



HAL
open science

Habilitation (Summary of main contributions to the field and some general perspectives on future research)

Ivan Varzinczak

► **To cite this version:**

Ivan Varzinczak. Habilitation (Summary of main contributions to the field and some general perspectives on future research). Artificial Intelligence [cs.AI]. Université d'Artois, 2019. tel-03048085

HAL Id: tel-03048085

<https://theses.hal.science/tel-03048085>

Submitted on 9 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Habilitation

Summary of main contributions to the field and
some general perspectives on future research¹

Ivan Varzinczak

CRIL, CNRS UMR 8188
Université d'Artois, Lens, France
varzinczak@cril.fr
<https://ijv.ovh>

“Habilitation à diriger des recherches” defended on 26 November 2019 at Université d'Artois before the following jury:

- Franz Baader (TU Dresden, Germany)
- Stéphane Demri (CNRS, France)
- Hans van Ditmarsch (CNRS, France)
- Sébastien Konieczny (CNRS, France)
- Pierre Marquis (Université d'Artois, France)
- Marie-Laure Mugnier (Université de Montpellier, France)
- Odile Papini (Université Aix-Marseille, France)
- Leon van der Torre (Université du Luxembourg, Luxembourg)

¹For a more detailed description of my recent research, the reader is referred to the companion monograph titled *Defeasible Description Logics*, which is the main scientific document for this “Habilitation à diriger des recherches”.

Contents

1	Introduction and a bit of context	1
2	My personal view on non-monotonic reasoning	7
3	My contributions from 2007 to 2019	11
3.1	Reasoning about actions	11
3.1.1	A solution to the frame and ramification problems	11
3.1.2	Foundations of regression for modal logics	13
3.1.3	Design principles for action theory specification	14
3.2	Belief change	15
3.2.1	Action theory revision	17
3.2.2	Horn contraction	18
3.2.3	Propositional typicality logic	19
3.3	Modal logics of defeasibility	21
3.3.1	Beyond propositional defeasible consequence	21
3.3.2	Beyond defeasibility of conditionals	23
3.4	Rationality in formal ontologies	25
3.4.1	Defeasible subsumption and representation results	26
3.4.2	Rational closure for defeasible DLs	27
3.4.3	Typicality of concepts and relations at the object level	27
3.4.4	Contextual defeasibility	28
4	Perspectives and future directions of investigation	30
	Bibliography	36

Chapter 1

Introduction and a bit of context

The success of modern artificial intelligence (AI) relies heavily on the realisation of general-purpose autonomous systems capable of behaving rationally in a given environment. Such systems are usually called (artificial) agents and may be of various types such as application software (e.g. a semantic search engine), or a hardware artefact (such as a network router), or even both (a robot). The environment in question can be a physical space, or a virtual realm such as the Web, or even some hybrid combination thereof.

Rational behaviour is commonly understood as the general ability of an autonomous system to make well-informed deliberations followed by appropriate interactions with the environment or with other agents within the system.¹ That is, rational autonomous systems have to be responsive to the environment, interact with it and with other programs or human users the way a human would normally do or sometimes in ways that are better than a human would do it. Such processes involve and are usually driven by goals to be achieved, questions posed by humans or by other artificial agents, or a need for adaptation when changes in the environment occur, amongst others.

What allows agents to be responsive to the intricacies of their environment is the fact they have an *internal representation* of the relevant aspects of the restricted part of the world they act upon. It is by consulting and manipulating this internal representation, which stands as a substitute for the real system, that agents are able to deliberate and interact with the environment and the human users. Such a representation constitutes the agents' *knowledge* about themselves, the environment and other agents, and should be as accurate as possible in capturing the agents' surroundings.

We expect agents to behave rationally (or at least sensibly) in their environment. It turns out that an agent's competence is largely influenced by its knowledge. Indeed,

¹Different notions of appropriateness of actions can be defined. For instance, in game theory, actions deemed as appropriate are those maximising an agent's expected utility.

how successful an agent is in exhibiting rationality will depend, to a large extent, on how the agent can deal with and use the information in the internal representation at its disposal.² Such information comprises both the knowledge the agent started off with as well as new information that is acquired through interaction with the environment and other agents. In particular, it is worth noting that the information we refer to here may be of a much more complex nature than that stored in standard data bases [1]: It is not only factual, but it can also be *prescriptive* [110] (telling *what* the agent should do in certain circumstances); it can be *taxonomic* [4] (telling the agent how the entities in its environment are related to each other); it can be *causal* [133, 134] (providing a description of the *effects* a given course of action will lead to), or several combinations thereof. Therefore the level of complexity of an agent’s internal representation goes far beyond that of standard data base systems.

In order to perform rationally in their environment, agents should be able to manipulate not only the information that is explicitly available to them, but also the information that implicitly follows from that which is explicitly stored. In other words, to be effective in taking appropriate decisions and interacting with its peers, an agent should be able to draw conclusions from what it knows. This amounts to the ability to perform *reasoning*, a faculty that goes beyond information retrieval from a data base. Therefore, the understanding and formalisation of human reasoning is an endeavour of great importance to AI and related disciplines.

Several approaches have been proposed for the representation of an agent’s internal knowledge and to allow it to perform reasoning on that representation [4, 83]. Some of them are based on data structures mimicking the inner workings of the human brain — the stance usually adopted in neural networks and, more recently, in machine learning, while others focus on rule-based procedural specifications — an approach commonly found in expert systems. Common to both approaches is the lack of a *semantics* [4], i.e., a shared and unambiguous understanding of the intuition behind the representation which is independent of a specific application domain. In that respect, ad hoc data structures and tailor-made algorithms are not general enough to serve alone as solid foundations for the diverse patterns of reasoning that are required.

Formal (mathematical) logic [111, 129], as it has been studied for the past millennia, and in particular as developed during the 20th century, provides us with solid foundations on which to base the study and formalisation of reasoning. This stems from at least three important features of logic: (*i*) its well-founded *formal semantics*, rendering it clear of ambiguity, (*ii*) its great degree of *generality*, making it applicable to different

²Of course, there are other factors that directly or indirectly contribute to the realisation of this goal, such as the accuracy of the agent’s sensors and actors (if it is a physical agent), amongst others. Here we are interested in the *cognitive* aspects of these agents, and therefore we do not investigate issues of the aforementioned nature, which are dealt with by the robotics community.

contexts, and (*iii*) its amenability to *implementation* in computer systems, allowing for reasoning not only in the physical world but also in the digital one. In this respect, the study of formal logic and its associated tools and techniques is of potential utility to numerous research fields, ranging from applied mathematics, philosophy and law to cognitive sciences, artificial intelligence and economics, to name but a few. As a result, the field of logic has established itself as one of the foundational areas of modern science and technology, in particular in knowledge representation and reasoning (KR&R).

In logic-based KR&R, the internal representation of an autonomous system's knowledge is materialised in the form of knowledge bases [104]. A knowledge base (KB) consists of a set of statements specified in some formal, machine-processable, language. Such a language is based on mathematical logic, which, as alluded to above, provides us with a well-defined formal semantics and an associated method for performing reasoning that can be automated in computer systems. The information in the KB is intended to capture or represent knowledge about some particular domain of interest in a concise and unambiguous way [4]. It is usually specified with the help of a domain expert. The prototypical logical formalisms behind knowledge representation languages in AI are classical (propositional and first-order) logic [24, 83, 115, 125]. Therefore, the idea behind the use of knowledge bases is for the agent to start off with a finite explicit representation of knowledge about the domain and then, by performing logical reasoning, derive pieces of knowledge that are implicit, i.e., consequences following logically from the explicit knowledge in the KB, allowing the agent to make informed decisions [104].

The type of language one chooses to represent knowledge is driven by the nature of the knowledge one wants to represent and the type of application one has in mind. State-of-the-art research in KR&R has a strong focus on the development of general-purpose logic-based representation languages and therefore issues related to expressivity and complexity are at the core of current open problems in the area [4, 24, 83].

All of this has to be done taking into account the *human user* (e.g. a knowledge engineer) who will encode the knowledge in the agent's KB. This means that knowledge representation languages should be designed according to common principles of software ergonomics [120], in other words they should be neat and elegant. Moreover, technical intricacies such as the complexity of the semantics or the reasoning mechanism itself should be hidden from the user. These issues have a direct impact on the usability of a logical formalism in the implementation of real-world applications.

A core aspect of rational behaviour is the ability to perform *commonsense reasoning* [107, 115], i.e., the type of reasoning (quite often tentative and even unsound from a logical point of view) that humans do in everyday life. The term commonsense reasoning is actually a quite broad notion encompassing many different yet interrelated reasoning tasks. It involves for instance the capacity to deal with exceptions or special cases [31, 74, 75], the ability to adapt one's knowledge in the presence of new,

possibly conflicting information [2, 58, 81], and, fundamentally, the aptitude to make decisions under incomplete or uncertain information about some aspect of the environment [99, 102], amongst others. These features are all examples of the fundamental component of commonsense reasoning commonly known as *defeasible reasoning*, the type of reasoning in which conclusions can be defeated, or withdrawn. Completing the picture, it is worth noting that as finite entities, agents usually have limited resources, both in time and space. Knowledge being a fundamental resource of an intelligent agent, it is obviously limited, too. This means that in many situations an agent will have to make decisions under incomplete or uncertain information about some aspect of its environment. Of course, ‘guesses’ of this type may be wrong — they are defeasible, but an agent with the ability to retract knowledge would not collapse when making incorrect assumptions. Hence the ability to make provisional, tentative conclusions depends on the capacity of handling exceptions and retracting previous knowledge as mentioned above. Without a proper formalisation of defeasible reasoning, the understanding and eventual realisation of the broader notion of commonsense reasoning in machines will never see the light of the day.

Historically, logical formalisms have shed light on how to reason correctly, i.e., how one ought to reason in ideal circumstances. However, it is equally important to understand how we ought to reason when these ideal circumstances are not met, notably in human quotidian contexts. It turns out that classical reasoning is not flexible enough when directly applied to certain aspects of human decision making. This is so because of properties and principles of classical reasoning that, in spite of making perfect sense in an ideal mathematical setting, are hard to justify or to deal with in everyday life. These comprise, for instance, the law of non-contradiction, the law of excluded middle, the so-called ‘paradoxes’ of relevance, and monotonicity, to mention just a few. Among these, monotonicity, which in its most basic form states that deductions are always incremental (if $\alpha \models \beta$, then $\alpha \wedge \gamma \models \beta$), has been acknowledged in the literature as the major stumbling block in a proper understanding and formalisation of commonsense reasoning [73, 108, 114, 115, 124]. As such, the eschewal of monotonicity, or rather a careful neutralisation thereof, is the main focus of the area of qualitative analysis of uncertainty in reasoning, also known as *non-monotonic reasoning*.³

The field of non-monotonic reasoning is quite broad and comprises several research domains. In this document, I will focus for the most part on the following areas: Reasoning about Actions and Change, Belief Revision, and Defeasible Consequence Relations. These are the areas in which I have mostly been active over the past twelve

³Sometimes the term *non-monotonic logics* is also used, although it seems to be less consensual. Here I shall use *non-monotonic reasoning*, which is broader in that it also includes methods of belief revision [81] and argumentation theory [12].

years and the purpose of the present memoir is precisely to provide a summary of my contributions to the area along that period. Before doing so, it is worth emphasising that when it comes to research methodology I have relied on the following fundamental assumptions [115]:

1. I assume, along with most cognitive scientists, that reasoning involves the use of *representations* and computational processes that operate on those representations;
2. I assume, along with researchers in symbolic artificial intelligence, that these representations are of a *symbolic nature*;
3. I assume, along with researchers in logic-based AI, that knowledge is best represented *declaratively* rather than procedurally, and
4. I have used as basic building blocks the declarative languages of *modal logics* and *description logics*, as well as fragments and extensions thereof, and their associated reasoning services.

Research in knowledge representation and reasoning usually focuses on two fundamental aspects, namely the *representational* dimension, i.e., that of designing suitable representation languages, and the *reasoning* dimension, i.e., that of devising correct algorithms and techniques for reasoning both with and about the representation in knowledge bases. In that respect, I have worked on these two intertwined dimensions by investigating representation languages with appropriate expressive power as well as methods for automating different forms of reasoning in such languages.

Organisation of the document The present memoir is organised as follows: In Chapter 2, I summarise my personal perspective on (logic-based) non-monotonic reasoning, its main aspects and the different levels at which it has to be dealt with. Chapter 3 contains a summary of my scientific contributions to non-monotonic reasoning for the period spanning from 2007 up until now and that I deem as the most relevant in the context of the present “habilitation”. Chapter 4 closes the memoir by highlighting some of the challenges in future developments of defeasibility in reasoning and specifying the directions I plan to follow in pursuing them. The appendices contain the required administrative documents, namely a complete CV and copies of the relevant publications for the period under consideration (2007–2019).

A Disclaimer The purpose of the present document is to provide a general overview of the work I have done over the past twelve years and its significance, as well as how its

different pieces fit together as a hotbed for further investigation into the area. Technical details can be found in the relevant papers referenced along the text and made available in the appendix. The presentation style here adopted is semi-technical, although formal. The companion monograph titled *Defeasible Description Logics* presents with a much higher level of detail my technical contributions to the development of defeasible reasoning for formal ontologies and on which I have worked from 2011 to 2019.

Chapter 2

My personal view on non-monotonic reasoning

New logics can be created in several different ways. We can add new constructs to a given language or restrict its grammar. We can enrich the semantic structures or impose restrictions on them. We can alter the underlying entailment relation, obtaining the so-called infra- or supra-classical logics. We can even restrict or augment the corresponding proof theory.

In a certain sense, non-monotonic reasoning shares some of the aims of traditional non-classical logics such as paraconsistent [121], intuitionistic [91], relevance [60] and adaptive [8] logics inasmuch as it gets to grips with one or more features of classical logic seen as problematic. Contrary to them, though, logics of non-monotonicity are not to be taken as replacements for or opposed to the classical one. They do not entirely reject classical consequence as incorrect; instead they build on it to get consequence relations that are at the same time stronger and more flexible. In that sense, logics of non-monotonicity are *not* to be seen as a family of non-classical logics [105]. They are (in general) *supra-classical*.¹

Designed to be stronger while more versatile than classical logics, non-monotonic formalisms usually show the two fundamental intertwined aspects below:

- **Ampliative aspect:** augmenting the underlying reasoning by allowing more conclusions. In practical contexts, this amounts to the ability to make inferences that venture beyond the scope of the premises, somehow in an ‘unsound’ but justifiable way. Prominent examples are (i) *default reasoning* [124]: jumping to conclusions deemed as plausible ‘by default’, i.e., in the absence of information to the contrary, such as applying negation as failure [50] or adopting the closed-world assumption [123], and (ii) inductive and abductive reasoning [64]: taking

¹There are also examples of infra-classical logics exhibiting non-monotonic features [33, 126].

chances in drawing conclusions that implicitly call for further scrutiny or tests by empirical observations, as in making inductive hypotheses in scientific theories or finding abductive explanations in forensics.

- **Defeasible aspect:** curtailing the underlying reasoning either by disregarding or by disallowing some conclusions that somehow ought not to be sanctioned. In practice, this amounts to the ability to backtrack one’s conclusions or to admit exceptions in reasoning. Some examples of this are (i) *retractive reasoning* [2]: withdrawing conclusions that have already been derived, such as in belief contraction [81] or in negotiation, and (ii) *preemptive reasoning* [96, 108]: preventing or blocking the inference of some conclusions by disallowing their derivation in the first place, as in dealing with exceptional cases in multiple inheritance networks [132] and in regulation specification [117].

Hence, put simply, the central research question in non-monotonic reasoning revolves around the problem of determining for a given context how one can sanction *more* conclusions as well as how to sanction *fewer* of them.

It turns out that the two aforementioned aspects of non-monotonicity manifest themselves in logical approaches, to a greater or lesser extent, at various different *levels* in the following way:

1. **Object level**, i.e., that of the *logical symbols* of the language: in this case the *connectives* can behave non-monotonically, as for instance a non-monotonic version of material implication \rightsquigarrow , showing an ampliative behaviour ($\alpha \rightsquigarrow \beta$ holds even if $\alpha \rightarrow \beta$ doesn’t) and a defeasible one ($\alpha \rightsquigarrow \beta$ is the case but not $\alpha \wedge \gamma \rightsquigarrow \beta$).
2. **Entailment level**, i.e., that of sanctioned *inferences* or *reasoning*: in this case it is the (classical) entailment relation \models (or the deducibility one \vdash , if one prefers) that has a non-monotonic counterpart ‘ \approx ’ (alternatively ‘ \sim ’), with both an ampliative behaviour ($\{\alpha\} \approx \beta$ even if $\{\alpha\} \not\models \beta$) and a defeasible one ($\{\alpha\} \approx \beta$ and $\{\alpha, \gamma\} \not\approx \beta$ do not conflict). (This level has been the one at the origin of research in non-monotonic reasoning and probably the most extensively studied.)
3. **Meta-reasoning level**, i.e., that of reasoning *about* the sanctioned inferences: here non-monotonicity happens ‘outside’ the logic since additions and removals of conclusions are dealt with at a level beyond that of entailment. Prominent examples of this level are the theory of belief revision [2, 81] and argumentation theory [12]. In both cases, the entailment relations that are reasoned with remain Tarskian (monotonic). Non-monotonicity shows up in the passage from one classical theory to a new one created through a series of revision cycles (additions

and removals of theorems)² or through a dialectical process of evolving possibly conflicting knowledge bases.

Over the past 40 years, efforts have been put into the study and definition of logical formalisms within which the aforementioned aspects of non-monotonic reasoning could adequately be captured in one or another of the above levels [24, 25, 67, 68, 73, 81, 83, 105]. These include, in particular, default logic [124] (*entailment* level), autoepistemic logic [114] (*entailment* level), AGM belief revision [2] (*meta-reasoning* level), circumscription [108, 109] (*entailment* level), abstract argumentation frameworks [59] (*meta-reasoning* level), dynamic epistemic logic [58] (at the *object* level) and preferential logics [68, 96, 102, 128] (both *object* and *entailment* levels), to name but a few. These have mainly been developed by philosophers and computer scientists, but of course this does not affect their mathematical interest or value.

Despite the importance and relevance of the foregoing approaches, all of them have somehow fallen short of providing a general and comprehensive account of non-monotonic reasoning. This is so because either they have not dealt with all the aforementioned levels of non-monotonicity, or they have only partially explored the two aspects of non-monotonicity within a given level, or both. We make this claim more precise in the following paragraphs.

At the object level, there have been many proposals of non-monotonic versions of conditional-like connectives [23, 34, 35, 46, 51, 55, 96]. These allow for exception-tolerant reasoning in the object language by blocking the conditional's *consequent* in non-typical situations. The usual example is that birds normally fly but penguins normally don't [96]. In spite of the success in characterising the intuitive behaviour of such non-monotonic connectives [96, 102], research in non-monotonicity at the object level remains somewhat limited in that it is restricted to non-monotonic *implication*, even in logics with more expressive power than the propositional one [27, 34, 35, 46, 101, 122]. Except for a few recent preliminary attempts [5, 19, 28, 37, 38, 131, 138], existing approaches have somehow left open the question of whether other logical connectives could also exhibit some aspect of non-monotonic behaviour.

At the level of entailment, which is the most thoroughly investigated level in non-monotonic reasoning, there has been a great deal of work done on how to define entailment relations with a built-in mechanism for allowing the jump to conclusions while still permitting the withdrawal of previously sanctioned inferences [3, 30, 99, 102, 106, 108, 109, 124]. In particular, the notions of *prototypical reasoning* (or *rational closure*) [102]

²In this level, the ampliative aspect is more 'conservative' than in the others in that the new piece of information must follow *classically* from the resulting theory, i.e., it is not necessarily regarded as being just provisional. This stems from the principle of *primacy of new information* [53], usually adopted in the field.

and *presumptive reasoning* [99] have shown to be quite promising in a propositional setting. Recent developments in the field have seen quite a few extensions of rational closure beyond propositional logic [27, 34, 35, 36, 46, 77, 78, 79], while those of presumptive reasoning are on their way. There seems to be an agreement, though, that both prototypical reasoning (delivering a more ‘conservative’ entailment relation) and presumptive reasoning (giving rise to more ‘venturous’ consequences) represent rather extreme cases and that there may be a whole spectrum of equally interesting consequence relations between the two. There is therefore a need to explore other types of non-monotonic entailment, in particular in the context of more expressive logical languages like those of *modal logic* [15, 49] and fragments of first-order logic such as *description logics* [4].

At the meta-reasoning level, an extensive body of research has been erected in the context of theory change (notably following the AGM paradigm [63, 69, 81]), abstract argumentation frameworks [12, 59] and their intersection [9, 13, 48]. While argumentation frameworks usually disregard the underlying logical language (thereby not addressing non-monotonicity at the object level), AGM-style revision assumes the formalism of (full) classical propositional logic. This stands as one of the historical limitations of the AGM paradigm, since there are good reasons to investigate the problem of belief change in logics other than the propositional one, such as modal logic, first-order logic and fragments thereof. It is only recently that theory change has been explored in less expressive logics such as Horn logic [18, 22, 56], and in some application-oriented fragments of modal and first-order logics [65, 113, 136]. Moreover, an outstanding issue of both argumentation frameworks and belief revision remains the imposed restriction to underlying entailment relations that are *monotonic*. It is quite reasonable to envisage theory change and argumentation frameworks based on logics that do not have a Tarskian consequence relation.

Given the discussion above, we can conclude that there are still many relevant issues across the three levels at which uncertainty shows up in reasoning that remain to be properly studied and formalised. Therefore, new and more general formal theories of non-monotonic reasoning, supplementing the existing ones in some crucial aspects, are called for.

In the next chapter, I will provide an overview of how I have dealt with the different aspects and levels of non-monotonic reasoning over the past twelve years in some application scenarios that are relevant from the standpoint of AI.

Chapter 3

My contributions from 2007 to 2019

In the following, I present a summary of the contributions I have made to the area of non-monotonic reasoning over the past twelve years. These have resulted in several papers in international conferences and articles in journals. The most significant of these publications are mentioned in the following sections.

3.1 Reasoning about actions

The area of reasoning about actions and change is central to artificial intelligence not least because it studies the foundations of intelligent agents interacting with the environment and with other agents. In order to behave rationally, agents have to make plausible decisions and perform appropriate *actions*. To achieve that, artificial agents therefore have to be equipped with an accurate description of which actions are available and what their effects and preconditions for executability are. It turns out that specifying such descriptions, as well as reasoning with them, is a hard task, in particular due to the *frame problem* and the *ramification problem*.

3.1.1 A solution to the frame and ramification problems

The so-called frame problem was first identified in the late 60's and is concerned with the things that are not affected by the execution of an action, i.e., the action's 'non-effects'. Indeed, when an action is performed, those aspects of the world that are, in a sense, irrelevant to the action remain unchanged through the action's execution. Having an accurate description of which aspects of the environment are not affected by a given action is crucial for intelligent agents. The problem is that for each action there are too many things that it does not affect. Hence, defining a correct and parsimonious way of specifying such information is of paramount importance.

The ramification problem, identified in the 80's, is, in a sense, the 'complement' of the frame problem. Actions may have effects that are directly observable (and therefore relatively easy to specify), but they may also produce a domino-like chain of effects, called 'ramifications' or indirect effects. Specifying and deriving *all* the relevant effects that a given action may produce as outcome is the core of the ramification problem.

Both the frame and the ramification problems received a lot of attention from the AI community in the 90's. In particular it has been observed that a satisfactory approach to the ramification problem should provide an appropriate account of the notion of *causality*. Nevertheless, two main criticisms can be made about the solutions that were proposed in that period. First, for the most part they were based on full first-order (predicate) logic, which besides being *undecidable* has a syntax that is not particularly user-friendly. Second, existing solutions to the frame and the ramification problems either could not cope with or did not deliver the expected results in more complex domains, notably scenarios involving actions with non-deterministic indirect effects. All of this questions the suitability and generality of the mentioned approaches for real-world applications.

I have contributed to this area by defining a new formalism for reasoning about actions and change with a solution to both the frame and the ramification problems. This formalism, called *LAPD*, the *logic of actions and plans with a ternary dependence relation*, is based on an expressive fragment of propositional dynamic logic (PDL) and therefore does not suffer from the decidability issues of first-order logic. (Indeed, I devised a sound and complete tableaux-based proof procedure for the new logic, which makes it amenable to implementation.) Moreover, the syntax of *LAPD* is much simpler than that of first-order based approaches, making it a suitable formalism in which to specify dynamic domains.

In line with contemporaneous approaches, *LAPD* is also equipped with a notion to express causality. The main difference, though, is that in *LAPD* the latter is formalised as a dependence relation between *actions*, their *effects* and execution *contexts*. This formalisation of causality allows for a solution to the frame problem via the assumption that facts which do not depend on an action in a given context remain unchanged. A solution to the ramification problem is ensured by the properties of the dependence relation: it can be semi-automatically generated from the set of direct effects for each action. Finally, *LAPD* does not suffer from the same counter-examples that plague other causality-based approaches because with dependence relations the link between indirect effects and the action originating a chain of changes is not broken.

Significance of the contribution In summary, the significance of *LAPD* is two-fold: (*i*) it challenged and refuted one of the widely-accepted assumptions in the community, namely that indirect effects do not depend on the actions triggering them, and

(ii) it provided a formalism which allows for the representation of and correct reasoning within complex scenarios that other approaches cannot cope with. As a result, I have chosen \mathcal{LAPD} as the underlying formalism for the specification of action domain descriptions and their evaluation in the contribution reported in Section 3.1.3, which appeared as an article in the journal Artificial Intelligence [90].

3.1.2 Foundations of regression for modal logics

An elegant solution to the frame and the ramification problems in domains with deterministic actions arose from a series of papers in the early 90's. In the context of the Situation Calculus (a dialect of full first-order logic), one can specify the effects of actions by stating a number of formulas of a special kind, called *successor-state axioms*. Roughly, each successor-state axiom says that a given property of the world changes its truth value if a specific action takes place under certain preconditions, and remains the same in any other context, in particular if some other action is executed. In other words, successor-state axioms allow for a quite concise specification of the effects and non-effects of actions in many scenarios.

With a domain specified in terms of successor-state axioms, it becomes possible to decompose any query about a deterministic domain into an equivalent simpler one by applying a rewriting procedure, called *regression*. Roughly, regression consists in reducing a given first-order formula with special symbols (such as actions) into another, equivalent one with just *propositional constants*. As a consequence, checking for validity of a first-order formula can then be done by checking validity of its regressed propositional counterpart, which is computationally more efficient.

With regression we can substitute a first-order formula by an equivalent sentence in propositional logic in which all action symbols (and other first-order constructs) have been adequately dismissed. The immediate advantage of this equivalence-preserving transformation is that when performing reasoning with the dynamic description of the world, one can reduce first-order satisfiability check to the problem of satisfiability in propositional logic, of which the computational cost is much lower.

It turns out that the 90's have seen a move towards logics that have a cleaner syntax and more appealing computational properties than first-order logic, notably modal logic, description logics, and combinations thereof. This raised the obvious question on whether regression techniques could also be developed for these formalisms. As it turns out, I contributed to the establishment of the foundations for regression in modal logics [57] by introducing a rewriting technique for reducing entailment of complex formulas specified in a fragment of quantified modal logic to entailment of plain Boolean sentences in classical propositional logic.

Significance of the contribution This work is significant to the reasoning about actions community in that it has shown that regression in formalisms other than the situation calculus is not only viable but also more elegant and fruitful in many contexts beyond that of reasoning about actions. In particular, this is witnessed by several spin-off papers by peer researchers on various extensions of modal logics of actions to deal with knowledge and beliefs. The body of research around this topic stands in support of the significance of this contribution to the field and of modal logic as a suitable core formalism for reasoning about actions, knowledge, beliefs and related notions, as well as several combinations thereof. As a result, I have chosen modal regression as the underlying proof procedure for the verification of specifications of deterministic scenarios in the contribution reported in Section 3.1.3, which appeared as an article in the journal Artificial Intelligence [90].

3.1.3 Design principles for action theory specification

Knowledge bases formalizing the specification of a dynamic domain are pieces of software applications and, as such, they should be designed according to good principles of software development. Notwithstanding the obvious need for a set of guidelines and practical tools to aid the specification of action domain descriptions, the reasoning about action community lacked formal principles on which to base the conception of action theories and the pinpointing of design errors in them.

In addressing this problem, I have investigated what makes a formal specification of a dynamic domain suitable for reasoning, error discovering and maintenance. As an approach towards answering the question, I studied and proposed a set of formal guidelines, borrowed from the well-established area of software engineering, to the specification of dynamic environments in a logic-based language. Going beyond the mere notion of logical consistency of a knowledge base, I proposed a set of postulates (meta-properties) that every domain description in reasoning about actions ought to satisfy. I have then devised algorithms that not only check whether an action theory satisfies the postulates but that also provide hints to the human user (a knowledge engineer in charge of debugging and maintaining the knowledge base) on why that is not the case. (These algorithms have been proved correct with respect to the set of postulates.) The main outcome of this work is an article in the top AI journal Artificial Intelligence [90].

The design principles that were put forward in this work have since been employed in many follow-up articles by researchers world-wide. In particular, the results of this investigation carried over to other areas of logic-based knowledge representation and reasoning such as general game playing as well as ontology engineering and debugging.

Significance of the contribution The significance of this contribution to the field resides in the fact that it has established formal links between reasoning about actions and software engineering by exploiting the formalisation in an action context of principles commonly adopted in software specification. It has provided the theoretical foundations of verification tools for knowledge engineers to evaluate logical formalisations of dynamic environments. Furthermore, the definitions for action theory evaluation provided in this contribution turned out to be fruitful in the definition of revision operators for action laws, summarised in Section 3.2.1. Indeed, domain descriptions abiding by the design principles I have put forward are more amenable to changes. The main publications resulting from this contribution have received 120+ citations altogether according to Google Scholar.

3.2 Belief change

When deploying knowledge bases for real-world applications, the use of the reasoning capabilities of the underlying logical formalism is not limited to deductions of what follows from the knowledge base at the object level (reasoning *with* the knowledge represented therein). A particularly important class of reasoning services associated with knowledge bases is that of knowledge base *evolution* and *maintenance*, which are examples of reasoning *about* the KB itself and that are defined more precisely in the paragraphs below.

During their lifetime, knowledge bases may need to be modified, and that may be due to different specific reasons. For instance, it may be the case that the world that the KB models has evolved since its first formalisation and therefore the knowledge base does not describe it accurately anymore. This is an example where the KB needs to be *updated* with new information. It might also be the case that the initial representation of the knowledge in the base was incomplete or incorrect with respect to the real world, or even inconsistent on its own or with respect to some new incoming piece of information. In that case, the knowledge base should go through a *revision* process.

Furthermore, during the lifetime of a knowledge base, there may also be a need to change the *language* in which it is formalised, i.e., one might eventually discover that the language used to formalise the domain is either not expressive enough to allow for expressing some intended subtleties, or it is too expressive, making the underlying reasoning unnecessarily more complex. Indeed, even in the case when the logical substratum is not meant to be changed, specific, localised, changes in the underlying signature of some constructs may occur (think of a 2-argument function that should now get 3 arguments when invoked). Finally, it might also be the case that different knowledge bases have to be somehow integrated in such a way for them to be operable

together. In this scenario, possible conflicts between the merging KBs should be identified and removed, and information with equivalent meaning in different KBs should not be considered as different in the bigger system. This is the problem of knowledge base *integration* or *alignment*.

All the aforementioned operations are instances of the problem of *evolving* knowledge bases. The most prominent approach in the literature on (classical) knowledge base revision and update is the so-called AGM paradigm [2, 69]) and extensions thereof. Issues related to knowledge base integration have been dealt with via *merging* operators [94]. These approaches essentially deal with knowledge bases specified in classical propositional logic.

Given that knowledge bases may evolve over time, there is a need for an automated way of storing, retrieving, controlling and articulating their several different versions. This is what we call knowledge base *maintenance*. Therefore knowledge engineers also need tools for maintaining and controlling existing different versions of the same knowledge base [66]. That also means that knowledge engineers need appropriate methods and tools for evaluating different versions of a KB and ensuring that each of them has good properties with respect to its intended applications [87, 89, 90, 134]. Needless to say, all these operations should be carried out in an efficient way, both in terms of computational and data complexities. They should also allow for both human supervised and unsupervised functioning.

Despite the obvious importance of maintenance frameworks for KBs, they are relatively neglected objects of investigation. This matter has only recently been given some attention in the ontology community [95, 116] and in the case of classical (propositional) knowledge bases [66]. As a consequence of being a novel research area, a prominent and widely accepted paradigm is still lacking in this field.

As already argued before, it turns out that, according to specific needs, many interesting applications require languages that are *more* expressive than classical propositional logic, yet having attractive computational properties. It may also be the case that one needs languages that are *less* expressive than propositional logic. Examples of the former are modal logics [14] and description logics (DLs) [4]. The most prominent example of the latter in computer science is Horn logic. There may also be good reasons for using extensions of first-order logic, or even higher-order logics.

As expected, the problems of evolution and maintenance also make sense for knowledge bases specified in logical formalisms other than standard propositional logic: As an example, consider ontologies specified in some DL. In that setting, revision cycles are normally the rule: information is added and removed frequently (ontology update, debugging and repair). Moreover, when developing or maintaining an ontology collaboratively, different simultaneous (possibly conflicting) versions thereof might exist at the same time. That can happen due to many reasons, such as different teams working

on different modules of the same ontology in parallel, or different developers having different views of the domain, among many others.

A further example is that of changing an agent's beliefs regarding the expected behaviour of actions and their effects [61, 85, 135, 136]. After attempting to perform an action, the agent may discover that the action in question actually has an outcome that is different from the one that it has always believed that action had. Alternatively, it might be the case that the action under consideration is no longer executable (due to some unforeseen reason), or that now it has unexpected side effects. Moreover, specific versions of an action theory can also be seen as *configurations* in which the agent might be according to specific *contexts*. Selecting and maintaining these configurations so that the agent is able to switch between them at run time in an efficient way becomes an issue of paramount importance.

All of this would be good and well if only the traditional AGM paradigm and related techniques could be directly applied to logics other than classical propositional logic. It turns out that logics that are *more* expressive than propositional logic, such as modal logics and DLs, have more structure, and therefore belief change for these logics needs to take this into account. Indeed, as shown by recent work [65, 136], existing methods for propositional theory change do not take care of the additional structure in these logics and therefore will give either incomplete solutions or unintuitive results when deployed in the evolution of more expressive KBs.

Similarly, limitations of the underlying language may also constrain the applicability of traditional methods for theory change. For instance, the fact that a given logic is a fragment of classical propositional logic does not necessarily mean that the existing theory change paradigm for full propositional logic applies to it: the result of the change operator might not be expressible in the underlying object language [18, 21, 56], compromising the completeness and therefore the reliability of these approaches.

From the above it follows that there is a clear need for new paradigms of KB evolution for logics that are either extensions or fragments of classical logic. It is also reasonable to expect this to hold for knowledge bases formalised in some specific infra- or supra-classical logical formalism.

I have made investigations into the problem of belief change for both logics that are more expressive and less expressive than classical propositional logic. The results are summarised in what follows.

3.2.1 Action theory revision

In a paper which appeared at the prestigious KR conference [135] and of which I am the sole author, I addressed the problem of changing a domain description for reasoning about actions specified in dynamic logic. At that time this was a novel and ill-explored

area of research. In the referred paper, I have formalised the notion of *contraction* of action laws from an action domain description.

More specifically, I have defined semantic operations on the set of models of the underlying theory capturing the intuition of the changes that ought to be performed in order for an action law (effect law, executability law and static law) to be removed from an agent’s set of beliefs about the actions under consideration. I have also provided the algorithmic counterpart of these semantic operations and shown their correspondence. An important point of the contribution is the role of a notion of minimality: modifications should be carried out in a minimal way, i.e., one should neither remove too much information nor add irrelevant new knowledge. I have shown that the proposed approach provides for a suitable notion of minimality when changing action domain descriptions.

This work was subsequently extended also to the case of action theory *revision*. I have defined the semantic operations corresponding to the notion of minimally revising action laws and I have provided the corresponding algorithms that effectively perform the expected changes at the level of the knowledge base. As done for the case of contraction, I have shown a correspondence result between the semantic constructions and the algorithms.

All the aforementioned results, together with the detailed proofs, appeared as an article in the prestigious Journal of Artificial Intelligence Research (JAIR), of which I am the sole author [136].

Significance of the contribution This contribution is significant in that it was the first to have defined contraction and revision operators for action laws in a reasoning about actions context, in particular for a modal-based formalism. Furthermore, the proposed approach stands as a bridge between standard AGM revision and existing approaches to ontology debugging and repair in description logics, which are closely related to modal logics. The main publications resulting from this contribution have received 40+ citations altogether according to Google Scholar.

3.2.2 Horn contraction

Horn logic is a fragment of classical propositional logic of great importance in computer science and artificial intelligence, notably because it is the backbone of logic programming. As a result, the ability to make modifications to Horn knowledge bases is a relevant matter. As it turns out, just applying the standard AGM approach to fragments of propositional logic is not enough: the expected result of revision is not always expressible in the target language of Horn logic. Quite surprisingly, it was only recently that this problem started drawing attention from the community [56].

In a paper which appeared at IJCAI [18], together with my collaborators I identified issues with existing preliminary construction methods for Horn contraction, namely they were not complete. In order to overcome the problem, we have proposed a new, more fine-grained construction method for contracting Horn clauses from Horn theories, with an accompanying set of AGM-style postulates, and we have shown the corresponding representation result. Furthermore, we have introduced a new form of contraction allowing for the contraction of the *disjunction* of sentences, which is not directly expressible in Horn logic.

In a follow-up paper which appeared in the Journal of Artificial Intelligence Research (JAIR) [22], we investigated the problem of Horn *belief base* contraction and compared it with the aforementioned proposal of ours. The main conclusion of this work is that the approach to contraction known as belief base contraction is more appropriate in the Horn case than the one based on belief sets.

Significance of the contribution These contributions are significant in the sense that they have surpassed those of previous approaches and have opened an avenue for exploration of belief change at the intersection of both belief base and belief set-based approaches. Moreover, they have strengthened the interest of the community for the problem of belief revision in fragments of propositional logic. This is witnessed by the number of publications having followed the aforementioned ones. The main publications resulting from this contribution have received 90+ citations altogether according to Google Scholar.

3.2.3 Propositional typicality logic

A fundamental, albeit tacit, semantic notion in belief revision, in particular of the AGM approach, is that of *normality* (alias *plausibility* or *typicality*). Indeed, semantically the result of revising a (propositional) knowledge base with a new sentence should correspond to the most *normal* (or *plausible* or *typical*) models of the knowledge base satisfying the new incoming piece of information. The same observation applies to the so-called KLM approach to defeasible reasoning [96, 102], where defeasible conditionals of the form $\alpha \sim \beta$ are read as “the most normal (alias typical or plausible) α -worlds are β -worlds.” Given the correspondence between the AGM and the KLM approaches [70], it is known that one can seamlessly switch between the two when reasoning non-monotonically.

Curiously enough, neither the AGM approach nor the KLM one allow for referring explicitly to a notion of typicality, be it at the object language or at the meta-level. This has important implications from the standpoint of knowledge representation and reasoning. In many cases, one wants to be able to capture the fact a given situation

is typical in the object language and use it in more complex constructs. Moreover, as proven in the DL case [74, 137], an explicit notion of typicality comes in handy when modelling exceptions in a given domain of application (also see Section 3.4.3).

In a paper which appeared at the JELIA conference [19] and of which the extended version appeared as a book chapter [20], we have filled precisely this gap. We have introduced an explicit operator to talk about typicality which intuitively allows us to single out those most typical states of affairs in which a given formula holds. We have shown that the resulting framework, called propositional typicality logic (PTL), delivers a more expressive language than that of standard KLM-style conditionals and, moreover, is also expressive enough to embed AGM belief revision.

In a paper which appeared at IJCAI [16] and in its more recent, expanded version, published in the journal Artificial Intelligence [17], we investigated different notions of (non-monotonic) entailment for PTL, each one based on the definition of rational closure as introduced by Lehmann and Magidor for KLM-style conditionals [102], and constructed using different notions of minimality. These results have shed light on the interplay between the expressive power of the underlying language and the associated notion of rational closure and constitute a springboard with which to investigate rationality in e.g. description logics (see Section 3.4).

Significance of the contribution PTL combines two important features, namely (i) it is a simple framework since it is plain propositional logic with an extra typicality operator added, and (ii) it is more expressive than the KLM approach and powerful enough in order to capture full AGM belief revision. Given the latter, PTL can also serve as a formal tool with which to analyse and compare different belief revision operators. Furthermore, in spite of being a simple syntactic superset of propositional logic, PTL shows that even minimal extensions such as the addition of a typicality operator widens the choices for an appropriate notion of defeasible entailment. Indeed, in the basic KLM framework, different semantic notions of minimality collapse to the same form of entailment, namely rational closure. In the case of PTL, different semantic constructions correspond to different definitions of entailment, some of which extending rational closure and some being incomparable to it. This makes PTL a good framework within which to investigate the properties of different definitions of defeasible entailment and to assess which ones are more appropriate in a given context. The main publications resulting from this contribution have received 70+ citations altogether according to Google Scholar.

3.3 Modal logics of defeasibility

For the most of AI’s existence, propositional logic was sufficient for the implementation of knowledge-based systems, in particular for the non-monotonic reasoning community. That used to be the case at least until the early 90’s. Nevertheless, modern KR&R applications require languages that are more powerful than the propositional one and with the ability to express e.g. the effects and preconditions of *actions* [125, 127]; an agent’s *beliefs* [58, 62]; a system’s *regulations* [110]; *ontologies* [4], and several combinations thereof. Notions such as these cannot properly be captured within the framework of propositional logic. This is because propositional languages are not expressive enough to represent the different types of knowledge mentioned above.

On the other side of the spectrum is First-Order Logic (FOL) [7, 111], within which all the notions mentioned above can be expressed, but that also turns out to be too expressive for our purposes, making it hard for users to understand and apply it in practice. To make things even worse, full FOL has bad computational properties — it is undecidable, and therefore it is not suitable as the core formalism of real-world applications. Even if full FOL is not assumed, the complexity of reasoning can easily become intractable [4].

One of the best alternatives to propositional and first-order logic has been the use of different systems of *modal logics* [11, 14, 49]. Modal logic allows for more expressive power than propositional logic without being hampered by the complexity and decidability issues of FOL. Amongst its features are a simplified syntax, an intuitive semantics and its amenability to implementation (especially via tableaux systems [80]).

Despite its philosophical origins, modal logic turned out to meet a variety of applications in artificial intelligence. Indeed, well-established formalisms for dealing with these notions in the AI literature are mostly variants of modal logic: The above mentioned examples illustrate applications of dynamic logic [82] and logics of action [47, 57, 90, 140, 141, 142], epistemic logic [62, 58], and deontic logic [110]. This is because modalities are quite versatile and can be interpreted as actions, beliefs and obligations, to name but a few.

Given the aforementioned reasons, modal logics stand as the foundation of a very important class of knowledge representation formalisms in artificial intelligence. Nevertheless, modal logics do not have all the ingredients that are necessary to give a satisfactory account of defeasible reasoning.

3.3.1 Beyond propositional defeasible consequence

As classical formalisms, standard modal logics do not provide an adequate framework within which the different aspects and levels of non-monotonic reasoning described

in Chapter 2 can be captured. Therefore enriching modal logics with non-monotonic reasoning capabilities constitutes a natural step in their development.

Although the KLM approach can directly be transplanted to a certain number of logics, it turns out that those constructions suffer from the limitation that they are largely *propositional* in nature — they do not cater for the additional expressivity of target languages other than the propositional one. Indeed, research on preferential reasoning has reached maturity only in a propositional context, whereas many logics of interest, like modal logic, have more structure in both the syntax and the semantics. If one wants to be able to capture the different forms of defeasible reasoning, then one has to fully move beyond propositional preferential consequence.

There has by now been a number of attempts to incorporate defeasible reasoning in logics other than propositional logic. After a first tentative exploration of preferential predicate logics by Lehmann and Magidor [101] and of a preference semantics for defaults in description logics by Quantz and Royer [122], some more recent investigations have attempted to define notions of defeasibility in deontic logics [117], and of defeasible subsumption relations in DLs [31, 46, 76]. One of the main obstacles in moving beyond the propositional setting is the lack of a generally accepted formal semantics which appropriately generalizes the propositional structures of Kraus et al.

The aforementioned extensions of the KLM approach have mostly been driven by either extending only the syntax [46, 76, 122] or just the underlying preferential semantics [31, 101]. An intuitive semantics, with a corresponding representation result, was still missing until my collaborators and myself provided a semantic foundation for extending defeasible consequence relations to modal logic [34, 36].

Our first contribution was the definition of a preferential semantics by enriching Kripke models with preference relations on possible worlds [34]. Such a semantics characterises appropriately the class of preferential modal defeasible conditionals. In particular, we have shown a representation result *à la* KLM linking the semantic constructions to a set of properties that basic modal conditionals deemed appropriate in a non-monotonic setting ought to satisfy. Besides that, we have also lifted the notion of rational closure as defined by Lehmann and Magidor in the propositional case [102] to modal logics, thereby providing a first account of this construction in formalisms with more structure than the propositional one.

In a spin-off paper [36], we investigated how the set of KLM-style properties can be extended in order to make use of the extra expressive power that a modal language provides. We have introduced one such a property, defined the corresponding semantic constraint to be imposed on preferential Kripke models, and shown new representation results linking the extended set of properties with the new semantics.

Significance of the contribution Besides providing the first account of KLM-style defeasible reasoning for modal logics, these results also paved the way for further exploration of different forms of defeasibility in modal and description logics, in particular at the object level (see Section 3.3.2). The publications resulting from this contribution have received 30+ citations according to Google Scholar.

3.3.2 Beyond defeasibility of conditionals

Defeasible reasoning, as traditionally studied in the literature on non-monotonic reasoning, has focused mostly on two aspects of defeasibility that we encountered in Chapter 2, namely exceptionality and information withdrawal. In particular, exceptionality, i.e., defeasibility at the object level, usually takes the form of an ‘argument’: in a statement (or a ‘conditional’, as it is also referred to) of the form $\alpha \sim \beta$, α plays the role of a ‘premise’, while β is its (defeasible) ‘conclusion’. Such is the case in the KLM approach [96, 102] and related frameworks [23, 51].

Nevertheless, there are many other appealing and equally useful aspects of defeasibility besides that of arguments or conditionals. These include notions such as typicality [19, 20, 23, 54, 76, 100], concerned with the most typical cases or situations (or even the most typical representatives of a class), degrees of belief [5, 6], relating to the most plausible epistemic possibilities held by an agent, vague notions [138] such as ‘generally’, ‘rarely’ or ‘most’, and relative normality [37, 38], amongst others. For instance, one may want to formalise the expected effect of an action, which may still fail; or plausible beliefs, which might be wrong; or certain forms of obligation that may be overridden.

It turns out that with KLM-style defeasible statements one cannot capture these aspects of defeasibility. This has to do partly with the syntactic restrictions imposed on \sim , namely no nesting of conditionals, but more fundamentally, it relates to where and how the notion of normality is used in such statements. The semantics of a conditional $\alpha \sim \beta$ says, roughly, that “the *normal* α -situations are β -situations”.¹ This sounds like the normality spotlight is somewhat put on α , as though normality was a property of the premise and not necessarily of the conclusion. Whether the β -situations alone are normal or not plays no role in the reasoning that is carried out. Furthermore, normality is assumed to be a property relating to the premise as a whole, and not of its sub-formulae. Technically this means one cannot refer directly to normality of a sentence in the scope of other logical operators. (This is also the case in all the aforementioned extensions of the KLM approach to logics that are more expressive than the propositional one.)

¹Friedman and Halpern [67] propose the alternative reading “the $\alpha \wedge \beta$ -situations are more normal (or more plausible) than the $\alpha \wedge \neg\beta$ -situations”.

The ability to capture forms of defeasibility beyond that of arguments is another fundamental challenge in the definition of a comprehensive theory of defeasible reasoning. In a recent work [37, 38] we investigated notions that we refer to as defeasible *modes* of inference. These amount to preferential versions of the traditional notions of knowledge, obligations and actions as studied in classical modal logics. For instance, in an action context, one may want to say that the outcome of a given action a is usually α [98], i.e., in the most normal situations resulting from a 's execution, α holds. This is notably different from saying that in the most normal worlds, the result of performing action a is *always* α . In the latter case, the normality of the situation, or state, before the action takes place is assessed, whereas in the former the relative normality of the situation is assessed against all possible outcomes [38]. Here we are interested in the formalisation of the former type of statement, where it becomes important to shift the notion of normality from the premise of an inference to the effect of an action, and, importantly, use it in the scope of other logical constructors.

The importance of defeasibility in modes of reasoning is also illustrated by the following example [37]. Although one may envisage a situation where the velocity of a sub-atomic particle in a vacuum is greater than c (the speed of light in a vacuum), it is in a sense known that c is the highest possible speed. We are then entitled to derive factual consequences of this scientific theory that also will be ‘known’. This venturesome version of knowledge, which patently differs from belief, provides for a more fine-grained notion of knowledge that may turn out to be wrong but that is not of the same nature as suppositions or beliefs.²

Scenarios such as the ones depicted above require an ability to talk about the normality of effects of an action, normality of knowledge or obligations, and so on. While existing modal treatments of epistemic reasoning and of preferential reasoning can express preferential semantics syntactically as modalities [23, 32, 76], they do not suffice to express defeasible modes of inference as described above.

Our contribution filled this gap by introducing non-standard modalities encapsulating a notion of defeasibility which allows us to talk about relative normality, i.e., about what is normal or expected relative to a given situation (or possible configuration of the world). These operators were defined within the extension of our preferential modal semantics [34, 36]. Given a possible world in a Kripke model, the normality (or exceptionality) of its accessible worlds are determined by a preference order on possible worlds, which unlike in the plausibility models of Baltag and Smets [5, 6], does not define an agent’s knowledge or beliefs. Rather, our notion of preference is part of the semantics of the background ontology described by the theory or knowledge base at

²Our proposal is not aimed at challenging the position of knowledge as indefeasible, justified true belief [72, 103], but rather provides an extension to epistemic modal logics to allow for reasoning with a modality that we shall, arguably for lack of a more suitable term, refer to as “defeasible knowledge”.

hand. As such, it is used in providing a meaning for defeasible actions, which can fail in their outcome, or defeasible knowledge, which may not hold in exceptional accessible worlds, in that the preference relation alters the classical semantics of these modalities. This allows for the definition of a formalism in which defeasible modes of inference can be expressed, and which can be integrated with existing non-monotonic modal logics.

Significance of the contribution For a long time defeasibility had been confined to conditionals. This contribution is significant in that it explored defeasibility also in other logical connectives and established a general semantic framework within which they can be studied and further extended. These allow for the exploitation of several nuances of quotidian reasoning at the object level. The main publications resulting from this contribution have received 60+ citations altogether according to Google Scholar.

3.4 Rationality in formal ontologies

For the past decades, there has been an increasing interest in formal ontologies and their role in the design of ‘smart’ applications. That has been the case not only in AI but also in other areas of computer science such as databases and semantic technologies. Description Logics (DLs) provide us with solid foundations on which to base the study and formalisation of ontologies as well as on reasoning and querying methods associated to them.

As it turns out, DLs still allow for meaningful, decidable extensions, as new knowledge representation requirements are identified. A case in point is the need to allow for exceptions and defeasibility in reasoning over logic-based ontologies. The different DL-based formalisms that have been proposed in the literature provide us with a wide choice of constructors in the object language. Nevertheless, these are intended to express only classical, unquestionable knowledge, and, as a consequence, they fall short of adequately capturing the different aspects of uncertainty and vagueness that often show up in everyday life. Examples of these comprise the various guises of exceptions, typicality (and atypicality), approximations and many others, as usually encountered in the different forms of human quotidian reasoning. A similar argument can be put forward when moving to the level of entailment, i.e., that of the sanctioned conclusions from an ontology. DL systems provide for a variety of standard and non-standard reasoning services, but the underlying notion of logical consequence remains classical and therefore, depending on the application one has in mind, DLs inherit most of the criticisms raised in the development of the so-called non-classical logics.

In this regard, endowing DLs and their associated reasoning services with the ability to cope with defeasibility is a natural step in their development. Even if some of the issues related to uncertainty in reasoning have been studied using probabilistic

approaches and statistical methods, their qualitative computational nature remains a large avenue for exploration.

The past two decades have witnessed the surge of some attempts to introduce non-monotonic reasoning capabilities in a DL setting, ranging from preferential approaches to circumscription-based ones, amongst others. Each of them investigated particular constructions and DL variants of existing approaches to defeasibility in propositional logic. Nevertheless, a *comprehensive* study of the formal foundations of (preferential) defeasible reasoning in DLs was still missing until we published a series of papers at logic and AI conferences [28, 29, 35, 39, 40, 41, 43] and journals [42, 137], in which we have filled precisely this gap.

Roughly, in these papers we have (i) defined a general and intuitive *semantics*; (ii) shown that the corresponding *representation results* linking our semantic constructions with a set of desirable properties hold, and (iii) presented an appropriate analysis of *entailment* in the context of ontologies with defeasible information with associated decision procedures that are implementable. Furthermore, we have also made a case for a number of additional defeasible constructs at the object level enriching the underlying DL concept language, to which we have given a corresponding intuitive semantics and devised associated reasoning methods. We have shown that these extensions do not negatively affect decidability or complexity of reasoning for an important class of DLs, and that many nuances of defeasible reasoning in natural language can be expressed in terms of our new constructs.

In what follows, I summarise the specific contributions to this area and their respective significance. More details, including the technical constructions, can be found in the companion document titled *Defeasible Description Logics*.

3.4.1 Defeasible subsumption and representation results

My first technical contribution to the development of defeasible description logics was the definition of a simple and intuitive semantics for the notion of defeasible subsumption [35]. In particular, together with my peers, we have provided a characterisation of two important classes of defeasible subsumption relations, namely preferential and rational subsumption, via the respective representation results, evidencing the fact that our semantic constructions are appropriate. In a follow-up publication [29], we revisited our semantics by making it simpler and more intuitive. The corresponding representation results carry over to the new semantics and their proofs are available in an article to appear in the ACM Transactions on Computational Logic [26].

Significance of the contribution These results are significant in that they show our basic framework is general enough to serve as a hotbed for (i) the definition of

the notion of *rational closure* for DLs (see below), and (ii) the introduction of further non-monotonic extensions to DLs, as witnessed by the results in the sections below. The main publications resulting from this contribution have received 100+ citations altogether according to Google Scholar.

3.4.2 Rational closure for defeasible DLs

Once the semantic foundations of defeasibility in DLs have been established, we carried out a thorough investigation of what an appropriate notion of defeasible entailment in a defeasible DL context means. In particular, we have analysed the suitability of a good candidate, namely rational entailment. Moreover, we have established the formal connection between rational entailment, the notion of rational closure of a defeasible knowledge base, and we have provided an algorithm for its computation. The proofs establishing the correctness of our approach are available in an article to appear in the ACM Transactions on Computational Logic [26].

Significance of the contribution These results are important in that they stand at the foundation of defeasibility at the *entailment level* and can be extended to more expressive formalisms, as witnessed by the contribution reported in Section 3.4.4. Furthermore, the main advantages of our approach are as follows: (i) it relies completely on classical entailment, i.e., entailment checking over defeasible ontologies can be reduced to a number of classical entailment checks over a rewritten ontology, (ii) it has computational complexity that is no worse than that of entailment checking in the classical underlying DL, and (iii) it is amenable to implementation, e.g. as a Protégé plugin³. Indeed, the referred implementation by a Ph.D. student I have co-supervised allows us to make experiments with large ontologies containing defeasible subsumption statements and its performance has been shown to scale well in practice [45]. The main publications resulting from this contribution have received 90+ citations altogether according to Google Scholar.

3.4.3 Typicality of concepts and relations at the object level

In another journal publication as the single author [137], I have introduced \mathcal{ALCH}^\bullet , a description logic allowing for an explicit notion of typicality that can be applied to both concepts and roles and of which the intuition is to capture the most typical instances of, respectively, classes and relations. In particular, I showed that reasoning w.r.t. \mathcal{ALCH}^\bullet knowledge bases is decidable through the definition of a tableau-based

³<https://github.com/kodymoodley/defeasibleinferenceplatform>

decision procedure that I have shown to be sound and complete w.r.t. the underlying preferential semantics.

When compared with other existing approaches to non-monotonicity in DLs, the novelty of \mathcal{ALCH}^\bullet resides in the provision of a framework for typicality of both classes and relations and that can serve as the foundation for further extensions of defeasible DLs of increasing expressivity, with non-monotonicity at the level of concepts as well as that of roles.

Significance of the contribution The significance of this contribution is supported by the fact that it was the recipient of the first Louis Couturat Logic Prize (France, 2018). It was then presented at the Universal Logic Contest at UNILOG 2018 in Vichy (<https://www.uni-log.org/vichy2018>) and subsequently won the first Universal Logic Prize.

3.4.4 Contextual defeasibility

Next, I have also carried out an investigation of defeasibility beyond concept subsumption along two main lines: First by introducing defeasible versions of value and existential restrictions [28, 39], thereby bringing defeasibility also to the object level; and second by introducing a (primitive) role-informed notion of context, which gives rise to multiple orderings on objects and the ability to model different, possibly incompatible, contexts of normality, making typicality a relativised construct [41].

Technically, this contribution addresses an important limitation in previous defeasible extensions of description logics, namely the restriction in the semantics of defeasible concept inclusion to a single preference order on objects. Semantically, it also answers the question of the meaning of multiple preference orders, namely that they reflect different contexts. Moreover, I have defined a tableau-based algorithm for checking preferential consistency of contextual defeasible knowledge bases [43], and I showed its soundness and completeness w.r.t. the underlying preferential semantics. These results are significant in that they are a central piece in the definition of other forms of contextual defeasible reasoning over ontologies.

Since the notion of entailment assumed in the above contribution is monotonic, there was a need to lift the definition of rational closure to a notion of *contextual rational closure*. We have achieved that by providing a semantic construction and a method for the computation of such a notion of closure, along with a correspondence result between the two [42].

Significance of the contribution These results are significant because they provide a more fine-grained version of the constructions reported in Section 3.4.2 and moreover

they have the potential to be applied to a wide class of real-world ontologies in which context plays an important role in modelling the knowledge.

Finally, I have also defined a meaningful extension of the highly expressive description logic \mathcal{SROIQ} with defeasible reasoning constructs in the concept language, in both concept and role inclusions, and in role assertions, together with an intuitive preferential semantics [40]. In particular, I have shown a translation of the entailment problem w.r.t. $d\mathcal{SROIQ}$ knowledge bases to concept satisfiability relative to DTBoxes and RBoxes only. I have also devised a terminating, sound and complete tableau-based algorithm for checking concept satisfiability w.r.t. $d\mathcal{SROIQ}$ knowledge bases. This logic can serve as the theoretical foundation for a defeasible version of the Web Ontology Language OWL.⁴

⁴<https://www.w3.org/OWL>

Chapter 4

Perspectives and future directions of investigation

The study of the theoretical foundations of uncertainty in reasoning is of paramount importance for both the understanding of human rationality and the design of fully-automated decision-making systems. In spite of all the progress obtained by the non-monotonic reasoning community over the past 40 years, including the work done by myself during the last 12 years, there still remain several open questions in the area and technical challenges to be addressed.

In particular, there are still many relevant issues across the three levels at which non-monotonicity shows up in reasoning (see Chapter 2) that remain to be more deeply studied and formalised, or even integrated with emerging non-symbolic approaches in AI. Indeed, specifications in defeasible knowledge bases can be learned from commonsense knowledge repositories via machine learning and help knowledge engineers to fine tune the formalisation of a given domain of application. Therefore, newer and even more general formal theories of non-monotonic reasoning, supplementing the existing ones in many crucial aspects, are still called for.

In that respect, for the upcoming five years I plan to pursue further investigations into the nature of non-monotonic reasoning and ways in which its two distinct aspects can be dealt with computationally at *all* its three different levels. In order to achieve the stated goal, I intend to bring to fruition the following specific objectives:

- **Objective 1:** At the **object level**, define new non-monotonic logical operators beyond standard non-monotonic implications or conditionals, endowed with an intuitive formal semantics together with a corresponding study of their formal properties and potential practical usefulness, in particular in capturing the various nuances of natural language. This point might benefit from e.g. a machine-learning component helping in identifying such constructs and their various syn-

onyms as they are usually used in making utterances. Furthermore, as alluded to above, the specification of knowledge bases, in particular of ontologies, a complex task that in the standard case requires many iterations and revisions by domain experts, can be made more efficient and less error prone by deploying tailor-made machine learning tools.

- **Objective 2:** At the **entailment level**, define new non-monotonic consequence relations that are neither limited to propositional languages nor to classical connectives, equipped with both an intuitive semantic characterisation and accompanying proof methods. Given that the formalisms I will take as point of departure, namely modal and description logics, are PSPACE-complete, an integration with the latest developments around the satisfiability problem (SAT) is envisaged in order to improve performance of the associated reasoning tools.
- **Objective 3:** At the **meta-reasoning level**, devise a comprehensive family of belief revision operators to logics of various expressivity and with a possibly non-monotonic underlying consequence relation.¹ This is admittedly the most ambitious part of the whole project. In a recent paper published at the KR conference [44], we have made the first steps towards this direction.

To bring about Objectives 1–3 set out above, I will pursue the activities motivated and laid out below. (The research tasks that are mentioned do not, by any means, constitute an exhaustive list of the activities that will be carried out. Rather, they are mentioned here because all of them are, to some degree, an important component of the proposed research. For the sake of presentation, technical details shall be omitted for now, especially since most if not all of them still remain to be worked out.)

- At the **object level**: I will investigate aspects of non-monotonicity in logical operators other than implication, both in propositional as well as in more expressive logics. For instance, I plan to investigate ampliative and defeasible versions of the propositional connectives (negation, conjunction and disjunction), of modalities (necessity and possibility), and of quantifiers in general first-order logic and fragments thereof, giving rise to exception-tolerant versions of these connectives. I will also examine the extent to which there is room for new defeasible connectives having no classical counterpart [19, 20].
- At the **entailment level**: I intend to explore new types of non-monotonic entailment relations, both for propositional and for less or more expressive logical

¹For now I will not address abstract argumentation frameworks, but will leave open for argumentation systems based on the above mentioned new non-monotonic consequence relations to be defined.

languages. For example, I will investigate different propositional supra-classical non-monotonic entailment relations, and assess their scalability beyond the propositional setting, namely in modal and predicate languages. In particular, I want to study how to define the respective proof methods for each of these consequence relations and how these can benefit from latest results from the SAT community. Finally, I will also explore non-monotonic consequence relations for the logical languages enriched with non-monotonic connectives as those mentioned above.

- At the **meta-reasoning level**: I will investigate the definition of new approaches and construction methods for belief revision in *(i)* logics that are more expressive than the propositional one, notably the different systems of modal logic and fragments of first-order logic; *(ii)* logics enriched with non-monotonic connectives but with an underlying classical entailment relation; *(iii)* logics with a non-Tarskian consequence relation, and finally *(iv)* logics with both non-monotonic connectives and an underlying non-monotonic entailment relation.

In order for me to realise the foregoing tasks, I will use formal methods to *(i)* design and specify the new logical languages, their underlying consequence relations and corresponding revision methods; *(ii)* endow them with an appropriate formal (unambiguous) semantics; *(iii)* design algorithms for the corresponding proof methods, and *(iv)* give mathematical proofs of the consistency of each formalism and correctness of the corresponding algorithms. Moreover, I will also analyse the computational complexity of the entailment problem for each new formalism. Whenever appropriate and scientifically significant, implementation of the results will be realised and tested in real-world applications.

It is worth pointing out that my agenda is not to diverge from the existing body of work in non-monotonic reasoning, but rather to build on it when striving for the stated objectives. The answer to the question of which existing approaches to non-monotonicity are more appropriate to serve as the substratum for the investigations that I have in mind is not cast in stone. A promising starting point is to consider those non-monotonic formalisms that are general enough in the propositional case, the most successful of which being the preferential approach [96, 128] and the AGM paradigm [2, 69, 81]. The main reason for doing so is that such formalisms satisfy well-established abstract properties which should also be present in general theories of uncertainty in reasoning.

The question as to which logical framework should be the backbone of the new representation languages and reasoning methods will mostly be guided by current trends in logic-based knowledge representation and reasoning [4, 83]. This means that we shall strive for a good trade-off between expressive power on the one hand and computational tractability on the other [24]. The prototypical logical formalisms usually considered

are classical propositional and first-order logic. It turns out that propositional logic is not expressive enough for the needs of modern KR&R, whereas (full) first-order logic is too expressive and has bad computational properties, the worst of which being semidecidability of the entailment problem. In this regard, from a KR&R perspective, the choice for modal logics and description logics as the core representation languages on which to initially base our general framework sounds appealing. These are elegant and well-established logical frameworks with a good balance between expressivity and computational complexity. This will provide the exploratory phase of our research with a strong initial boost, and will ensure that tangible results are obtained much sooner than would otherwise have been the case.

An important aspect of the conceptual framework to be deployed here is the fact that all tasks specified above will be carried out with an eye on the potential users of the new logical formalisms. This means that we shall design our new representation languages and reasoning methods in accordance with commonly adopted principles of software ergonomics [120]; in other words, besides being significant and mathematically sound, the formalisms I intend to develop should be neat, elegant and amenable to implementation in order to be applied in practice.

It is legitimate to raise the question on whether the research here described could be based on probabilities instead of logic. Indeed probabilistic reasoning [119] is also non-monotonic in that any inference of probability less than 1 can in principle fail (defeasible aspect). There are results showing that non-monotonic reasoning can be given a probability-based account. Examples of these are Pearl’s system Z [118], Hawthorne’s conditionals [84], and the possibilistic approach [10], which give essentially the same results as the (non-probabilistic) preferential approach [96]. In that respect, one could in principle base our investigations on a probabilistic setting. The requirement for that would be the availability of appropriate probabilistic (and possibilistic) treatments of the three levels we have encountered above on which to initially base our work. At the moment it is not clear whether all this machinery is available and therefore the logic-based route seems to be more viable. Since probabilistic and logical approaches to reasoning are not antagonistic but rather complementary, possible combinations of both is certainly a point that will be investigated in more depth during the exploratory phase of this research.

Finally, it is worth emphasising that my goal is not to propose yet another family of disconnected formalisms for non-monotonic reasoning, each one dealing with a single aspect or level of non-monotonicity. Instead, I aim at developing a unifying comprehensive approach to all different aspects and levels of non-monotonicity, benefitting from existing results in the literature that show the links among frameworks sitting at different levels. For instance, there is a close connection between classical knowledge base revision [92] and KLM-style rational consequence relations [102] in the sense that

each one can be defined in terms of the other [70]. A similar result established the link between KLM rational consequence and belief revision in PTL (see Section 3.2.3). It is reasonable to expect results like these to provide the basic springboard with which to develop the general overarching approach I have in mind.

Once all of this is in place, further research questions will arise. From a knowledge representation and reasoning perspective, when one deals with knowledge bases, issues related to modularisation [52, 71, 88, 89], consistency checking [86, 87, 97, 139], knowledge base integration [112], maintenance and repair [85, 135, 136] as well as versioning [66, 116] show up. These are tasks acknowledged as important by the community in the classical case [90, 93, 113, 130, 134] and that also make sense in a non-monotonic setting, such as the one here envisaged. Nevertheless, when moving to a defeasible approach, such tasks have to be reassessed and specific methods and techniques re-designed. This constitutes an extra challenge in my research endeavour and definitely an avenue worthy of exploration.

The innovation of the present research project relies mostly in the investigation of nuances of reasoning hitherto largely unexplored and in the quest for a comprehensive framework for multifarious non-monotonic reasoning. By pushing the frontiers of formal logic in the ways described above, I hope to advance the body of knowledge in several disciplines and at the same time provide more solid foundations for the realisation of a new variety of relevant real-world computer-based applications. The following are good examples of the latter:

- Application in **artificial intelligence**: the definition of several forms of consequence relations for a variety of languages together with more powerful methods for belief revision will allow for the design of more fine-grained exception-tolerant intelligent systems capable of adapting themselves to an ever-changing environment, whether real or virtual.
- Application in **computer science** and **software engineering**: the definition of non-monotonic constructs in the object language will allow for the design of methods for program verification and software specification capable of dealing with foreseeable exceptions.
- Application in **diagnosis**: the formalisation of several ampliative consequence relations will provide the basis for the generation of explanations in failure diagnosis based on the formal specification of the expected behaviour of a given device or system or process.
- Application in **reasoning about actions** and **planning**: the definition of constructs capturing the notions of defeasible actions, defeasible plans and retractable goals will allow for the design of systems that can recover when facing ‘hidden’ preconditions preventing an action or a whole plan from succeeding.

- Application in **general game playing**: an analysis of defeasibility at the three different levels will uplift the conception of robust intelligent systems. It will allow for the deployment of agents that can be reprogrammed to play new games and that do not collapse when exceptions to the rules are met or when other players bend the rules.
- Application in **information systems** and the **semantic web**: by enriching description logics with multifarious non-monotonic reasoning capabilities, we will provide for more accurate representations of exceptions in formal ontologies together with methods to better deal with the dynamics of ontologies.
- Other applications: the outcomes of this project will also meet potential applications in **machine vision** (default conclusions about the rest of a scene one cannot see) and in **natural language understanding** (default conclusions about presuppositions, or how a given sentence or story or dialogue will continue). Of course, this will have to be complemented by current advances in learning-based AI.

Non-monotonic reasoning as I intend to investigate here may prove useful not only in real-world applications like the above mentioned ones, but also in other theoretical domains beyond artificial intelligence and mathematical logic. These comprise, for instance, cognitive science and philosophical logic, since I will analyse and formalise linguistic constructs and reasoning patterns not yet fully explored.

Bibliography

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [2] C. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [3] O. Arieli and A. Avron. General patterns for nonmonotonic reasoning: From basic entailments to plausible relations. *Logic Journal of the IGPL*, 8:119–148, 2000.
- [4] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2 edition, 2007.
- [5] A. Baltag and S. Smets. Dynamic belief revision over multi-agent plausibility models. In W. van der Hoek and M. Wooldridge, editors, *Proceedings of LOFT*, pages 11–24. University of Liverpool, 2006.
- [6] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT’08)*, number 3 in Texts in Logic and Games, pages 13–60. Amsterdam University Press, 2008.
- [7] D. Barker-Plummer, J. Barwise, and J. Etchemendy. *Language, Proof and Logic*. University of Chicago Press, 2 edition, 2011.
- [8] D. Batens. *Adaptive Logics and Dynamic Proofs*. Forthcoming, 2011.
- [9] R. Baumann, G. Brewka, W. Dvorák, and S. Woltran. Parameterized splitting: A simple modification-based approach. In E. Erdem, J. Lee, Y. Lierler, and D. Pearce, editors, *Correct Reasoning — Essays on Logic-Based AI in Honour of Vladimir Lifschitz*, LNCS, pages 57–71. Springer, 2012.
- [10] S. Benferhat, D. Dubois, and H. Prade. Possibilistic and standard probabilistic semantics of conditional knowledge bases. *Journal of Logic and Computation*, 9(6):873–895, 1999.
- [11] J. van Benthem. *Modal Logic for Open Minds*. Center for the Study of Language and Information, 2010.
- [12] P. Besnard and A. Hunter. *Elements of Argumentation*. University of Michigan Press, 2008.
- [13] P. Bisquert, C. Cayrol, F. Dupin de Saint-Cyr, and M.-C. Lagasquie-Schiex. Duality between addition and removal: A tool for studying change in argumentation. In S. Greco, B. Bouchon-Meunier, G. Coletti, M. Fedrizzi, B. Matarazzo, and R.R. Yager, editors, *Advances on Computational Intelligence — 14th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, pages 219–229. Springer, 2012.

- [14] P. Blackburn, J. van Benthem, and F. Wolter. *Handbook of Modal Logic*. Elsevier North-Holland, 2006.
- [15] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.
- [16] R. Booth, G. Casini, T. Meyer, and I. Varzinczak. On the entailment problem for a logic of typicality. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2805–2811, 2015.
- [17] R. Booth, G. Casini, T. Meyer, and I. Varzinczak. On rational entailment for propositional typicality logic. *Artificial Intelligence*, 277, 2019.
- [18] R. Booth, T. Meyer, and I. Varzinczak. Next steps in propositional Horn contraction. In C. Boutilier, editor, *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 702–707. AAAI Press, 2009.
- [19] R. Booth, T. Meyer, and I. Varzinczak. PTL: A propositional typicality logic. In L. Fariñas del Cerro, A. Herzig, and J. Mengin, editors, *Proceedings of the 13th European Conference on Logics in Artificial Intelligence (JELIA)*, number 7519 in LNCS, pages 107–119. Springer, 2012.
- [20] R. Booth, T. Meyer, and I. Varzinczak. A propositional typicality logic for extending rational consequence. In E.L. Fermé, D.M. Gabbay, and G.R. Simari, editors, *Trends in Belief Revision and Argumentation Dynamics*, volume 48 of *Studies in Logic – Logic and Cognitive Systems*, pages 123–154. King’s College Publications, 2013.
- [21] R. Booth, T. Meyer, I. Varzinczak, and R. Wassermann. A contraction core for Horn belief change: Preliminary report. In *Proceedings of the 13th International Workshop on Nonmonotonic Reasoning (NMR)*, 2010.
- [22] R. Booth, T. Meyer, I. Varzinczak, and R. Wassermann. On the link between partial meet, kernel, and infra contraction and its application to horn logic. *Journal of Artificial Intelligence Research (JAIR)*, 42:31–53, 2011.
- [23] C. Boutilier. Conditional logics of normality: A modal approach. *Artificial Intelligence*, 68(1):87–154, 1994.
- [24] R. Brachman and H. Levesque. *Knowledge Representation and Reasoning*. Morgan Kaufmann, 2004.
- [25] G. Brewka, V.W. Marek, and M. Truszczynski, editors. *Nonmonotonic Reasoning: Essays Celebrating its 30th Anniversary*. King’s College Publications, 2011.
- [26] K. Britz, G. Casini, T. Meyer, K. Moodley, U. Sattler, and I. Varzinczak. Principles of klm-style defeasible description logics. *ACM Transactions on Computational Logic*, 22(1), 2020.
- [27] K. Britz, G. Casini, T. Meyer, K. Moodley, and I. Varzinczak. Ordered interpretations and entailment for defeasible description logics. Technical report, CAIR, CSIR Meraka and UKZN, South Africa, 2013.
- [28] K. Britz, G. Casini, T. Meyer, and I. Varzinczak. Preferential role restrictions. In *Proceedings of the 26th International Workshop on Description Logics*, pages 93–106, 2013.

- [29] K. Britz, G. Casini, T. Meyer, and I. Varzinczak. A KLM perspective on defeasible reasoning for description logics. In C. Lutz, U. Sattler, C. Tinelli, A.-Y. Turhan, and F. Wolter, editors, *Description Logic, Theory Combination, and All That. Essays Dedicated to Franz Baader on the Occasion of His 60th Birthday*, volume 11560, pages 147–173. Springer, 2019.
- [30] K. Britz, J. Heidema, and W. Labuschagne. Entailment, duality, and the forms of reasoning. Technical Report OUCS-2007-01, Department of Computer Science, University of Otago, 2007.
- [31] K. Britz, J. Heidema, and T. Meyer. Semantic preferential subsumption. In J. Lang and G. Brewka, editors, *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 476–484. AAAI Press/MIT Press, 2008.
- [32] K. Britz, J. Heidema, and T. Meyer. Modelling object typicality in description logics. In A. Nicholson and X. Li, editors, *Proceedings of the 22nd Australasian Joint Conference on Artificial Intelligence*, number 5866 in LNAI, pages 506–516. Springer, 2009.
- [33] K. Britz, J. Heidema, and I. Varzinczak. Constrained consequence. *Logica Universalis*, 5(2):327–350, 2011.
- [34] K. Britz, T. Meyer, and I. Varzinczak. Preferential reasoning for modal logics. *Electronic Notes in Theoretical Computer Science*, 278:55–69, 2011. Proceedings of the 7th Workshop on Methods for Modalities (M4M’2011).
- [35] K. Britz, T. Meyer, and I. Varzinczak. Semantic foundation for preferential description logics. In D. Wang and M. Reynolds, editors, *Proceedings of the 24th Australasian Joint Conference on Artificial Intelligence*, number 7106 in LNAI, pages 491–500. Springer, 2011.
- [36] K. Britz, T. Meyer, and I. Varzinczak. Normal modal preferential consequence. In M. Thielscher and D. Zhang, editors, *Proceedings of the 25th Australasian Joint Conference on Artificial Intelligence*, number 7691 in LNAI, pages 505–516. Springer, 2012.
- [37] K. Britz and I. Varzinczak. Defeasible modes of inference: A preferential perspective. In *Proceedings of the 14th International Workshop on Nonmonotonic Reasoning (NMR)*, 2012.
- [38] K. Britz and I. Varzinczak. Defeasible modalities. In *Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 49–60, 2013.
- [39] K. Britz and I. Varzinczak. Introducing role defeasibility in description logics. In L. Michael and A.C. Kakas, editors, *Proceedings of the 15th European Conference on Logics in Artificial Intelligence (JELIA)*, number 10021 in LNCS, pages 174–189. Springer, 2016.
- [40] K. Britz and I. Varzinczak. Toward defeasible *SROIQ*. In *Proceedings of the 30th International Workshop on Description Logics*, 2017.
- [41] K. Britz and I. Varzinczak. Rationality and context in defeasible subsumption. In F. Ferrarotti and S. Woltran, editors, *Proceedings of the 10th International Symposium on Foundations of Information and Knowledge Systems (FoIKS)*, number 10833 in LNCS, pages 114–132. Springer, 2018.
- [42] K. Britz and I. Varzinczak. Contextual rational closure for defeasible *ALC*. *Annals of Mathematics and Artificial Intelligence*, 2019.
- [43] K. Britz and I. Varzinczak. Preferential tableaux for contextual defeasible *ALC*. In S. Cerrito and A. Popescu, editors, *Proceedings of the 28th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods (TABLEAUX)*, number 11714 in LNCS. Springer, 2019.

- [44] G. Casini, E. Fermé, T. Meyer, and I. Varzinczak. A semantic perspective on belief change in a preferential non-monotonic framework. In M. Thielscher and F. Toni, editors, *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 220–229. AAAI, 2018.
- [45] G. Casini, T. Meyer, K. Moodley, U. Sattler, and I. Varzinczak. Introducing defeasibility into OWL ontologies. In M. Arenas, O. Corcho, E. Simperl, M. Strohmaier, M. d’Aquin, K. Srinivas, P.T. Groth, M. Dumontier, J. Heflin, K. Thirunarayan, and S. Staab, editors, *Proceedings of the 14th International Semantic Web Conference (ISWC)*, number 9367 in LNCS, pages 409–426. Springer, 2015.
- [46] G. Casini and U. Straccia. Rational closure for defeasible description logics. In T. Janhunen and I. Niemelä, editors, *Proceedings of the 12th European Conference on Logics in Artificial Intelligence (JELIA)*, number 6341 in LNCS, pages 77–90. Springer-Verlag, 2010.
- [47] M.A. Castilho, A. Herzig, and I. Varzinczak. It depends on the context! A decidable logic of actions and plans based on a ternary dependence relation. In *Proceedings of the 9th International Workshop on Nonmonotonic Reasoning (NMR)*, 2002.
- [48] C. Cayrol, F. Dupin de Saint-Cyr, and M.-C. Lagasquie-Schiex. Revision of an argumentation system. In J. Lang and G. Brewka, editors, *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 124–134. AAAI Press/MIT Press, 2008.
- [49] B. Chellas. *Modal logic: An introduction*. Cambridge University Press, 1980.
- [50] K. Clark. Negation as failure. In H. Gallaire and J. Minker, editors, *Logics and Databases*, pages 293–322. Plenum Press, 1978.
- [51] G. Crocco and P. Lamarre. On the connections between nonmonotonic inference systems and conditional logics. In R. Nebel, C. Rich, and W. Swartout, editors, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 565–571. Morgan Kaufmann Publishers, 1992.
- [52] B. Cuenca Grau, B. Parsia, E. Sirin, and A. Kalyanpur. Modularity and web ontologies. In P. Doherty, J. Mylopoulos, and C. Welty, editors, *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 198–208. Morgan Kaufmann, 2006.
- [53] M. Dalal. Investigations into a theory of knowledge base revision: preliminary report. In R. Smith and T. Mitchell, editors, *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI)*, pages 475–479. Morgan Kaufmann Publishers, 1988.
- [54] J. Delgrande. A first-order logic for prototypical properties. *Artificial Intelligence*, 33:105–130, 1987.
- [55] J. Delgrande. An approach to default reasoning based on a first-order conditional logic: Revised report. *Artificial Intelligence*, 36:63–90, 1988.
- [56] J. Delgrande. Horn clause belief change: Contraction functions. In J. Lang and G. Brewka, editors, *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 156–165. AAAI Press/MIT Press, 2008.
- [57] R. Demolombe, A. Herzig, and I. Varzinczak. Regression in modal logic. *Journal of Applied Non-Classical Logics*, 13(2):165–185, 2003.

- [58] H. van Ditmarsch, W. van der Hoek, and B.P. Kooi. *Dynamic Epistemic Logic*. Springer, 2007.
- [59] P.M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and nDperson games. *Artificial Intelligence*, 77(2):321–357, 1995.
- [60] J.M. Dunn and G. Restall. Relevance logic. In D.M. Gabbay and F. Günthner, editors, *Handbook of Philosophical Logic*, volume 6, pages 1–128. Kluwer Academic Publishers, 2002.
- [61] T. Eiter, E. Erdem, M. Fink, and J. Senko. Updating action domain descriptions. In L. Kaelbling and A. Saffiotti, editors, *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 418–423. Morgan Kaufmann Publishers, 2005.
- [62] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
- [63] E. Fermé and S.O. Hansson. AGM 25 years. *Journal of Philosophical Logic*, 40(2):295–331, 2011.
- [64] P. Flach and A. Kakas, editors. *Abduction and induction: Essays on their relation and integration*. Kluwer Academic Press, 2000.
- [65] G. Flouris, D. Plexousakis, and G. Antoniou. Generalizing the AGM postulates. In *Proceedings of the 10th International Workshop on Nonmonotonic Reasoning (NMR)*, 2004.
- [66] E. Franconi, T. Meyer, and I. Varzinczak. Semantic diff as the basis for knowledge base versioning. In *Proceedings of the 13th International Workshop on Nonmonotonic Reasoning (NMR)*, 2010.
- [67] N. Friedman and J.Y. Halpern. Plausibility measures and default reasoning. *Journal of the ACM*, 48(4):648–685, 2001.
- [68] D.M. Gabbay and K. Schlechta. Roadmap for preferential logics. *Journal of Applied Non-Classical Logic*, 19(1):43–95, 2009.
- [69] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, 1988.
- [70] P. Gärdenfors and D. Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65(2):197–245, 1994.
- [71] J. Garson. Modularity and relevant logic. *Notre Dame Journal of Formal Logic*, 30(2):207–223, 1989.
- [72] E. Gettier. Is justified true belief knowledge? *Analysis*, 23(6):121–123, 1963.
- [73] M. Ginsberg, editor. *Readings in nonmonotonic reasoning*. Morgan Kaufmann, 1987.
- [74] L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. Preferential description logics. In N. Dershowitz and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, number 4790 in LNAI, pages 257–272. Springer, 2007.
- [75] L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. Reasoning about typicality in preferential description logics. In S. Hölldobler, C. Lutz, and H. Wansing, editors, *Proceedings of the 11th European Conference on Logics in Artificial Intelligence (JELIA)*, number 5293 in LNAI, pages 192–205. Springer, 2008.
- [76] L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. $\mathcal{ALC} + T$: a preferential extension of description logics. *Fundamenta Informaticae*, 96(3):341–372, 2009.

- [77] L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. A minimal model semantics for non-monotonic reasoning. In L. Fariñas del Cerro, A. Herzig, and J. Mengin, editors, *Proceedings of the 13th European Conference on Logics in Artificial Intelligence (JELIA)*, number 7519 in LNCS, pages 228–241. Springer, 2012.
- [78] L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. A non-monotonic description logic for reasoning about typicality. *Artificial Intelligence*, 195:165–202, 2013.
- [79] L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. Semantic characterization of rational closure: From propositional logic to description logics. *Artificial Intelligence*, 226:1–33, 2015.
- [80] R. Goré. Tableau methods for modal and temporal logics. In M. D’Agostino, D.M. Gabbay, R. Hähnle, and J. Posegga, editors, *Handbook of Tableau Methods*, pages 297–396. Kluwer Academic Publishers, 1999.
- [81] S.O. Hansson. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer Academic Publishers, 1999.
- [82] D. Harel, J. Tiuryn, and D. Kozen. *Dynamic Logic*. MIT Press, 2000.
- [83] F. van Harmelen, V. Lifschitz, and B. Porter, editors. *Handbook of Knowledge Representation*. Elsevier Science, 2008.
- [84] J. Hawthorne. Nonmonotonic conditionals that behave like conditional probabilities above a threshold. *Journal of Applied Logic*, 5(4):625–637, 2007.
- [85] A. Herzig, L. Perrussel, and I. Varzinczak. Elaborating domain descriptions. In G. Brewka, S. Coradeschi, A. Perini, and P. Traverso, editors, *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI)*, pages 397–401. IOS Press, 2006.
- [86] A. Herzig and I. Varzinczak. Domain descriptions should be modular. In R. López de Mántaras and L. Saitta, editors, *Proceedings of the 16th European Conference on Artificial Intelligence (ECAI)*, pages 348–352. IOS Press, 2004.
- [87] A. Herzig and I. Varzinczak. Cohesion, coupling and the meta-theory of actions. In L. Kaelbling and A. Saffiotti, editors, *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 442–447. Morgan Kaufmann Publishers, 2005.
- [88] A. Herzig and I. Varzinczak. On the modularity of theories. In R. Schmidt, I. Pratt-Hartmann, M. Reynolds, and H. Wansing, editors, *Advances in Modal Logic*, 5, pages 93–109. King’s College Publications, 2005.
- [89] A. Herzig and I. Varzinczak. A modularity approach for a fragment of \mathcal{ALC} . In M. Fisher, W. van der Hoek, B. Konev, and A. Lisitsa, editors, *Proceedings of the 10th European Conference on Logics in Artificial Intelligence (JELIA)*, number 4160 in LNAI, pages 216–228. Springer-Verlag, 2006.
- [90] A. Herzig and I. Varzinczak. Metatheory of actions: beyond consistency. *Artificial Intelligence*, 171:951–984, 2007.
- [91] A. Heyting. *Intuitionism: An Introduction*. North-Holland, 3 edition, 1971.
- [92] H. Katsuno and A. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 3(52):263–294, 1991.

- [93] B. Konev, D. Walther, and F. Wolter. The logical difference problem for description logic terminologies. In A. Armando, P. Baumgartner, and G. Dowek, editors, *Proceedings of the 4th International Joint Conference on Automated Reasoning (IJCAR)*, number 5195 in LNAI, pages 259–274. Springer-Verlag, 2008.
- [94] S. Konieczny, J. Lang, and P. Marquis. DA^2 merging operators. *Artificial Intelligence*, 157:49–79, 2004.
- [95] R. Kontchakov, F. Wolter, and M. Zakharyashev. Can you tell the difference between DL-Lite ontologies? In J. Lang and G. Brewka, editors, *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 285–295. AAAI Press/MIT Press, 2008.
- [96] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
- [97] J. Lang, F. Lin, and P. Marquis. Causal theories of action – a computational core. In V. Sorge, S. Colton, M. Fisher, and J. Gow, editors, *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1073–1078. Morgan Kaufmann, 2003.
- [98] N. Laverny and J. Lang. From knowledge-based programs to graded belief-based programs, part I: on-line reasoning. *Synthese*, 147:277–321, 2005.
- [99] D. Lehmann. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15(1):61–82, 1995.
- [100] D. Lehmann. Stereotypical reasoning: logical properties. *Logic Journal of IGPL*, 6(1):49–58, 1998.
- [101] D. Lehmann and M. Magidor. Preferential logics: the predicate calculus case. In *Proceedings of TARK*, pages 57–72, 1990.
- [102] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992.
- [103] K. Lehrer and T. Paxson. Undeclared justified true belief. *The Journal of Philosophy*, 66(8):225–237, 1969.
- [104] H. Levesque and G. Lakemeyer. *The Logic of Knowledge Bases*. MIT Press, 2001.
- [105] D. Makinson. *Bridges from Classical to Nonmonotonic Logic*, volume 5 of *Texts in Computing*. King’s College Publications, 2005.
- [106] D. Makinson. How to go nonmonotonic. In *Handbook of Philosophical Logic*, volume 12, pages 175–278. Springer, 2 edition, 2005.
- [107] J. McCarthy. Programs with common sense. In *Proceedings of the Teddington Conference on the Mechanization of Thought Processes*, pages 75–91, 1959.
- [108] J. McCarthy. Circumscription, a form of nonmonotonic reasoning. *Artificial Intelligence*, 13(1-2):27–39, 1980.
- [109] J. McCarthy. Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence*, 28(1):89–116, 1986.

- [110] P. McNamara. Deontic logic. In D.M. Gabbay and J. Woods, editors, *Handbook of the History of Logic*, volume 7: Logic and the Modalities in the Twentieth Century, pages 197–288. Elsevier North-Holland, 2006.
- [111] E. Mendelson. *Introduction to Mathematical Logic*. D. Van Nostrand Company, 4 edition, 1997.
- [112] T. Meyer, K. Lee, and R. Booth. Knowledge integration for description logics. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 645–650. AAAI Press/MIT Press, 2005.
- [113] K. Moodley. Debugging and repair of description logic ontologies. M.Sc. thesis, University of KwaZulu-Natal, Durban, South Africa, 2011.
- [114] R. Moore. Semantical considerations for nonmonotonic logic. *Artificial Intelligence*, 25:75–94, 1985.
- [115] E. Mueller. *Commonsense Reasoning*. Morgan Kaufmann, 2006.
- [116] N. Noy and M. Musen. PromptDiff: A fixed-point algorithm for comparing ontology versions. In R. Dechter, M. Kearns, and R. Sutton, editors, *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI)*, pages 744–750. AAAI Press/MIT Press, 2002.
- [117] D. Nute, editor. *Defeasible Deontic Logic*, volume 263 of *Synthèse Library*. Kluwer Academic Publishers, 1997.
- [118] J. Pearl. System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Proceedings of the 3rd Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 1990.
- [119] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2 edition, 2009.
- [120] R. Pressman. *Software Engineering: A Practitioner’s Approach*. McGraw-Hill, 1992.
- [121] G. Priest. Paraconsistent logic. In D.M. Gabbay and F. Günthner, editors, *Handbook of Philosophical Logic*, volume 6, pages 287–393. Kluwer Academic Publishers, 2nd edition, 2002.
- [122] J. Quantz and V. Royer. A preference semantics for defaults in terminological logics. In *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 294–305, 1992.
- [123] R. Reiter. On closed world data bases. In H. Gallaire and J. Minker, editors, *Logics and Databases*, pages 119–140. Plenum Press, 1978.
- [124] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1-2):81–132, 1980.
- [125] R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, 2001.
- [126] G. Restall. Relevant and substructural logics. In D.M. Gabbay and J. Woods, editors, *Handbook of the History of Logic*, volume 7: Logic and the Modalities in the Twentieth Century, pages 289–398. Elsevier North-Holland, 2006.
- [127] M. Shanahan. *Solving the frame problem: a mathematical investigation of the common sense law of inertia*. MIT Press, 1997.

- [128] Y. Shoham. *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press, 1988.
- [129] H.C.M. de Swart. *Logic: Mathematics, Language, Computer Science and Philosophy*, volume 1. Peter Lang, 1993.
- [130] M. Thielscher. A unifying action calculus. *Artificial Intelligence*, 175(1):120–141, 2011.
- [131] L. van der Torre and Y. Tan. The many faces of defeasibility in defeasible deontic logic. In D. Nute, editor, *Defeasible Deontic Logic*, volume 263 of *Synthese Library*, pages 79–121. Kluwer, 1997.
- [132] D. Touretzky. Implicit ordering of defaults in inheritance systems. In G. Shafer and J. Pearl, editors, *Readings in Uncertain Reasoning*, pages 668–671. Morgan Kaufmann, 1990.
- [133] I.J. Varzinczak. Causalidade e dependência em raciocínio sobre ações (“Causality and dependency in reasoning about actions”). M.Sc. thesis, Universidade Federal do Paraná, Curitiba, Brazil, 2002.
- [134] I.J. Varzinczak. *What is a good domain description? Evaluating and revising action theories in dynamic logic*. PhD thesis, Université Paul Sabatier, Toulouse, 2006.
- [135] I.J. Varzinczak. Action theory contraction and minimal change. In J. Lang and G. Brewka, editors, *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 651–661. AAAI Press/MIT Press, 2008.
- [136] I.J. Varzinczak. On action theory change. *Journal of Artificial Intelligence Research (JAIR)*, 37:189–246, 2010.
- [137] I.J. Varzinczak. A note on a description logic of concept and role typicality for defeasible reasoning over ontologies. *Logica Universalis*, 12(3-4):297–325, 2018.
- [138] P. Veloso, S. Veloso, J. Viana, R. de Freitas, M. Benevides, and C. Delgado. On vague notions and modalities: a modular approach. *Logic Journal of the IGPL*, 18(3):381–402, 2009.
- [139] D. Zhang, S. Chopra, and N. Foo. Consistency of action descriptions. In M. Ishizuka and A. Sattar, editors, *Proceedings of the 7th Pacific Rim International Conference on Artificial Intelligence (PRICAI): Trends in Artificial Intelligence*, number 2417 in LNCS, pages 70–79. Springer, 2002.
- [140] D. Zhang and N. Foo. EPDL: A logic for causal reasoning. In B. Nebel, editor, *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 131–138. Morgan Kaufmann Publishers, 2001.
- [141] D. Zhang and N. Foo. Interpolation properties of action logic: Lazy-formalization to the frame problem. In S. Flesca, S. Greco, N. Leone, and G. G. Ianni, editors, *Proceedings of the 8th European Conference on Logics in Artificial Intelligence (JELIA)*, number 2424 in LNCS, pages 357–368. Springer-Verlag, 2002.
- [142] D. Zhang and N. Foo. Frame problem in dynamic logic. *Journal of Applied Non-Classical Logics*, 15(2):215–239, 2005.