adopts a two-step interpretation too. SDRT *proper* characterizes what a discourse representation is, describing the logical language of *Segmented Discourse Representation Structures* (SDRSs), their syntax and their semantics in model-theoretic terms. The *Glue Logic* deals with the construction of the SDRS from a given text or dialogue.

SDRT improves DRT in that it takes into account the fact that a several sentence long discourse is not simply a single structure (one “big box”) but has an internal hierarchical structure. This move is done integrating other approaches to discourse than dynamic semantics [MT88, Hob85, Pol88, GS86], of which the most influential one prob-

ably is Mann & Thompson’s Rhetorical Structure Theory (RST) [MT88]. A discourse (and an SDRS) recursively divides up into *segments* which can range from basic clauses (represented as DRSs) to complex discourse stretches (SDRSs), and are connected by *discourse* or *rhetorical relations*, for example, *Narration*, *Elaboration*, *Result*, *Explanation*, or *Contrast*.⁴⁶ To give a graphical illustration of SDRSs, the previous example (8), which would be a case of *Narration*, would be represented in full by the first graph below; the second graph focuses only on the discourse (or rhetorical) structure of this SDRS.⁴⁷



By establishing a link with the previous discourse, discourse relations represent the relevance of each single segment within the discourse structure, and by contributing their various semantics to discourse content, they make explicit the specific communication intention apparently motivating the utterance of each such segment.⁴⁸ In the example

⁴⁶I will not get into the details of each relation here (but see at the end of this section for a few of them), referring the reader to the SDRT literature. SDRT stresses the need to have a clear semantic characterization for each relation and that distinct relations must contribute differently to the SDRSs in truth-conditional terms. As a result, SDRT adopts a limited number of relations (a few tens). In other theories a large variety of relations have been adopted, with different names for similar relations, and in numbers varying from just two to an unlimited one. See [HM93] for a review.

⁴⁷Here again, this is not the place for a crash-course on SDRT that would explain the details of these graphs. Just three points though. In a SDRS, segments are labelled, π_a, π_b , etc; more on this below in Section 4.3. I omit here the topic required by the Narration relation, cf. Section 4.4. On the second graph, the nodes are sub-SDRS labels, and the lines represent attachment with discourse relations. Of course, SDRT provides a formal definition of SDRSs, not just a graphical representation, together with a proper model-theoretic semantics. See [AL03], and [AL08] for the updated definition of dialogue SDRSs.

⁴⁸SDRT does not belong to the mentalist tradition in dialogue theories and presupposes no mind-reading. SDRT is rather a conventionalist theory, so the intentions behind each speech act are those that any linguistically competent addressee or hearer could interpret on the basis of the discourse only. Such intentions could be different from the actual intentions of the speaker.

above, *Narration* contributes a precedence relation between the main eventualities of the segments related, thus indirectly conveying the DRS condition $e1 < e2$ of (8).

This provides an account of the *coherence* of a discourse. Not any sequence of sentences makes up an acceptable discourse: a discourse is coherent if each new segment is connected to the previous context by a discourse relation. This of course comes on top of the mere interpretability of each segment in context and in particular the resolution of any underspecification (e.g., anaphors), an aspect of coherence sometimes called *cohesion* [HH76]. But as shown at length by SDRT studies, the hierarchical structure of discourse in fact governs anaphora resolution, so both aspects of coherence are intertwined.⁴⁹

Explaining possible and impossible anaphora resolution is one important source of empirical validation for the SDRT approach. Other achievements include lexical disambiguation [AL95], bridging anaphora [AL98a], presupposition [AL98c], and, first of all, the correct account of the temporal structure of discourse, where DRT found its limits [LA93]. SDRT accounts for the temporal inclusion between events by the *Elaboration* involved in (10) and in (11), and for the temporal reversal of events in example (1) repeated below as (12)a-b, by contrasting a *Narration* and an *Explanation* case.

- (10) *L'été de cette année-là vit plusieurs changements dans la vie de nos héros (a). François épousa Adèle (b). Jean-Louis partit pour le Brésil (c) et Paul s'acheta une maison à la campagne (d).*

(The summer of that year saw several changes in the lives of our heroes. François married Adèle. Jean-Louis left for Brazil and Paul bought a house in the countryside. – From [KR83])

- (11) *L'an dernier, Jean escalada le Mont Cervin. Le premier jour, il monta jusqu'à la cabane H. Il y passa la nuit. Ensuite il attaqua la face Nord. Douze heures plus tard il arriva au sommet.*

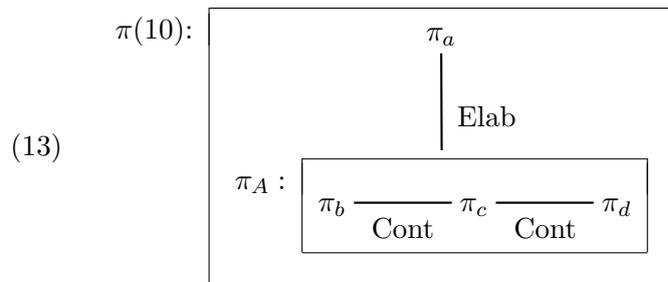
(Last year, Jean climbed the Matterhorn. The first day, he went up to the cabin H. He spent the night there. Then he tackled the north face. Twelve hours later he reached the summit. – From [KR83])

- (12) a. Paul fell. Max helped him up.
 b. Paul fell. Max pushed him.
 c. Paul fell. Then Max helped him up.
 d. Paul fell because Max pushed him.
 e. Paul fell. Then Max pushed him.

The Glue Logic takes as input the SDRS representing the previous discourse and the formula representing the new clause to be “attached”. Such formulas can be obtained using Montague’s Grammar lambda-calculus framework (which is what I assume here),

⁴⁹The hierarchy in the discourse structure, i.e., the dominance between segments results from the embedding between segments and sub-segments as well as the *subordinating* character of some discourse relations, as opposed to the *coordinating* character of others. On graphs, subordinating relations are depicted as vertical or oblique lines, while coordinating relations are horizontal lines (cf. the graph in (13)). See Section 4.2 below. Note that contrarily to what assumed in RST [MT88], in Polanyi’s LDM [Pol88] and in D-LTAG [WKJ01, FRWJ06], the discourse structure in a SDRS is not a tree but a graph.

adapted to take underspecification into account [Rey93]. Updating the context with the new clause consists in finding an open node in the SDRS, inferring the discourse relations concretizing the attachment via rules of the Glue Logic and resolving underspecifications such as anaphors. The open nodes are determined by the “right frontier” [PS84, Pol88] of the hierarchical structure (see Section 4.2 below). For instance, on the graph for (10), represented in (13), only π_a , π_d and π_A are on the right frontier and open.⁵⁰



Sometimes a discourse marker signals in a definite way a relation, as in (12c-d) with *then* (*Narration*) and *because* (*Explanation*), but in many cases the relations are left implicit, as in (12a-b) which are discursive equivalents of (12c-d). However, implicit relations can be ruled out by a contradictory explicit marker ((12e) clearly differs from (12b)) so inferences are done in a non-monotonic logical framework.

The strength of SDRT stems from its ability to integrate formal semantics, formal pragmatics and lexical semantics and account for their interactions, getting closer to the human processing of linguistic interpretation. A great attention to cognitive plausibility, especially as far as computational complexity is concerned, is paid in [AL03]. On the other hand, the complexity of SDRT makes its presentation not straightforward. In addition, this ambitious theory has been evolving and is still work-in-progress since in particular neither is a complete set of discourse relations fully investigated yet nor are the associated triggering rules of the Glue Logic fully spelled out.⁵¹

In an attempt to overcome these presentation difficulties, I have contributed in several occasions, via seminars, tutorials and articles to make more accessible the original proposal in [Ash93] and emphasize the progress made in [AL03]. For instance, [BVA01] provides a simplified and more systematic account of SDRT, with some technicalities (especially the update procedure) corrected. I also contributed to point out methodological

⁵⁰It is possible to attach to complex segments like π_A only with some relations, like *Contrast*, *Result* and *Explanation*. For many relations, e.g., *Narration* and *Elaboration*, it is not possible to attach to π_A .

⁵¹The rules of the Glue Logic are based on a set of varied clues based on the lexicon and commonsense knowledge whose range and dimension are unknown. To make progress and keep the interaction between the rules manageable, these rules are divided into three categories. First of all, discourse markers give rise to undefeasible rules that directly infer the relation indisputably marked. Other undefeasible rules, in large numbers, simply conclude to the presence of clues of some kind. Then, defeasible rules jump to the conclusion of the most plausible discourse relation in a given context on the basis of such families of clues. The latter rules, in reduced numbers thus manageable, control the interaction between possibly contradictory clues. Obviously, regularities in the lexicon and relevant knowledge on which the second category of rules rests need to be investigated to engage towards a full theory. A first attempt at this is made in [AL03, Chap.6].

issues that still remain unsolved [PV08].

More concretely, I have worked on clarifying several fundamental aspects which remained underspecified or unclear, contributing to develop some missing features. To these aspects and features I now turn.

4.2 Subordinating and coordinating relations

Most theories of discourse structure assume that hierarchy in discourse results from a distinction within discourse relations, on top of the embedding between segments and subsegments.⁵² The intuitive motivation is that some segments of a discourse play a subordinate role relative to previous, separate, segments they are connected to, while others are considered on a par. SDRT along with Hobbs calls this distinction subordinating vs. coordinating, RST calls it nucleus/satellite vs. multinuclear, and Grosz and Sidner reduce all discourse relations to just this distinction calling their two relations dominance and satisfaction-precedence. Since the hierarchical structure of discourse is based on it, this distinction has significant consequences. Indeed, the analysis of discourse coherence exploits the “right-frontier constraint” [PS84, Pol88], the principle that stipulates that a new basic segment (a new clause) can be attached only to segments appearing on the *right frontier* of the graph representing the discourse, where segment embedding and subordinating relations yield top-down edges⁵³ while coordinating relations appear as left-right edges, and the last basic segment determines the bottom end of the right frontier. This principle explains the incoherence of texts like (14a-b-c) and the coherence of its variant (14a-b'-c):

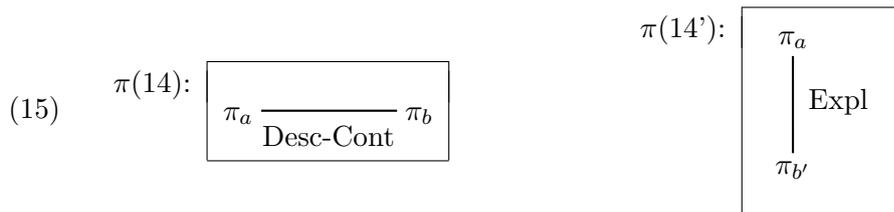
- (14) a. Arts-and-Crafts jewels tend to be elaborate.
 b. They are often mass-produced.
 b'. Ornateness was the fashion at the turn of the century.
 c. However, this jewel is simple in form.
 (From [KOOM01])

(14c) clearly needs to be attached by *Contrast* (marked by the adverb *however*) to (14a) to make sense. The relation between (14a) and (14b), being coordinating (a *Description-continuation*⁵⁴), pushes (14a) to the left out of the right frontier; the unavailability of (14a) for attaching (14c) explains the incoherence. This doesn't happen with (14b') which attaches to (14a) with a subordinating relation (an *Explanation*). The difference between these two graphs is illustrated in (15).

⁵²An exception is D-LTAG [WKJ01, FRWJ06] whose trees take only embedding into account.

⁵³To avoid confusion between the two kinds of domination sources, I favour a hybrid graphical notation which I adopt here and in which the embedding appears as inclusion within larger boxes and not as edges. See, e.g., the graph on (13).

⁵⁴See [PVA09] for reasons to introduce this relation in SDRT, as distinct from *Continuation*.



The grounds for categorizing discourse relations as coordinating or subordinating were quite vague and varied in the literature, and in SDRT, the type of each relation was attributed on an intuitive basis. In [AV01, AV05] we proposed a set of four converging⁵⁵ tests checking whether a given relation occurs in specific linguistic patterns (e.g. patterns involving anaphora) or not. This strategy avoids relying on the semantics of the relations to define a non-semantic feature. The subordinating-coordinating distinction indeed manifests itself at the “information-packaging” level, i.e., the way the speaker presents the information, which in SDRT is grasped by the discourse structure. Our strategy rather relies on discourse (counter)examples that make such manifestations explicit like in (15), exploiting fundamental and established structural principles like the right-frontier constraint.

The tests were proved sound with respect to the widely assumed hypothesis that *Narration* and *Elaboration* respectively are prototypes of coordinating and subordinating relations. But these tests had unexpected consequences. They showed that the coordinating character of a relation like *Result* should be taken as a default only, since further information-packaging tools like coordination markers and punctuation can force them to be subordinating on some occasions [AV05]. This confirms that the subordinating/coordinating distinction is not strictly dependent on the semantics of the relations. The tests also showed that a relation like *Background* which in SDRT [Ash93, LA93, AAB⁺95, AL03] was assumed to be coordinating⁵⁶ actually is subordinating (proof independently developed in [VP04] and [FHRSB05]), in line with the hypothesis of RST that the corresponding relations of *Background* and *Circumstance* are of the nucleus/satellite type. Other authors have applied these tests to new SDRT relations, e.g., *Afterthought* in [AK08].

4.3 The rhetorical nature of discourse relations

Information-packaging (or structural) features like those we’ve just examined show that discourse relations are not purely semantic. For instance, *Narration* cannot be reduced to temporal precedence between the events described by the segments related, and *Result*

⁵⁵In practice the four tests all aim at checking whether the attachment site still belongs to the right-frontier after the attachment. They can be considered as variants of a same global test, and in fact, we never found any divergence.

⁵⁶In [AAB⁺95], observation of *Background* examples regarding temporal structure and the accessibility of referents already showed that *Background* wasn’t a simple coordinating relation. However, instead of concluding to its subordinating nature, we introduced a quite ad-hoc and complex topic structure called the Focus-Background Pair or FBP. [VP04] compares the two solutions and opts for the subordinating solution. See though Section 4.5.2 below for subtleties regarding the reverse version of *Background*.

and *Elaboration* are not just causation and parthood between events. This is why SDRT models the relationships between discourse relations and their “semantic effects” as implications and not equivalences. For *Result* and *Explanation* such rules are (slightly simplified):

$$\begin{aligned} \textit{Result}(\alpha, \beta) &\rightarrow \textit{cause}(e_\alpha, e_\beta) \\ \textit{Explanation}(\alpha, \beta) &\rightarrow \textit{cause}(e_\beta, e_\alpha) \end{aligned}$$

where α and β are variables for segments and e_α and e_β the “main eventualities” referred to in these segments.

In fact, if one would assume such a reduction, *Explanation* would simply be a *Result* with switched arguments, but the two discourses in (16a-b) are not equivalent. They generate very different graph structures (*Explanation* is subordinating while *Result* here is coordinating), which is proved by the necessary change in where pronouns appear and by the continuation (16c) referring to the pushing, possible only with (16a).

- (16) a. Paul fell. Sue pushed him. / Paul fell because Sue pushed him.
 b. Sue pushed Paul. He fell. / Sue pushed Paul. As a result, he fell.
 c. But it wasn’t on purpose.

A further reason for not identifying discourse relations with their semantic effects is their rhetorical nature. Relations relate *utterances* or *speech acts*, and they characterize the rhetorical role of a speech act in discourse, just as assumed in RST [MMT92]. They reflect the rhetorical intentions of the speaker, his or her options in presenting the content of the discourse, his or her choices to focus on some links rather than others, in other words, additional claims or *public commitments*⁵⁷ on top of those associated with the isolated utterance of an assertion [AL08]. This is why segments in SDRT are not simply dynamic predicate logic propositions. Each dynamic proposition is *labelled* by a different label which is the trace of a speech act. In SDRT’s terminology, segments are represented as *labelled constituents*, the constituents being occurrences of propositions. Discourse relations then do not simply relate propositions and add extra semantic conditions on their referents (as they would if they were reduced to their semantic effects). They relate utterances, i.e., the labels, and characterize the rhetorical role of these utterances. The speech act so characterized in a relation $R(\alpha, \beta)$ is β , the segment which is attached to some previous segment α . This is why discourse relations are asymmetric: $R(\alpha, \beta)$ cannot be identical with any $R'(\beta, \alpha)$ which would be characterizing the speech act of

⁵⁷SDRT adopts a non-mentalistic approach to discourse and dialogue, so as to avoid relying on impossible mind-reading, and rather focusses on what a discourse objectively conveys. Beyond this stance, [GHL06] has shown that the classical account of speech acts in terms of beliefs and intentions of the speaker and the addressee runs into problems when sincerity is abandoned. The authors, although mainly focussed on dialogue, advocate to use public commitments [Ham70, WK95], or grounded information [Tra94], instead, a position now adopted in [AL08]. I show later in this section that public commitments are an essential component of rhetorical relations in monologue as well. I will use indistinctly the terms “commitment” or “public commitment”, thus ignoring the so-called “dark-side” commitments, i.e., private commitments, introduced in [WK95].

α .⁵⁸

This view on discourse relations is unfortunately not widely adopted, but is quite clear in most writings on SDRT. However, it has surely not been emphasized enough since several authors using SDRT, including Asher and Lascarides themselves, have made proposals incompatible with this view. For instance, [BLDA06] proposed to define the relations of *Result* and *Weak-Result* in terms of causal relations on eventualities. Asher and Lascarides also proposed in [AL03, Chap.6] along with [Dan06] that causative verbs and those Danlos calls “discourse verbs” like *to precede*, should be given a semantics in terms of discourse relations, thus considering the two discourses in (17) as equivalent.

- (17) a. Ted left. Then Sue arrived.
b. Ted left. This preceded Sue’s arrival.

Such proposals amount to mix the discourse and propositional (or clause) levels,⁵⁹ and in fact it turns out to be inadequate. Continuations again show that, unlike the *Narration* in (17a), the relation involved in (17b) is subordinating (likely a *Commentary*): (17b) can be continued by *And it followed Max’s finishing up the wine*, while a continuation of (17a) with *Max had finished up the wine* is unable to give the same reading, i.e., the same temporal order.

Beyond such structural evidence, the linguistic phenomenon of “blocking” that I investigated in [Vie07] brings the rhetorical role of discourse relations to light. Blocking in discourse occurs when the semantics of a linguistic marker blocks the inference to discourse relations that would hold in its absence. For instance, we have shown in [BLDV01a] that the adverb *puis* (roughly equivalent to *then*) blocks *Result*:

- (18) a. *L’acide tomba dans le liquide. Une explosion se produisit.*
(The acid fell into the liquid. An explosion happened.)
b. *L’acide tomba dans le liquide. Puis une explosion se produisit.*
(The acid fell into the liquid. Then an explosion happened.)

The *Result* relation, present in (18a), is absent from (18b). The reading obtained in this second variant is, like in (18a), that the two events occurred in sequence,⁶⁰ but, unlike in (18a), that the speaker presents these events without any commitment regarding their causal relationship. In other words, the speaker conveys something like “I don’t want to claim that the two events are causally related”, and commits not to claim anything in this respect.⁶¹ However, the discourse in (18b) may still truthfully describe

⁵⁸This observation applies to independent clauses; when relations attach subordinate clauses to their main clauses, this may no longer hold.

⁵⁹Defining discourse relations in terms of their semantic effects makes the discourse structure collapse, leaving only propositional content. Introducing discourse relations where there are none makes distortions in the discourse structure and alters the associated commitments.

⁶⁰*Puis* marks *Narration*, whose semantic effects include precedence between events —see Section 4.5.1—, and causation between events also implies precedence.

⁶¹Indeed **Puis le mélange réagit en explosant* (Then the mixture reacted by exploding) is contradictory, and continuing (18b) with *En fait, l’explosion fut provoquée par le mélange* (Actually, the explosion was caused by the mixing) is possible but requires a marker of *Contrast* (*en fait* / actually). Of course,

a world in which the two events do happen to be causally related, thus it is not the case that $\neg Result(a, b) \rightarrow \neg cause(e_a, e_b)$, a further proof against the equivalence. There is a significant difference between conveying (or even asserting) “I don’t want to claim that the two events are causally related” and asserting “the two events are not causally related”, and even not saying anything in this respect.

This difference clearly reveals the rhetorical role of discourse relations. And this is the case not just for those relations that clearly involve the speaker’s mental attitudes or their speech acts in their semantics, as “cognitive-level” relations (e.g. *Question-Answer-Pair* and other dialogue relations [AL98b, LA09b]) and “metatalk” relations (e.g. *Result**)⁶² do, but applies to the standard veridical “content-level” (or “semantic”) relations like *Result*. The relation *Result* conveys not simply causation between events but the public commitment of the speaker regarding such a causation. Disputes in dialogue of course challenge the commitments of the other speakers, and dialogue modelling requires a proper account of the evolution of commitments [LA09a]. It is interesting to observe here with the blocking phenomenon that even for monologues involving standard content-level relations only, in which sincerity is not questioned, ignoring the speaker’s public commitment involved in discourse relations by reducing them to their semantic effects puts the theory into trouble.

Note that the case of *puis* and *Result* is not isolated. Similar blocking effects have been observed in [GT03] for the conjunction *and*, and in [Tab06] for the adverb *anyway*.

Assuming the separation between the rhetorical level and the propositional level implies that discourse relation markers and all direct markers for discourse structure cannot be found inside clauses, but externally to them, for instance as clause modifiers, i.e., IP-adjuncts. This is clearly the case of standard connectives like *puis* and *mais*, and also the case of the frame locating adverbials that will be examined in 4.6.1. More evidence in favour of this hypothesis is given by the fact that causal verbs, whose semantics can be seen as equivalent to the semantic effects of discourse relations, are *not* markers for discourse relations, contrary to what is put forward in [Dan06] and as seen above. But this still is a coarse hypothesis, and because of the self-referential nature of language, it is to be expected that some clause elements indeed can refer to the discourse structure. Exploring the limits of this hypothesis and doing full justice to the rhetorical nature of discourse relations is part of my project for future work described in Chapter 7.

not only rhetorical relations involve commitments, any assertion includes a commitment to the truth of its propositional content. But the blocking phenomenon shows that a single commitment operator over a whole discourse as proposed in [AL08] cannot do. Blocking has access to the commitment operator within discourse by stacking a commitment on top of the negation of a commitment to ϕ , and this is something different from simply not committing anything, i.e., not saying anything about ϕ . So, in writing $\neg Result(\alpha, \beta)$, the negation operator should modify a commitment operator, that is, commitment is involved *in* the interpretation of the discourse relation.

⁶²*Result** appears in *It’s cold. Please close the window*, for the state described in the first sentence causes the request, i.e., the *speech act* of the second sentence. Such relations are called “pragmatic” in [San97, Kno01], as opposed to “semantic” ones; see also [JR01].

4.4 Discourse topics and forward-looking segments

An essential feature of SDRT is to take into account implicit discourse relations between segments, i.e. relations that are not explicitly marked with a connective (*because*, *but*, *then*, etc.). This is in line with most discourse theories and with the observation that 60 to 70% of discourse links are unmarked [Tab06]. What is more distinctive of SDRT is that it also introduces in the discourse representations implicit segments, called “discourse topics”.

Topic is a word that refers to various notions, well-studied and well-defined as far as sentence topic is concerned, but quite vague and much less consensual for discourse topic. When trying to make this latter notion more precise, three major views have been adopted: discourse topic could be a *question* under discussion, an *entity* under discussion, or a *proposition* summarizing a discourse segment (see [Ash04, Ste04] and the other papers on discourse topic in the same issue). Some authors have proposed that topics play a central role in discourse [vK95], while in SDRT, their use is motivated by some specific phenomena only.⁶³ Since [Ash93] SDRT argues for the need to include propositional topics within the SDRSs to account for some semantic and structural phenomena related to the relation *Narration*.⁶⁴ Semantically, *Narration* requires not only temporal sequence between two events (or actually a more complex temporal relation, see 4.5.1 below), but that the two events belong to a same “story”: they share some common theme or topic, which in SDRT is represented as a basic clause, i.e., a labelled simple constituent. This common topic may be explicitly given in the discourse by a previous segment elaborated by those related by *Narration* as in (19a), or it might be left implicit as in (19b).

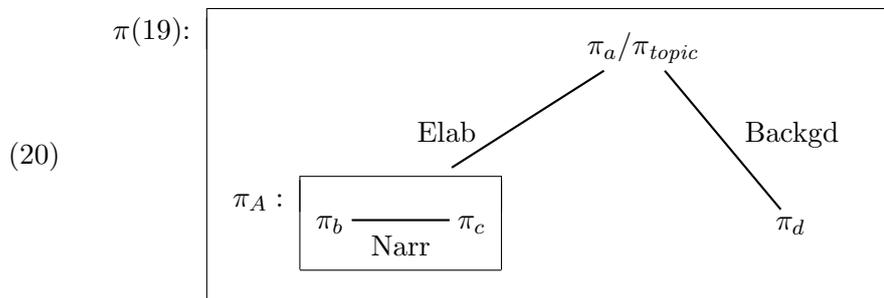
- (19) a. Alice met friends last night (a). She had dinner with Lea (b). Then, she went to the pub with Bill (c). They had a nice time (d).
b. Alice had dinner with Lea (b). Then, she went to the pub with Bill (c). They had a nice time (d).

Structurally, evidence from a variety of phenomena, for instance plural anaphora (structural parallelism is another case, see [Ash04]), shows that an implicit topic requires to be inserted as a constituent in the SDRS: In (19b), to be able to obtain the reading in which *they* refers to *Alice + Lea + Bill*, a segment similar in content to the first sentence of (19a) needs to be recovered and the last sentence (d) attached to it (otherwise, the attachment is made to the previous sentence (c), the only open node, and *they* resolves to *Alice + Bill*). This common discourse structure of (19a) and (19b) makes this point

⁶³Actually, one may consider that all segments dominating others via a subordinating relations have a topic-like role, so the phenomenon could be very much widespread. Those dominating others via an *Elaboration* relation are clearly acknowledged as such in SDRT.

⁶⁴*Background* is another relation for which the insertion of some sort of discourse topic (the “focus-background pair”) has been proposed when it was still considered as coordinating [AAB⁺95]. Even though a semantic constraint of having a common “topic” remains, with a subordinating account there is no longer a structural need to insert a new constituent, at least for the standard *Background* case. See Section 4.5.2 however for the reversed *Background* case.

clear:⁶⁵



Discourse topics may appear as theoretical artefacts generated by attachment and anaphora resolution principles intrinsic to SDRT. Indeed, they rather induce skepticism (see the papers discussing [Ash04] in that same issue, e.g., [Keh04]). Although this will probably not convince those skeptics, I have contributed to show that yet another use of discourse topics as subsuming a question and its answers brings further technical advantages to correctly account for the structure of dialogue [PMDV02, PV08].

Linguistically more compelling evidence comes from a quite different use of implicit discourse topics that I and other colleagues introduced in order to handle anticipation phenomena in discourse. Implicit *Narration* topics and *Question-Answer-Pair* topics are all constructed summarizing two segments when these are attached together; such topics are elaborated by *past* discourse (although possible continuations yet to come may trigger an update). However, topic changes are often signalled in discourse with “forward-looking” markers. This is the case of locating adverbials in IP-adjunct position as introduced in Section 3.2 (see example (6) on p.34), and analyzed in Section 4.6.1. We will see there that these locating adverbials mark a topic change and introduce a “frame” that locates a new discourse segment, possibly encompassing more clauses than the one modified by the adverbial. In other words, the discursive role of locating adverbials like *in 1983* or *that summer* when preposed to a clause as in example (6a) roughly is: “I’m going to tell you something new that happened in the summer of that year”. The semantic and structural aspects related to both the introduction of a new topic and the distribution of a locative property over a (possibly complex) segment to come point at a new use of discourse topics in SDRT, which I call “forward-looking” discourse topics. Such topics are similar to topics used for summarizing past discourse: they are propositional topics that summarize the (possibly complex) segment attached to them via *Elaboration*. The difference is that when the topic segment is first introduced, it is waiting for a segment—that is not yet present in discourse—to be attached to it and which it will summarize, and its propositional content (that “something new”) is still unspecified. Like all implicit topics, this propositional content is updated after the

⁶⁵I consider myself the relation of *Topic*, often used in SDRT to link a topic to the segment it dominates, to be simply the relation of *Elaboration*, since they have the same semantic effects. That an implicit topic should be updated when the segment that elaborates it grows actually is a property of the topic segment itself—it comes from the fact that this segment has no proffered content—and not of the discourse relation that links the segment it subsumes to it.

attachment of each new basic segment elaborating it.

This strategy proved adequate also to deal with the forward-looking (or reversed) *Background* relation, that is, the one in which the background is described before the foreground, as will be seen in Section 4.5.2. Yet other forward-looking markers have been studied by Myriam Bras and colleagues in the category of “structural markers” like *d’abord* (*firstly*), which are frame adverbials too, but not locating ones [Bra08].

4.5 Study of some Discourse Relations

Discourse relations are fundamental ingredients of SDRT and other discourse theories, and yet, they have been rarely deeply investigated for their own sake; they are usually peripheral to the study of discourse markers and connectives, which on the contrary are sometimes quite thorough [Sch87, JR99]. Most often, for each discourse relation, a few characterizing features, a textual description and some examples are given [SSN92, HM93].⁶⁶ Some works are exceptions in this picture, trying to formally characterize either the semantics of the relations themselves or their triggering rules (i.e., how they are signalled and recognized in texts), or both: the seminal work by Hobbs [Hob79, Hob85] and Kehler’s work [KKRE08] focus on the triggering rules and the inference process, while SDRT covers both aspects. However, the proposals made for the discourse relations used in SDRT are not all based on dedicated linguistic investigations. *Contrast* and *Parallel* have been studied in details [AHB01], but for most relations, we gradually discovered that the early proposals required refinements and improvements (see, e.g., the first attempts to correct *Background* and *Narration* in [AAB⁺95]). My colleagues and I have therefore specifically worked on some relations, starting from the basic temporal relations of *Narration* and *Background*.

4.5.1 Narration

Narration has been characterized from the start in SDRT through two aspects: its semantics effects of temporal precedence and its semantic and structural requirement of a shared topic [LA93]. The first enhancement concerned taking into account spatial effects on top of temporal ones. We showed in [AAB⁺94, AABV95, AAB⁺95] that *Narration* implies that time moves on but space doesn’t in the absence of indications to the contrary (spatial locating adverbials, motion verbs, or a trajectory context). We next showed in [BLDV01b, BLDV01a] that the temporal consequences themselves were to be strengthened. *Narration*(α, β) implies that there is no *significant* temporal gap between the two sequential events e_α and e_β , i.e., e_β needs to be conceived as “the next step” in the story told. The length of the interval between the two is not important, the point is that no event that is relevant to the story occurs in between. Relevance here is shown to be a question of interference with either the poststate of e_α or the prestate of e_β . Thus a same condition demanding that the poststate of e_α be compatible with the prestate of e_β and that they both hold during the whole temporal interval captures at the

⁶⁶In the Penn Discourse Treebank project, there is even no attempt at all to characterize discourse relations, as they simply use discourse connectives instead of relations [PMD⁺07].

same time the spatial and temporal effects.⁶⁷ This condition is quite strong, and it turns out that a weaker relation of Narration, that we called *weak Narration* implying only temporal precedence and with no common topic requirement, can be involved in some coherent texts where the standard Narration or *strong Narration* is blocked.⁶⁸ These two relations probably are the two extremes of the range spanned by *Narration* seen as a global scalar relation. The scalarity of *Narration*, essentially based on the specificity of the contents of the narrative topic, is advocated for in [AL03] although not yet accounted for in the theory.⁶⁹

As far as triggering conditions are concerned, the study in [BLDV01a, BLDV01b], also reported in [BBLD⁺04], proved that *puis* (roughly equivalent to *then*, but more restricted) is a marker of *strong Narration* thus contributing a non-defeasible triggering rule. The same work, especially [BLDV01b], discusses what are *connectives* (like *puis*), whose semantics is directly in terms of discourse relations, as opposed to some adverbials for instance temporal locating adverbials indicating sequence (even anaphoric ones, like *un peu plus tard* (a bit later)), whose semantics primarily contribute to the propositional content of discourse segments. This distinction is grounded in syntax: discourse markers like *puis* or *mais* (*but*) necessarily are in IP-adjunct position which corresponds to their role at the sentence level, while locating adverbials primarily act as VP modifiers.⁷⁰

4.5.2 Background

This relation relates an event-describing clause, the foreground, and a state-describing clause which presents the scene in which the event occurs, its background. The original temporal overlap semantics and aspect-based triggering rules of *Background* didn't change much since [LA93]. In this respect, the only addition has been to make spatio-temporal the temporal overlap involved in its semantics.⁷¹ However, *Background* has been the subject of much changes regarding the overall discourse structure in which it

⁶⁷The spatial effect is obtained on the basis of the straightforward assumption that the temporal overlap between the poststate and the prestate implies a spatial one when the two states share participants. Note that the simple temporal overlap between poststate and prestate as proposed in [AAB⁺95] and taken up in [AL03] is too weak to guarantee the “no significant temporal gap” constraint: the poststate of e_α should run (at least) till the start of e_β and the prestate of e_β needs to hold when e_α ends.

⁶⁸We showed that this blocking is triggered by temporal locating adverbials denoting sequence like *un peu plus tard* (a bit later). A serious hypothesis would be to generalize the blocking of *strong Narration* to all locating adverbials in IP-adjunct position. This would make unnecessary our earlier proposal (done in [AABV95, AAB⁺95] and taken up in [AL03]) of applying a temporal shift function contributed by the adverbial to the poststate of e_α in the semantic effects of *strong Narration*. Indeed, with a simple precedence between events for *weak Narration*, there is no poststate to shift. As far as the shifting effect of spatial locating adverbials is concerned, since the spatial overlap was a consequence of the temporal overlap, there is no need to take it into account either. With *weak Narration*, the only semantic effect is temporal sequence, with no “space doesn't move on” effect, as observed.

⁶⁹The “Maximize Discourse Coherence” constraint as defined in [AL03] does not take into account the contents of topics, although the authors show that this would be needed.

⁷⁰As seen below Section 4.6.1, this does not mean that such adverbials have no discourse structuring role when they are in an IP adjunct position.

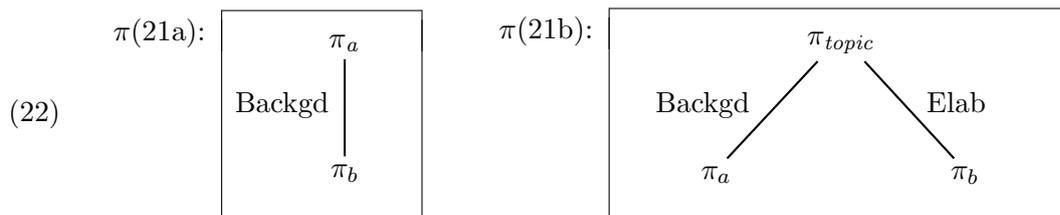
⁷¹Contrarily to *Narration*, the spatial dimension of the overlap needs to be asserted because it isn't restricted to cases in which the event and the state share participants.

appears. *Background* was assumed from the start to be coordinating, while, applying the tests presented above and in accordance with the way RST considers analogue relations, one shows that it actually is subordinating [VP04]. Problems due to failing to recognize the subordinating nature of *Background* lead to the introduction of a complex ad hoc topic structure (the Foreground-Background Pair) in [AAB⁺95], something still appearing in [AL03].

Background, or rather *Background*₁ [AL03] or *Background*_{backward} [APV07], applies to discourses in which the foreground is presented before the background as in (21a). But the opposite order of presentation also occurs as in (21b), which in SDRT, unlike in RST, cannot be dealt with by a mere inversion of arguments since the order of arguments reflects the order of attachment and constrains among other things the availability of referents for anaphora resolution.⁷²

- (21) a. Mary came home (a). It was pouring with rain (b).
 b. It was pouring with rain (a). Mary came home (b).

The situation in which the foreground is presented before the background, called *Background*₂ or *Background*_{forward}, requires attention too but has been neglected till recently. When a state (or a sequence of states) is described with no event in the context to anchor that state, there is an expectation that a foreground (possibly constituted by a sequence of events) is to come in the stage thus set. Indeed, the first sentence of (21a), but not that of (21b), could stand alone and make a very short discourse.⁷³ This is exactly the kind of forward-looking phenomena for which the introduction of a “forward”-topic structure has been proposed (see Section 4.4 above), and this solution is adopted in [APV07]. In fact, this solution avoids the introduction of a second *Background* relation: the structure generated is constituted by a topic, to which both the state-describing segment is attached by the usual *Background* relation and the expected foreground is to be attached by *Elaboration*. Here is an illustration of the discourse structure of (21a) and (21b):



While with a mere inversion of arguments, (a) in (21b) cannot serve as background for several successive clauses, this is allowed by the forward-looking topic structure. In addition, a controlled “percolation” of some referents from the background to the

⁷²To simplify annotation, only one relation of *Background* has been considered in [RDA⁺07] and [ANN08] although the two cases are recorded differently. The topic solution is not enforced in the annotation as it is an SDRT-dependent theoretical construct; see Section 4.7.

⁷³For states described with Imparfait in French, this incompleteness phenomenon is to be related to the anaphoricity of Imparfait [Vet80], which in these cases gives rise to a cataphor.

topic can be allowed to account for some anaphora patterns, while this is impossible with the inversion. We show in [APV07] that this proposal gives the right structural properties regarding attachment, referent scope and anaphora resolution, for uses of *Background_{forward}* in narratives as well as for its presuppositional uses (*Background* is the relation to attach presuppositions with, see [AL98c]). In Section 4.6.1, we will see how this proposal combines with other forward-topic triggering patterns.

4.5.3 Entity-Elaboration

Another important SDRT relation used from the start is *Elaboration* whose semantics involves a parthood relation between the main eventualities as in (10), (11) or (19). I contributed to stress some facts related to the semantics of parthood, for instance that in addition to implying temporal inclusion, it implies spatial inclusion, or actually spatio-temporal inclusion.⁷⁴ There are no markers for *Elaboration*, and only lexical clues together with word knowledge can (defeasibly) signal *Elaboration*.⁷⁵

We faced difficulties with *Elaboration* in discourse annotation tasks, such as the one carried in the project described below in Section 4.7. Just as in other projects [WG05, RDA⁺07], annotators obtain low agreement regarding *Elaboration*, which is repeatedly confused with *Background*, for instance. One reason for this problem is that some other relation is missing, which is somehow close to both. The relation of *Entity-Elaboration* that we introduced in [PVA09] “provides more details” on an entity described through an NP, rather than on the main eventuality of the clause as for *Elaboration*. It appears in (23) between the clauses (b) and (d), as well as between (b) or (d) and the non-restrictive appositions (a) or (c). This relation is called “Attribute/object Elaboration” in [MT87] and [KOOM01], and “Individual Elaboration” in [FHB01], but wasn’t considered in SDRT before [PVA09].

- (23) Pierre Vinken, {61 years old (a)}, will join the board as a nonexecutive director Nov. 29 (b). Mr. Vinken is chair man of Elsevier N.V., {the Dutch publishing group (c)} (d). (From [PMD⁺07])

No connectives signal this relation. *Entity-Elaboration*, especially in news texts, frequently appear with specific syntactic constructions: non-restrictive appositions and non-restrictive subordinated clauses, the former being perhaps more frequent in Romance languages than in English. And the syntactic position of appositions, being attached (“hosted by” in Marandin’s terms) to the NP or to the sentence (IP) —the

⁷⁴What exactly is the spatial extension of an eventuality remains an open problem in the literature. We usually assume that it includes the extension of its main participants, those that are true arguments of the verb describing the eventuality.

⁷⁵[KV09] criticizes the analysis of *Elaboration* in SDRT, especially that in [Ash93] which indeed was way too complex. The authors tend to confuse triggering rules and semantic effects, but they are right on one point. Triggering rules based on the lexical relation of subtype —or subsumption— cannot capture all *Elaboration* cases, since subtype signals parthood essentially when the event elaborated is a plural event, as in (10) or (19). Some Glue Logic rules are missing for the frequent cases in which the elaborating events describe phases of a single larger event, as in (11).

two positions being distinguished by possible or impossible movements in the sentence [Mar98]— is an helpful clue: *Entity-Elaboration* only appears when attached to the NP, thus making the distinction between (24a), an *Explanation* and/or *Background* case, and (24b), and *Entity-Elaboration* case.

- (24) a. *Marie, toute agitée de soupçons, demanda à rencontrer le commissaire.*
Toute agitée de soupçons, Marie demanda à rencontrer le commissaire.
Marie demanda à rencontrer le commissaire, toute agitée de soupçons.
 (Marie, very troubled with suspicions, asked to meet the superintendent.
 —From [Mar98])
- b. *Marie, mère de 3 enfants, demanda à rencontrer le commissaire.*
**Mère de trois enfants, Marie demanda à rencontrer le commissaire.*
**Marie demanda à rencontrer le commissaire, mère de trois enfants.*
 (Marie, mother of three, asked to meet the superintendent.)

Giving segment status to something which is not a clause, like an apposition, is debatable. However, non-restrictive appositions introduce secondary predications, that, like other parentheticals,⁷⁶ are “external” to the syntactic structure of the sentence [Mar98], so there are syntactic grounds to consider them as segments. On the semantic side, [Com98] shows the discursive role of non-restrictive appositions, something we have observed in many examples too.

Considering that non-restrictive appositions introduce segments and thus are equivalent to full clauses doesn’t contradict the hypothesis, put forward at the end of Section 4.3, that discourse relations are triggered by phenomena appearing only at the inter-clausal level. However, as seen above for the *Entity-Elaboration* cases, some appositions are “hosted” by the NP and not the sentence. Accounting for this fact would jeopardize somewhat my hypothesis.⁷⁷ Some authors, especially in “macro syntax”, emphasize though that parenthetical constructions show the limits of classic grammar [Deu03]; so investigating what exactly is the relation between an apposition and the NP that “hosts” it still remains to be done to take a definite stance on the syntax-discourse interface.

4.6 The spatio-temporal structure of discourse

Since [LA93] (and [BA94] for French which displays more subtle tense effects) it is known that the rhetorical structure of discourse, i.e., its hierarchical segmentation and the discourse relations holding between the segments, interacts with what is called the temporal structure, i.e., the temporal relations between the various discourse referents—especially the eventualities—involved in the propositional content of the segments.

⁷⁶See [Ash00] for an account of clausal parentheticals in SDRT.

⁷⁷And it is not clear at all what an “attachment to a discourse referent” would mean in SDRT. The problem is though that currently, two appositions attached to two different NPs in the same clause introduce two segments attached to the same discourse segment, thus ignoring the inherent separate links with the two different discourse referents. The solution we put forward in [PVA09] is to introduce another notion of discourse topic, that of an “entity under discussion” (cf Section 4.4), but their need and use should be better examined.

I contributed, together with Nicholas Asher, Michel Aurnague and Myriam Bras, to extend those analyses to the spatio-temporal dimension in [AABV93, AAB⁺94, AAB⁺95]. The standard content-level relation of SDRT such as *Narration*, *Background* and *Elaboration*, beyond their obvious temporal semantic effects, have spatio-temporal semantic effects, i.e., they imply spatio-temporal relations between the main eventualities of the segments they relate (see Section 4.5 above). As a result, the spatial properties of the eventualities denoted by lexemes such as motion verbs, spatial prepositions and spatial locating adverbials, and encoded in the propositional content of the segments interact with the inference to discourse relations in the Glue Logic, e.g., blocking some inferences (see [AAB⁺95] for a detailed treatment of an example). So, in particular, the semantics of temporal and spatial locating adverbials taken into account within the propositional content of segments has an indirect effect on discourse structure.

It is only later that we realized that the syntactic position of adverbials mattered (in [ABVA01], see Section 3.2 above), and that preposed locating adverbials in IP-adjunct position had a direct effect, introducing specific patterns within discourse structure [VBAA05]. In this section, I will first (4.6.1) present the new SDRT account we proposed to account for several discourse phenomena involving locating adverbials that have been observed in different fields.

Following our first work on the spatio-temporal structure of discourse, we had observed that locating adverbials, be they temporal or spatial, could have a spatio-temporal meaning in route-describing texts. We analyzed this phenomenon in [AABV95] as a semantic ambiguity of locating adverbials. This wasn't satisfactory since it was clear that the phenomenon was a pragmatic one due to the discourse context. However, we hadn't noticed that this phenomenon occurred only when the adverbial was in IP-adjunct position. It turns out that our account of IP-adjunct locating adverbials provides a simple explanation of the spatio-temporal reading of locating adverbials, without assuming that they are ambiguous, as we showed in [VBAA05]. I will review this next (4.6.2).

The spatio-temporal structure of discourse is especially relevant in directions-giving dialogues, whose main objects are route descriptions. Such dialogues have been studied by Laurent Prévot in his PhD, which I contributed to supervise [Pré04] (see also [PMDV02]).

4.6.1 Discourse effects of locating adverbials in IP-adjunct position

As seen above in Section 3.2, being clause modifiers, locating adverbials in IP-adjunct position are best accounted for at the discourse level rather than within the propositional content of the basic segment representing the clause. Indeed, linguistic work in both syntax and semantics [Kun75, Cha84, Cha97, Cha03] has shown that in this position, locating adverbials take on a role of “frame adverbials”, possibly having scope over a larger segment than just the clause modified and “setting the scene” for this whole segment (cf. example (6a) on p.34). [LDPW03] has examined the locating (also called “indexing”) role of temporal frame adverbials on a corpus, observing in particular that the scope boundary of the adverbial is fuzzy in many occasions.

Separate work in descriptive and psycholinguistic discourse studies [BV00] (see also

[Dic01]) has shown that preposed adverbials have a role in discourse segmentation, this time not so much because of their scope over larger segments, but because of their role as markers of discourse topic shifts. We were not aware of that literature when we wrote [VBAA05]. In our work based on the above-mentioned corpus study [LDPW03] from which example (25) below is taken, we actually rediscovered the marker of topic shift role of preposed IP-adjunct locating adverbials.

For instance, in (25), *en 1932* marks a “discourse pop”, i.e., the elaboration of the story of the previous resistance network is closed off, and a new sub-topic of the main topic of the whole text (Churches and nazism in Germany) is introduced, the setting-up of the group the German Christians. And similarly for *en septembre 1933* and *quelques semaines plus tard*.

(25) [...] *Mais lorsque la Gestapo démantela ce réseau de résistance en 1943, Bonhoeffer fut arrêté et déporté en camp de concentration. Il fut pendu le 9 avril 1945.*

En 1932 [In 1932] *se forma un groupe protestant national-socialiste, les “Chrétiens allemands”, qui réclamèrent après l’arrivée au pouvoir de Hitler la formation d’une Église du Reich, structurée selon le “Führerprinzip” et rejetant les juifs, ce qui se réalisa quelques mois plus tard. L’Église protestante, désormais dirigée par les “Chrétiens allemands”, était devenue un instrument entre les mains de Hitler. En septembre 1933* [In September 1933] *fut organisé le “synode brun” ; la majorité des responsables ecclésiastiques s’y rendirent en uniforme nazi. Il fut décidé, malgré l’opposition des adversaires des “Chrétiens allemands”, que les pasteurs qui n’étaient pas aryens seraient exclus de l’Église du Reich ; 70 responsables ecclésiastiques suivirent l’exemple du pasteur Koch et quittèrent alors la salle en signe de protestation.*

Quelques semaines plus tard [Some weeks later], *le pasteur Martin Niemöller appela les pasteurs hostiles à ces mesures antisémites à s’unir au sein d’une nouvelle organisation, le “Pfarrernotbund”, la “Ligue d’urgence des pasteurs”, qui respecterait les principes de tolérance énoncés par la Bible et la profession de foi réformatrice. [...]*

To account for both these roles, introducers of new frames and markers of topic shift, we proposed in [VBAA05] to extend the use of discourse topics in SDRT to a “forward-looking” use, as described above in Section 4.4. An adverbial like *in 1932* in preposed IP-adjunct position introduces a new topic—in the technical SDRT sense—that is, a new basic segment whose associated speech act is “I’m now going to tell you something new that happened in 1932”. This forward-looking topic with initial propositional content “something that happened in 1932” is waiting to be updated and elaborated by a segment to come, including the clause modified by the adverbial and possibly more.

Such a proposal explains why example (6a) on p.34 gets the same meaning as (10) on p.40, while (6b) is incoherent. Contrary to *Narration* in (19b) on p.47, the relation *Continuation* holding between the clauses of (6a) is not able in (6b) to generate the

topic⁷⁸ which is explicit in (10) and triggered by the adverbial in (6a).

This proposal is also in line with yet other evidence coming from formal semantics. [dS99] emphasized the effect of the position of the adverbial on the information partition into focus/topic structure. The relationship between discourse topic and the information partition of a sentence is not a straightforward one (see [GT97]). However, in the case at hand, information partition brings some light on the contents of the new discourse topic introduced by the preposed adverbial in IP-adjunct position. The information partition is illustrated by the related question under discussion as in (26): in (26b), the semantics of the QUD directly corresponds to the contents of the topic introduced by the adverbial, at least before it gets updated by the summarization process once the segment to come is attached to it.

- (26) a. John arrived at 5. / When did John arrive?
 b. At 5, John arrived. / What happened at 5?

Technically, the contents of this new topic segment is obtained by the standard procedure of existential closure [Die92] on the semantic contribution of the adverbial seen above in (7) on p.33, which in IP-adjunct position is not expecting to be combined with a VP. We thus obtain $\exists P \exists x (P(x) \wedge Q(x))$ (where $Q(x)$ is the specific contribution of a particular adverbial) which, as contents of a new SDRS π waiting to be elaborated by another SDRS to come π' , yields (27).

$$(27) \quad \lambda \pi' \exists \pi (\pi : [\exists P \exists x (P(x) \wedge Q(x))] \wedge \text{Elaboration}(\pi, \pi'))^{79}$$

The vexing problem of the apparent change in semantics of temporal locating adverbials according to the aspect of the eventuality⁸⁰ also troubled us in [VBAA05], in which we suggested using an general principle to the effect that “states perdure”. Such a principle is quite plausible from an aspectual point of view, but its overall impact should be assessed before claiming it is adequate. On the other hand, this proposal was

⁷⁸Actually *Continuation* cannot be triggered in (6b). *Continuation* requires some already existing dominance structure. As its name suggests, this coordinating relation marks the continuation of a previous, subordinate, rhetorical link—here *Elaboration*.

⁷⁹In [VBAA05], the formula actually is altered. First, an atom marks that P is to be updated by the topic summarization process every time π' is updated. This is an unnecessary addition, as the standard topic update procedure should simply apply, provided SDRT uses some general device to signal π as constructed topic, along with the other implicit topics. Second, the contribution of the adverbial itself, $Q(x)$ actually is taken outside π 's contents to be “distributed” over π' , by default. This quite ad-hoc move was proposed to account for the phenomena of “frame-exiting” without “topic-exiting” that may occur when the topic-shift role predominates over the framing role. (This happens for instance in the segment dominated by *in 1932* in (25); in that segment, whose topic is the setting up of the German Christians group, some events actually happen in 1933 after Hitler came to power; see [LDPW03].) We thus wanted to account for the locating contribution of the adverbial as a default only, but this solution is not satisfactory as this is a pragmatic phenomenon that should be accounted as such and not through an arbitrary manipulation of the semantics of the adverbial. This is left for future work, probably eased when a large annotated corpus will be available (see Section 4.7).

⁸⁰For adverbial like *in 1932*, there appears to be a reversal of the temporal inclusion relation according to the aspect of the eventuality it combines with: an event is located within the time 1932, while a state includes—or simply overlaps—that time. See, e.g., [KR93].

done without considering the interaction between the forward-looking topic structure triggered by the adverbial in IP-adjunct position and the relation of *Background* often present with states. In fact, as seen above in 4.5.2, forward-looking *Background* gives rise to similar topic structures as IP-adjunct locating adverbials. In [APV07] we also examined the interaction between forward-looking adverbials and *Background*, to the effect that in (28), a variant of example (21b) on p.51, a single topic is introduced and the same graph as that of $\pi(21b)$ in (22) is obtained.

(28) Yesterday evening, it was pouring with rain (a). Mary came home (b).

As shown in [APV07], to solve this issue, we simply need to assume that “non-spanning” locating adverbials require that the eventuality they locate be an event, in contrast with “spanning” locating adverbials that require states, i.e., those built with the prepositions *depuis* (since), *jusque* (until) or *de ... à* (from ... to) [BM93]. Most locating adverbials discussed here are non-spanning, whether “punctual” (e.g., *at 5, two hours later*) or not (e.g., *on Monday, during 1983*). Such an adverbial in IP-adjunct position will not be able to combine with a stative clause as the clause elaborating the topic, for an ontological constraint of homogeneity is imposed by the parthood relation in the semantics effects of *Elaboration*. Instead, this stative clause will attach with *Background* to the topic, thus yielding an overlap with the topic event, and will be indirectly located by the adverbial. This means that the topic still waits for some event clause to elaborate it, which predicts that the first sentence of (28) cannot stand alone in null context.⁸¹

Interestingly, this proposal also explains the change in location meaning of a present or past perfect clause with adverbials in VP-adjunct or IP-adjunct positions, as seen in example (5) on p.33. I standardly (see e.g., [KR93, MdSV⁺04]) assume that a perfect tense introduces both a past event and its resulting state, the resulting state being focused on and thus, in a SDRT account, being by default the main eventuality. When in VP-adjunct position, the composition process enables the adverbial expecting an event to ignore the default focus and pick the event also provided by the VP to modify it. When in IP-adjunct position, the topic segment generated by the adverbial cannot combine (for the expected *Elaboration* attachment) with the main clause’s segment that has been separately built and where the main eventuality standardly is a state. The main clause is thus attached to the adverbial topic segment by *Background*, with the expected semantic effects.

Topic shifts are grasped through the introduction new topics, that is, only their forward-looking aspects are considered here. The closing-off of the previous topic, a “discourse pop” in SDRT terms, is not yet accounted for (the new topic could be attached to the last constituent and constitute a sub-topic of the previous one). There are no standard means to directly force a pop-up in SDRT, it usually is a secondary effect of some semantic incompatibility, so this is left for the future. In any case, we should first of all examine if the pop-up is systematic, and if not, in which cases it occurs.

⁸¹If no event candidate to be located by the adverbial ever comes, we can hypothesize that there is an existential closure yielding, say, an implicit perception event, as often observed with the French imperfect tenses.

Future work should also reconsider the semantic comparison made in [BLDV01b, BBLD⁺04] between connectives like *puis* (then) and locating adverbials like *un peu plus tard* (a bit later), in which we didn't take into account the structural effects due to IP adjoining. In [BLDV01a], we had observed in our corpus discourse organizing phenomena due to *puis* that can indeed be described as topic shifts. However, since *puis* has no locating character (see 4.5.1 above), it obviously has no framing role. This difference may be in line with psycholinguistic studies ordering discourse markers along a continuity-discontinuity signalling line [BV95].

4.6.2 Spatiotemporal interpretation of locating adverbials

In [AABV95], we showed that both temporal and spatial locating adverbials could take what we called a “spatio-temporal interpretation” in route-description contexts, e.g., in narratives describing journeys. For example, in (29e), *dix minutes plus tard* (ten minutes later) and *dix kilomètres plus loin* (ten kilometres further) can be substituted one for the other, with the same spatio-temporal reading: it started thundering at some point after Cordes along the trajectory *and* after the time at which the rain became hail. The same phenomenon is involved in (d): *à Cordes* (at Cordes) in this context is equivalent to “[Jean] having reached Cordes”.

- (29) *Jean arriva à la hauteur d'Albi (a). On ne voyait rien du paysage (b). Il pleuvait à verse depuis Toulouse (c). À Cordes, la pluie se transforma en grêle (d), et, dix minutes plus tard / dix kilomètres plus loin, le tonnerre se mit à gronder (e).*

Jean reached Albi. One couldn't see anything of the landscape. It had been pouring from rain from Toulouse. At Cordes, the rain became hail, and ten minutes later / ten kilometers further, there were rumbles of thunder.

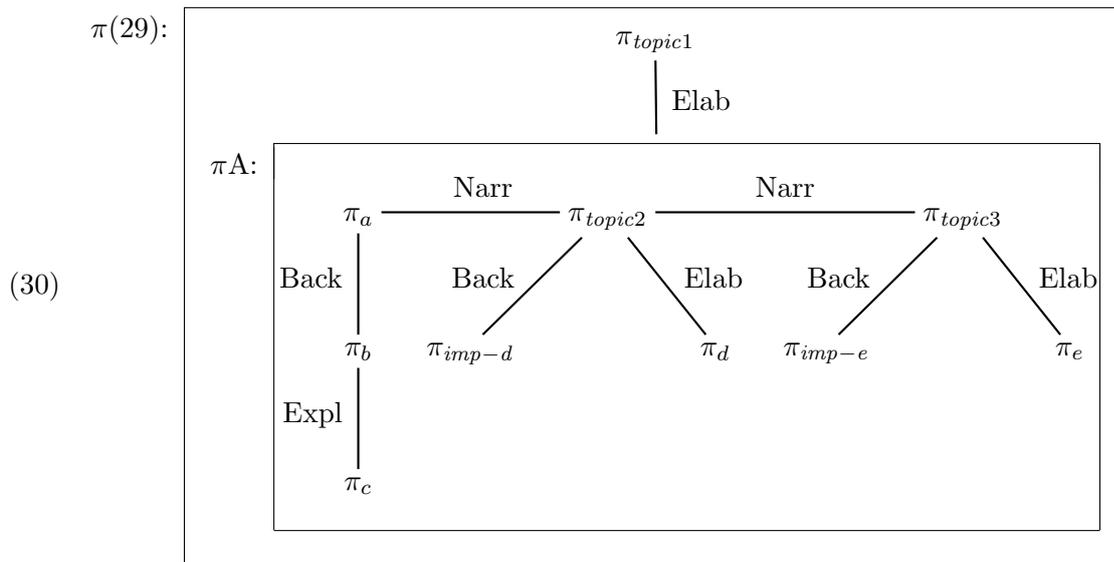
We analyzed this phenomenon in [AABV95] as a semantic ambiguity of locating adverbials. But the spatio-temporal interpretation clearly is a pragmatic phenomenon since it arises only in some specific contexts and can be overridden; in [VBAA05] we accordingly proposed to account for it as an implicature⁸² that the objects involved in the trajectory have continued their movement up to being located themselves by the adverbial. We also observed in [ABVA01, VBAA05] that this implicature is triggered (in a route-description context) by the adverbial only when it is in IP-adjunct position, something we missed in [AABV95].

Both the IP-adjunct position and the implicature converge to introduce a forward-topic structure, assuming that implicatures, like presuppositions [AL98c, APV07], are

⁸²In that paper, we used the term presupposition instead of implicature. The phenomenon observed is insensitive to negation, but that isn't really an argument in favour of projection since the adverbial in IP-adjunct position already is outside the clause. Embedding the whole sentence in reported speech, as in *On annonça à la radio qu'à Cordes la pluie se transformait en grêle* doesn't licence projection, but this test isn't really conclusive because of the necessary aspectual change (*transforma / transformait*). Context sensitivity more clearly points at an implicature and not a presupposition. I nevertheless assume the implicature is attached with *Background*, as a (local) presupposition would.

in SDRT treated as introducing segments related by a forward-looking *Background* (see Section 4.5.2). The implicated segment, whose propositional content includes a motion event and a resulting state of being located by the adverbial (the resulting state being the main eventuality) is attached to the topic by *Background*. The main clause is attached to the topic by *Elaboration* locally yielding graphs similar to that of $\pi(21b)$ in (22) p.51. Because *Background* entails spatio-temporal overlap (see Section 4.5.2), we obtain the observed spatio-temporal locating effect on the topic event, which here is identical to that of the main clause.

The whole SDRS graph for (29) is depicted in (30).



4.7 Towards the next SDRT: a more empirical approach

One could describe the work on SDRT done in Toulouse (a significant part of which has been done by Myriam Bras and myself in collaboration with other colleagues at IRIT and ERSS, and very often with Nicholas Asher), as empirically grounded on detailed lexical and compositional semantics studies. In the last decade, the attention to corpus data has grown (see, e.g., [BLDV01a, PMDV02, VBAA05]), which helped us discover unexpected phenomena and generate more general analyses than when working exclusively on constructed examples.

Corpus-based work in discourse semantics is however impaired by the difficulty of searching for instances of a given phenomenon within the corpus, because more often than not, no specific marker for a given discourse relation is used [Tab06]. This difficulty points at the need for a corpus annotated with discourse relations. More generally, such an annotated corpus is required to search for large discourse structures displaying a complex hierarchy, something invisible through a search only based on syntax and the lexicon. In addition, it would certainly ease the search for semantic patterns (e.g., lexical) associated with a given discourse marker to study the corresponding discourse relations.

This latter need has motivated the creation of the Penn Discourse TreeBank, which has been made available in 2008 (a preliminary version in 2006) [PMD⁺07, PDL⁺08].⁸³ This large corpus has the advantage to be based on the same corpus as the syntactically annotated Penn Treebank and the semantically annotated (predicate-argument structures) Propbank. Nevertheless, its use is quite restricted by the fact that the annotation is essentially an annotation of discourse markers, rather than discourse relations, thus leaving a large part of discourse links uncovered.⁸⁴ This originates in the DLTAG approach to discourse [WKJ01], and despite the explicit theoretically neutral standpoint, it is not clear whether other features of the PDTB are not DLTAG dependent (e.g., complex segments appear to be more restricted than in SDRT).

Recently, we have started the project ANNODIS funded by ANR (cf. <http://w3.erss.univ-tlse2.fr/annodis> [PWAE⁺09]), partially on the basis of the previous experience of an annotation project conducted at Austin [RDA⁺07]. The project aims at obtaining an annotated corpus integrating two views on discourse. There is first the annotation of segments and discourse relations in a “bottom-up” way, that is, from basic segments to larger ones, looking into the semantic contents of each basic segment. This is the sort of annotation I referred to above, and it is inspired by discourse theories like SDRT, Hobb’s theory [Hob85], RST [MT88] or D-LTAG [WKJ01, FRWJ06], while trying to remain as neutral as possible to be able to use the corpus to test specific hypotheses underlying each theory.⁸⁵ The second view is rather a “top-down” one. It is inspired by theories of discourse organization like [Gou96], and is based in particular on visual discourse cues in texts, like headings and text divisions, which are essentially ignored in the first view. Integrating both views in a single annotated corpus will gather all possible discourse structure information, and will allow the confrontation of the two views, for instance, checking the compatibility of the macro-level discourse structures obtained in both views.

The ANNODIS corpus will thus have a much richer annotation of discourse structure than the PDTB corpus; another difference is that it will be composed of French texts (where tense and aspect in particular give rise to different discourse phenomena).

The project currently is at a middle stage: the annotation manuals have been finalized and the systematic annotation of the whole corpus is now starting. Writing the manual

⁸³Other efforts have been done towards the annotation of discourse relations in corpora, although, to my knowledge, no other has been made widely available [WG05].

⁸⁴There recently has been a move in the PDTB towards extending the annotation to unmarked discourse relations (actually assuming implicit discourse markers, which is not really the same thing, see footnote 86 below), but this is limited to adjacent clauses and doesn’t cover discourse pops (as with (d) in example (19) on p.47, represented in (20)).

⁸⁵Theory neutrality actually is somewhat unrealistic, and annotation with no theoretical background may on the contrary yield annotated corpora which are difficult to use as in [WG05]. In fact, to guarantee precise annotations, a theory should support the definition of elementary discourse units (basic discourse segments) to segment properly and the characterization of discourse relations to distinguish clearly between them. What one can aim more realistically at is maximizing expressivity. SDRT arguably is more expressive than RST and DLTAG. As a result, annotations in ANNODIS which are made compatible with SDRT are less constrained than in other corpora. Efforts to remove non consensual features like constructed topics used in SDRT have nevertheless been made in ANNODIS.

for the “bottom-up” part and running many annotation experiences, already made us confront our view of discourse relations coming from SDRT with randomly selected real data. As a result, we faced in particular the need to enlarge the number of relations that was first considered. For instance, we were forced to add the relation of Entity-Elaboration [PVA09] (see Section 4.5.3 above).

Once the corpus will be there to be exploited, two types of uses can be conceived. The first is to use the data to check many hypotheses, from the fundamental assumptions of a given discourse theory, especially SDRT, to the proposals already done regarding some specific phenomena, like the behaviour of adverbials in discourse; obviously this sort of use can also support the study of discourse phenomena that have not been addressed yet. The second is to bite the statistical bullet and use machine-learning methods to design “discourse parsers”, thus departing altogether from the logically-based theoretical approach.⁸⁶ I am not myself fond of statistically-based computational linguistics, but of course, I would be very interested to see to what extent a complex structure, a significant part of which is based on inferences using a variety of knowledge sources, can be uncovered relying only on simple surface cues. I am eager though to exploit the corpus to put a number of hypotheses to the test, and I expect SDRT to take a significant turn thereafter.

By way of conclusion of this chapter, here is a list of theoretical issues in SDRT that I believe can be addressed and hopefully resolved through extensive testing on annotated corpora. Many of these issues are related to the fundamental right-frontier principle on which SDRT rests; I have already started examining (without a large empirical ground, though) some of these in [PV05, PV08].

- Being coordinating or subordinating was long considered as a permanent feature of discourse relations, but evidence has shown that it is best considered as a default [AV05] (see Section 4.2). Exploiting the tests based on attachment information that we proposed, we can check which relations admit variations from the default, with which frequency, and in which context. A better understanding of the phenomena that interact at the information-packaging level would result from this. It is expected that punctuation as well as syntactic coordination and subordination have a significant impact, although SDRT hasn’t fully exploited these resources yet.
- Demonstrating the need for implicit discourse topic segments on real data would perhaps convince the skeptics about topics (see Section 4.4). In the ANNODIS project, to preserve theoretical neutrality, there is no annotation scheme for implicit discourse topics, for they are not consensual. Besides, detecting the presence of implicit segments is a difficult task to train the annotators for. Nevertheless, where there would be an obvious discourse pop requiring the attachment to a

⁸⁶ Efforts in this direction have already started, using automatic annotation with discourse markers as definite clues for some discourse relations, in the absence of a manually annotated corpus. The conclusion of [SL08] is though that such automatic annotation is of limited use due to the difference in patterns in which a relation is marked or not, confirming the need for a real annotated corpus.

constructed topic in SDRT, as for (19b) on p.47 (whose structure is shown on (20)), we instruct the annotators to use as attachment point the complex segment that would elaborate the topic. There will thus be some data to test “backward” topics. There will be also means to check the importance and role of “forward” topics, as the annotators are instructed to segment preposed IP-adjunct adverbials, their extra-clausal role being consensual. And although the annotators will use a single relation of *Background*, the forward cases will be identifiable. More generally, the corpus will certainly serve as a basis to study all “forward-looking” discourse structures, whose importance is just starting to be uncovered.

- SDRT has explicitly allowed for attachments with multiple relations, in contrast with other theories. Multiple attachments of a same segment to different attachment points surely are more peculiar still. They are not forbidden by any SDRT principle, provided all the attachment points are on the right frontier, yet it is not a situation deliberately allowed for, and it is not clear how it interacts with the “continuing discourse patterns” principle. However, on several occasions, we did encounter examples that required double attachments, especially with two coordinating relations giving rise to “diagonal” patterns [BLDV01a, BLDV03]; [APV07] presents another case of diagonal attachments, this time with one coordinating and one subordinating relation. The corpus will provide the means to check for the occurrences of this phenomena and study it, because the annotators are instructed to annotate all the discourse links that they believe are present.
- As seen above, the annotation process itself, with the possible difficulties encountered by the annotators and their possibly low agreement in places, may point at problems with the chosen set of discourse relations itself. This will help focus future studies on the relations that require attention because they are semantically not characterized enough or they are substitute for some missing relation, as was the case for Entity-Elaboration (see Section 4.5.3). We should eventually obtain a stable set of all-purpose discourse relations, a quite elusive goal [HM93].
- Finally, without necessarily using machine-learning techniques, an obvious use of the corpus will be to extract triggering rules by generalizing over patterns in which a given relation appears, with the help of a rich enough lexicon connected to an ontologically well-founded knowledge base. This could be the basis to the building a Glue Logic with large coverage.

5 Ontology of concrete entities

As explained in Chapter 1, and seen all along the previous chapters, my approach to “representing content” is to anchor all representations in a well-founded ontological base, in other words, a foundational ontology. This means in particular not to study a given domain in isolation, but to pay attention to the global ontological commitment of any proposal, describing how a category of entities depends on the existence of others and is related to them.

Because of my original and long-standing interest in the lexical semantics of space and time, this chapter examines the category of *concrete entities*, that is, those that are located in time and possibly in space (as opposed to abstract entities, that are out of time), or in other words, all the objects⁸⁷ that surround us and all the events⁸⁸ that involve them. There are proposals in metaphysics that argue to either reduce all to events (or processes), e.g. [Sei97], and others that reduce all to objects (or successions of instantaneous objects) e.g. [Chi73]; the distinction between objects and events and is nevertheless central in most foundational ontologies, which are mainly dedicated to the ontology of concrete entities (see also [GM09]).

The first section (5.1) of this chapter discusses the ontological issues to be considered when selecting an ontological scheme, and presents the one I have adopted in my work on the semantics of space and time.

I then (5.2) focus on my contribution to the representation of space. As explained in Section 2.3, I adopt a relational approach to space, which motivates the development of theories based on *extended* spatial entities, as opposed to standard geometry based on abstract points, lines and figures. This is the reason why my work belongs to *mereotopology*, a research area at the intersection of formal ontology and artificial intelligence that developed considerably in the 90s.

As its name tells, mereotopology is based on *mereology*, a fundamental theory (or family of theories) in formal ontology. Mereology is the study of the relation of *parthood*, seen as unique general relation. As will be seen below, I generally adopt a restricted interpretation of the primitive *P* (for part) used in mereology, since I simply use it as spatio-temporal “inclusion”. This is not to say that I disregard what distinguishes parthood from spatio-temporal inclusion. On the contrary, I developed a proposal to account for this difference, which in language and cognition actually comes into a variety of cases. The last section (5.3) describes my contribution to the representation of parthood relations.

⁸⁷Objects (or “endurants” or “continuants”) in metaphysics cover material entities from piles of sand, clouds, microorganisms, and people to cars and orchestras, as well as immaterial entities in space like holes and cracks.

⁸⁸In metaphysics, events (or “perdurants” or “occurrents”) are all concrete entities that *happen* in time. The existence of events, once seen as dubious entities, is nowadays widely accepted, following the linguistic and ontological arguments of Davidson [Dav67]. Depending on the authors, this may or may not cover both “events proper” and states, i.e., perfective and imperfective “eventualities”, as we called them in the previous chapter. See, e.g., [CV96, HPV00].

5.1 Concrete entities in a foundational ontology

An ontology is said to be “foundational” when it is formal, philosophically well-founded, and focuses on general and basic concepts axiomatizing them in a rich enough way. In knowledge representation, this is a guarantee for having large scope and being generic and transparent enough to be highly reusable [BM09].⁸⁹

When I started searching for general schemes in which to ground my representations of “content” in the early 90’s, there was no foundational ontology to rely upon in the field of knowledge representation. The closest to this were the first proposals by Pat Hayes to develop his own “naive physics” program [Hay78, Hay85b, Hay85a]. Although I hadn’t really familiarized myself with that literature at the time, in metaphysics, relatively few contemporary philosophers had proposed comprehensive ontological theories either, as influential philosophers have argued [LZ03]. In fact, still today no general scheme proposed in philosophy has really influenced foundational ontologies developed in knowledge representation, perhaps because of their high-level of abstraction. Rather, various metaphysical studies on various fundamental issues guide and support the existing foundational ontologies put forward by computer scientists. I will now very briefly review such issues (of course without doing justice to their complexity at all) and situate the position I have adopted with respect to them. The resulting scheme is presented in several papers (especially [AV93b, AAB⁺95, AVB97, VA07]; see also [VBM08] for a methodological discussion), however only partially as none of them is focused on just such ontological foundation.

5.1.1 Properties and individuals

An age-old question in metaphysics is to decide whether *properties* are allowed as citizens of the ontology on top of *individuals* [MO97]. Some, e.g., Quine, have said “never”, some, like Russell, have said “of course”. Properties are distinguished from individuals in that properties but not individuals can have other entities (or tuples of entities) as *instances*, i.e., the distinction is made through a formal relation of instantiation.

The standard account in metaphysics is to see properties as abstract entities called universals [Arm89]. I cover in this chapter only concrete entities, which I take to be entities that are at least in time and possibly in space-time, so such a view is not developed here.⁹⁰ A usual implementation of properties as universals is to consider them as predicates, but since for quite another reason I stick to a first-order logical framework (see Section 1.3), I anyway do not admit properties as predicates in the domain of quantification.

⁸⁹The basic categories of entities in an ontology are now commonly said in knowledge representation to form a “top-level”, referring to the usual graphical representation of category taxonomies as (inverted) trees. “Top-level ontology” is sometimes used instead of “foundational ontology”, but the former doesn’t imply generality, philosophical well-foundedness, nor deep axiomatization.

⁹⁰I do not exclude abstract entities from my domain of interest once and for all, but these would be limited to certain subclasses of abstract objects (e.g., certain propositions or states of affairs, to avoid ill-founded objects and cardinality issues).

A major alternative to properties as universals is the trope approach, in which properties are seen as concrete entities [Dal97].⁹¹ DOLCE adopts a related account of reified properties to represent *qualities* like colours [MBG⁺03, MB05], that I will apply to my theory of artefacts in Section 6.2. I have also developed a theory of *concepts*, reified properties with a concrete aspect due in particular to their dependence on creators and societies that adopt them (see Section 6.1).

However, the ontological scheme I adopt at a first approximation and I present in this section is limited to individuals. Of course, properties and relations have an important role in any ontology, but I take them into account as predicates not belonging to the domain of quantification. The three basic and fundamental relations in formal ontology I make use of are: **parthood**, the basic relation of *Mereology* that relates two individuals one of which is a part of the other [Les31, Sim87]; **dependence**, especially existential dependence, the binary relation specifying that the existence of an individual implies that of another [Fin95, Sim87]; **constitution**, the special case of dependence in which the two entities coincide without being identical, as occurs e.g., with a person dependent on her body, the body constituting the person and both being colocated [Rea97].

In the next two sections of this chapter, I will develop on mereological relations, which gave rise to considerable formal developments. Dependence assumes a characterization of what *existence* is and most often uses a modal necessity operator; I will make use of a several types of dependence in Section 5.3, as well as in next chapter. Constitution is much more elusive, and formal accounts rely only on its relationships with dependence, parthood, and location in space-time; it nevertheless is a crucial relation to distinguish categories of entities, as will be seen shortly.

5.1.2 Space and time

Without surprise, the nature of space and time and the constraints bearing on the location of concrete individuals in time and space is the next important issue. It is important in general, because it is the very basis of the definition of concrete individuals and because it has impact on all of the above: parthood, existence and constitution are all usually seen as being temporalized relations and as being closely related to spatial location. Of course, it is important for my purposes, as it underlies any representation of space and time.

Many philosophers take a realist stance on space and time, that is, they consider that there are such things as “container” space and time, made of absolute locations existing prior to having any concrete entity located at them. The opposite view is to consider that we can refer to and conceive of such locations because there are concrete, especially material, entities around that display spatial or temporal relations between them, and that time and space are just abstractions over those relations. This opposition is usually described as the absolute vs. relational views on space and time, referring back to, respectively, Newton’s and Leibniz’s positions (see, e.g., [Ner03]). My personal taste

⁹¹For a review of alternative views on properties and how to implement them in knowledge representation, see [BM09].

is to apply parsimony and adopt the relational view, in absence of linguistic evidence supporting absolutism.⁹² I have therefore focussed on relations between concrete entities which means between *extended* entities.⁹³ Anyway, abstract places and times can be recovered as some sort of equivalence classes on relations between concrete entities. And whether or not such equivalence classes eventually are bound to prior spatial or temporal substrates, it might be very convenient to refer to them as extended locations in space or time, also called *regions* and *intervals* (or periods). However, assuming those locations are sets of points or instants, as standardly done in Cartesian geometry and continuous models of time, would be adding a much stronger ontological commitment, for sets, and points and instants themselves, are all mathematical abstractions. The structure of space-time induced by relations on regions actually is subject of more interesting developments than the absolute/relational issue. We will see that in the next section (5.2).

The most controversial issue about space-time in metaphysics really lies elsewhere. It doesn't concern time and space themselves, but how concrete entities *are in* time and space, in relation to their mereological structure [Rea97]. Puzzles readily arise when one assumes on the one hand mereological extensionality (the principle by which an entity is identical to the sum of its proper parts, see next section) and on the other hand the persistence of identity through change, especially through the loss of a part (e.g., the cat Tibbles which remains Tibbles after losing its tail). Solutions suggest to distinguish Tibbles from the matter that constitutes it (Tibbles would persist, but the matter that constitutes it could be different at different times—this is the “standard account”), to use a temporalized parthood, or to assume that any concrete entity is extended in time and has temporal parts (thus eliminating the idea of persistence altogether). This last view is called four-dimensionality (or perdurantism) for it considers time as yet another dimension added to space; any concrete entity is located at some spatio-temporal “worm”⁹⁴ [Qui60, Lew86, Sid03]. It contrasts with three-dimensionality (or endurantism) which is the view, arguably more commonsensical, that ordinary objects are “wholly present” at any time, i.e., they don't miss any (past or future) part [Bak07]. Three-dimensionality has the advantage of directly supporting the distinction between objects and events, events being the concrete entities that do have temporal parts; it nevertheless has more difficulties in solving the puzzles of identity through time. Four-dimensionality usually consider there is no ontological distinction between objects and events; many assume that all concrete entities simply *are* spatio-temporal worms, thus also rejecting the possibility of coincidence or colocation and having no need for a constitution relation.

⁹²When we talk of places and times in language, we are *not* referring to absolute locations: places are all tied to portions of earth surface, to buildings, to seats on trains, etc. and, similarly, times are anchored to events or to calendars, but calendars themselves are defined on the basis of some original event and the earth's rotation around the sun, the pulsation of some crystal, etc.

⁹³There arguably are no 0-dimensional perceptible entities. One could object that only material entities are physically necessarily extended. But it can be held that immaterial concrete entities such as a hole in a sheet are all dependent on material ones, and can therefore be ignored in characterizing space. A similar argument is proposed in [Kam79] regarding the construction of time based on only relations between events (proper), states being assumed to be dependent on events.

⁹⁴Or “histories” in the application done in [Hay85a].

My own view combines pieces of standard solutions. It goes without saying that language clearly distinguishes objects from events, so I adopt this distinction. But I consider it might be done through other properties than just the having or not of temporal parts: I may distinguish someone from her life because only the person has a weight while only her life can be hectic; more generally, objects *participate in* events, not the other way around, and participation is not parthood. So, I am not a strong four-dimensionalist, i.e., one that admits only space-time worms in its domain of concrete entities.⁹⁵ However, I do consider that being in time characterizes all concrete entities. Objects are in time as well as in space, i.e., they are in space-time, so they are arguments of temporal and spatial relations. In fact, I deem it essential to escape Tibbles’s and other puzzles to temporalize parthood for objects, i.e., to specify when an object is part of another. A very close alternative [Mas09], with a stronger ontological commitment though, is to use a binary parthood on “temporal slices” of entities, assuming that any concrete entity can be temporally “sliced” with respect to the temporal extension of some other entity, e.g., an event.⁹⁶ I actually do not have strong preferences towards one or another (see Section 5.2.2). This approach, which can be seen as a weak four-dimensionalism (as it accepts temporal slices of objects but refuses to reduce all to spatio-temporal worms), is the one I have adopted in most of my papers on the semantics of space, explicitly [AV93b, AAB⁺95], or implicitly when temporal issues are ignored [AV93a, AVB97].⁹⁷

Finally, I also accept the existence of different coincident objects related by the relation of constitution.⁹⁸ For instance, I distinguish a person, its body, and the matter making it up at a given time, because these three entities have different *essential properties*: to exist, the person needs to be capable of intending, the body to have a specific shape and structure, and the amount of matter to have certain parts (for these cannot change across time) but no particular shape. I thus adopt a multiplicativist position [VBM08]. Like the distinction between objects and events, the distinction between an ordinary material object (e.g. the body) and the amount of matter that constitutes it is reflected in language. In fact, it has been advocated by many philosophers of language studying the count/mass distinction and taken up in formal semantics (see e.g., [Lin83], and for an excellent survey [PS03]). Such works influenced my choices quite a lot, although my attention to the category of amounts of matter doesn’t stem from

⁹⁵This would in addition contradict the view that concrete entities are prior to space and time.

⁹⁶Strong four-dimensionalists, who are absolutists, slice concrete entities directly with respect to times. The slice of entity x during t is the maximal temporal part of x having t as temporal extension.

⁹⁷In my papers, I sometimes used a function *stref* (for “spatio-temporal referent”) for restricting concrete entities to their spatio-temporal properties. Other views would be that this function would pick either the (spatio-temporal) region at which the entity is located in an absolutist framework, or some representing element of the equivalence class of an entity in a relational approach. In any case, for arguments of temporal and spatial relations, I used terms of the form *stref*(x), with *stref*(x) = *stref*(y) meaning that x and y are spatio-temporally colocated. Today, I would formulate this differently, simply using the concrete entities themselves and a colocation relation.

⁹⁸I do not characterize, contrary to many, coincidence and constitution in terms of sharing the same parts, but in terms of being spatio-temporally colocated. Having the same parts is a stronger property (though not for strong four-dimensionalists). See [Rea97], and [VBM08] for some notes on this point.

the semantics of noun phrases. It is necessary for the semantics of the expression of parthood relations in language, as will be seen in Section 5.3.⁹⁹

Additional categories of entities are motivated by linguistic evidence, as we have seen for the semantics of space in Section 2.3, without finding much support in metaphysics nor philosophy of language. There are first the immaterial objects, that I have also called “space portions”, like holes, containers’ insides, interiors, and entities described by phrases like “the space beneath the table”, “the gap between the two buildings”.¹⁰⁰ These actually did get some attention in formal ontology [CV94]. As explained above, they are not regions of absolute space and they depend on their “hosts”: the inside of a cup depends on the cup and is bound to move along with it.

Language also distinguishes “locations” (see Section 2.3), which again, are not regions of absolute space. In a first approximation, they are material objects that are fixed with respect to some larger whole object, conventionally taken as frame of reference. For instance, geographical locations are areas of land, at a fixed position on earth. Locations in addition are similar to containers: they determine a space portion in which objects located at those locations are placed (a space portion of a suitable thickness above ground, in the previous example). Differently from ordinary containers, language displays a systematic ambiguity such that locations may directly, i.e., without the use of a spatial preposition, refer to those associated space portions: e.g., a room and a garden can be spacious. On the basis of my and especially Michel Aurnague’s work [Aur95c, Aur04], it has been recently proposed that the “dot type” construction (see Section 1.1.1) would be appropriate for locations: they would be in the category (i.e., type) “ordinary material object • immaterial object”, thus grasping both aspects at once [Ash07, Ash10] (see also [Don04] that models space portions and locations in a single theory of “relative places”).

Finally, discourse theories like DRT and SDRT rely on the use of two subcategories of events, or rather, eventualities: perfective events, or events proper, and imperfective ones, or states [KR93]. As was clear in the previous chapter, my own work also assumes this distinction which presents many similarities to the count/mass distinction [Bac86a]. It has been argued that durative perfective eventualities, i.e., accomplishments in Vendler’s terminology, are constituted by imperfective eventualities, i.e., processes or activities [MS88]. Like mass nouns, imperfective eventuality verbs have the cumulative property, and to some extent, the dissective property [Dow79].

⁹⁹The account of mass terms I used builds on Parson’s [Par70, Par75]. This implies that in addition to being based on amounts of matter and the relation of constitution (the one above, not exactly Parson’s), it makes use of a further category of entities, *substances* (the “name” use of mass nouns like *water* in sentences like *water is widespread* and *water is a liquid*), and a relation of “being an amount of a substance”. I did not take up the abstraction operator proposed by Parsons, yet considering substances as concrete entities is quite controversial. I have not developed further work on substances after my PhD. Additional investigations on the ontological status of a category of substances, as well as a category of *kinds* necessary to account for some uses of generics [Car77], are required to analyse again how they fit the whole picture, as they are currently not accounted for in DOLCE.

¹⁰⁰Contrary to what these phrases may suggest, immaterial objects do not only denote empty space areas. They survive being occupied by a material entity, as e.g., with the inside of a filled cup.

5.1.3 Categories and number

To sum up, in my approach, concrete entities divide into events and objects. Objects further divide into (countable) ordinary material objects, amounts of matter, and immaterial objects.¹⁰¹ Events divide into perfective (“events” proper) and imperfective (states) ones.¹⁰² Locations form a complex category constructed on ordinary material objects and immaterial objects. We will see in next chapter that mental and social aspects introduce further important distinctions in both categories of objects and events.

Number complexifies this scheme further by adding a new dimension, orthogonal to the categorial divisions just seen (see, e.g., [AV93b, AVB97, VA07]). Number, i.e., the property that distinguishes singular from plural entities is of course particularly relevant in linguistics. It has an ontological import as soon as one wants to account for collective predication, either on plurals, as in *Lea and Luc brought the piano upstairs* and *the lions surrounded their prey* or on singular collections, as in *the orchestra played a sonata*, for what exactly brought the piano upstairs, surrounded the prey or played the sonata? Surely some concrete entity and so not a set, which is abstract. But since those predicates are not distributive (i.e., it is not the case that a single lion surrounds the prey), referring to each of the singular entities Lea and Luc, to each lion, and to each musician cannot do, something else is needed. Proposals based on mereology have been thoroughly developed in formal semantics (see e.g. [Lin83, Bac86a, Lan89, Lin97]) but it is only recently that the issue of plural reference has become subject of debate in philosophy [Ray07]¹⁰³ and in I have included a theory of plurals and collections (based on Link’s and Bach’s, with some variations) in my scheme, again not for a general concern on the semantics of noun phrases, but because collective entities are used in spatial expressions and induce meaning variations with respect to singular entities (for instance, you can put sugar *in your strawberries*, but (usually) not *in a single strawberry*). They also give rise to specific parthood relations (see Section 5.3). As done by several authors, I distinguish mere plurals from collections (also called groups in [Sim87, Lan89]) because the latter have some *unity*, a unifying property characterizing them, while the former doesn’t. In fact each collection is constituted by a plural entity, and just as with ordinary object and the amount of matter that constitutes it, the plural entity constituting the collection may change over time, something frequent for venerable orchestras and football teams.¹⁰⁴

¹⁰¹Two further subcategories for substances and kinds might be included here; see footnote 99.

¹⁰²Although I didn’t study time nominals, I expect that the need for some sort of “time portions”, the temporal dual of immaterial objects or “space portions”, exists. This would add a third sub-category of temporal entities.

¹⁰³Actually, the trend there is to opt for a logic of plural quantification without any ontological commitment rather than an ontology admitting of collective entities (see, e.g., [Yi06], and for a generalization of this approach to mass terms as well [Nic08]). This won’t convince the ontologist who already needs to commit to a rich domain structured by mereological relations and constitution, and is hardly a solution for those who chose standard first-order logic for practical reasons.

¹⁰⁴Collections whose unifying property is social may be called organizations. There are non-social collections, e.g., a forest; in this case (but not all) unity is based on spatial proximity.

5.1.4 Towards a comprehensive foundational ontology

The scheme presented above is presented in more or less details in my papers on the semantics of space (e.g., [AV93a, AV93b, AAB⁺95, AVB97, VA07]). As explained in Section 2, it has been successfully used to represent the semantics of a variety of spatial expressions. In fact, although its main features find support by one or another proposal in philosophy, the choice for such features has been largely motivated by linguistic observations. Searching for yet more empirical evidence, [AHV05, AHV07a] reports of a joint project with psycholinguists aimed at testing the cognitive plausibility of a category scheme such as the present one (see Section 2.4).

Apart from terminological issues, this scheme has many similarities with two of the best-known foundational ontologies, DOLCE [MBG⁺03, BM09] (see also <http://www.loa-cnr.it/DOLCE.html>) and BFO [GS03, Gre03, MBG⁺03] (see also <http://www.ifomis.org/bfo>). DOLCE, which aims at linguistic and cognitive adequacy,¹⁰⁵ is particularly close to my choices. It distinguishes, among objects, amounts of matter, physical objects (what I called ordinary material objects) and non-physical objects (what I called immaterial objects); regarding events, DOLCE separates events proper (perfective eventualities) from stative (or imperfective) eventualities. More importantly, it adopts a similar set of formal relations: general temporalized parthood, dependence and constitution.

A few differences are worth noting though. Both DOLCE and BFO adopt an absolutist approach to space-time, and accordingly include categories for spatial regions and for temporal regions. In DOLCE, spatial and temporal regions constitute subcategories of abstract entities, which in principle would include other subcategories, like facts and sets although those have not been axiomatically characterized. I plan to investigate such subcategories in the future (see Section 7.2). DOLCE doesn't deal with number, and has no category for pluralities, groups or collections. BFO, which does not aim at linguistic adequacy but at realism, has a unique category of "aggregates" for both amounts of matter and plural objects. It also collapses in a unique category of "sites" immaterial objects (space portions) and locations.¹⁰⁶ On the other hand, it focuses on unity, and thus splits ordinary material objects into "substances" or bona fide objects and "flat parts", those material objects that do not have a unity of their own and are arbitrarily identified within a bona fide object. It distinguishes "processes" from instantaneous "events", but apparently ignores perfectivity and imperfectivity, in fact, it is not even clear if states are considered among processes. Both DOLCE and BFO include in their schemes *qualities*, reified properties akin to tropes. In DOLCE, a quality, e.g., the colour of a given rose, is not a fixed colour hue, e.g., carmine, but an individual that possibly changes its location in a *quality space* over time [MB05]. I recently used this feature of DOLCE to model the notion of capacity of an artefact (see Section 6.2), so it seems reasonable to extend my scheme with qualities as well.

The position I currently adopt though is instead to join my efforts with those of the

¹⁰⁵Its name stands for "Descriptive Ontology for Linguistic and Cognitive Engineering".

¹⁰⁶A peculiar position is that it explicitly considers all sites, including cavities inside a human body, to be independent entities.

LOA and contribute to the future revision and extension of DOLCE. This is being done in several directions: concepts and roles (see Section 6.1), pluralities and collections, as well as different parthood relations (see Section 5.3), actions (see Section 6.3), and artefacts (see Section 6.2). As mentioned above, substances and kinds, which arguably are not socially defined concepts, are also needed; these may be seen modelled as further categories of reified properties. The extension of DOLCE will be done in parallel with the development of general methodological guidelines on how to introduce new categories in a foundational ontology; the clarification of why and when reifying properties and why and when multiplying colocated entities was already initiated in [VBM08].

5.2 Region-based representations of space

My first interest in representations of space was quite practical: I needed some framework to represent the geometric level in the semantics of spatial expressions (see Section 2.2).¹⁰⁷ Gradually, my interest widened to ontological issues and made me search for an ontology of space, or rather, of space-time, adopting a relational approach based on extended entities, as explained in the previous section. This focus led me to participate at an early stage to the field of qualitative spatial reasoning in artificial intelligence which developed in the 1990's, and the subsequent area at the intersection of artificial intelligence and formal ontology called "mereotopology" (for reviews of the field, see e.g., [Vie97, CH01, Gal09]).

This area takes as foundation a family of theories building on Lesniewski's work, usually called classical or standard mereology [Les31]. Classical mereology was introduced by Lesniewski to serve as foundation for mathematics, when set theory was still in trouble with Russell's paradox. In contemporary formal ontology studies, mereologies are taken to be theories of general parthood [Sim87]. When focusing on concrete entities having an extension in space, parthood clearly implies spatial "inclusion".¹⁰⁸ Moreover, for those adopting an absolutist approach to space, parthood between spatial regions just amounts to spatial inclusion. Mereology was thus an obvious candidate to serve as a departure point for representations of space, either relational ones, necessarily based on extended entities, or absolutist ones having chosen extended regions (also called solids in early writings [dL22, Nic23, Tar56a] contemporary with Lesniewski's) as primitive entities often for cognitive adequacy reasons.

Mereology basically is a partial order, whose unique primitive relation is called either P for parthood or PP for proper parthood.¹⁰⁹ What distinguishes mereology (in particular from linear orders) is some degree of "supplementation" to guarantee, in presence of a proper part, the existence of another, separate, part that makes the difference between the proper part and the whole. In addition, mereologies most often include axioms to allow the introduction of operators for sum, product and difference. A complement operator derives from difference operator and the assumption of the existence

¹⁰⁷In this search, I investigated most existing approaches, which gave rise to my survey paper [Vie97].

¹⁰⁸This term is not used here in its restricted set-theoretical sense.

¹⁰⁹Some authors choose to take *overlap* O as primitive instead, while O is defined in the previous case as the sharing of some common part.

of a “universe” of which anything is a part, something usually done. Instead of such operators, classical mereology has a powerful general “fusion” operator (similar to the comprehension axiom of set theory, requiring an axiom schema) from which such binary operators can be defined.¹¹⁰ For a presentation of mereology, see e.g., [Sim87, Var96] and the dated but readable [LG40]; in my own papers, see, e.g, [MV99]. Mereology in itself is insufficient to characterize space or space-time. Some spatial aspects are grasped through additional topological notions, leading to *mereotopologies*, while fuller accounts are found in *mereogeometries*.

In the rest of this section, general issues in mereotopology are presented first, focusing in particular on the specific problems of dealing with atomic representations, a very relevant case for finiteness implies atomicity. I then address extensions of mereotopology, evoking approaches to mereogeometries and space-time.

5.2.1 Mereotopology

Theories combining mereological notions and topological ones have been called *mereotopologies*. Topology is a mathematical theory based either on the notion of open set or on that of accumulation point, two arguably non-commonsensical notions because based on abstract notions like that of a point. Mereotopology, however, uses primitives arguably closer to commonsense like “being connected with” (ie, sharing a part or a boundary), “being an interior part of” (ie, being a part entirely surrounded by the whole) or “being self-connected” (ie, being one-piece).

It is believed that the earliest attempts at mereotopology have been made by Whitehead [Whi19, Whi20], who proposed to use a single primitive of “being extensionally connected with”, from which the mereological relation of part is defined. Whitehead’s work inspired many others [dL22, Nic23, Tar56a], some of them clearly motivated by cognitive adequacy concerns [Nic23], but was turned into an axiomatic theory based on a primitive of connection only much later by Clarke [Cla81, Cla85]. The Region Connection Calculus or RCC [RCC92b], probably the most famous mereotopology in AI, builds on Clarke’s work.¹¹¹ Actually, a whole family of connection-based theories emerged from this approach, among which the one I have adopted and developed in [AV93a, AV93b, AV95b, AVB97], which differs from Clarke’s [Cla81] by the use of unary and binary operators only (no fusion operator), the choice of atomicity, and some extensions (see below).

In [AV95a], I proved the first representation theorem of the field,¹¹² thus clarifying the models of Clarke’s theory [Cla81] in terms of classical point-set theory.¹¹³ This also

¹¹⁰The fusion operator makes a real difference only for infinite domains.

¹¹¹A change in the definition of complement was included though, so that RCC is significantly different. In particular, RCC is bound to be non atomic.

¹¹²Earlier but partial results were proven in [Grz60], although these remained unnoticed for quite a while in AI.

¹¹³I haven’t stressed anywhere that the models proposed in [AV95a] could be stripped of the elements aimed at grasping the notion of (weak) contact (see below) to yield models of a version of Clarke’s theory without the fusion operator, with a still valid completeness proof. Unfortunately, this has confused some authors regarding the importance of the result that has been taken to be limited to atomic domains,

corrects a problem in [Cla85], independently spotted in my PhD [Vie91], in [BG91] and in [Esc94].¹¹⁴ The representation theorem in [AV95a] shows that the regions in Clarke’s family of mereotopologies are *regular*, i.e., in point-set topological terms, that the closure of their interior is equal to their closure, and conversely, that the interior of their closure is equal to their interior. This property corresponds to the commonsensical fact that extended regions corresponding to concrete objects have no holes and cracks of the size of a point or of a lesser-dimensional region (e.g., portions of lines or planes in a 3D space), nor attached such elements, like dangling linear spikes.

Building on [Sim87, Var96] I then contributed with Claudio Masolo to the effort of systematizing the field.¹¹⁵ In [MV99], to be able to analyse the issue of the characterization of atomicity and divisibility in mereotopologies (see below) we formally compared several families of mereotopologies, mainly in terms of formulas that are theorems or not but also in terms of models. There are at least two strategies in mereotopologies. In the first, one extends a mereology with an added topological primitive (see e.g., [CV94, BGM96, Ste00b]). In the second strategy, adopted by Whitehead, Clarke, RCC and myself, both mereological and topological notions are axiomatized with a unique primitive. The authors just mentioned all take a primitive of connection, but other proposals exist (see e.g., [Got96] in mereotopology, and [dL22] for mereology, topology and geometry all at once). The mereotopologies based on a unique primitive of connection are particular cases of theories based on both a mereological primitive and a connection primitive, so in [MV99], the first strategy is used to facilitate the comparison of theories. These two strategies correspond to different ways of relating mereological to topological notions. The link can be weak if parthood just implies that anything connected to the part is connected to the whole, or strong with the converse implication as well, in which case parthood could be defined on the basis of a unique primitive of connection.

The choice of an implication or an equivalence can be interpreted as different degrees of “spatiality” one wants to ascribe to the mereological parthood. In strong mereotopologies, what occurs at the boundaries of the entities, i.e., how they are externally connected to others, contributes to their identity (because parthood is antisymmetric). This choice interacts with the adoption of a general point of view on space and concrete entities

something applying only to the contact part [CH01].

¹¹⁴In [Cla85], Clarke introduced points as filters in the object language itself (thus going higher order), following the Russel-Wiener construction. Like most early authors (Whitehead, de Laguna, Nicod, Tarski) Clarke aimed at proving that the abstract points used in classical topology could be defined as constructs in terms of regions (more on this in Section 5.2.2). Unfortunately, Clarke’s theory [Cla81] extended with the definition of points as sets of regions of [Cla85] has the unexpected result of having the same models as classical mereology, i.e., the topological dimension collapses. This can be explained by the fact that the construction used only captures interior points. It is still possible to give a more complex definition of points that preserves topological properties, as used in the proof of the representation theorem in [AV95a] (see also [AV93b], where I use a higher-order logic in Clarke’s style). This fact was not noticed in [BG91, Esc94], nor by other authors [CH01, DV06] who have confused the problem in the construction of points of [Cla85] with a problem in the mereotopology itself [Cla81].

¹¹⁵Other efforts have concentrated on the expressivity of the connection relation [Got94, CV98] or on algebraic versions of some mereotopologies [Ste00b, DW05], following Tarski’s work on the equivalence between classical mereology and Boolean algebra [Tar56b].

discussed in the previous section. Absolutists may restrict the application of topological notions to spatial regions, so even if using a general parthood relation (which of course will not be defined in terms of connection in this case), the equivalence and the “spatiality” of parthood are restricted to regions. This is the position adopted in [CV99, Chapter 7], for instance. Relationalists that are strong four-dimensionalists obviously do not have any problem with a spatial, or rather spatio-temporal, parthood. Relationalists that, on the other hand, accept the existence of different collocated entities, like someone and his life, or a statue and the clay it is made of (which is the position I adopt), run into a problem with the spatiality or spatio-temporality of a general parthood relation. There are two ways out. First, the mereological P can be interpreted as pure inclusion and the identity in the antisymmetry of P replaced with a spatio-temporal equivalence, ie, collocation. Such a purely spatial or spatio-temporal relation can be considered to come in addition to a general parthood, that is another mereological relation, or one can consider that there is no need for a general parthood since specific parthood relations may be defined on the basis of P and additional constraints, as done in Section 5.3. This is the strategy I have usually adopted. In the second solution, to instead keep a general parthood without an additional spatio-temporal inclusion relation, one needs to renounce to the strong link between mereology and topology. A theorem might well surprise those who would like to keep a general reading of the mereological relation in mereotopologies: if one wishes to characterize the fact that entities are spatially extended by stating that they have an interior part (they are not “all boundary”), a constraint weaker than regularity, then a strong link between mereology and topology is necessarily established (cf. theorem T5 in [MV99]).

Further variety among theories is brought by the choices of a definition for complement and that of distinguishing between open and closed regions (or having only one such kind), but less variety than expected, since interactions and constraints make some distinctions collapse, as we showed in [MV99]. The topological notions of open and closed regions, i.e., regions without or with their boundaries seen as infinitely thin elements, have often been pointed out as non commonsensical, and in addition, having this distinction forces either to accept that the difference between a closed region and its interior exists therefore boundaries, i.e., non-extended entities,¹¹⁶ are in the domain, or to abandon the strong version of mereological supplementation (see above). All options have been developed (e.g., in RCC there is no open / closed distinction, and in [Gal96, Got96, Smi96] boundary individuals are admitted), but not all are compatible with other desirable properties, like the atomicity and the connectedness of the space [MV99]. My own position is to require the extension of all individuals, ie, refuse infinitely thin boundaries, renounce to mereological extensionality to allow for atomic spaces, and accept topological operators to allow for a self-connected space. This option corresponds to Clarke’s family of theories.

Atomicity is especially important if one seeks to build a bridge between spatial rea-

¹¹⁶2D surfaces in a 3D space are in some sense extended. What I call here non-extended are all individuals of a lower dimension than the whole space, e.g., in a 3D space, 0D points, 1D lines and 2D surfaces and any combination of such entities.

soning and spatial databases approaches in areas like vision or GIS. It is also the only option for a cognitive approach.¹¹⁷ Atomicity is also unavoidable when building space from spatial entities and relations described in a text, in the fashion of the construction of time done in [Kam79]: linguistic descriptions of everyday space do not imply the existence of an infinite domain. In [MV99] we examined the particular case of tessellations (discrete partitions of a connected space into “tiles”), motivated by the fact that spatial data as pixel images or country maps are all tessellations of a portion of the plane. In a tessellation, tiles are atoms, they are extended, and they are connected to their neighbours yielding a self-connected space. [MV99] shows a negative result: it is impossible to characterize tessellations in a mereotopology. This result is in line with works in tiling theory [GS87] that show that in a space of two or more dimensions, discreteness may be defined only using a metrics, therefore not in mereotopologies. The systematic analysis carried out in [MV99] has nevertheless precisely characterized a number of notions like atomicity, divisibility and density in mereotopologies and shown a number of theorems which can guide the choice for a theory in function of specific representational needs.

The use of the open / closed topological distinction may be seen as not commonsensical, yet, it is useful to characterize the commonsensical distinction between mereotopological connection and everyday notion of contact. Connection implies the sharing of a boundary and is appropriate to represent the relation between two joined parts of a one-piece material object, e.g., the connection between a hand and a wrist. Two separated material objects in contact arguably do not share any boundary, as in e.g., the contact between a hand and a pen held in the hand. In [AV93a, AV93b, AV95a] I proposed to define, in an atomic domain, contact between two objects as the impossibility to find another object in between, something characterized in terms of open regions. More generally, I have tried to explore how mereotopologies can be used for commonsensical notions that do not directly correspond to the abstract concepts of topology. In particular, we do refer to boundaries in natural languages, viz. the terms “surface”, “edge”, “point”, but the entities such terms denote arguably are extended just as the material objects they bound and depend of: surfaces can be scratched and points can be blunt (for a detailed analysis of the commonsense concept of surface and related paradoxes see [Gal07]). Together with Michel Aurnague, I have therefore proposed to analyse such “limit” entities in atomic spaces as regions one-atom thick [AV93a, AV93b, AVB97].

The study of mereotopologies and qualitative spatial reasoning is very rich, and the above certainly doesn’t exhaust all relevant issues. There are of course many works dedicated to reasoning on mereotopological representations. A most-used approach is that of “transitivity” or composition tables originally proposed for temporal reasoning [VKvB89, Fre92], adopted in [RCC92a] yielding the so-called RCC-8 calculus,¹¹⁸ whose

¹¹⁷Since our perceptual and cognitive capacities are finite and finiteness implies discreteness, if cognitive space contains some notion of density, it is probably a density in intension.

¹¹⁸The RCC-8 calculus is significantly less expressive than the RCC mereotopology presented in [RCC92b]. In fact, in most mereotopologies (among which Clarke’s theory) the same set of jointly exhaustive and pairwise disjoint relations called RCC-8 can be defined and the formulas encoded in the compositional tables are theorems. Note too that while the calculus defined by the composition table is decidable, most mereotopologies aren’t.

decidability and complexity has been largely explored (see, e.g., [RN99], and [CH01, Gal09] for an overview). This approach is systematized in the study of relational algebras [Dün05], although relational algebras impose a strong reading of compositional tables, therefore restricting their expressivity.¹¹⁹

5.2.2 Beyond mereotopology: geometry and time

Mereogeometries. Whitehead, de Laguna, Nicod and Tarski’s aims were not limited to topological aspects of space. They all wanted to characterize geometry in region-based theories, nowadays called mereogeometries (see [BM10] for a detailed analysis of this area).

Mereogeometries can be approached in two ways. The first approach takes two steps: after defining points as sets of regions in a mereology or mereotopology, one axiomatizes Euclidean geometry on those points for instance with a distance primitive. The construction of points as sets of regions, often with filters [Nic23, Cla85, BG91, AV93b, Roe97], imposes a higher-order language if geometrical relations are to be predicated on points. In the second approach, geometrical relations bear directly on regions, as with the “can connect” ternary relation of [dL22]. Only this second approach is suitable to keep the theory first-order.¹²⁰

One of the most famous work in mereogeometries, Tarski’s [Tar56a], actually is a hybrid method. Tarski uses a properly axiomatized mereological primitive and a geometrical primitive of “sphere”, not axiomatized at all, from which he constructs points as sets of concentric spheres. He then axiomatizes a ternary equidistance primitive relation on points to characterize Euclidean geometry, and at the same time indirectly constrain the sphere primitive.¹²¹ Efforts were made to turn Tarski’s proposal into a theory of the second type, without any point construction in the theory itself. In [BGM96], spheres are defined within a mereotopology extended by a morphological primitive of congruence, which is axiomatized adapting Hilbert’s system. [Ben01] directly axiomatizes Tarski’s primitive of sphere and a proves that the resulting theory is complete with respect to standard \mathbf{R}^n metric space.¹²²

Grounding geometrical relations on the existence of perfect regions like spheres, almost as abstract as points, arguably isn’t cognitively adequate. In my work together with Michel Aurnague [AV93b, AVB97], I have attempted to put forward a theory of the second type without assuming regions having a perfect shape. To the mereotopology

¹¹⁹For instance, the composition of PP with PP yielding PP simply corresponds to the transitivity of proper part ($\forall x, y, z((PP(x, y) \wedge PP(y, z)) \rightarrow PP(x, z))$) in the standard, weak, reading of compositional tables, but in relational algebras, this corresponds to the axiom $\forall x, z(\exists y(PP(x, y) \wedge PP(y, z)) \leftrightarrow PP(x, z))$, thus buying transitivity and density at the same time.

¹²⁰This is of course not guaranteed. For instance, [Ger90] chooses to base mereogeometry on a mereology and a real binary distance function, which is not qualitative, not cognitive, and obviously not first-order axiomatizable.

¹²¹In [Nic23], the same hybrid approach is most radically adopted, as neither his mereological primitive nor his “conjugaison” quaternary geometrical primitive are axiomatized.

¹²²Such a theory is not completely first-order as [Ben01] also aimed at proving categoricity, and the reals are not first-order definable.

described in the previous section, we added orientation primitives adapted from Allen’s interval relations on a line. A different set of relations is required for each relevant axis, corresponding to an implicit projection on those axes. Axes, or straight lines, are abstract entities that are thus added to the domain, but these have an undeniable linguistic relevance. For instance, the three frontal, lateral and vertical axes associated to most oriented objects (e.g., humans) are used in a form or another in all semantic analyses of the so-called projective prepositions (in front of, behind, left of, right of, above, below). We then added a ternary qualitative distance relation “closer to ... than to ...” adapted from [vB83], initially on constructed points [AV93b] and then directly on extended individuals [AVB97]. This proposal proved able to represent the semantics of prepositions and internal location nouns and to support a few useful theorems.

The same approach was pursued further in [ADV97], where instead of introducing straight lines as individuals, an indefinite number of orders was introduced (these are “linear” orders between extended entities, which in one dimension behave like intervals). The theory was compatible with the hypothesis of a finite, thus atomic, domain. On the whole though, this direction of investigation remained unaccomplished, like other proposals in mereogeometry (e.g., [CBGG97] based on a morphological convex closure relation), since no representation theorem has been proved to understand what the models of the theory obtained are.

However, in addition to the significant [Ben01], two more works have brought important results in the field of mereogeometry. First, Maureen Donnelly achieved in her PhD (that I partially supervised) the axiomatization of de Laguna’s ternary “can connect” relation, grasping in a single relation mereology, topology and geometry [Don01]. Such a theory, equipped with a representation theorem to prove it has Euclidean geometry as models, almost satisfies the cognitive requirements we discussed above. It doesn’t assume a higher-order construction of points, nor is founded on regions of perfect shape contrary to [Ben01]. It is though non-atomic, but this is obviously unavoidable if one aims at grasping the density of Euclidean geometry. Second, the extensive analysis of existing mereogeometries done in [BM10] compares the expressivity of the disparate primitives chosen in the various theories (unary *sphere*, binary *congruence*, binary *convex closure*, ternary *can connect*, ternary *closer*, quaternary *conjugaison*). Taking the intended interpretation in the standard \mathbf{R}^n metric space as a basis for comparison (without proving representation theorems though), the authors conclude that all such choices, except the convex closure which is weaker, in principle have the same expressivity and may be axiomatized so as to grasp Euclidean geometry.¹²³

How mereogeometries can be suitably weakened into first-order theories describing atomic domains still is a remaining issue.

¹²³This is by no means trivial though. Early authors like Nicod or Tarski renounced to such an axiomatization. More recently, van Benthem notes in [vB83] that the axiomatics of the closer primitive he provides is too weak to grasp all properties of Euclidean space. [BM10] suggests that such axiomatization can be obtained on the basis of definitions of a primitive in terms of spheres they provide and the axiomatization of geometry in terms of spheres of [Ben01], which is proved to be complete.

Space-time. As expounded in Section 5.1.2, the entities involved in the representation of space I use actually are extended in space-time. This means first of all that temporal relations can be predicated on entities of the same domain of quantification than that of the mereotopology presented above. I have chosen to use Kamp’s temporal theory, exploited in DRT and SDRT and presented in [Kam79]. It is a theory originally proposed for a domain of eventualities (entities extended in time), based on two primitive relations: an overlap for mereological structure and a precedence for partial order.¹²⁴ In a spatio-temporal domain, the relation of connection of the mereotopological theory used also gets a spatio-temporal interpretation. Additional axioms for linking this spatio-temporal connection and the mereological sum operator to the temporal primitives are therefore introduced in [AV93b].

In such an approach there is also the need to temporalize the parthood relation, and all other derived between objects. Parthood relations between objects often are not permanent, e.g., most of the cells of my body when I was born are now gone. The temporalization of parthood is unnecessary for events: one event is wholly part of another, not during a limited time (such a situation would simply be seen as a proper overlap). As evoked in Section 5.1.2 (see also [Mas09]), there are two expressively similar ways (but with different ontological commitments) to temporalize parthood between objects, and I have used both in my papers. One may add a third argument z , a concrete individual,¹²⁵ yielding a ternary $P(x, y, z)$ where z specifies the temporal extent during which x is part of y . One may instead use a “slice” operator and simply relate the temporal parts of the objects with a binary relation as in $P(x@z, y@z)$ where $x@z$ denotes the slice of x during the temporal extent of z . This second option guarantees a uniform binary parthood across the domain of concrete entities (and possibly beyond), but increases the domain of quantification of objects with large numbers of object slices, an ontological commitment unpalatable to many.

In a theory of space-time, the difficult issue of continuity¹²⁶ can be addressed. Continuity is particularly relevant to the approach I have chosen, in which issues of persistence and identity through time are solved by taking a (weak) four-dimensionalism stance. In this approach, identity through time simply is determined by the interpretation function. A constant, say, Caesar, denotes an entity from its coming into existence to its end. So, deciding if two temporally distant entities are the “same” entity, e.g. Caesar, considered at different times, boils down to a simple mereological question, whether or not these entities are temporal parts of Caesar. In a domain with unrestricted mereological sums, the difficult question then is to isolate the sub-domain of entities corresponding to ordinary objects and events. Any two temporally distant entities are temporal parts of a same entity, their sum. But not all sums of (temporal slices of) objects make up

¹²⁴To properly recover Allen’s temporal “meets”, a mereotopological temporal connection is needed, instead of the temporal overlap. This is developed in [Mul98c, Mul98a], a PhD I supervised and which extends my spatio-temporal scheme for increased spatio-temporal expressivity, regarding continuity (see just below) and motion.

¹²⁵Object or event, it doesn’t matter. All are in time and implicitly characterize a temporal extent. Of course, absolutists would use instead directly a temporal interval.

¹²⁶This term refers here to smooth motion and not to the cardinality of the reals.

objects. Ceasar’s left leg during the year 50 BC plus the Eiffel tower during 1950 is not an object. Nor is the temporally connected Ceasar’s left leg during 50 BC plus Ceasar’s right leg during 49 BC. Spatiotemporal “jumps” are not admitted within objects.¹²⁷

Such a notion of continuity within four-dimensional entities can be analyzed in terms of temporal slices and self-connectedness. Philippe Muller defined continuity this way in his PhD, which I supervised [Mul98c, Mul98b, Mul98a]. This notion of continuity also yields a partial account of the continuity of movement, recovering the continuous transitions within “conceptual neighborhoods”¹²⁸ developed for the RCC-8 set of mereotopological relations [CCR92]. The mereotopological approach to continuity developed by Philippe Muller doesn’t suffice though to handle deeper issues in the continuity of movement. Antony Galton has extensively studied different forms of continuity in space-time, extending in particular the conceptual neighborhoods with the notion of “dominance” and arguing for the need of a representation of time with both instants and intervals (see, e.g., [Gal95, Gal97, Gal00, Gal01]).

Philippe Muller has also applied his work on spatio-temporal representation to a spatio-aspectual characterization of classes of motion verbs [MS99].

5.3 Parthood relations

Several authors in formal ontology (e.g., [Sim87, CV99]) have looked at whether and how the general parthood of mereology should be specialized into various less general relations. In fact, mereology has been criticised. Transitivity in particular has been considered as not applying indiscriminately across all domains; variants of the argument in (36) below have repeatedly appeared, suggesting mereology interacts with the notion of collection. This by itself was enough to motivate my interest in the variety of parthood relations. In addition, the analysis of the semantics of spatial prepositions (see Chapter 2) has shown that beyond their common use to locate an object with respect to a separate one, they can be used to situate a part within a whole. I therefore needed to study what exactly are the parthood relations referred to in language.

Constructions referring to parthood relations are quite frequent and diverse in language: general “is part of” expressions (31), genitives (32) and possessive constructions (33), as well as relational nouns (34) and specific verbs (35) in large numbers:

(31) *Ma main fait partie de mon corps* (My hand is part of my body)

(32) *la poignée de la porte* (the handle of the door)

¹²⁷Most objects are not perfectly spatio-temporally “smooth” though: the loss of a hair, a tail or a plank, is not considered to alter the identity of a human, a cat or a ship, while topologically it involves the same sort of discontinuity as cutting a ship into halves. Some form of spatio-temporal continuity certainly is involved in the solution of difficult puzzles like that of the ship of Theseus, but vagueness is here an issue to be addressed. The relative size of the parts lost or gained is to be considered, and this requires a spatio-temporal mereogeometry.

¹²⁸Conceptual neighborhoods have been introduced in [Fre92] for reasoning over temporal relations. They have become popular in qualitative temporal and spatial reasoning, see e.g. [Euz95] for an application to solve granularity issues.

- (33) *Cet atome a un électron* (This atom has one electron)
 (34) *le haut de l'armoire* (the top of the cupboard)
 (35) *Ce plat est constitué de poivrons et de morue* (This dish is made up of pepper and cod)

A semantic analysis of such structures shows that language makes sharp distinctions among parthood relations: a construction cannot be substituted for any other in any contexts at will [ILE88, Aur02, VA07]. It has also been pointed out long ago [Lyo77, Cru86] that there are frequent cases of “intransitivity” of parthood occurrences in language. For instance,

- (36) *Marguerite's tail is part of Marguerite the cow, and Marguerite is part of the herd, but Marguerite's tail is not part of the herd;*

similarly,

- (37) even if *The house has a door* and *The door has a handle* are true, *The house has a handle* cannot describe the same state of affairs (it evokes at best a toy house with a carrying handle on top).

Facing such evidence, and especially concerned by the intransitivity issue, several authors have explored the hypothesis of a multiplicity of parthood relations (e.g., [WCH87, ILE88]), thus departing from the mereological approach and its unique, universal, and transitive relation of (proper)¹²⁹ parthood. [WCH87] postulated that each relation taken separately is transitive, while a combination of two different relations blocks transitivity. [WCH87] has had a large impact on the literature in computational linguistics and in knowledge representation (even without linguistic concern), and many authors aware of the varied behaviour of parthood use schemes of relations that are variations of the set of six relations introduced in that work.

The focus on multiple relations as opposed to the universal mereological relation conceals another important distinction between individual-level parthood and type-level meronymic/holonymic relations. Studies of parthood stemming from mereology address parthood relations between two individuals like in *this wheel is part of this car*. On the other hand, the meronymic/holonymic relations in lexical resources such as WordNet, the “has-part” and “part-of” relations of many computational ontologies, as well as the relations most often studied by cognitive psychologists (e.g., those of [WCH87]) represent generic, often typical, information on lexemes, concepts or kinds that refer to classes of individuals (I will here use the term “type” to globally refer to these), like in *a house has a roof*. Such information can be represented in quantified formulas involving individual parthood (here generically noted P , of course to be adapted with any specific parthood relation): either the meronymic “has-part” $\forall x(A(x) \rightarrow \exists y(B(y) \wedge P(y, x)))$ ¹³⁰ in which

¹²⁹Parthood expressions in natural languages would denote proper parthood rather than the reflexive parthood most often taken as primitive in mereology.

¹³⁰These formulas are very simplified. They omit the actual existence or “presence” predicate required

the whole type A depends on the part type B as in *a car has wheels*, or the holonymic “part-of” $\forall y(B(y) \rightarrow \exists x(A(x) \wedge P(y, x)))$ in which the part type B depends on the whole type A as in *a finger is part of a hand*.¹³¹

My original work on parthood relations initiated in my PhD [Vie91] and has been partially presented in [AV93a, AV93b, AVB97, VA07]. In this work, I showed that the multiplicity of parthood relations is largely explained by the different ontological nature of the arguments. A mereological relation together with the categorization (see Section 5.1.3) into ordinary material objects, events, portions of matter, and substances, and the number distinction into singular entities and plural or collection ones, properly axiomatized, suffice to define four relations: “Member-collection” (e.g., a cow/a herd), “Subcollection-collection” (e.g., Benelux/European Union, missing from [WCH87, ILE88]), “Portion-whole” (e.g., a slice/a cake) and “Substance-whole” (e.g., meat/this stew). This account respects the inferential behavior regarding transitivity, which actually is more complex than what postulated in [WCH87].

The prevalent parthood relation, *functional parthood*, also called “Component-integral whole” or “functional component” (e.g., my hand/my body, the handle/the door), couldn’t be accounted this way though. Indeed, functional parthood presents some famous transitivity puzzles ((38a) would be a transitive occurrence while (38b) an intransitive one) and involves an elusive notion of function.

- (38) a. The jacket has a sleeve, the sleeve has a cuff, and the jacket has a cuff.
 b. The house has a door and the door has a handle but the house does not have a handle.

Only few previous formal accounts of functional parthood were available, and only two of them attempted to explain its various (in)transitivity patterns [Cru86, Mol97]. Some simply assume it is transitive (e.g., [WCH87]) while others assume there are no regularities but a large number of specific functional parthood relations, each of which may be transitive or intransitive (e.g., [CV99]). On the basis of joint work with Michel Aurnague who examined in details the issue of dependencies associated with parthood relations in French and Basque [Aur98, Aur01, Aur02], I completed my analysis of parthood relations with a new account of functional parthood, presented in [VA07, Vie06]. I also show in [Vie06] that the two existing attempts [Cru86, Mol97] at explaining the transitive and intransitive patterns of functional parthood are inadequate because they leave typical cases unexplained.

The new account of functional parthood, an individual-level relation, draws on de-

to correctly account for the generic existential dependence involved [Fin95, Sim87]. They omit the temporal dimension of both the existence and the parthood relations. They also ignore typicality: statements like *a house has a roof* that actually present exceptions are represented just as necessary ones like *an atom has a nucleus*.

¹³¹In some pairs of types both generic dependences are involved, as in *a hand has fingers* and *a finger is part of a hand*, but the two formulas are not equivalent, i.e., a meronymic relation between two types is not the inverse of a holonymic one. This is correctly accounted for in WordNet for instance, but not in FrameNet which doesn’t systematically include both type-level versions for each individual-level parthood relation.

pendence relations akin to those underlying type-level meronymies and holonymies. In fact, the normativity of the notion of function requires to take into account the lexical types used to describe the entities, distinguishing, e.g., *lit* (bed) from *meuble* (piece of furniture). The expression of functional parthood is sensitive to the object descriptions used: although one says *la tête du lit*, *la tête du meuble* is awkward at best. Such lexical types are not considered as ontological categories, but as “social concepts”, dependent on a specific language and the linguistic society that produced it (see Section 6.1 below). As social objects, we can quantify over them (for simplification, they are superficially distinguished here using uppercase letters, e.g., *X*, although I generally do not use a sorted logic).

Functional parthood does not imply that the objects involved are functional and functioning. An engine can be part of a car even when the car is parked and even if the engine is broken. The notion of function required to model functional parthood is quite weak, and amounts to a relationship between the function of the part and that of the whole. This is grasped through a notion of *generic functional dependence*, based on the primitive property of “being functioning as a *X* at time *t*”.¹³² Building on the literature on dependence in formal ontology [Sim87, Fin95], several notions of “functional dependence” are defined: generic functional dependence between lexical types and individual functional dependence between specific entities under a description. The direction of the dependence, from the part to the whole (holonymic) or from the whole to the part (meronymic), and the direct use in the description of the type with which the generic functional dependence is encoded in the lexicon or the indirect use of a subtype,¹³³ yield the distinction of four functional parthood relations.

On the basis of these definitions and the axioms characterizing functional dependence, I show in [Vie06] that only some combinations of the four relations yield transitivity theorems, thus explaining the various patterns observed, for example those in (38). In [VA07], we also show that this analysis predicts the linguistic behaviour of determinative compound nouns based on parthood, explaining why *a car engine* makes sense while *a flower petal* doesn’t.

¹³²For expository reasons, I used times in [VA07, Vie06], but for coherence of the whole scheme, I would rather use a concrete entity instead for its implicit temporal extent as explained in Section 5.2.2.

¹³³This distinction is involved in (38): the type ‘handle’ generically functionally depends on the type ‘object that can be moved or used by hand’ (any handle functioning as a handle requires the existence of a movable object functioning as a movable object), of which ‘door’ only is a subtype, while the type ‘cuff’ generically functionally depends on the type ‘sleeve’ itself.

6 Ontology of social reality

Most application areas of knowledge representation, and therefore of foundational ontologies, concern a reality where humans, groups of humans or societies, the product of their activities, and their conventions are central. Of course, this is the case of any area centered on human interaction, like dialogue modelling, and by itself this motivates my interest in the ontology of social reality. However, this is much more pervasive, and domains ranging from spatial representation for geographic information systems to function modelling for computer aided manufacturing processes are based on concepts depending on social norms.

Issues in ontology of social reality¹³⁴ have received much less attention in formal ontology than the basic problems in the ontology of concrete entities seen in the previous chapter. However, the topic has recently become a hot one. Philosophers [Sea95, Tuo95, Bra99] now meet and work together with cognitive psychologists [CC95] and researchers in artificial intelligence [DL06, LLGH09]. Their work is often anchored to the better explored domain of mental attitudes and the philosophy of mind; in particular, individual and collective action and agency is a subtopic that has been largely addressed for long in logic-based artificial intelligence. We can expect that in the near future such a work on social reality will be reflected in foundational ontologies, and this is indeed the objective that together with the LOA I have been pursuing in the work reported in this chapter.¹³⁵

In agreement with what anticipated in the first version of DOLCE (although not developed), I consider social entities to be concrete entities. Social entities depend on a given society. They are created, adopted, rejected by (members of) a society: they are in time. Therefore, the present chapter actually examines a subcategory of the category of objects seen in the previous chapter. Of course, the focus in the previous chapter was on basic properties and relations applying across all sorts of concrete entities, especially spatio-temporal and material, and these do not suffice for the characterization of social reality. This chapter is dedicated to the properties and relations specific to social objects among all objects. In reality, the area is not yet mature enough to draw a map of general issues and solutions as could be done in the previous chapter. We are approaching the ontology of social reality by addressing a series of specific cases, trying to generalize step by step, at least from a methodological point of view [VBM08]. Such a work will contribute with a significant part of the extension of DOLCE currently under development for a next release.

I have been involved in three topics, arguably significant enough, that I present in this chapter. The first section is devoted to roles and socially defined concepts. Then I

¹³⁴Ontology opposes to epistemology in that metaphysical studies aim at objective representations. This is not in contradiction with the study of social reality, and even mental reality: subjective concepts can be accounted for from an external, objective, point of view.

¹³⁵I'm strongly convinced that this can be developed only on the basis of interdisciplinary collaborations. This motivated the creation of the virtual laboratory ILIKS (Interdisciplinary Laboratory on Interacting Knowledge Systems <http://www.loa-cnr.it/iliks>) that I coordinate together with Nicola Guarino and that gathers French and Italian researchers in philosophy, logic, computer science, cognitive psychology, economics and linguistics. The projects developed within ILIKS involved the two PhD students, one of whom in co-tutorship, that I recently directed (see Section 6.3).

address artefacts and function. The last section deals with action and agency.

6.1 Roles and concepts

The first element of social reality we examined at LOA was that of *relational roles* (henceforth called roles) of objects [MVB⁺04, MG⁺05]. Typical examples of roles are socially relevant notions such as *student*, *president* or *customer*, but *catalyzer* is an example in another domain. As will be seen shortly, roles, and more generally concepts, are properties belonging to social reality. They are treated as social objects, in other words, they are reified properties.

Following the main literature on this topic (see [Ste00a]), roles are considered as dynamic, anti-rigid, and relationally dependent properties. The first two aspects (dynamism and anti-rigidity) regard, respectively, the temporal and modal nature of the relation between roles and their *players*. Entities could play a role only during a specific time interval (in a possible world or set of possible worlds). The third aspect, the relational dependence, characterizes the fact that a property depends via a *pattern of relationships* [Sow88] on additional “external” properties.¹³⁶ Accordingly, roles are standardly taken to be predicates defined on the basis of a relation whose arguments are characterized by specific properties. For example, the role of ‘being a customer’ can be defined as: “a customer is a person that (repeatedly) buys (something) from a company”. In this case, the (unary) predicate ‘being a customer’ is defined on the basis of ‘buy’, ‘being a person’, and ‘being a company’.

The novelty of the approach we introduced in [MVB⁺04] consisted in taking seriously into consideration two additional aspects of roles, their *intensional* and *conventional* nature. The intensional nature relies on the fact that the previous definition not only specifies the extension of the role ‘customer’ but defines what a customer is. While in classical logic two co-extensional predicates are necessarily indistinguishable, one would like to consider that two co-extensional roles are different if they are defined in different ways. The conventional nature implies an existential dependence on some society that produced the conventions, sometimes described as context-dependence. For instance, the role of president (of a country) depends on the existence of that country, but also on the existence of some sort of constitutive text defining what ‘being a president’ means in that country: constraints on who can be player, ways in which players are appointed, norms constraining what the player may or may not do, etc. In addition, this role can be dated: it has been created at some point, and so exists in time.

Staying first-order, these latter aspects can be captured by reifying the roles and the social conventions or contexts that define them, so that the definition and dependence relationships can be expressed. Making roles, and more generally socially defined properties, called here *concepts*, part of the domain of quantification as social objects also enables making justice to their temporal dimension, i.e., their concrete character.

¹³⁶A property X (generically) depends on an external property Y if, necessarily, for every instance x of X there exists an instance y of Y which is an *entity external* to x . The notion of “external entity” is not straightforward. Note though that notions like part, constituent, and quality typically refer to entities that are “internal” to other entities.

The approach we followed in [MVB⁺04] is based on a clear distinction between (i) the properties in the *ground ontology* (in our case, DOLCE), represented as predicates and therefore assumed as static, rigid, extensional, and not explicitly defined or linked to a social context (e.g., the primitive predicates of the theory); and (ii) the properties (called “concepts”) reified at the object level, that are not necessarily static, rigid, and extensional and for which it is possible to explicitly describe some aspects of the social contexts that define them (called “descriptions”). The extension of this approach to social relations appears straightforward.¹³⁷

An important characteristic of roles consists in the possibility of introducing ‘new’ attributes or of hiding attributes of the players: students but not persons have a registration number, passengers but not persons have a flight number, customers but not persons have a code or a purchase number, persons but not customers have weights... Let us consider, for example, the ‘customer code’ attribute. Luc can be (simultaneously) the customer of different companies and therefore he can have different codes, one for each company he is customer of. But if *code* is a function of the person playing the customer role, then Luc can have only one code value. In the case of customer codes, it is possible to solve this issue by modelling *code* as a function with two arguments, the person and the company. However, the problem becomes more serious in the case of a classical puzzle, the *counting passengers problem* [Gup80]. The same person can fly several times the same company, and in this case should be counted as distinct passengers. To solve the problem, one may consider adding a temporal argument for distinguishing passengers. And in other cases, still other arguments.

The approach we followed in [MVB⁺04] and developed in [MGV⁺05] is more general and arguably more powerful. It is based on two originally unrelated proposals in computer science [Ste00a] and in philosophy [Fin82]. This approach is a multiplicative one. It assumes that when an object starts playing a role, this generates a new dependent entity called *qua-entity* (following the terminology in [Fin82]). For instance, “Luc-qua-customer of company A”, “Luc-qua-customer of company B”, “Nixon-qua-Republican”, or “Nixon-qua-Quaker” are separate individuals existentially dependent on Luc or Nixon, and *inherent* in them.¹³⁸ It is such qua-entities that are the bearers of the role attributes, and not the role players.

As we discuss it in [VBM08], using qua-entities may appear excessively multiplicative, and one may wonder whether alternative solutions could be found. For instance, to solve the counting problem we could think to count events instead of qua-entities. Indeed, in the passenger example, the problem disappears counting “carrying events” because

¹³⁷In a linguistically or cognitively motivated ontology, the dividing line between properties and relations considered as *ground*, i.e., objective, and those considered *social*, i.e., dependent on social conventions, may seem very difficult to draw. One could aim at making explicit all social dependences within the ontology itself, and thus model as much as possible as social concepts and relations. Establishing what is the minimum ground level still remains to be done. At the very least, this ground level includes the properties and relations to characterize what a social object is, and to relate a concept to its “instances” and a social relation to its “arguments”.

¹³⁸The inherence relation is typically addressed in trope theory. It is an asymmetrical relation that specializes existential specific dependence. Inherence differs from constitution because the dependence is specific: every time the president gets elected, a new qua-entity is born (see the discussion in [VBM08]).

they have a one-to-one correspondence with qua-entities. However, counting events is not enough when the same person participates in the same single event (or more generically in the same kind of events) with several roles. For instance, let us suppose that Berlusconi participated to some industrial meeting both as Italian Prime Minister and as the President of Mediaset. If we want to count the *representatives* present at the meeting we cannot count persons (just one Berlusconi) and we cannot count events (just one meeting). That qua-entities do participate as such in events appears compelling in examples of this kind and shows that is not so obvious to avoid qua-entities.

6.2 Artefacts

Taking advantage of my experience with the notion of function in analyzing functional parthood (see Section 5.3), I recently addressed the domain of artefacts together with Stefano Borgo [BV09, BV06].

The notion of *physical artefact* is a slippery one even when limiting ourselves to non-agentive artefacts: we all think we know what we mean when talking about these objects and we can provide good examples of them. When asked to name artefacts, most of our examples point to manufactured items. Indeed, the recognition of physical manipulation on an item gives us a strong indication that that entity is an artefact. Nonetheless, physical manipulation is not a key element for artefacts: tons of manipulated entities are definitely not counted as artefacts (e.g., sawdust, cut-off hair, mowed grass) while a moment of thought suffices to find physically unaltered objects that actually make up artefacts (e.g., a pebble used as paperweight, the unworked shells used as money in the past).

The fundamental element to single out artefacts is intentionality [Dip93, Bak04]: we intentionally select objects in order to use them for a purpose perhaps physically modifying them to suit our tasks. Intentionality then is part and parcel of the process of attributing functions to objects, i.e., of the process in which artefacts are created. These observations lead to consider artefacts as ontologically separated from other physical entities like water and trees, and therefore to multiplying entities: the paperweight is not the pebble, it is co-located with it and constituted by it (see Section 5.1.2). Indeed, it can be argued that the pebble does not depend on any creation event, nor on any agent, that it is not meant to hold papers, and that it is older than the paperweight.

The class of artefacts we addressed in [BV09] is the most studied in the philosophical literature [Bak04, KM06, Eld07, Tho07]. It collects entities constituted by either amounts of matter (olive oil, pieces of glass) or physical objects (statues, boats, microchips).¹³⁹ In our approach, artefacts have an ontological status and are essentially the result of an intentional act of an agent (or group of agents) called creator. This intentional act, the creation, is considered as an act of *intentional selection* of the entity to constitute the artefact. Intentional selection is not enough, though. In the example of the paper-weight made of a pebble, the artefact is the result of some agent intentionally

¹³⁹The notion of artefact, though, arguably covers individuals constituted by entities of yet other categories: agentive physical objects (robots), features (speed bumps, folds in a skirt), perdurants (judgements, performances, wars) and more abstract ones like pieces of music, laws or social institutions.

selecting the pebble and *attributing* to it some *capacities* (holding paper without damaging it, being easily grasped by hand, being firm etc). Of course, the artefact might turn out not to have the capacities the agent attributed to it, as it could be flawed or malfunctioning, but that does not affect the existence of the artefact itself.¹⁴⁰

The notion of capacity we adopted building on [Cum75] characterizes the dispositions [Mum98] or behaviors a physical endurant is able to express, independently of any agent, even in the specific case of artefacts. Capacities are a type of individual *qualities*. Individual qualities in DOLCE are reified properties (e.g., the colour of this rose) akin to tropes that are each mapped to a value in the quality space (e.g., the space of colors, the space of times, etc.) that characterizes how they are structured [MBG⁺03, BM09]. The notion of *capacity space* is quite complex and not yet well understood, yet we can say that the value corresponding to the capacity quality of an entity at a given time is a region of the capacity space collecting all the various dispositions the entity is able to express at that time. For instance, the capacity of this pen has now the value of writing finely in black when swept on paper, fitting in one's hand when grasped, making a certain noise when struck on the table...

The *attributed capacity* is a distinct individual quality of physical artefacts only that maps to the *same* space as the capacity quality, and characterizes the purpose or function of the artefact as determined by its creator. The pen above certainly has the attributed capacity to write finely in black when swept on paper and to fit in one's hand when grasped, but most probably not to make a certain noise when struck on the table. The fact that (actual) capacities and attributed capacities are elements of a same space has a number of advantages since it allows to define malfunctioning and reconcile the physical and mental nature of artefacts [KM06].

The formalization of physical artefact proposed in [BV09] makes use of the qualities *capacity* and *attributed capacity* as well as a primitive relation of *intentional selection* which characterizes the events that we called *creations*. We then characterize the subclass of *social* physical artefacts, those whose attributed capacity is socially recognizable and thus depend on a given society, making them social objects. Social recognition goes through the recognition of the intention of the creator to make an artefact of a given type. Types of artefacts are social concepts, seen in Section 6.1. The subclass of social artefacts gathers most of the artefacts that we commonly encounter, bottles, books, pens, cars etc. Actually, some argue that only these really deserve the name of artefact [Dip93]. Finally, products are the special kind of social artefacts that are meant to be exchanged, typically in a commercial transaction [BV06].

¹⁴⁰Establishing the conditions of destruction of an artefact is a delicate issue, since they are not related to disappearing capacities. We assume the artefact persists while the material object constituting it does (this is based on properties like shape and internal structure), but the artefact might see its constituting object change, as when a car gets one wheel substituted. An additional condition of spatio-temporal continuity (see Section 5.2.2), with an appropriate degree of vagueness, is required to solve infamous puzzles like that of the ship of Theseus [BV09, VBM08].

6.3 Action

Action is not a social concept *per se*. The study of action rather belongs to the ontology of the mental realm, as what fundamentally distinguishes actions among events are the intentions of its agent. However, all socially relevant events are actions, and most social relations such as commitment, trust or delegation are related in some way to actions. Obviously, a particular category of social actions, speech acts, is of interest to me. But before being able to address social acts and interaction, a good understanding of what action is is required.

The philosophy of action is a rich topic (see, e.g., [Mel97]), but unfortunately, little has percolated into formal ontology. Action also is a crucial notion for a variety of application domains, e.g., planning, robotics, dialogue systems, and multiagent systems of any kind. Many theoretical accounts have been proposed, but these proposals are often unrelated. A correlate is that no well-developed ontology of action and agency has been incorporated into a foundational ontology till now.

This situation motivated the work developed in two PhD theses of the University of Trento that I directed simultaneously, one completely, that of Robert Trypuz [Try07], and one in co-tutorship with Andreas Herzig, that of Nicolas Troquard [Tro07].

This work focused initially on the relationship between *agency* and *action*, mostly studied separately. Agency is the relationship between an agent (or a group of agents) and the states of affairs it can bring about, without referring to how this is done, i.e., the actions performed. This view yielded a wealth of proposals in modal logic, e.g., [Pör74, Che92, BPX01, Pau02b]. The family of STIT (“Seeing To It That”) logics [Che92, BPX01] is one of the most suitable logical systems dealing with agency, both in terms of expressivity and formal properties.¹⁴¹ As a result, it currently enjoys a large popularity in logic-based AI for modelling multi-agent systems. The semantics of STIT captures the intrinsic link between choices agents make and the causal structure of the flow of time. Choices constrain the possible futures: in the branching time structure underlying STIT models, a choice prunes some branches of the tree.

In STIT and more generally in logics of agency, action remains only implicit. Reducing the ontological commitment is of course positive in principle, but here expressivity is very limited. The only possibility to distinguish actions is on the basis of their actual effects. But different actions can have the same effect (e.g., taking buses of different lines that end up in the same terminal). Moreover, actions may abort before completion or simply fail. So the essence of an action cannot lie in the actual effects obtained. In fact, a large part of the philosophy of action is devoted to other features of actions. Two essential dimensions of actions are missing in logics of agency, the eventive dimension and the intentional dimension.

Davidson’s argumentation in favour of the reification of events, widely accepted in formal semantics, originally concerned actions [Dav67]. Actions are a subclass of events, and as such enjoy all the properties events have. They have spatio-temporal properties,

¹⁴¹One such property is that STIT is more expressive than logics of *ability* that come from computer science and social software, namely ATL and CL [BHT05, BHT06].

notably a duration, something ignored in STIT logic. They are involved in causal links which are not necessarily reduced to actions causing states of affairs to hold, but also actions causing other events. They have participants, which include the agent but are not necessarily reduced to it. They have an internal mereological structure, in particular temporal parts that may or may not be seen as individual actions. They also have a “mode of action”, something we usually specify in language through verbs and adverbs to categorize actions. Just like for events in general, such distinctions are essential to individuate actions, at least for the multiplicativist that distinguishes the spatio-temporally coincident rotation of the earth and its cooling off [PV00]. The categorization of actions also allows us to state that an action of a given type (e.g., making a mayonnaise) was executed while it failed (the mixing didn’t properly result in an emulsion). All this strongly suggests that we need to refer to actions as individuals in order to predicate over them a large variety of properties and relations.

The second missing dimension is the intentional one, well-studied in the philosophy of action (see e.g., [Bra87]). Clearly, not all events are actions. Actions are events in which an *agent* participates. In the least committing sense, an agent is a physical object capable of mental attitudes, among which intentions. Beyond this mere category interpretation, “agent” also evokes the physical causal role of this entity in the event, something that can be grasped as agency. Among events involving an agent, the actual intentions of the agent are what distinguishes actions from reflex movements and unintended accidents: a deliberate jumping is not an accidental falling. Crucially, the intentions of the agent have a causal role in its actions. We distinguish three types of intentions: “outcome intention-that” whose object is a proposition the truth of which is sought, “intention-to-do” whose object is an action intended to be done, and “intention-in-action” whose object is an action being executed.¹⁴²

Intention being standardly modelled in modal logics (see e.g., [CL90]), extensions of agency logics with epistemic and intentional operators have naturally emerged. Similar extensions of Propositional Dynamic Logic (PDL), dynamic logics being a family of logics introduced for reasoning about programs but now also applied to model actions (see e.g., [Mey00]), have been proposed.

The eventive dimension therefore is what appears the most challenging to the agency logic approach. In agency logics, one could model complex actions as sequences of choices, something which could also be used to model action duration. This is the solution adopted in [TV06], where we developed a combination of STIT logic and PDL, PDL-like operators being used to individuate actions. But for actions that are homogeneous processes, like gliding or stirring, decomposing them into basic atomic actions to account for their duration is artificial. In addition, it can be argued that there is no separate intention-to-do for those atomic actions. We introduced dependent action “continuations” in [TV06] to avoid this last problem, but on the whole the approach appears very artificial, and is committed to a discrete time to avoid assuming the agents

¹⁴²These three types result from marrying a logical or AI distinction between outcome intention-that (or intention-to-be) and intention-to-do [CL90, HL04, Woo00] and a philosophical distinction between prior (future-directed) intention to do and intention-in-action (present-directed) [Sea01, Bra87].

are continuously busy deliberating about the actions they do. In any case, the individuation of actions through PDL-like operators does not allow the predication of the variety of properties and relations required by their eventive nature. Moreover, seeing actions as operators prevents them to be properly integrated in a foundational ontology (e.g., DOLCE) in which events are fully acknowledged as citizens (see Chapter 5). Therefore, the development of a first-order theory in which agents and actions are individuals of the domain of quantification cannot be avoided. An advantage of such a first-order account is of course the clarification within the object language of the properties of such individuals, hidden in the semantics with the modal logic approach (see Section 1.3).

The approach I chose with Nicolas Troquard and Robert Trypuz was to start from a first-order order of agency, by axiomatizing in a first-order language the models of STIT. The methodology adopted has thus been to first express the ontological assumptions of STIT in a first-order theory called *OntoSTIT* [TTV06]; it is proved in [Try07] that *OntoSTIT* share its models with STIT, and can therefore be seen as a real equivalent of STIT.¹⁴³ The next step was to extend this theory by enlarging its language and its domain of interpretation. *OntoSTIT+* includes the full predicative description of the states of affairs (instead of simple propositional variables as in STIT) and of course actions as a kind of events [TTV06, Tro07, Try07]. Contrary to most action logics, *OntoSTIT+* has a large expressivity: actions are considered as having a non-null duration and as being non-deterministic, i.e., actions may be suspended before ending or may simply fail. In fact, actions in *OntoSTIT+* are what have been called “attempts” in other theories [TLV07]. Of course, successful actions can be characterized on the basis of their associated outcome-intention.

Indeed, in a final step, the theory has been enriched with the three types of agents’ intentions and beliefs. This enables the characterization of action types by their associated intended outcome [Try07, TLV07]. The overall theory was proved expressive enough to solve the so-called Bratman’s video-game puzzle [Bra87] regarding intentions and actions, a puzzle that makes clear that, even under the assumption that agents are rational, the intention to do an action whose expected outcome is ϕ cannot be reduced to the intention that ϕ obtain [TLV07]. The theory also supports a subcategorization of action events that can be seen as a new ontological counterpart of *aktionsarten*, the aspectual classes of verbs according to linguistic properties (see Section 1.1.1) [TV07].

This theory has been developed so as to be integrated with DOLCE, which takes an absolutists stance with respect to time. Clearly, the branching time structure on which this theory of action is based doesn’t fit the relational approach I have adopted elsewhere as argued in Section 5.1.2. Developing an equivalent theory of action within a relational framework in which time is described only through temporal relations on events would require a non-trivial extension to deal with the implicit alethic modality embedded in the alternative futures of the branching time structures. This modality is the basis of the account of agency within the theory and cannot be dispensed with.

¹⁴³An error appeared in [TTV06]. *OntoSTIT* is proved in [Try07] to be complete with respect to the “bundle trees” built out of the branching tree structures, and not directly the STIT structures themselves, due to a cardinality issue analyzed in [Zan02].

7 Perspectives: The mental and social content of discourse

The two sides of my work, the linguistically-oriented one and the ontologically-oriented one, have up to now met mainly within the realm of space-time and material entities. The perspectives that I envision at this point are to make my linguistic and my ontological investigations meet in other realms. First of all, the specific sub-realm of concrete entities that I have started to study, namely social reality. Next, the mental realm on which I also touched with the intentional aspects of action, and on which social reality depends. Finally, the abstract realm which is little explored in formal ontology-inspired knowledge representation, and which is also required for some areas of social reality.

Those domains still call for ontological studies, especially to integrate them within the larger picture (i.e., within DOLCE). For instance, accounting for the logical structure of certain mental and abstract entities is currently a challenge in first-order foundational ontologies. Both to gain useful empirical data for this task, and to improve our account of linguistic phenomena, I will focus on the needs and the means to enrich an already quite comprehensive theory of discourse semantics like SDRT with mental and social concepts. In addition, I will envisage what a discourse theory can tell to ontology analysis and what applications without linguistic concerns it could find in knowledge representation.

7.1 Enriching SDRT with an explicit account of communicative actions

As seen in Chapter 4, SDRT is a theory that lies on the assumption that discourse has a rhetorical nature; each utterance is acknowledged as being a speech act and not a mere proposition, and discourse relations relate those speech acts, embodying part of their associated public commitment. Classical work on speech acts builds on the fact that, like all actions, they are events with a mental (intentional) component. In addition, they have a remarkable social dimension, that can be best described in terms of public commitments or grounded information [WK95, GHL06] (see Section 4.3). Public commitments associated with speech acts play an important role in [LA09a, LA09b]’s new account of dialogue in SDRT. The evolution of each speaker’s commitment-store is tracked using different SDRS for each, with the hypothesis that an SDRS constitutes in itself a public commitment (to the truth of its assertions, its discourse relations’ semantic effects, etc.). We have also seen in Section 4.3 that accounting for the rhetorical nature of discourse relations, still in terms of public commitments, is important to correctly model monologue as well.

However, that an SDRS is a public commitment and that rhetorical relations involve commitments remains implicit in SDRT. [AL08] does clarify how the semantics of SDRT for dialogue can be extended to include a logic of public commitments. But what is considered in this work is that there is an implicit commitment operator applied to the content of the SDRS as a whole before interpretation. We have seen instead that commitments need to be accounted for within rhetorical relations themselves, with the objective to fully characterize them, and to open the possibility to apply negation *over* them to account for the blocking phenomenon.

In addition, the speech acts themselves appear in SDRSs as constants called “labels”.

These constants are characterized just by their contents and by how they are related by discourse relations to each other. That they represent actions only appears through the dynamic effects captured in the general framework of context-change-potential semantics. They are not considered as events, at least not of the same sort of events as those described within discourse, with temporal properties and participants among other things.

My first general objective is thus to explore how SDRT could be enriched towards an explicit account of communicative actions and their public commitments. This will be done by considering both the eventual nature of speech acts, and the mental and social characteristics of those events.

Speech acts currently are taken into account in SDRT in two different ways. The speech acts underlying any utterance, the main speech acts, are, as just recalled, represented as constituent “labels” and used as arguments of discourse relations. The speech acts described in reported speech (e.g., in *Lea **said** that it’s sunny outside*) or by performative verbs (e.g., in *I’m **telling** you that it’s sunny outside*) are instead represented as Davidsonian events, on which many properties (e.g., temporal, causal) can be predicated. Doing full justice to the main speech acts means considering them as actions and therefore first of all as events, thus unifying these two views. The next step is to take into account what makes these events communicative actions, that is, their mental aspect (speaker’s goals and intentions) and their causal aspect (agency and action effects), the latter having both a mental character (effects on the addressee’s beliefs) and, more importantly, a social character (public commitments), as argued in Section 4.3. This will lead to the full characterization of discourse relations (possibly through definitions), on the basis of a representation of the public commitments associated to speech acts.

7.1.1 Empirical evidence

Taking into account explicitly in SDRT the eventual and rhetorical nature of the speech acts on which the discourse structure is built certainly isn’t straightforward. But such a move would improve the account of some linguistic phenomena, and conversely, studying these will guide the analysis. I will thus first examine the following issues, in particular relying on empirical data from the ANNODIS corpus (see Section 4.7). I of course plan to carry out this work in collaboration with Myriam Bras and Nicholas Asher and my other colleagues of the ANNODIS project.

Metatalk relations. Modelling metatalk relations, also called pragmatic relations, is the most straightforward issue. Such relations involve, instead of the main eventuality of a constituent, the speech act of the current utterance or that of the point of attachment. For instance in (39) the state described in (a) causes the speech act in (b) and is a *Result** case, while in (40), where (b) explains why (a) is claimed, we have an *Explanation** (also sometimes called *Evidence*).

- (39) a. It’s cold.
b. Please close the door.

- (40) a. It has rained,
 b. (because/ \emptyset) the ground is wet.

Metatalk relations are quite frequent in both dialogues and monologues. They have specific markers (e.g. *puisque* (since)) but are often signalled by the same markers as standard semantic relations (e.g., *alors* (then), *donc* (therefore), *car* (because)) [Kno01, JR01]. Being able to refer to the utterance events in SDRT makes the modelling of pragmatic relations quite simple: an alteration of the semantic effects of semantic relations would suffice. Further studies are required though to understand which relations are able to take a pragmatic reading, in which context, and with which argument position being affected.¹⁴⁴ Checking whether the pragmatic reading affects the hierarchical structure of discourse (e.g., altering the default subordinating or coordinating nature of their corresponding semantic relations) is also needed.

Utterance adverbs. Some sentence adverbs expressing a speaker's attitude, like *franchement* (frankly) or *sérieusement* (seriously), are utterance modifiers, i.e., they modify a main speech act event [Mit77].¹⁴⁵ Just like metatalk relations, these adverbs require an SDRT extended with reference to the main speech acts as events. Being sentence adverbs (or IP-adjuncts), when in a preposed position they display in addition the same framing potentialities as locating adverbials in such a position, i.e., they are able to modify a series of speech acts. After analyzing their occurrences in corpus, to model their semantics I will therefore also consider if the topic-based technique used in Section 4.4 is appropriate. In particular, it is not clear that the topic thus inserted should be a standard constituent since the event modified is a main speech act, and at this point it is too early to assume main speech acts may enter standard constituents.

Discourse verbs. Pushing further the study of lexemes that have a discursive semantics, but looking now *inside* the propositional content of clauses, I will examine the role of declarative verbs (or verbs of communication) that make speech acts explicit, like in (41b) to reinforce the simple (41a).

- (41) a. I'm upset
 b. I'm telling you that I'm upset

A closely related class of verbs that I will call rhetorical, make in addition rhetorical links explicit, especially metatalk relations as in (42), a *Result** case.

- (42) The ground is wet. I conclude it has rained.

¹⁴⁴It seems for instance that causal metatalk relations, *Result** and *Explanation**, can only substitute the effect event with a speech act, and not the cause event.

¹⁴⁵This is not always the speech act of the utterance in which the adverb appears. In a question, e.g. *Franchement, que veux-tu faire ?* (Frankly, what do you want to do?) the speech act modified is that of the expected answer [Mar98, BGKM04].

Declarative and rhetorical verbs can be correctly called “discourse verbs” because they genuinely refer to speech acts, as opposed to verbs like *to precede* or *to cause* which denote a temporal or causal relation between events of any type (see the discussion of Danlos’s proposal [Dan06] in Section 4.3).

Discourse verbs can be used as performatives or in reported speech. In the performative case (as the previous two examples), the extended SDRT should make clear what is the relationship between the speech act corresponding to the whole utterance, and the speech act described within the propositional content of this utterance. In a first approximation it is identity, thus relying on language’s self-referring capacity, and considering the matrix clause to be simply redundant. Pragmatic differences like the reinforcement effect and structural effects (in subordinating / coordinating terms) should be examined though in details on corpus.

A different approach is needed for reported speech, which is very frequent in news, as in (43).

- (43) President Obama said nations of the world are failing to work together to confront the most pressing challenges of today.

Following the tradition in RST (see e.g., [RE06]), reported speech is specifically annotated in ANNODIS with an *Attribution* relation attaching the embedded clause to its matrix clause. This is done without distinguishing the standard intensional use of reported speech, as in (43), and its “evidential” use, as in (44) [HARD06].

- (44) Lea said that it has rained. Yet, the ground is dry.

In principle, the relation *Attribution* is reserved for the intensional use, whereas the evidential use is modelled with a relation of *Evidence* —actually *Explanation**. In fact, the very existence of the *Attribution* relation may be questioned if one considers that the main eventuality of the matrix clause (the speech act described) is identical to the speech act of the embedded clause. The embedded clause is not part of the main discourse and doesn’t play any rhetorical role with respect to it: it specifies only the propositional content of the speech act described in the matrix clause, and just as with belief attributions, it is embedded in an opaque context.¹⁴⁶ On the other hand, in the evidential use, the reported speech clearly interacts with the structure of the main discourse. In (44), the rhetorical role of the second sentence is to mark a contrast with the content of the reported speech, not with the matrix clause of the first sentence. In this case, the reported speech act appears within the main discourse structure, and is attached to the previous context instead of the matrix clause. The matrix clause plays a secondary role; it is attached to the reported speech with *Explanation**, for it constitutes the reason why the content of the reported speech is in some sense endorsed by the speaker of the main discourse.

¹⁴⁶There however still is the need to analyse the internal discourse structure of the reported speech that may span over several segments.

Although questionable, annotating all reported speech cases with an *Attribution* relation will enable studying these two cases thoroughly. The boundary between the two cases is not clear-cut, and in a long reported speech, the two perspectives alternate, making the analysing of discourse structure and speech acts themselves a challenge. Complementing the extensive descriptive literature on speech acts and enunciation, it is expected that such a corpus analysis will bring some light on the status of speech acts in discourse, on how exactly the main speech acts should be accounted for within the SDRSs, and on the interactions between different levels of enunciation in discourse.

7.1.2 Theoretical account

From a theoretical point of view, the generalization of such empirical studies should bring me to examine two general issues, and, of course, to consider how the extension of SDRT to a new account of speech acts could be done.

Self-referring discourse and forward-looking phenomena. Studies on the interaction between different discourse levels, the main discourse and discourses described within the propositional content, especially the self-referencing cases with performatives, should serve as a basis to generalize the analysis of forward-looking phenomena in discourse, like locating adverbials in IP-adjunct position (see Section 4.6.1) or forward-looking *Background* (see Section 4.5.2). Indeed, other forward-looking phenomena make use of much more explicit means to refer to future discourse, but without explicit performatives. For instance, full clauses can describe the structure of a discourse segment to come, as in “initiation clauses” for enumeration structures like in (45), studied by [BPVC08, Bra08].

- (45) *Deux conséquences importantes en découlent :*
 - *Premièrement, ... ;*
 - *Deuxièmement,*

I will for this rely on a close collaboration with Myriam Bras, who is currently studying structural markers in discourse, especially adverbials (e.g. *d’abord, premièrement, dans un premier temps* in enumerative structures). A challenge is of course to identify when there is a real anticipation in discourse as in (45) (which would be incoherent with only the first clause), and when we simply have an *Elaboration* structure as in (11) on p.40; indeed, example (10) seems to have an intermediary status. Detecting initiation clauses is much more subtle than detecting the successive items, because they are not marked by sentence adverbials but for instance by plural determiners.

Locus of structural and rhetorical markers. The baseline hypothesis formulated at the end of Section 4.3 that direct discourse markers cannot be found inside the clauses, but externally, for instance among clause modifiers, finds its limits with the “initiation clauses” just seen and with declarative or rhetorical verbs used as performatives or in reported speech. Such “discourse verbs” constitute a bridge between the propositional

level (they characterise the main eventuality of the matrix clause) and the discourse level (they characterise the speech act of the embedded clause which in some cases plays a role in the main discourse –see above). To generalize those cases and refine the baseline hypothesis, paying attention to controversial syntax problems regarding what is external to the clause and what is internal to it (see e.g., “incident” elements of [Mar98] and the “macrosyntax” approach of [Deu03]; see also Section 4.5.3) is needed.

In contrast with discourse verbs, verbs like *to precede* or *to cause* do not denote speech acts and in principle have no effect at the discourse level (see previous section and my criticism of [Dan06] in Section 4.3). Matters are probably more complex though since, as shown in [Bra08], there is a continuum between the description of a fact that an event causes another with purely propositional means and the implication of the same fact through the semantics effects of a causal discourse relation:

- (46) a. *Le départ de Marie a causé celui de Pierre.*
 b. *Marie est partie, ce qui a causé le départ de Pierre.*
 c. *Marie est partie. Cela a causé le départ de Pierre.*
 d. *Marie est partie. A cause de cela, Pierre est parti.*
 e. *Marie est partie. De ce fait, Pierre est parti.*

The hypothesis above will thus be reconsidered taking into account both syntax and lexical semantics, with a focus on clarifying the range of linguistic means to mark discourse structure (see previous issue) and rhetorical relations, the latter having been initiated in [BLDV01a, BLDV01b] (see Section 4.5.1) and being currently pursued by Myriam Bras [Bra08]. This will be done of course in collaboration with Myriam Bras, and building on the conclusions of Section 4.3. It is essential to distinguish and pay attention to the three contributions of a discourse relation: its semantic effects at the propositional level, its subordinating / coordinating nature that determines the graph’s right frontier, and its rhetorical character with its embodying of a public commitment.

Representing speech acts and their commitments within SDRSs. Considering the main speech acts in SDRT as events means representing their action character not in an implicit way, as, e.g., through their effects only in action logics (or dynamic logics), but in an explicit way, as e.g., in a first-order theory of action (see Section 6.3).¹⁴⁷ This means that I will need to investigate, with Nicholas Asher and Alex Lascarides, how the dynamic effects of speech acts, now accounted for by a variant of public announcement logic in the semantics of dialogue SDRSs [AL08], can be made explicit at the language level, and also for monologues. In fact, we have seen above that those effects, essentially

¹⁴⁷I will not consider here all aspects of actions. In particular, I will ignore action progress and duration. This might be necessary to account for dialogues or multilogues in which several speakers talk simultaneously; however, this would require an account of the incremental interpretation of sentence fragments, which would take me too far from my present competences (see e.g., [SS09] on this issue). Similarly, as I focus on monologue texts (the ANNODIS corpus doesn’t include spoken dialogues) I will not consider the fact that speech act may fail, as, e.g., when the message is not even perceived or is misunderstood.

in terms of public commitments, need to appear *within* the SDRSs, especially *within* discourse relations.

The formal ontology of action I have begun to develop (see Section 6.3) will have to be supplemented by a representation of commitments. It is not yet clear whether it will be possible to avoid the very difficult issue of how to represent within a first-order language entities such as mental attitudes or public commitments, maintaining enough inference power. In any case, I plan to start from an adequate commitment logic, either to directly use a commitment operator within the SDRS language, or to use it as a starting point for developing a commitment theory (applying the same strategy as for action in [TTV06]). An adequate commitment logic still needs to be found. Indeed, since discourse makes the difference between the blocking of a relation and not expressing anything in this respect (see footnote 61 on p.45), a modal logic with Euclidean axiom 5, i.e., where not committing to ϕ (not expressing anything) is equivalent to committing to not committing to ϕ (e.g., using a marker of relation blocking), is too strong for modelling the public commitments associated with an assertion. The proposals in [GHL06] and [AL08] need to be weakened.¹⁴⁸

Commitment is of course related with other mental and social notions, and in fact, [GHL06] relates commitment with the social notion of public groundedness and the mental notion of (private) belief. In the larger picture, social notions like delegation, trust and social roles probably need to be involved as well. For instance, one can expect that rhetorical links involving requests are related to action delegation, and that those involved in suggestions imply a public assumption that the addressee trusts the speaker for the topic at hand. My colleagues at both IRIT and LOA-ISTC have significant experience in this area, with complementary logical and ontological approaches and I myself have already worked on roles (see Section 6.1). I expect to work in the future on how this area could constitute another possibility of extension of SDRT.

7.2 Contribution of discourse semantics to knowledge representation

In parallel to work described above, I will investigate interactions in the opposite direction, i.e., from discourse semantics to knowledge representation.

I will first consider what discourse analysis can tell to the ontology of abstract entities, as part of the more general aim of extending DOLCE for its next release (see the introduction to Chapter 6).

In another line of work, I plan to start applying SDRT to knowledge representation in an e-government project.

Extending DOLCE to abstract entities. As shown in [MVB⁺04] (see Section 6.1), concepts, i.e., properties seen as entities belonging to social reality, are defined through descriptions. Descriptions (e.g., scientific theories characterizing a scientific concept,

¹⁴⁸In particular, commitment cannot be defined, as done in [GHL06], as the public grounding of a belief if one assumes that both public grounding and belief are KD45 modalities, since the Euclidean character of commitment then comes as a theorem.

constitutions defining institutional roles, novels describing kinds of fictional entities, (implicit or not) dictionary entries defining a lexeme of some language) also are entities belonging to social reality. They are created at some point, they can be socially adopted or refuted, they are thus in time, and not abstract at all. On the other hand, descriptions do have abstract *content*. Such content allows inferences, comparisons, and translations. For instance, it is on the basis of this abstract content that the concept of democracy used by the ancient Greeks can be compared to the concepts of democracy adopted in different current societies.

Introducing such abstract entities in a foundational ontology like DOLCE is thus necessary, not only for the sake of obtaining a comprehensive account of what there is, but even to satisfactorily characterize social entities, thus contributing together with my colleagues at LOA to the current effort to update and extend DOLCE. The difficulty here is not so much to analyse the nature of abstract entities and situate them within the large picture, relating them to other realms. It is rather to account for the logical structure of description content, required to perform the formal comparisons and inferences alluded to above, while remaining within a first-order framework.

To address this difficult issue, I plan to rely on the one hand on [Ash93] which shows that some abstract objects can be modelled as the content of an SDRS, thus accounting for part of its structure. On the other hand, to represent logical propositions in a first-order language, I will consider making use of underspecification techniques nowadays widespread in linguistics, from syntax to semantics and discourse. The Glue Logic of SDRT exploits such a language (see Section 4.1) building on Underspecified DRT [Rey93]. I will first evaluate which of these languages (e.g., those developed in SDRT, Constraint Language for Lambda Structures [EKN01], or Minimal Recursion Semantics [CFPS05]; see [Bun08] for a comparison) could be adequate to my needs. Then I should assess what sort of operations on representations of description content are possible, and how this supports reasoning about social entities.

Discourse relations for document workflow in e-government. Finally I will investigate how SDRT can be applied in a domain for which discourse semantics doesn't seem relevant at first sight, e-government. In collaboration with Nicola Guarino, I will work on a project of management of administrative processes, their associated document workflow, and the dossiers collecting such documents.

Managing large administrative dossiers (e.g., for healthcare, law, public inquiries) is very costly. At present, document workflow softwares only handle multiple accesses and document versioning and retrieving issues. We assume that dossier management and the associated public services and administrative processes will improve if the administration could perform integrity checks on the contents of dossiers, so as to know at any time whether some issues remain unsettled (e.g., a placed query have not been replied to yet), whether there are controversies (e.g., two expert reports argue for opposite conclusions), and in general what are the semantic dependencies between documents. This would improve document retrieval, on the basis of arguments rather than only keywords, and possibly support reasoning.

The approach we propose is then to handle dossiers as if each document constituted a speech act, and represent rhetorical relations between them. To the best of our knowledge, rhetorical relations *across* documents have not been studied. We expect most such relations to be of dialogue and argumentation type: question/answer pair (many forms and reports are either queries or answers to queries), correction, acknowledgement, etc.

The LOA is currently a contractor in a project of computer science for law. The applications we envision concern trial dossiers, jurisprudence, and norm consolidation processes.

References

- [AAB⁺94] Nicholas Asher, Michel Aurnague, Myriam Bras, Pierre Sablayrolles, and Laure Vieu. Computing the spatiotemporal structure of discourse. In H. Bunt, R. Muskens, and G. Rentier, editors, *International Workshop on Computational Semantics*, pages 11–20. ITK, Tilburg, 1994.
- [AAB⁺95] Nicholas Asher, Michel Aurnague, Myriam Bras, Pierre Sablayrolles, and Laure Vieu. De l’espace-temps dans l’analyse du discours. *Sémiotiques*, 9:11–62, 1995.
- [AABV93] Nicholas Asher, Michel Aurnague, Myriam Bras, and Laure Vieu. Space, time and discourse. In Franck Anger, Hans W. Guesgen, and Johan van Benthem, editors, *IJCAI-93 Workshop on Spatial and Temporal Reasoning*, pages 219–235, 1993.
- [AABV95] Nicholas Asher, Michel Aurnague, Myriam Bras, and Laure Vieu. Spatial, temporal and spatio-temporal locating adverbials in discourse. In Pascal Amsili, Mario Borillo, and Laure Vieu, editors, *Time, Space and Movement. Meaning and Knowledge in the Sensible World, Workshop Notes of the 5th International Workshop TSM’95*, pages A107–119. IRIT, Toulouse, 1995.
- [ABVA01] Michel Aurnague, Myriam Bras, Laure Vieu, and Nicholas Asher. The syntax and semantics of locating adverbials. *Cahiers de Grammaire*, 26:11–35, 2001.
- [ACV⁺07] Michel Aurnague, Maud Champagne, Laure Vieu, Andrée Borillo, Philippe Muller, Jean-Luc Nespoulous, and Laure Sarda. Categorizing spatial entities with frontal orientation: The role of function, motion and saliency in the processing of the french internal localization nouns avant/devant. In Michel Aurnague, Maya Hickmann, and Laure Vieu, editors, *Categorization of Spatial Entities in Language and Cognition*, pages 153–175. John Benjamins, Amsterdam, 2007.
- [ADV97] Nicholas Asher, Maureen Donnelly, and Laure Vieu. A mereotopologically-based commonsense geometry with directions and relative distance. In *IJCAI’97*, 1997. poster.
- [AER04] Jennifer Austin, Stefan Engelberg, and Gisa Rauh. *Adverbials. The interplay between meaning, context, and syntactic structure*. John Benjamins, Amsterdam, 2004.
- [AGDS08] Nathalie Aussenac-Gilles, Sylvie Despres, and Sylvie Szulman. The TERMINAE method and platform for ontology engineering from texts. In Paul Buitelaar and Philipp Cimiano, editors, *Bridging the Gap between Text and Knowledge: Selected Contributions to Ontology Learning and Population from Text*, pages 199–223. IOS Press, 2008.
- [AH85] James Allen and Patrick Hayes. A common-sense theory of time. In Aravind Joshi, editor, *Proceedings of the 9th IJCAI*, volume 1, pages 528–531. Morgan Kaufmann, 1985.
- [AHB01] Nicholas Asher, Daniel Hardt, and Joan Busquets. Discourse parallelism, ellipsis and ambiguity. *Journal of Semantics*, 18(1):1–25, 2001.
- [AHV05] Michel Aurnague, Maya Hickmann, and Laure Vieu. Les entités spatiales dans la langue: étude descriptive, formelle et expérimentale de la catégorisation. In Catherine Thinus-Blanc and Jean Bullier, editors, *Agir dans l’espace*, pages 217–232. Editions de la Maison des sciences de l’homme, Paris, 2005.

- [AHV07a] Michel Aurnague, Maya Hickmann, and Laure Vieu, editors. *Categorization of Spatial Entities in Language and Cognition*. Human Cognitive Processing 20. John Benjamins, Amsterdam, 2007.
- [AHV07b] Michel Aurnague, Maya Hickmann, and Laure Vieu. Introduction: Searching for the categorization of spatial entities in language and cognition. In Michel Aurnague, Maya Hickmann, and Laure Vieu, editors, *Categorization of Spatial Entities in Language and Cognition*, pages 1–32. John Benjamins, Amsterdam, 2007.
- [AK08] Maria Averintseva-Klisch. German right dislocation and afterthought in discourse. In Anton Benz and Peter Kühnlein, editors, *Constraints in Discourse*, pages 225–247. John Benjamins, 2008.
- [AL95] Nicholas Asher and Alex Lascarides. Lexical disambiguation in a discourse context. *Journal of Semantics*, 12(1):69–108, 1995.
- [AL98a] Nicholas Asher and Alex Lascarides. Bridging. *Journal of Semantics*, 15(1):83–113, 1998.
- [AL98b] Nicholas Asher and Alex Lascarides. Questions in dialogue. *Linguistics and Philosophy*, 21(3):237–309, 1998.
- [AL98c] Nicholas Asher and Alex Lascarides. The semantics and pragmatics of presupposition. *Journal of Semantics*, 15(3):239–300, 1998.
- [AL03] Nicholas Asher and Alex Lascarides. *Logics of Conversation*. Cambridge University Press, Cambridge, 2003.
- [AL08] Nicholas Asher and Alex Lascarides. Commitments, beliefs and intentions in dialogue. In *Proceedings of the 12th Workshop on the Semantics and Pragmatics of Dialogue (Londial)*, pages 35–42, London, 2008.
- [All84] James Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23(2):123–154, 1984.
- [ANN08] ANNODIS, French ANR project on discourse annotation. <http://w3.erss.univ-tlse2.fr/textes/pagespersos/annodis/ANNODISen.html>, 2008.
- [APV07] Nicholas Asher, Laurent Prévot, and Laure Vieu. Setting the background in discourse. *Discours*, 1, 2007. On-line journal, available at <http://discours.revues.org/index301.html>.
- [Arm89] David Armstrong. *Universals: An opinionated introduction*. Westview Press, 1989.
- [Ash93] Nicholas Asher. *Reference to abstract objects in discourse*. Kluwer Academic Publishers, Dordrecht, 1993.
- [Ash00] Nicholas Asher. Truth-conditional semantics for parentheticals. *Journal of Semantics*, 17(1):31–50, 2000.
- [Ash04] Nicholas Asher. Discourse topic. *Theoretical Linguistics*, 30(2-3):161–201, 2004.
- [Ash07] Nicholas Asher. Objects, locations and complex types. In Michel Aurnague, Maya Hickmann, and Laure Vieu, editors, *The Categorization of Spatial Entities in Language and Cognition*, pages 337–361. John Benjamins, 2007.

- [Ash10] Nicholas Asher. *Lexical Meaning in Context: A web of words*. Cambridge University Press, 2010.
- [Aur93] Jean-Christophe Aurisset. *Intégration de facteurs pragmatiques dans la synthèse déclarative d'images dans le cadre du projet VILAIN*. PhD thesis, Université Paul Sabatier, 1993.
- [Aur95a] Michel Aurnague. L'expression de l'espace en basque : à propos du génitif et de l'inessif. *Linguisticae Investigationes*, 19(1):15–55, 1995.
- [Aur95b] Michel Aurnague. Orientation in French spatial expressions: Formal representations and inferences. *Journal of Semantics*, 12(3):239–267, 1995.
- [Aur95c] Michel Aurnague. Some remarks about the notion of location. Manuscript, 1995.
- [Aur96] Michel Aurnague. Les noms de localisation interne : Tentative de caractérisation sémantique à partir de données du basque et du français. *Cahiers de Lexicologie*, 69(2):159–192, 1996.
- [Aur98] Michel Aurnague. Basque genitives and part-whole relations: Typical configurations and dependences. *Carnets de Grammaires 1*, ERSS, Université de Toulouse 2, 1998.
- [Aur01] Michel Aurnague. *Entités et relations dans les descriptions spatiales : l'espace et son expression en basque et en français*. Habilitation à diriger les recherches, Université de Toulouse-Le Mirail, 2001.
- [Aur02] Michel Aurnague. Relations de partie à tout, configurations typiques et dépendances : Analyse sémantique de quelques constructions génitives du basque. *Revue de Sémantique et Pragmatique*, 11:69–85, 2002.
- [Aur04] Michel Aurnague. *Les structures de l'espace linguistique : Regards croisés sur quelques constructions spatiales du basque et du français*. Peeters, Leuven, 2004.
- [AV93a] Michel Aurnague and Laure Vieu. A three-level approach to the semantics of space. In Cornelia Zelinski-Wibbelt, editor, *Semantics of Prepositions: From Mental Processing to Natural Language Processing*, Natural Language Processing 3, pages 393–439. Mouton de Gruyter, Berlin, 1993.
- [AV93b] Michel Aurnague and Laure Vieu. Towards a formal representation of space in language: A commonsense reasoning approach. In Franck Anger, Hans W. Guesgen, and Johan van Benthem, editors, *IJCAI-93 Workshop on Spatial and Temporal Reasoning*, pages 123–158, Chambéry, 1993.
- [AV95a] Nicholas Asher and Laure Vieu. Toward a geometry of common sense: A semantics and a complete axiomatization of mereotopology. In Chris Mellish, editor, *Proceedings of the 14th IJCAI*, pages 846–852. Morgan Kaufmann, San Mateo, CA, 1995.
- [AV95b] Michel Aurnague and Laure Vieu. A theory of space-time for natural language semantics. In Kepa Korta and Jesús María Larrazabal, editors, *Semantics and Pragmatics of Natural Language: Logical and Computational Aspects, Proceedings of the Toulouse-Donostia Workshop*, ILCLI Series 1, pages 69–126. ILCLI, Donostia-San Sebastian, 1995.

- [AV01] Nicholas Asher and Laure Vieu. Subordinating and coordinating discourse relations. In *First International Workshop on Semantics, Pragmatics and Rhetoric (SPR'01)*, pages 57–64. ILLI San Sebastian, 2001.
- [AV05] Nicholas Asher and Laure Vieu. Subordinating and coordinating discourse relations. *Lingua*, 115(4):591–610, April 2005.
- [AVB97] Michel Aurnague, Laure Vieu, and Andrée Borillo. La représentation formelle des concepts spatiaux dans la langue. In Michel Denis, editor, *Langage et cognition spatiale*, Collection Sciences Cognitives, pages 69–102. Masson, Paris, 1997.
- [BA94] Myriam Bras and Nicholas Asher. Le raisonnement non-monotone dans la construction de la structure temporelle de textes en français. In *Actes de RFIA '94*, volume 2, pages 223–234. AFCET, Paris, 1994.
- [Bac86a] Emmon Bach. The algebra of events. *Linguistics and Philosophy*, 9:5–16, 1986.
- [Bac86b] Emmon Bach. Natural language metaphysics. In Ruth Barcan Marcus, Georg J.W. Dorn, and Paul Weingartner, editors, *Logic, Methodology and Philosophy of Science*, volume VII, pages 573–595. North Holland, Amsterdam, 1986.
- [Bak04] Lynne Rudder Baker. The ontology of artifacts. *Philosophical explorations*, 7(2):99–111, 2004.
- [Bak07] Lynne Rudder Baker. *The Metaphysics of Everyday Life: An Essay in Practical Realism*. Cambridge University Press, 2007.
- [Bar95] Chris Barker. *Possessive Descriptions*. Lecture Notes. CSLI, 1995.
- [BB03] Patrick Blackburn and Johan Bos. Computational semantics. *Theoria*, 18(1):27–45, 2003.
- [BB05] Patrick Blackburn and Johan Bos. *Representation and Inference for Natural Language. A First Course in Computational Semantics*. Studies in Computational Linguistics. CSLI Publications, Stanford, 2005.
- [BBLD⁺04] Andrée Borillo, Myriam Bras, Anne Le Draoulec, Laure Vieu, Arie Molendijk, Henriëtte de Swart, Henk Verkyul, Co Vet, and Carl Vettters. Tense, connectives and discourse structure. In Henriëtte de Swart and Francis Corblin, editors, *Handbook of French Semantics*, pages 309–348. CSLI, Stanford, 2004.
- [BCM⁺03] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.
- [Ben75] David Bennett. *Spatial and temporal uses of English prepositions: An essay in stratificational semantics*. Longman, 1975.
- [Ben01] Brandon Bennett. A categorical axiomatization of region-based geometry. *Fundamenta Informaticae*, 46(1-2):145–158, 2001.
- [BG91] Loredana Biacino and Giangiacomo Gerla. Connection structures. *Notre Dame Journal of Formal Logic*, 32(2):242–247, 1991.
- [BGKM04] Olivier Bonami, Danièle Godard, and Brigitte Kampers-Manhe. Adverb classification. In Francis Corblin and Henriëtte de Swart, editors, *Handbook of French Semantics*, pages 143–183. CSLI Publications, 2004.

- [BGM96] Stefano Borgo, Nicola Guarino, and Claudio Masolo. A pointless theory of space based on strong connection and congruence. In L. Carlucci Aiello and S. Shapiro, editors, *KR'96, Principles of Knowledge Representation and Reasoning*, pages 220–229. Morgan Kaufmann, San Mateo (CA), 1996.
- [BHR95] John A. Bateman, Renate Henschel, and Fabio Rinaldi. Generalized Upper Model 2.0: documentation. Technical report, GMD/Institut für Integrierte Publikations- und Informationssysteme, Darmstadt, Germany, 1995.
- [BHS09] Franz Baader, Ian Horrocks, and Ulrike Sattler. Description logics. In Steffen Staab and Rudi Studer, editors, *Handbook on Ontologies*, pages 21–43. Springer, 2nd edition, 2009.
- [BHT05] Jan Broersen, Andreas Herzig, and Nicolas Troquard. From Coalition Logic to STIT. In Wiebe van der Hoek, Alessio Lomuscio, Erik de Vink, and Mike Wooldridge, editors, *Int. Workshop on Logic and Communication in Multiagent Systems (LCMAS 2005)*, Electronic Notes in Theoretical Computer Science. Elsevier, 2005.
- [BHT06] Jan Broersen, Andreas Herzig, and Nicolas Troquard. Embedding Alternating-time Temporal Logic in strategic STIT logic of agency. *Journal of Logic and Computation*, 16(5):559–578, 2006.
- [Bie88] Manfred Bierwisch. On the grammar of local prepositions. In Manfred Bierwisch, Wolfgang Motsch, and Ilse Zimmermann, editors, *Syntax, Semantik, und Lexicon*, number 29 in *studia grammatica*, pages 1–65. Akademie Verlag, Berlin, 1988.
- [BLDA06] Myriam Bras, Anne Le Draoulec, and Nicholas Asher. Evidence for a scalar analysis of Result in SDRT from a study of the French connective *alors*. In *Proc. of the SPRIK Conference on Explicit and Implicit Information in Text - Information Structure accross Languages*, pages 75–79, Oslo, 2006.
- [BLDV01a] Myriam Bras, Anne Le Draoulec, and Laure Vieu. French adverbial *puis* between temporal structure and discourse structure. In Myriam Bras and Laure Vieu, editors, *Semantic and Pragmatic Issues in Dialogue: Experimenting with Current Theories*, CRiSPI, pages 109–146. Elsevier, Oxford, 2001.
- [BLDV01b] Myriam Bras, Anne Le Draoulec, and Laure Vieu. Temporal information and discourse relations in narratives: The role of French connectives ‘puis’ and ‘un peu plus tard’. In ACL, editor, *Proceedings of the ACL'2001 Workshop on Temporal and Spatial Information Processing*, pages 49–56, Toulouse, July 2001.
- [BLDV03] Myriam Bras, Anne Le Draoulec, and Laure Vieu. Connecteurs et temps verbaux dans l’interprétation temporelle du discours : Le cas de *puis* en interaction avec l’imparfait et le passé simple. *Cahiers Chronos*, 11:71–97, 2003.
- [BM93] Myriam Bras and Frédérique Molinès. Adverbials of temporal location: Linguistic description and automatic processing. In Józef Darski and Zygmunt Vetulani, editors, *Proceedings of 26th Colloquium of Linguistics: Language, Computation, Communication*, Linguistische Arbeiten 293, pages 137–146. Niemeyer, Tübingen, 1993.
- [BM09] Stefano Borgo and Claudio Masolo. Foundational choices in DOLCE. In Steffen Staab and Rudi Studer, editors, *Handbook on Ontologies*, pages 361–381. Springer, 2nd edition, 2009.

- [BM10] Stefano Borgo and Claudio Masolo. Full mereogeometries. *The review of symbolic logic*, in press, 2010.
- [BMCVW92] Laurent Buisson, Roger Martin-Clouaire, Laure Vieu, and Jean-Luc Wybo. Artificial intelligence and environmental protection: A survey of selected applications in France. In R. Aiken, editor, *Education and Society - Information Processing 92, Proceedings of IFIP'92*, volume II. Elsevier, Amsterdam, 1992.
- [Bob84] Daniel Bobrow. Qualitative reasoning about physical systems: An introduction. *Artificial intelligence*, 24(1-3):1–5, 1984.
- [BPVC08] Myriam Bras, Laurent Prévot, and Marianne Vergez-Couret. Quelle(s) relation(s) de discours pour les structures énumératives ? In J. Durand, B. Habert, and B. Laks, editors, *Actes du Colloque Mondial de Linguistique Française (CMLF'08)*, pages 1945–1964. Institut de Linguistique Française, 2008.
- [BPX01] Nuel Belnap, Michael Perloff, and Ming Xu. *Facing the future: Agents and choices in our indeterminist world*. Oxford University Press, Oxford, 2001.
- [Bra87] Michael E. Bratman. *Intentions, plans, and practical reason*. Harvard University Press, Cambridge, 1987.
- [Bra99] Michael E. Bratman. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge University Press, Cambridge, 1999.
- [Bra08] Myriam Bras. *Entre relations temporelles et relations de discours*. Habilitation à diriger les recherches, Université de Toulouse, 2008.
- [BTF07] John Bateman, Thora Tenbrink, and Scott Farrar. The Role of Conceptual and Linguistic Ontologies in Discourse. *Discourse Processes*, 44(3):175–213, 2007. Special Issue on Dialogue Modelling: Computational and Empirical Approaches.
- [Bun08] Harry Bunt. Semantic underspecification: Which technique for what purpose? In Harry Bunt and Reinhard Muskens, editors, *Computing Meaning*, volume 3, pages 55–85. Springer, 2008.
- [BV95] Yves Bestgen and Wietske Vonk. The role of temporal segmentation markers in discourse processing. *Discourse Processes*, 19:385–406, 1995.
- [BV00] Yves Bestgen and Wietske Vonk. Temporal adverbials as segmentation markers in discourse comprehension. *Journal of Memory and Language*, 42:74–87, 2000.
- [BV06] Stefano Borgo and Laure Vieu. From physical artefacts to products. In Roberta Cuel and Roberta Ferrario, editors, *Formal Ontology Meets Industry (FOMI 2006)*, pages 85–99, Trento, 2006. 14-15 decembre.
- [BV09] Stefano Borgo and Laure Vieu. Artefacts in formal ontology. In Anthonie Meijers, editor, *Handbook of Philosophy of Technology and Engineering Sciences*, pages 273–308. Elsevier, 2009.
- [BVA01] Joan Busquets, Laure Vieu, and Nicholas Asher. La SDRT : Une approche de la cohérence du discours dans la tradition de la sémantique dynamique. *Verbum*, XXIII(1):73–101, 2001.
- [Car56] Rudolf Carnap. *Meaning and Necessity: A study in semantics and modal logic*. The University of Chicago Press, 1956.

- [Car77] Gregory Carlson. *Reference to kinds in English*. PhD thesis, University of Massachusetts at Amherst, 1977.
- [CBGG97] Anthony Cohn, Brandon Bennett, John Gooday, and Nicholas Mark Gotts. Qualitative spatial representation and reasoning with the region connection calculus. *Geoinformatica*, 1:275–316, 1997.
- [CC95] Rosaria Conte and Cristiano Castelfranchi. *Cognitive and Social Action*. University College of London Press, London, 1995.
- [CCR92] Zhan Cui, Anthony G. Cohn, and David A. Randell. Qualitative simulation based on a logical formalism of space and time. In *Proceedings of AAAI'92*. AAAI/MIT Press, Cambridge, MA, 1992.
- [CFPS05] Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan Sag. Minimal recursion semantics: An introduction. *Research on Language and Computation*, 3:281–332, 2005.
- [CGN89] Lewis Creary, Mark Gawron, and John Nerbonne. Reference to locations. In Julia Hirschberg, editor, *Proc. of the 27th Annual Meeting of the Association for Computational Linguistics (ACL-89)*, pages 42–50. ACL, 1989.
- [CH01] Anthony G. Cohn and S. M. Hazarika. Qualitative spatial representation and reasoning: An overview. *Fundamenta Informaticae*, 43(1):2–32, 2001.
- [Cha84] Wallace Chafe. How people use adverbial clauses. In C. Brugman and M. Macaulay, editors, *Proceedings of the Tenth Annual Meeting of the Berkeley Linguistic Society*, volume 10, pages 437–449, 1984.
- [Cha97] Michel Charolles. L'encadrement du discours : univers, champs, domaines et espaces. *Cahier de Recherche Linguistique*, 6:1–73, 1997.
- [Cha03] Michel Charolles. De la topicalité des adverbiaux détachés en tête de phrase. *Travaux de Linguistique*, 47:11–51, 2003.
- [Che92] Brian F. Chellas. Time and modality in the logic of agency. *Studia Logica*, 51(3-4):485–518, 1992.
- [Chi73] Roderick Chisholm. Parts as essential to their wholes. *Review of Metaphysics*, 26:581–603, 1973.
- [Chi94] Gennaro Chierchia. Intensionality and context change. *Journal of Logic, Language and Information*, 3:141–168, 1994.
- [CL90] Paul Cohen and Hector Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.
- [Cla81] Bowman L. Clarke. A calculus of individuals based on “connection”. *Notre Dame Journal of Formal Logic*, 22(3):204–218, 1981.
- [Cla85] Bowman L. Clarke. Individuals and points. *Notre Dame Journal of Formal Logic*, 26(1):61–75, 1985.
- [CMSV09] Philipp Cimiano, Alexander Mädche, Steffen Staab, and Johanna Völker. Ontology learning. In Steffen Staab and Rudi Studer, editors, *Handbook on Ontologies*, pages 245–267. Springer, 2nd edition, 2009.
- [Com98] Bernard Combettes. *Les constructions détachées en français*. Ophrys, Gap, 1998.

- [Cov92] Kenny Coventry. *Spatial prepositions and functional relations: The case for minimally specified lexical entries*. PhD thesis, University of Edinburgh, 1992.
- [CPSR01] Kenny Coventry, Mercè Prat-Sala, and Lynn Richards. The interplay between geometry and function in the comprehension of *over*, *under*, *above*, and *below*. *Journal of Memory and Language*, 44:376–398, 2001.
- [CR06] Philipp Cimiano and Uwe Reyle. Towards foundational semantics. Ontological semantics revisited. In Brandon Bennett and Christiane Fellbaum, editors, *Proceedings of the Fourth International Conference on Formal Ontology in Information Systems (FOIS 2006)*, pages 51–62. IOS Press, 2006.
- [CRCL99] Laura A. Carlson-Radvansky, Eric S. Covey, and Kathleen M. Lattanzi. “What” effects on “where”: Functional influences on spatial relations. *Psychological Science*, 10:516–521, 1999.
- [CRI93] Laura Carlson-Radvansky and David Irwin. Frames of reference in vision and language: Where is above? *Cognition*, 46:223–244, 1993.
- [CRR96] Laura A. Carlson-Radvansky and Gabriel A. Radvansky. The influence of functional relations on spatial term selection. *Psychological Science*, 7:56–60, 1996.
- [Cru86] D.A. Cruse. *Lexical semantics*. Cambridge University Press, Cambridge, England, 1986.
- [Cum75] Robert Cummins. Functional analysis. *Journal of Philosophy*, 72:741–765, 1975.
- [CV94] Roberto Casati and Achille C. Varzi. *Holes and Other Superficialities*. A Bradford Book. The MIT Press, Cambridge, Massachusetts, 1994.
- [CV96] Roberto Casati and Achille Varzi, editors. *Events*. Dartmouth, Aldershot, 1996.
- [CV98] Anthony Cohn and Achille Varzi. Connection relations in mereotopology. In Henri Prade, editor, *Proceedings of ECAI’98*, pages 150–154. Wiley, 1998.
- [CV99] Roberto Casati and Achille Varzi. *Parts and Places - The Structures of Spatial Representation*. MIT Press, Cambridge, MA, 1999.
- [Dal97] Chris Daly. Tropes. In D.H. Mellor and Alex Oliver, editors, *Properties*, pages 140–159. Oxford University Press, 1997.
- [Dan06] Laurence Danlos. “Discourse verbs” and discourse periphrastic links. In Candy Sidner, John Harpur, Anton Benz, and Peter Kühnlein, editors, *Second Workshop on Constraints in Discourse (CID06)*, pages 59–65, Maynooth, Ireland, 2006.
- [Dav67] Donald Davidson. The logical form of action sentences. In N. Rescher, editor, *The Logic of Decision and Action*, pages 81–95. University of Pittsburgh Press, Pittsburgh, 1967.
- [Dav90] Ernest Davis. *Representations of Commonsense Knowledge*. Morgan Kaufmann, Palo Alto, 1990.
- [Deu03] Henri-José Deulofeu. L’approche macrosyntaxique en syntaxe : un nouveau modèle de rasoir d’Occam contre les notions inutiles ? *Scolia*, 16:77–95, 2003.
- [Dic01] Michael Dickey. *The processing of tense*. Kluwer, Dordrecht, 2001.
- [Die92] Molly Diesing. *Indefinites*. MIT Press, Cambridge (MA), 1992.

- [Dip93] Randall Dipert. *Artifacts, Art Works, and Agency*. Temple University Press, Philadelphia, 1993.
- [dL22] Theodore de Laguna. Point, line and surface, as sets of solids. *The Journal of Philosophy*, XIX(17):449–461, 1922.
- [DL06] Robert Demolombe and Vincent Louis. Norms, institutional power and roles: Towards a logical framework. In Floriana Esposito, Zbigniew W. Ras, Donato Malerba, and Giovanni Semeraro, editors, *Foundations of Intelligent Systems. Proc. of the 16th International Symposium ISMIS 2006*, volume 4203 of *Lecture Notes in Artificial Intelligence*, pages 514–523, 2006.
- [Don01] Maureen Donnelly. *An Axiomatization of Common-Sense Geometry*. Phd thesis, The University of Texas at Austin, Philosophy Dpt, 2001.
- [Don04] Maureen Donnelly. Relative places. In Achille Varzi and Laure Vieu, editors, *Formal Ontology in Information Systems, Proceedings of the third international conference (FOIS 2004)*, pages 249–260. IOS Press, Amsterdam, 2004.
- [Dow79] David Dowty. *Word Meaning and Montague Grammar. The Semantics of Verbs and Times in Generative Semantics and Montague’s PTQ*. Reidel, Dordrecht, 1979.
- [Dow82] David Dowty. Tenses, time adverbs and compositional semantic theory. *Linguistics and Philosophy*, 5(1):23–55, 1982.
- [dS91] Henriëtte de Swart. *Adverbs of quantification: A generalized quantifier approach*. PhD thesis, University of Groningen, 1991.
- [dS99] Henriëtte de Swart. Position and meaning: Time adverbials in context. In Peter Bosch and Rob van der Sandt, editors, *Focus. Linguistic, cognitive and computational perspectives*, pages 336–361. Cambridge University Press, Cambridge, 1999.
- [Dün05] Ivo Düntsch. Relation algebras and their application in temporal and spatial reasoning. *Artificial Intelligence Review*, 23(4):315–357, 2005.
- [DV06] Georgi Dimov and Dimiter Vakarelov. Contact algebras and region-based theory of space: A proximity approach – I. *Fundamenta Informaticae*, 74:209–249, 2006.
- [DW05] Ivo Düntsch and Michael Winter. A representation theorem for boolean contact algebras. *Theoretical Computer Science*, 347:498–512, 2005.
- [EKN01] Markus Egg, Alexander Koller, and Joachim Niehren. The constraint language for lambda structures. *Journal of Logic, Language and Information*, 10:457–485, 2001.
- [Eld07] Crawford Elder. On the place of artifacts in ontology. In Eric Margolis and Stephen Laurence, editors, *Creations of the Mind: Theories of Artifacts and Their Representation*, pages 33–51. Oxford University Press, Oxford, 2007.
- [ES07] Jérôme Euzenat and Pavel Shvaiko. *Ontology matching*. Springer, 2007.
- [Esc94] Carola Eschenbach. A note on Clarke’s points. Unpublished manuscript, 1994.
- [Euz95] Jérôme Euzenat. An algebraic approach to granularity in qualitative time and space representation. In *IJCAI’95*, pages 894–900. Morgan Kaufmann, San Mateo, CA, 1995. Montréal Août 1995.

- [Fel98] Christiane Fellbaum, editor. *WordNet. An Electronic Lexical Database*. MIT Press, Cambridge (MA), 1998.
- [FHB01] Cathrine Fabricius-Hansen and Bergljot Behrens. Elaboration and related discourse relations viewed from an interlingual perspective. SPRIK Report 13, University of Oslo, 2001. Also appeared in Proceedings of the Third Workshop on Text Structure, Carlota Smith (ed.), University of Texas at Austin, October 2000.
- [FHRSB05] Cathrine Fabricius-Hansen, Wiebke Ramm, Kåre Solfjed, and Bergljot Behrens. Coordination, discourse relations and information packaging – cross-linguistic differences. In M. Aurnague, M. Bras, A. Le Draoulec, and L. Vieu, editors, *First International Symposium on the Exploration and Modelling of Meaning*, pages 85–93, 2005.
- [Fin82] Kit Fine. Acts, events and things. In *Sixth International Wittgenstein Symposium*, pages 97–105, 1982. Kirchberg-Wechsel (Austria).
- [Fin95] Kit Fine. Ontological dependence. *Proceedings of the Aristotelian Society*, 95:269–90, 1995.
- [Fre92] Christian Freksa. Temporal reasoning based on semi-intervals. *Artificial Intelligence*, 54:199–227, 1992.
- [Fre03] Werner Frey. Syntactic conditions on adjunct classes. In Ewald Lang, Claudia Maienborn, and Cathrine Fabricius-Hansen, editors, *Modifying Adjuncts*, pages 163–209. Mouton de Gruyter, 2003.
- [FRWJ06] Katherine Forbes-Riley, Bonnie Webber, and Aravind Joshi. Computing discourse semantics: the predicate-argument semantics of discourse connectives in D-LTAG. *Journal of Semantics*, 23(1):55–106, 2006.
- [FS06] Nissim Francez and Mark Steedman. Categorical grammar and the semantics of contextual prepositional phrases. *Linguistics and Philosophy*, 29:381–417, 2006.
- [Gal90] Antony Galton. A critical examination of Allen’s theory of time. *Artificial Intelligence*, 42:159–188, 1990. Reprinted in Mani, I., Pustejovsky, J., and Gaizauskas, R., *The Language of Time: A Reader*, Oxford University Press, 2005.
- [Gal95] Antony Galton. A qualitative approach to continuity. In Pascal Amsili, Mario Borillo, and Laure Vieu, editors, *Time, Space and Movement. Meaning and Knowledge in the sensible world. Proceedings of TSM’95*, 1995.
- [Gal96] Antony Galton. Taking dimension seriously in qualitative spatial reasoning. In W. Wahlster, editor, *ECAI’96*, pages 501–505. John Wiley, Chichester, 1996.
- [Gal97] Antony Galton. Space, time and movement. In Oliviero Stock, editor, *Spatial and Temporal Reasoning*. Kluwer, Dordrecht, 1997.
- [Gal00] Antony Galton. Continuous motion in discrete space. In Anthony Cohn, Fausto Giunchiglia, and B. Selman, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Seventh International Conference (KR2000)*, pages 26–37. Morgan Kaufmann, 2000.
- [Gal01] Antony Galton. Dominance diagrams: A tool for qualitative reasoning about continuous systems. *Fundamenta Informaticae*, 46:55–70, 2001.

- [Gal07] Antony Galton. On the paradoxical nature of surfaces: Ontology at the physics/geometry interface. *The Monist*, 90(3):380–390, 2007.
- [Gal09] Antony Galton. Spatial and temporal knowledge representation. *Journal of Earth Science Informatics*, 2:169–187, 2009.
- [Ger90] Giangiacomo Gerla. Pointless metric spaces. *The Journal of Symbolic Logic*, 55(1):207–219, 1990.
- [GFC99] Simon Garrod, Gillian Ferrier, and Siobhan Campbell. In and on: Investigating the functional geometry of spatial prepositions. *Cognition*, 72:167–189, 1999.
- [GGMO03] Aldo Gangemi, Nicola Guarino, Claudio Masolo, and Alessandro Oltramari. Sweetening WordNet with DOLCE. *AI Magazine*, 24(3):13–24, 2003.
- [GHL06] Benoit Gaudou, Andreas Herzig, and Dominique Longin. A logical framework for grounding-based dialogue analysis. In Wiebe van der Hoek, Alessio Lomuscio, Erik de Vink, and Mike Wooldridge, editors, *International Workshop on Logic and Communication in Multi-Agent Systems (LCMAS 2005)*, volume 157-4 of *Electronic Notes in Theoretical Computer Science (ENTCS)*, pages 117–137. Elsevier, 2006.
- [GL96] John Gumperz and Stephen Levinson, editors. *Rethinking linguistic relativity*. Cambridge University Press, Cambridge, 1996.
- [GM09] Antony Galton and Riichiro Mizoguchi. The water falls but the waterfall does not fall: New perspectives on objects, processes and events. *Applied Ontology*, 4(2):71–107, 2009.
- [GOS09] Nicola Guarino, Daniel Oberle, and Steffen Staab. What is an *ontology*? In Steffen Staab and Rudi Studer, editors, *Handbook on Ontologies*, pages 1–17. Springer, 2nd edition, 2009.
- [Got94] Nicholas Gotts. How far can we ‘C’? Defining a ‘doughnut’ using connection alone. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Principles of Knowledge Representation and Reasoning. Proceedings of KR’94*, pages 246–257. Morgan Kaufmann, San Mateo, CA, 1994.
- [Got96] Nicholas Gotts. Formalizing commonsense topology: The inch calculus. In *Proceedings of the Fourth International Symposium on Artificial Intelligence and Mathematics - AI/MATH’96*, pages 72–75, Fort Lauderdale (FL), 1996.
- [Gou96] Dionysis Goutsos. A model of sequential relations in expository text. *Text*, 16(4):501–533, 1996.
- [Gre03] Pierre Grenon. Spatio-temporality in basic formal ontology: Snap and span, upper-level ontology, and framework for formalization. Technical Report 05-2003, Saarland University IFOMIS, 2003.
- [Gri75] H. Paul Grice. Logic and conversation. In Peter Cole and Jerry Morgan, editors, *Syntax and semantics 3: Speech Acts*, pages 41–58. Academic Press, New-York, 1975.
- [Gru93] Thomas Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.

- [Grz60] Andrzej Grzegorzczuk. Axiomatizability of geometry without points. *Synthese*, 12(2-3):228–235, 1960.
- [GS86] Barbara J. Grosz and Candace L. Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204, 1986.
- [GS87] Branko Grünbaum and G.C. Shephard. *Tilings and patterns*. Freeman, New York, 1987.
- [GS91] Jeroen Groenendijk and Martin Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–106, 1991.
- [GS03] Pierre Grenon and Barry Smith. SNAP and SPAN: Towards dynamic spatial ontology. *Spatial Cognition and Computation*, 4(1):69–103, 2003.
- [GT97] Isabel Gómez Txurruka. *La partición informacional en el discurso*. Phd dissertation, Universidad del País Vasco, 1997.
- [GT03] Isabel Gómez Txurruka. The natural language conjunction *and*. *Linguistics and Philosophy*, 26(3):255–285, 2003.
- [Gua95] Nicola Guarino. Formal ontology, conceptual analysis and knowledge representation. *International Journal on Human and Computer Studies*, 43(5–6):625–640, 1995.
- [Gua98] Nicola Guarino. Some ontological principles for designing upper level lexical resources. In A. Rubio, N. Gallardo, R. Castro, and A. Tejada, editors, *First International Conference on Language Resources and Evaluation*, pages 527–534. European Language Resources Association, 1998.
- [Gup80] Anil Gupta. *The Logic of Common Nouns: An investigation in quantified modal logic*. Yale University Press, New Haven and London, 1980.
- [GW04] Nicola Guarino and Chris Welty. An overview of OntoClean. In S. Staab and R. Studer, editors, *Handbook on Ontologies*, pages 151–159. Springer-Verlag, 2004.
- [Ham70] Charles Hamblin. *Fallacies*. Methuen, London, 1970.
- [HARD06] Julie Hunter, Nicholas Asher, Brian Reese, and Pascal Denis. Evidentiality and intensionality: Two uses of reportative constructions in discourse. In Candy Sidner, John Harpur, Anton Benz, and Peter Kühnlein, editors, *Proceedings of the Workshop on Constraints in Discourse*, pages 99–106, National University of Ireland, Maynooth, 2006.
- [Hay78] Patrick Hayes. The naive physics manifesto. In D. Michie, editor, *Expert systems in the micro-electronic age*. Edinburgh University Press, Edinburgh, 1978.
- [Hay85a] Patrick J. Hayes. Naive physics I: Ontology for liquids. In Jerry R. Hobbs and Robert C. Moore, editors, *Formal theories of the commonsense world*, pages 71–107. Ablex, Norwood (NJ), 1985.
- [Hay85b] Patrick J. Hayes. The second naive physics manifesto. In Jerry R. Hobbs and Robert C. Moore, editors, *Formal theories of the commonsense world*, pages 1–36. Ablex, Norwood (NJ), 1985.
- [HCG⁺10] Chu-ren Huang, Nicoletta Calzolari, Aldo Gangemi, Alessandro Lenci, Alessandro Oltramari, and Laurent Prévot, editors. *Ontology and the Lexicon. A Natural Language Processing Perspective*. Cambridge University Press, 2010.

- [Hei83a] Irene Heim. File change semantics and the familiarity theory of definiteness. In Rainer Bäuerle, Christopher Schwarze, and Arnim von Stechow, editors, *Meaning, Use and the Interpretation of Language*, pages 164–190. de Gruyter, Berlin, 1983.
- [Hei83b] Irene Heim. On the projection problem for presuppositions. In M. Barlow, D. Flickinger, and M. Wescoat, editors, *Proceedings of the second annual West Coast Conference on Formal Linguistics (WCCFL2)*, pages 114–125, 1983. Reprinted in: *Formal Semantics*, P. Portner and B. Partee (eds.), Blackwell, 2002.
- [Her86] Annette Herskovits. *Language and Spatial Cognition. An Interdisciplinary Study of the Prepositions in English*. Cambridge University Press, Cambridge, London, 1986.
- [HH76] Michael Halliday and Ruqaiya Hasan. *Cohesion in English*. Longman, 1976.
- [Hin81] Erhard Hinrichs. *Temporale Anaphora im Englischen*. Magisterarbeit, University of Tübingen, 1981.
- [Hin86] Erhard Hinrichs. Temporal anaphora in discourses of English. *Linguistics and Philosophy*, 9(1):63–82, 1986.
- [Hit93] Janet Hitzeman. *Temporal adverbials and the syntax-semantics interface*. PhD thesis, Univ. of Rochester, 1993.
- [Hit97] Janet Hitzeman. Semantic partition and the ambiguity of sentences containing temporal adverbials. *Natural Language Semantics*, 5(2):87–100, 1997.
- [HL04] Andreas Herzig and Dominique Longin. C&L intention revisited. In Didier Dubois, Chris Welty, and Mary-Anne Williams, editors, *Proceedings 9th Int. Conf. on Principles on Principles of Knowledge Representation and Reasoning(KR2004)*, pages 527–535. AAAI Press, 2004.
- [HM87] Steve Hanks and Drew McDermott. Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33(3):379–412, 1987.
- [HM93] Eduard H. Hovy and Elisabeth Maier. Parsimonious or profligate: How many and which discourse structure relations? Unpublished manuscript, available at <http://www.isi.edu/natural-language/people/hovy/papers/93discproc.pdf>, 1993.
- [Hob79] Jerry R. Hobbs. Coherence and coreference. *Cognitive Science*, 3:67–90, 1979.
- [Hob85] Jerry R. Hobbs. On the coherence and structure of discourse. Report CSLI-85-37, Center for Study of Language and Information, 1985.
- [Hof07] Thomas Hofweber. Innocent statements and their metaphysically loaded counterparts. *Philosophers' Imprint*, 7(1), 2007. Online Journal: <http://hdl.handle.net/2027/spo.3521354.0007.001>.
- [HPV00] James Higginbotham, Fabio Pianesi, and Achille C. Varzi, editors. *Speaking of Events*. Oxford University Press, 2000.
- [HR91] Otthein Herzog and Claus-Rainer Rollinger. *Text understanding in LILOG - Integrating computational linguistics and artificial intelligence*. Number 546 in Lecture Notes in Artificial Intelligence. Springer, Berlin, 1991.
- [Hus70] Edmund Husserl. *Logical investigations*. Routledge, London, 1970. Original German edition 1913–1921, Niemeyer.

- [ILE88] Madelyn Anne Iris, Bonnie E. Litowitz, and Martha Evens. Problems of the part-whole relation. In Martha Evens, editor, *Relational models of the lexicon: Representing knowledge in semantic networks*, pages 261–288. Cambridge University Press, Cambridge, 1988.
- [Jac83] Ray Jackendoff. *Semantics and Cognition*. MIT Press, Cambridge, MA, 1983.
- [Jac96] Ray Jackendoff. Semantics and cognition. In Shalom Lappin, editor, *The Handbook of Contemporary Semantic Theory*, chapter 20, pages 539–559. Blackwell, 1996.
- [Joh94] Michael Johnston. *The syntax and the semantics of adverbials adjuncts*. PhD thesis, University of California at Santa Cruz, 1994.
- [JR99] Jacques Jayez and Corinne Rossari. Pragmatic connectives as predicates. The case of inferential connectives. In Patrick Saint-Dizier, editor, *Predicative Forms in Natural Language and in Lexical Knowledge Bases*, pages 285–319. Kluwer Academic Publishers, 1999.
- [JR01] Jacques Jayez and Corinne Rossari. The discourse level sensitivity of consequence discourse markers in French. *Cognitive Linguistics*, 12:275–290, 2001.
- [Kam79] Hans Kamp. Events, instants and temporal reference. In Rainer Bäuerle, Urs Egli, and Arnim von Stechow, editors, *Semantics from different points of view*, pages 376–417. de Gruyter, Berlin, 1979.
- [Kam81] Hans Kamp. A theory of truth and semantic representation. In Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, editors, *Formal Methods in the Study of Language*, volume Part 1 of *Mathematical Centre Tracts 135*, pages 277–322. University of Amsterdam, Amsterdam, 1981.
- [Keh04] Andrew Kehler. Discourse topics, sentence topics, and coherence. *Theoretical Linguistics*, 30(2-3):227–240, 2004.
- [KKRE08] Andrew Kehler, Laura Kertz, Hannah Rohde, and Jeffrey Elman. Coherence and coreference revisited. *Journal of Semantics*, 25(1):1–44, 2008.
- [KM06] Peter Kroes and Anthonie Meijers. The dual nature of technical artefacts. *Studies in History and Philosophy of Science*, 37(1):1–4, 2006.
- [Kno01] Alistair Knott. Semantic and pragmatic relations and their intended effects. In Ted Sanders, J Schilperoord, and W. Spooren, editors, *Text representation: linguistic and psycholinguistic aspects*, pages 127–151. Benjamins, 2001.
- [KOOM01] Alistair Knott, J. Oberlander, M. O’Donnell, and C. Melish. Beyond elaboration: The interaction of relations and focus in coherent text. In T. Sanders, J. Schilperoord, and W. Spooren, editors, *Text representation: linguistic and psycholinguistic aspects*, pages 181–196. Benjamins, 2001.
- [KR83] Hans Kamp and Christian Rohrer. Tense in texts. In Rainer Bäuerle, Christopher Schwarze, and Arnim von Stechow, editors, *Meaning, Use and the Interpretation of Language*, pages 250–269. de Gruyter, Berlin, 1983.
- [KR93] Hans Kamp and Uwe Reyle. *From Discourse to Logic*. Kluwer Academic Publishers, Dordrecht, 1993.

- [Kun75] Susumu Kuno. Three perspectives in the functional approach to syntax. In R.E. Grossman, L.J. San, and T.J. Vance, editors, *Papers from the Parasession on Functionalism*, pages 276–336. Chicago Linguistic Society, University of Chicago, 1975.
- [KV09] Georges Kleiber and Hélène Vassiliadou. Sur la relation d’Elaboration : des approches intuitives aux approches formelles. *Journal of French Language Studies*, 19:183–205, 2009.
- [LA93] Alex Lascarides and Nicholas Asher. Temporal interpretation, discourse relations, and commonsense entailment. *Linguistics and Philosophy*, 16(5):437–493, 1993.
- [LA09a] Alex Lascarides and Nicholas Asher. Agreement, disputes and commitments in dialogue. *Journal of Semantics*, 26(2):109–158, 2009.
- [LA09b] Alex Lascarides and Nicholas Asher. The interpretation of questions in dialogue. In Arndt Riestler and Torgrim Solstad, editors, *Proceedings of Sinn und Bedeutung 13*, volume 5 of *SinSpeC. Working Papers of the SFB 732*, pages 17–30. University of Stuttgart, 2009.
- [Lan89] Fred Landman. Groups I, Groups II. *Linguistics and Philosophy*, 12:559–605 & 723–744, 1989.
- [LDPW03] Anne Le Draoulec and Marie-Paule Péry-Woodley. Time travel in text: Temporal framing in narratives and non-narratives framing in narratives and non-narratives. In L. Lagerwerf, W. Spooren, and L. Degan, editors, *Determination of Information and Tenor in Texts: Multidisciplinary Approaches to Discourse*, pages 267–275. Stichting Neerlandistiek VU & Nodus Publikationen, Amsterdam & Münster, 2003.
- [Lee69] Geoffrey Leech. *Towards a semantic description of English*. Longman, London, 1969.
- [Les31] Stanislaw Lesniewski. O podstawach matematyki [on the foundations of mathematics]. *Przegląd Filozoficzny [Philosophical Review]*, 30-34, 1927-1931. French translation: Sur les fondements de la mathématique, Hermès, Paris, 1989.
- [Lev93] Beth Levin. *English verb classes and alternations. A preliminary Investigation*. The University of Chicago Press, 1993.
- [Lew86] David K. Lewis. *On the Plurality of Worlds*. Blackwell, 1986.
- [LG40] Henry Leonard and Nelson Goodman. The calculus of individuals and its uses. *The Journal of Symbolic Logic*, 5(2):45–55, 1940.
- [LG90] Douglas Lenat and R. V. Guha. *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*. Addison-Wesley, 1990.
- [Lin83] Godehard Link. The logical analysis of plurals and mass terms: A lattice theoretical approach. In Rainer Bäuerle, Christopher Schwarze, and Arnim von Stechow, editors, *Meaning, Use and the Interpretation of Language*, pages 302–323. de Gruyter, Berlin, 1983.
- [Lin97] Godehard Link. *Algebraic Semantics in Language and Philosophy*. Lecture Notes. CSLI, 1997.

- [LLGH09] Emiliano Lorini, Dominique Longin, Benoit Gaudou, and Andreas Herzig. The logic of acceptance: grounding institutions on agents' attitudes. *Journal of Logic and Computation*, page to appear, 2009.
- [Lyo77] John Lyons. *Semantics*. Cambridge University Press, Cambridge, 1977.
- [LZ03] Michael Loux and Dean Zimmerman. Introduction. In Michael Loux and Dean Zimmerman, editors, *The Oxford Handbook of Metaphysics*, pages 1–7. Oxford University Press, 2003.
- [Mai95] Claudia Maienborn. Toward a compositional semantics for locative modifiers. In *Proceedings of Semantics and Linguistic Theory V*, pages 237–254, 1995.
- [Mar98] Jean-Marie Marandin. Grammaire de l'incidence. Unpublished manuscript, available at http://www.llf.cnrs.fr/Gens/Marandin/GrammaireDeL_incidence.doc, 1998.
- [Mas00] Claudio Masolo. *Critères de Comparaison et de Construction de Théories Axiomatiques pour la Représentation de la Connaissance : Ontologies de l'Espace et du temps / Criteri di Confronto e Costruzione di Teorie Assiomatiche per la Rappresentazione della Conoscenza: Ontologie dello Spazio e del Tempo*. Phd thesis, Université Paul Sabatier - IRIT / Università di Padova - Ladseb-CNR, 2000.
- [Mas09] Claudio Masolo. Parthood simpliciter vs. temporary parthood. In Gerhard Lakemeyer, Leora Morgenstern, and Mary-Anne Williams, editors, *Proceedings of the Ninth International Symposium on Logical Formalizations of Commonsense Reasoning (Commonsense 2009)*, pages 97–102, Toronto, 2009.
- [MB05] Claudio Masolo and Stefano Borgo. Qualities in formal ontology. In *Proceedings of the Foundational Aspects of Ontologies Workshop (FOnt 2005) at KI 2005*, Koblenz, September 2005.
- [MBG⁺03] Claudio Masolo, Stefano Borgo, Aldo Gangemi, Nicola Guarino, and Alessandro Oltramari. The WonderWeb library of foundational ontologies and the DOLCE ontology. WonderWeb (EU IST project 2001-33052) deliverable D18, LOA-ISTC-CNR, 2003.
- [McD82] Drew McDermott. A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6:101–155, 1982.
- [MCV92] Roger Martin-Clouaire and Laure Vieu. Spatial and qualitative reasoning in concentrated flow erosion simulation. In *Proceeding of the AAAI-92 Workshop on AI in Natural & Agricultural Resource Development*, San José (CA), 1992.
- [MdSV⁺04] Arie Molendijk, Henriëtte de Swart, Carl Vettters, Andrée Borillo, Myriam Bras, Anne Le Draoulec, Laure Vieu, Henk Verkyul, and Co Vet. Meaning and use of past tenses in discourse. In Henriëtte de Swart and Francis Corblin, editors, *Handbook of French Semantics*, pages 271–308. CSLI, Stanford, 2004.
- [Mel97] Alfred R. Mele, editor. *The philosophy of action*. Oxford Readings in Philosophy. Oxford University Press, Oxford, 1997.
- [Mey00] John-Jules Meyer. Dynamic logic for reasoning about actions and agents. In J. Minker, editor, *Logic-Based Artificial Intelligence*, pages 281–311. Kluwer, 2000.

- [MGV⁺05] Claudio Masolo, Giancarlo Guizzardi, Laure Vieu, Emanuele Bottazzi, and Roberta Ferrario. Relational roles and qua-individuals. In Guido Boella, James Odell, Leendert van der Torre, and Harko Verhagen, editors, *Roles, an Interdisciplinary Perspective: Ontologies, Programming Languages, and Multiagent Systems. Papers from the AAI Fall Symposium*, pages 103–112. AAAI Press, 2005.
- [Mit77] Anita Mittwoch. How to refer to ones’s own words: speech-act modifying adverbials and the performative analysis. *Journal of Linguistics*, 13(2):153–368, 1977.
- [MMT92] William Mann, Christian Matthiessen, and Sandra Thompson. Rhetorical Structure Theory and text analysis. In William Mann and Sandra Thompson, editors, *Discourse Description. Diverse linguistic analyses of a fund-raising text.*, pages 39–78. John Benjamins, Amsterdam, 1992.
- [MN78] David Marr and Keith Nishihara. Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London*, B 200:269–294, 1978.
- [MO97] D.H. Mellor and Alex Oliver, editors. *Properties*. Oxford University Press, 1997.
- [Mói01] Telmo Móia. Telling apart temporal locating adverbials and time-denoting expressions. In *Proceedings of the workshop on Temporal and spatial information processing - Annual Meeting of the ACL*, Morristown, NJ, USA, 2001. Association for Computational Linguistics.
- [Mol97] Friederike Moltmann. *Parts and Wholes in Semantics*. Oxford University Press, Oxford, 1997.
- [Mon73] Richard Montague. The proper treatment of quantification in ordinary english. In K.J.J. Hintikka, J.M.E. Moravcsik, and P. Suppes, editors, *Approaches to Natural Language*, pages 221–242. Reidel, Dordrecht, 1973.
- [Mon74] Richard Montague. *Formal Philosophy: Selected Papers of Richard Montague*. Yale University Press, New Haven, 1974.
- [Mor06] Leora Morgenstern. Knowledge representation and commonsense reasoning: Review of four books. *Artificial Intelligence*, 170:1239–1250, 2006.
- [MS88] Marc Moens and Mark Steedman. Temporal ontology and temporal reference. *Computational Linguistics*, 14(2):15–28, 1988.
- [MS99] Philippe Muller and Laure Sarda. Représentation de la sémantique des verbes de déplacement transitifs du français. *Traitement Automatique des Langues*, 39(2):127–147, 1999.
- [MT87] William C. Mann and Sandra A. Thompson. Rhetorical Structure Theory: A theory of text organization. Reprint Series ISI/RS-87-1190, Information Sciences Institute, 1987.
- [MT88] William C. Mann and Sandra A. Thompson. Rhetorical Structure Theory: Towards a fonctionnal theory of text organization. *Text*, 8(3):243–281, 1988.
- [Mul95] Philippe Muller. *Sémantique de l’espace : Formalisation des prépositions “avant” et “après”*. Mémoire de DEA, Université Paul Sabatier IRIT, 1995.
- [Mul98a] Philippe Muller. *Eléments d’une théorie du mouvement pour la formalisation du raisonnement spatio-temporel de sens commun*. Thèse de doctorat, Université Paul Sabatier IRIT, 1998.

- [Mul98b] Philippe Muller. A qualitative theory of motion based on spatio-temporal primitives. In Anthony G. Cohn, Lenhart Schubert, and Stuart C. Shapiro, editors, *Principles of Knowledge Representation and Reasoning (KR'98)*, pages 131–141. Morgan Kaufmann, San Francisco, CA, 1998. Trento June 1998.
- [Mul98c] Philippe Muller. Space-time as a primitive for space and motion. In Nicola Guarino, editor, *Formal Ontology in Information Systems - Proceedings of the First International Conference (FOIS'98)*, pages 63–76, 1998.
- [Mum98] Stephen Mumford. *Dispositions*. Oxford University Press, Oxford, 1998.
- [MV99] Claudio Masolo and Laure Vieu. Atomicity vs. infinite divisibility of space. In Christian Freksa and David Mark, editors, *Spatial Information theory. Proceedings of COSIT'99*, LNCS 1661, pages 235–250. Springer Verlag, Berlin, 1999.
- [MVB⁺04] Claudio Masolo, Laure Vieu, Emanuele Bottazzi, Carola Catenacci, Roberta Ferrario, Aldo Gangemi, and Nicola Guarino. Social roles and their descriptions. In Didier Dubois and Chris Welty, editors, *Proceedings of the 9th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2004)*, pages 267–277. AAAI Press, Menlo Park (CA), 2004. Whistler June, 2-5, 2004.
- [Ner03] Graham Nerlich. Space-time substantivalism. In *The Oxford Handbook of Metaphysics*. Oxford University Press, 2003.
- [Nic23] Jean Nicod. *La géométrie dans le monde sensible*. PhD thesis, Paris, 1923. Presses Universitaires de France, 1962.
- [Nic08] David Nicolas. Mass nouns and plural logic. *Linguistics and Philosophy*, 31:211–244, 2008.
- [Noy09] Natalya Noy. Ontology mapping. In Steffen Staab and Rudi Studer, editors, *Handbook on Ontologies*, pages 573–590. Springer, 2nd edition, 2009.
- [NP01] Ian Niles and Adam Pease. Towards a standard upper ontology. In Chris Welty and Barry Smith, editors, *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)*, pages 2–9, 2001.
- [NR04] Sergei Nirenburg and Victor Raskin. *Ontological semantics*. MIT Press, 2004.
- [OAH⁺07] Daniel Oberle, Anupriya Ankolekar, Pascal Hitzler, Philipp Cimiano, Michael Sintek, Malte Kiesel, Babak Mougouie, Stephan Baumann, Shankar Vembu, Massimo Romanelli, Paul Buitelaar, Ralf Engel, Daniel Sonntag, Norbert Reithinger, Berenike Loos, Robert Porzel, Hans-Peter Zorn, Vanessa Micelli, Christian Schmidt, Moritz Weiten, Felix Burkhardt, and Jianshen Zhou. DOLCE ergo SUMO: On foundational and domain models in the SmartWeb Integrated Ontology (SWIntO). *Web Semantics: Science, Services and Agents on the World Wide Web*, 5(3):156–174, 2007.
- [Par70] Terence Parsons. An analysis of mass terms and amount terms. *Foundations of language*, 6:363–388, 1970.
- [Par75] Terence Parsons. Afterthoughts on mass terms. *Synthese*, 31:517–521, 1975.
- [Par84] Barbara Partee. Nominal and temporal anaphora. *Linguistics and Philosophy*, 7:243–286, 1984.

- [Par90] Terence Parsons. *Events in the semantics of English: A study in subatomic semantics*, chapter 11 “Tenses and temporal adverbials”. The MIT Press, Cambridge, MA, 1990.
- [Pau02a] Sabina Pauen. Evidence for knowledge-based category discrimination in infancy. *Child Development*, 73(4):1016–1033, 2002.
- [Pau02b] Marc Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.
- [PB98] Barbara Partee and Vladimir Borschev. Integrating lexical and formal semantics: genitives, relational nouns, and type-shifting. In Robyn Cooper and T. Gamkrelidze, editors, *Second Tbilisi Symposium on Language, Logic, and Computation*, pages 229–241. Center on Language, Logic, Speech, Tbilisi, 1998.
- [PDL⁺08] Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltsakaki, Livio Robaldo, Aravind Joshi, and Bonnie Webber. The Penn Discourse Treebank 2.0. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC 2008)*, pages 2961–2968, 2008.
- [Pel79] Francis J. Pelletier, editor. *Mass terms: Some philosophical problems*. Reidel, Dordrecht, 1979.
- [PF97] Ian Pratt and Nissim Francez. On the semantics of temporal prepositions and preposition phrases. Technical Report LCL9701, Computer Science Department, University of Manchester, 1997.
- [PF01] Ian Pratt and Nissim Francez. Temporal prepositions and temporal generalized quantifiers. *Linguistics and Philosophy*, 24(2):187–222, 2001.
- [PMD⁺07] Rashmi Prasad, Eleni Miltsakaki, Nikhil Dinesh, Alan Lee, Aravind Joshi, Livio Robaldo, and Bonnie Webber. The Penn Discourse Treebank 2.0 annotation manual. IRCS technical report, Institute for Research in Cognitive Science, University of Pennsylvania, 2007.
- [PMDV02] Laurent Prévot, Philippe Muller, Pascal Denis, and Laure Vieu. Une approche sémantique et rhétorique du dialogue. Un cas d’étude : l’explication d’un itinéraire. *Traitement Automatique des Langues*, 43(2):43–70, 2002.
- [Pol88] Livia Polanyi. A formal model of the structure of discourse. *Journal of Pragmatics*, 12:601–638, 1988.
- [Pör74] Ingmar Pörn. Some basic concepts of action. In S. Stenlund, editor, *Logical Theory and Semantic Analysis*. Reidel, 1974.
- [PP02] Paul Portner and Barbara Partee, editors. *Formal Semantics*. Blackwell, 2002.
- [Pré04] Laurent Prévot. *Structures sémantiques et pragmatiques pour la modélisation de la cohérence dans les dialogues finalisés*. Thèse de doctorat, Université de Toulouse III, Paul Sabatier, 2004.
- [PS84] Livia Polanyi and Remko Scha. A syntactic approach to discourse semantics. In *Proceedings of COLING*, pages 413–419, 1984.
- [PS03] Francis J. Pelletier and Lenhart K. Schubert. Mass expressions. In *Handbook of philosophical logic*, volume 10, pages 249–335. Kluwer, 2nd edition, 2003. First edition in 1989.

- [Pus91] James Pustejovsky. The generative lexicon. *Computational Linguistics*, 17(4):409–441, 1991.
- [Pus95] James Pustejovsky. *The generative lexicon*. MIT Press, Cambridge (MA), 1995.
- [Pus98] James Pustejovsky. The semantics of lexical underspecification. *Folia Linguistica*, 32(3-4):323–347, 1998.
- [PV00] Fabio Pianesi and Achille Varzi. Events and event talk: An introduction. In James Higginbotham, Fabio Pianesi, and Achille Varzi, editors, *Speaking of Events*, pages 3–47. Oxford University Press, New York, 2000.
- [PV05] Laurent Prévot and Laure Vieu. The moving right frontier. In Claudia Sassen, Anton Benz, and Peter Kühnlein, editors, *Proceedings of the First International Workshop on Constraints in Discourse*, pages 135–141, Dortmund, 2005. Universität Dortmund.
- [PV08] Laurent Prévot and Laure Vieu. The moving right frontier. In Anton Benz and Peter Kühnlein, editors, *Constraints in Discourse, Pragmatics & Beyond* New Series, pages 53–66. John Benjamins, Amsterdam, 2008.
- [PVA09] Laurent Prévot, Laure Vieu, and Nicholas Asher. Une formalisation plus précise pour une annotation moins confuse: la relation d’élaboration d’entité. *Journal of French Language Studies*, 19(2):207–228, 2009.
- [PWAE⁺09] Marie-Paule Péry-Woodley, Nicholas Asher, Patrice Enjalbert, Farah Benamara, Myriam Bras, Cécile Fabre, Stéphane Ferrari, Lydia-Mai Ho-Dac, Anne Le Draoulec, Yann Mathet, Philippe Muller, Laurent Prévot, Josette Rebeyrolle, Ludovic Tanguy, Marianne Vergez-Couret, Laure Vieu, and Antoine Widlocher. ANNODIS: une approche outillée de l’annotation de structures discursives (poster). In *Actes de TALN’09*, 2009.
- [Qui48] Willard V. Quine. On what there is. *Review of Metaphysics*, 2(1):21–28, 1948.
- [Qui60] Willard V. Quine. *Word and object*. MIT Press, Cambridge, MA, 1960.
- [Ray07] Agustín Rayo. Plurals. *Philosophical Compass*, 2(3):411–427, 2007.
- [RCC92a] David Randell, Anthony Cohn, and Zhan Cui. Computing transitivity tables. In *Proceedings of CADE 11*. Springer, 1992.
- [RCC92b] David Randell, Zhan Cui, and Anthony Cohn. A spatial logic based on regions and connection. In B. Nebel, C. Rich, and W. Swartout, editors, *Principles of Knowledge Representation and Reasoning, Proceedings of the Third International Conference*, pages 165–176. Morgan Kaufmann, San Mateo (CA), 1992.
- [RDA⁺07] Brian Reese, Pascal Denis, Nicholas Asher, Jason Baldrige, and Julie Hunter. Reference manual for the analysis and annotation of rhetorical structure (v 1.0). Technical report, University of Texas at Austin, 2007.
- [RE06] Gisela Redeker and Markus Egg. Says who? On the treatment of speech attributions in discourse structure. In Candy Sidner, Anton Benz, John Harpur, and Peter Kühnlein, editors, *Proceedings of the Workshop on Constraints in Discourse (CID 2006)*, pages 140–146, 2006.
- [Rea97] Michael Rea, editor. *Material Constitution. A Reader*. Rowman and Littlefield Publishers, Lanham, MD, 1997.

- [Rei47] Hans Reichenbach. *Elements of Symbolic Logic*. Macmillan, New York, 1947.
- [Rei81] Tanya Reinhart. Definite NP anaphora and C-command domains. *Linguistic Inquiry*, 12(4):605–635, 1981.
- [Rei01] Ray Reiter. *Knowledge in action: Logical foundations for specifying and implementing dynamic systems*. MIT Press, 2001.
- [Rey93] Uwe Reyle. Dealing with ambiguities by underspecification: Construction, interpretation and deduction. *Journal of Semantics*, 10:123–179, 1993.
- [RN99] Jochen Renz and Bernhard Nebel. On the complexity of qualitative spatial reasoning: A maximal tractable fragment of the region connection calculus. *Artificial Intelligence*, 108(1-2):69–123, 1999.
- [Roe97] Peter Roeper. Region-based topology. *Journal of Philosophical Logic*, 26(3):251–309, 1997.
- [San97] Ted Sanders. Semantic and pragmatic sources of coherence. *Discourse Processes*, 24:119–147, 1997.
- [Sch87] Deborah Schiffrin. *Discourse markers*. Cambridge University Press, Cambridge, 1987.
- [Sea95] John Searle. *The Construction of Social Reality*. The Free Press, New York, 1995.
- [Sea01] John Searle. *Rationality in Action*. The MIT Press, 2001.
- [Sei97] Johanna Seibt. Existence in time: From substance to process. In Jan Faye, Uwe Scheffler, and Max Urchs, editors, *Perspectives on Time*, pages 143–182. Kluwer, 1997.
- [Sid03] Theodore Sider. *Four-Dimensionalism - An Ontology of Persistence and Time*. Clarendon Press Press, Oxford, 2003.
- [Sim87] Peter Simons. *Parts - A study in ontology*. Clarendon Press, Oxford, 1987.
- [SK07] Elizabeth Spelke and Katherine Kinzler. Core knowledge. *Developmental Science*, 10:89–96, 2007.
- [SL08] Caroline Sporleder and Alex Lascarides. Using automatically labelled examples to classify rhetorical relations: A critical assessment. *Natural Language Engineering*, 14(3):369–416, 2008.
- [Smi96] Barry Smith. Mereotopology: A theory of parts and boundaries. *Data and Knowledge Engineering*, 20(3):287–303, 1996.
- [Sow88] John F. Sowa. Using a lexicon of canonical graphs in a semantic interpreter. In M. W. Evens, editor, *Relational models of the lexicon*. Cambridge University Press, 1988.
- [Spe03] Elizabeth Spelke. What makes us smart? Core knowledge and natural language. In Dedre Gentner and Susan Goldin-Meadow, editors, *Language in Mind: Advances in the study of language and thought*, chapter 10, pages 277–311. The MIT Press, Cambridge, MA, 2003.
- [SS09] David Schlangen and Gabriel Skantze. A general, abstract model of incremental dialogue processing. In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pages 710–718. ACL, 2009.

- [SSN92] Ted Sanders, Wilbert Spooren, and Leo Noordman. Towards a taxonomy of coherence relations. *Discourse Processes*, 15:1–35, 1992.
- [Ste00a] Friedrich Steimann. On the representation of roles in object-oriented and conceptual modelling. *Data and Knowledge Engineering*, 35:83–106, 2000.
- [Ste00b] John Stell. Boolean connection algebras: A new approach to the Region-Connection Calculus. *Artificial Intelligence*, 122:111–136, 2000.
- [Ste04] Manfred Stede. Does discourse processing need discourse topics? *Theoretical Linguistics*, 30(2-3):241–253, 2004.
- [Tab06] Maite Taboada. Discourse markers as signals (or not) of rhetorical relations. *Journal of Pragmatics*, 38(4):567–592, 2006.
- [Tal83] Leonard Talmy. How language structures space. In Herbert Pick and Linda Acredolo, editors, *Spatial orientation: theory, research and application*, pages 225–282. Plenum, New-York, 1983.
- [Tar56a] Alfred Tarski. Foundations of the geometry of solids. In J. Corcoran, editor, *Logic, semantics, metamathematics*, pages 24–30. Oxford University Press, Oxford, 1956. First published in French in suppl. to *Annales de la Société Polonaise de Mathématique*.
- [Tar56b] Alfred Tarski. *Logic, semantics, metamathematics. Papers from 1923 to 1938*. Clarendon Press, Oxford, 1956.
- [Tho07] Amie Thomasson. Artifacts and human concepts. In Eric Margolis and Stephen Laurence, editors, *Creations of the Mind: Theories of Artifacts and Their Representation*, pages 52–73. Oxford University Press, Oxford, 2007.
- [TLG⁺06] Lynda Temal, Pascal Lando, Bernard Gibaud, Michel Dojat, Gilles Kassel, and Lapujade Anne. OntoNeuroBase: a multi-layered application ontology in neuroimaging. In *Proceedings of FOMI 2006*, 2006.
- [TLV07] Robert Trypuz, Emiliano Lorini, and Laure Vieu. Solving Bratman’s video game puzzle in two formalisms. In Xabier Arrazola and Jesús María Larrazabal, editors, *ILCLI International Workshop on Logic and Philosophy of Knowledge, Communication and Action (LogKCA-07)*, pages 411–426. The University of the Basque Country Press, Donostia, 2007.
- [Tra94] David Traum. *Computational theory of grounding in natural language conversation*. PhD thesis, Computer Science Departement, University of Rochester, 1994.
- [Tro07] Nicolas Troquard. *Independent agents in branching time: Towards a unified framework for reasoning about multiagent systems*. Phd dissertation, University Paul Sabatier & University of Trento, IRIT & LOA-ISTC-CNR, July 20th, 2007 2007.
- [Try07] Robert Trypuz. *Formal Ontology of Action: A Unifying Approach*. Phd dissertation, University of Trento, ICT School LOA-ISTC-CNR, December 19th, 2007 2007.
- [TTV06] Nicolas Troquard, Robert Trypuz, and Laure Vieu. Towards an ontology of agency and action. From STIT to OntoSTIT+. In Brandon Bennett and Christiane Fellbaum, editors, *Formal Ontology in Information Systems (FOIS 2006)*, pages 179–190. IOS Press, Amsterdam, 2006.

- [Tuo95] Raimo Tuomela. *The importance of us: A philosophical study of basic social notions*. Stanford Series in Philosophy. Stanford University Press, 1995.
- [TV06] Nicolas Troquard and Laure Vieu. Towards a logic of agency and actions with duration. In G. Brewka, S. Coradeschi, A. Perini, and P. Traverso, editors, *ECAI06 proceedings (17th conference)*, pages 775–776. IOS Press, Amsterdam, 2006.
- [TV07] Robert Trypuz and Laure Vieu. An ontology of the aspectual classes of actions. In Xabier Arrazola and Jesús María Larrazabal, editors, *ILCLI International Workshop on Logic and Philosophy of Knowledge, Communication and Action (LogKCA-07)*, pages 393–409. The University of the Basque Country Press, Donostia, 2007.
- [Tve96] Barbara Tversky. Spatial perspective in descriptions. In Paul Bloom, M.A. Peterson, L. Nadel, and M.F. Garrett, editors, *Language and space*, pages 463–491. The MIT Press, Cambridge, MA, 1996.
- [VA07] Laure Vieu and Michel Aurnague. Part-of relations, functionality and dependence. In Michel Aurnague, Maya Hickmann, and Laure Vieu, editors, *The Categorization of Spatial Entities in Language and Cognition*, pages 307–336. John Benjamins, Amsterdam, 2007.
- [Van86] Claude Vandeloise. *L’espace en français : sémantique des prépositions spatiales*. Travaux en Linguistique. Seuil, Paris, 1986. English translation: *Spatial Prepositions: A Case Study from French*, The University of Chicago Press, 1991.
- [Van88] Claude Vandeloise. Les usages spatiaux statiques de la préposition à. *Cahiers de lexicologie*, 53:119–148, 1988.
- [Var96] Achille Varzi. Parts, wholes, and part-whole relations: The prospects of mereotopology. *Data and Knowledge Engineering*, 20(3):259–286, 1996.
- [Var02] Achille Varzi. Words and objects. In A. Bottani, M. Carrara, and D. Giaretta, editors, *Individuals, Essence, and Identity. Themes of Analytic Metaphysics*, pages 49–75. Kluwer Academic Publishers, Dordrecht, Boston, and London, 2002.
- [Var07] Achille Varzi. From language to ontology: Beware of the traps. In Michel Aurnague, Laure Vieu, and Maya Hickmann, editors, *The Categorization of Spatial Entities in Language and Cognition*, pages 269–284. John Benjamins, 2007.
- [vB83] Johan van Benthem. *The logic of time*. Reidel, Dordrecht, 1983.
- [VBAA05] Laure Vieu, Myriam Bras, Nicholas Asher, and Michel Aurnague. Locating adverbials in discourse. *Journal of French Language Studies*, 15(2):173–193, 2005.
- [vBD83] Johan van Benthem and Kees Doets. Higher-order logic. In Dov Gabbay and Franz Guenther, editors, *Handbook of philosophical logic*, volume I, Elements of Classical Logic, pages 275–329. Reidel, 1983.
- [VBM08] Laure Vieu, Stefano Borgo, and Claudio Masolo. Artefacts and roles: Modeling strategies in a multiplicative ontology. In Carola Eschenbach and Michael Gruninger, editors, *Proceedings of the 5th International Conference on Formal Ontology in Information Systems (FOIS 2008)*, pages 121–134. IOS Press, Amsterdam, 2008.

- [VC98] Pepijn Visser and Zhan Cui. On accepting heterogeneous ontologies in distributed architectures. In P. Borst, V. Benjamins, A. Farquhar, and A. Gómez-Pérez, editors, *Proceedings of the ECAI'98 workshop on Applications of Ontologies and Problem-solving methods*, pages 112–119, 1998.
- [Ver04] Jean-Christophe Verstraete. Initial and final position for adverbial clauses in English: the constructional basis of the discursive and syntactic differences. *Linguistics*, 42(4):819–853, 2004.
- [Vet80] Co Vet. *Temps, aspects et adverbes de temps en français contemporain*. Droz, Genève, 1980.
- [Vie91] Laure Vieu. *Sémantique des relations spatiales et inférences spatio-temporelles : une contribution à l'étude des structures formelles de l'espace en Langage Naturel*. Thèse de doctorat, Université Paul Sabatier IRIT, 1991. 22 Avril 1991.
- [Vie93] Laure Vieu. A logical framework for reasoning about space. In Andrew Frank and Irene Campari, editors, *Spatial Information Theory - A Theoretical Basis for GIS, Proceedings of COSIT'93*, Lecture Notes in Computer Science 716, pages 25–35. Springer, Berlin, 1993.
- [Vie97] Laure Vieu. Spatial representation and reasoning in artificial intelligence. In Oliviero Stock, editor, *Spatial and Temporal Reasoning*, pages 5–41. Kluwer, Dordrecht, 1997.
- [Vie06] Laure Vieu. On the transitivity of functional parthood. *Applied Ontology*, 1(2):147–155, 2006.
- [Vie07] Laure Vieu. On blocking: The rhetorical aspects of content-level discourse relations and their semantics. In Michel Aurnague, Kepa Korta, and Jesús María Larrazabal, editors, *Language, Representation and Reasoning. Memorial Volume to Isabel Gómez Txurruka*, pages 263–282. The University of Basque Country Press, Bilbao, 2007.
- [vK95] Jan van Kuppevelt. Main structure and side structure in discourse. *Linguistics*, 33:809–833, 1995.
- [VKvB89] Marc Vilain, Henry Kautz, and Peter van Beek. Constraint propagation algorithms for temporal reasoning: A revised report. In D. S. Weld and J. de Kleer, editors, *Readings in Qualitative Reasoning about Physical Systems*, pages 373–381. Morgan Kaufmann, San Mateo, CA, 1989.
- [VMC94] Laure Vieu and Roger Martin-Clouaire. Spatial and qualitative reasoning for modeling physical processes. *AI Applications*, 8(1):61–74, 1994.
- [VP04] Laure Vieu and Laurent Prévot. Background in Segmented Discourse Representation Theory. In *Workshop Segmented Discourse Representation Theory, 11th conference on Natural Language Processing (TALN)*, pages 485–494, Fes, Maroc, 2004.
- [WCH87] Morton Winston, Roger Chaffin, and Douglas Herrmann. A taxonomy of part-whole relations. *Cognitive Science*, 11:417–444, 1987.
- [WdK90] Daniel S. Weld and Johan de Kleer, editors. *Readings in Qualitative Reasoning about Physical Systems*. Morgan Kaufmann, San Mateo (CA), 1990.

- [WG01] Chris Welty and Nicola Guarino. Supporting ontological analysis of taxonomic relationships. *Data and Knowledge Engineering*, 39(1):51–74, 2001.
- [WG05] Florian Wolf and Edward Gibson. Representing discourse coherence: A corpus-based study. *Computational Linguistics*, 31(2):249–287, 2005.
- [Whi19] Alfred N. Whitehead. *An enquiry concerning the principles of natural knowledge*. Cambridge University Press, Cambridge, 1919.
- [Whi20] Alfred N. Whitehead. *The concept of nature*. Cambridge University Press, Cambridge, 1920.
- [WK95] Douglas Walton and Erik Krabbe. *Commitment in Dialogue: Basic concepts of interpersonal reasoning*. State University of New York Press, 1995.
- [WKJ01] Bonnie Webber, Alistair Knott, and Aravind Joshi. Multiple discourse connectives in a lexicalized grammar for discourse. In Bunt, Muskens, and Thijsse, editors, *Computing Meaning*, volume 2, pages 229–245. Kluwer, Dordrecht, 2001.
- [Woo00] Michael Wooldridge. *Reasoning about Rational Agents*. MIT Press, Cambridge, 2000.
- [Yi06] Byeong-Uk Yi. The logic and meaning of plurals, I-II. *Journal of Philosophical Logic*, 34 & 35:459–506 & 239–288, 2005–2006.
- [Zan02] Alberto Zanardo. First-order and second-order aspects of branching-time semantics. In *Proceedings of HPLMC-02*, 2002.
- [Zwa97] Joost Zwarts. Vectors as relative positions: A compositional semantics of modified PPs. *Journal of Semantics*, 14:57–86, 1997.

Résumé

Le travail présenté se situe dans deux domaines de recherche, la sémantique formelle des langues naturelles et l'ontologie formelle pour la représentation des connaissances, tous les deux abordés par des méthodes logiques dans un cadre résolument interdisciplinaire (intelligence artificielle, linguistique, philosophie, psychologie cognitive).

Il se penche particulièrement sur leurs interactions, partant du principe qu'une approche référentielle de la sémantique des langues requiert une ontologie complète du monde décrit et qu'une approche cognitive de la représentation des connaissances requiert des données empiriques, pour lesquelles la langue est une source privilégiée, et sont par conséquent interdépendantes.

Parmi l'ensemble des phénomènes de nature sémantique, ce travail se focalise sur le "contenu", c'est-à-dire sur les éléments en langue comme en représentation des connaissances qui dénotent et caractérisent les entités dont on parle, leurs propriétés et leurs relations, par opposition aux éléments qui affectent exclusivement la forme logique des messages, et aux problèmes de langage de représentation et de raisonnement. Les aspects linguistiques couverts vont de la sémantique lexicale (notamment prépositions spatiales) à la sémantique compositionnelle (adverbiaux de localisation temporelle ou spatiale) et à la sémantique et pragmatique du discours, avec, pour cette dernière, une contribution à la Segmented Discourse Representation Theory (notamment la caractérisation des relations de discours et l'étude du rôle des topiques de discours).

Les domaines de représentation couverts vont des entités concrètes (espace, espace-temps et relations de partie-à-tout dans une approche méréologique) à la réalité sociale et mentale (concepts, rôles, artefacts et actions). Le point de vue le plus récent adopté est celui de contribuer au développement de l'ontologie fondationnelle DOLCE (Descriptive Ontology for Linguistic and Cognitive Engineering).