

Resolving Conflicts in Knowledge for Ambient Intelligence

MARTIN HOMOLA¹, THEODORE PATKOS², GIORGOS FLOURIS²,
JÁN ŠEFRÁNEK¹, ALEXANDER ŠIMKO¹, JOZEF FRTÚS¹, DIMITRA
ZOGRAFISTOU² and MARTIN BALÁŽ¹

¹*Comenius University in Bratislava, Mlynská dolina, 842 48 Bratislava, Slovakia*

E-mail: {homola, frtus, simko, sefranek, balaz}@fmph.uniba.sk

²*Foundation for Research and Technology Hellas (FORTH), Heraklion Crete, Greece*

E-mail: {patkos, fgeo, dzograf}@ics.forth.gr

Abstract

Ambient Intelligence (AmI) proposes pervasive information systems composed of autonomous agents embedded within the environment who, in orchestration, complement human activity in an intelligent manner. As such, it is an interesting and challenging application area for many computer science fields and approaches. A critical issue in such application scenarios is that the agents must be able to acquire, exchange and evaluate knowledge about the environment, its users, and their activities. Knowledge populated between the agents in such systems may be contextually-dependent, ambiguous, and incomplete. Conflicts may thus naturally arise, that need to be dealt with by the agents in an autonomous way. In this survey, we relate AmI to the area of Knowledge Representation and Reasoning (KR), where conflict resolution has been studied for a long time. We take a look at a number of KR approaches that may be applied: context modelling, multi-context systems, belief revision, ontology evolution and debugging, argumentation, preferences, and paraconsistent reasoning. Our main goal is to describe the state of the art in these fields, and to draw attention of researchers to important theoretical issues and practical challenges that still need to be resolved, in order to reuse the results from KR in AmI systems or similar complex and demanding applications.

1 Introduction

1.1 Background and Focus

The knowledge representation and reasoning research community contributed over the years a multitude of well-defined theoretical results, as well as practical solutions for engineering information systems tailored to the needs of diverse application domains that deal with knowledge. Recently, the domain of Ambient Intelligence (AmI) emerged as an important objective merging research trends from different disciplines; many of the problems relevant with respect to the handling of knowledge within AmI were studied in KR for a number of years. Although some solutions from this field already found their way to AmI systems, as intelligent environments move from the lab to the real world their behaviour becomes more sophisticated and the development of viable holistic approaches requires a thorough reconsideration of the applicable tools and methodologies that need to be seamlessly combined within them.

The field of AmI (Zhelka, 1998; ISTAG, 2013) studies information systems embedded within the environment, sensitive to the human presence, that are able to facilitate distributed and networked computing machinery with the aim to accommodate and support human users with their everyday activities and tasks. Application domains of AmI range from ambient assisted living and health-care monitoring, to smart home and office automation, transportation services, classroom and education

environments, smart shopping, and others (Cook et al., 2009; Chan et al., 2008; Rubel et al., 2004; Sadri, 2011).

The envisioned AmI applications materialize a long anticipated application objective for Artificial Intelligence (AI), and many of the subproblems studied within AmI can be addressed by AI methods, including: how to recognize activities, how to detect, anticipate and respond to users' needs and intentions, how to develop autonomous entities that can exhibit commonsense behaviour, how to conduct distributed reasoning, etc. AmI systems need to be able to process knowledge about the environment in which they are embedded, but also about the users' activities, goals, and tasks. As the knowledge of the environment may be imperfect and ambiguous and the goals of diverse users may be contradictory, one particular issue that has been long studied within KR becomes relevant also in this domain: the problem of *conflict resolution*. The autonomous entities involved in an AmI application need to be able to handle conflicting knowledge, and to find (a form of) mutual consensus in their actions, in order to serve their users well, in a consistent, and unobtrusive fashion. This issue was already recognized by the researchers within the AmI domain (Resendes et al., 2014).

In this survey, we focus on selected KR approaches and formalisms that address conflict resolution. Each of the surveyed approaches addresses the problem with a specific motivation, following certain use cases. We do believe that the research in KR has now advanced to a state where it is useful to consider also more broadly defined problems, rooted in challenging real-world applications. For this sake, we look towards AmI as a model domain that integrates a set of important features that need to be considered together, and not only in isolation: (a) distributed and decentralized nature of the system, with multiple autonomous reasoning entities (agents) and the need to resolve the conflicts reaching a certain consensus between the entities; (b) the need to recognize the context: the environment where the agent is placed, users within the environment, their preferences and needs; (c) being able to react appropriately to a possible change in the situation; (d) computational effectivity of reasoning and conflict resolution; (e) unobtrusiveness: the system should be able to work autonomously without calling for interference of the users.

While we could have selected other model domains, this combination of requirements together with some other AmI specifics, that we discuss later on, introduce complex interdependencies between agents and their interactions that are much more intensively emphasised within the AmI setting, and forbid many simplifying assumptions one may take in in other domains.

1.2 Goals and Audience

The goal of this survey is to review and evaluate the relevant approaches from KR that can be applied on the conflict resolution problem. In particular, we investigate the extent to which they can be applied to a complex setting, as framed by the requirements (a)–(e) above. That is, our goal is to identify to which extent the KR approaches may possibly be applied, and to pinpoint important issues that still need to be resolved in the respective subfields, in order to become applicable in such challenging domains as AmI.

Having this goal in mind, we believe that the survey can be especially useful to the researchers in KR, who will get an overview of research directions relevant to the problem of conflict resolution, they will be able to compare the various approaches within the field, contrasting their applicability and open issues. Researchers will be able to compare how analogous issues were addressed in different subfields, and also where and how different approaches need to be combined in order to meet the given goals.

The survey can also be useful to researchers who are looking for a suitable conflict resolution methodology for their application in AmI or a similar area. They will learn about the approaches coming from the KR area and about their current status and potential applicability.

Other similar areas which share some specifics with AmI and have strong needs for conflict resolution may include, e.g., open multi-agent systems, systems of agents interacting with Web services and processing information from the Web, etc.

1.3 Survey Scope and Overview

In Section 2, we first introduce the necessary background from AmI and discuss how KR is relevant to AmI, where in the architecture of AmI systems KR methods can be best applied. In Section 2.3, we take a closer look at the problem of conflict resolution, and we analyse the different kinds of conflicts that appear in AmI applications, especially from the point of view of the diversity of knowledge types that necessarily need to be processed by AmI systems.

Then, in Section 3, we survey a number of selected KR areas in which the problem of conflict resolution was pursued. In each area, we focus on the representational issues (i.e., what problems the approach allows to deal with), but also on computational issues (i.e., if the solutions can be computed, and how hard is it to compute them, if there are any known implementations, etc.), and as well on the relevance of the area to AmI and the known applications in this domain. The surveyed areas are as follows.

Context Modelling (Section 3.1) was considered for a long period of time an important issue in KR, and it is a central problem for AmI as well, where *context recognition* is equally important to context modelling. Apart from answering the question which information is needed to capture the current situation of an agent, and how this contextual model should be organized, reasoning agents are challenged with uncertainty and ambiguity of the data on which they need to build their contextual models, and they need to resolve conflicts that may thus arise (e.g., in sensory data, or between the sensory data and background knowledge). We will discuss both more traditional KR-based approaches where the uncertainty and ambiguity are captured symbolically, and data-driven approaches where they are captured numerically. Hybrid approaches try to combine the results of the former two.

Multi-Context Systems (Section 3.2) and similar approaches in the area of distributed KR focus on the problem of combining multiple knowledge sources for reasoning. The combination is achieved with so called bridge-rules which allow to transfer conclusions from local reasoning in one knowledge source into another one as facts. Particularly relevant to AmI is the assumption that the knowledge sources may be distributed and heterogeneous, e.g., each coming from a different agent that may possibly be placed in a different context, and may even use a different representation language. The knowledge derived from different sources may be conflicting, and multi-context systems provide methods to resolve such conflicts. The focus on distribution of knowledge is also relevant, as this is often required in real world applications.

Belief change (Section 3.3). Often referred to as *belief revision* is the problem of determining how to modify an agent's KB in the face of new, possibly contradictory information. The focus is on identifying and resolving problems before they actually creep into the KB. Belief change approaches could be used to prevent conflicts arising from conflicting sensor readings or from information provided by other agents that is conflicting with the context that the local agent understands. The conflicts considered by belief change approaches are logical inconsistencies, and many of the works in this area deal with the theoretical and philosophical aspects of the problem of updating a KB. Thus, the field is quite relevant for understanding the process of updating knowledge bases, and, consequently, the semantics that a rational agent should apply in order to prevent conflicts from creeping into its KB.

Ontology evolution (Section 3.4.1) refers to the process of modifying an ontology in response to a certain change in the domain or its conceptualization. Ontologies and ontology languages are being increasingly applied also by AmI applications, therefore this area is relevant. It is similar to belief change, in the sense that ontology evolution also tries to prevent conflicts from appearing in the KB. Ontology evolution has a more practical nature compared to belief change, in the sense that most approaches are dealing with the practical aspects of the problem of evolution, rather than understanding the evolution process per se. It deals with both the schema and the data of the ontology. It can thus serve to resolve conflict of various types, depending especially on the role the ontology is playing in the agent's knowledge.

Ontology debugging (Section 3.4.2), just like ontology evolution, deals with ontological languages at a practical level. The main difference is that ontology debugging is applicable after the conflicts have appeared in the KB, which can happen either because they were somehow allowed to appear, or because of reckless updating, or because the rules associated with the data had to be changed. Ontology debugging may take two different forms: direct approaches (*repairing*) actively try to restore the conflicts by removing conflicting elements, whereas indirect approaches (based on *consistent query answering*

methods) do not directly address the conflicts but provide appropriate inconsistency-tolerant semantics that allow consistent answers over a conflicting dataset, essentially “hiding” the conflict from the user.

Argumentation (Section 3.5) aims to understand the process of exchanging rational arguments. More specifically, argumentation studies how arguments relate to each other, and how one can take decisions in the presence of possibly conflicting arguments. Argumentation was successfully applied to conflict resolution, because the resolution of a conflict can be modeled as the process of deciding which part of the evidence (arguments) is acceptable, given a complex evidence set parts of which support (or attack) conflicting arguments.

In Section 3.6, we have a look at the existing body of work on the relation between argumentation and belief change, which we suggest as particularly interesting development w.r.t. AmI, as combining and revising argumentation systems will make them more applicable in complex and dynamic environments.

Preferential reasoning (Section 3.7). KR formalisms are employed to encode a problem in a formal language, and use reasoning capabilities of the formalism to compute the solutions to the problem. Often multiple solutions exist, e.g., due to the nature of a problem, or due to the use of general rules that are used to model the problem. Preferences are then used to select preferred solutions. Or preferences can be used to select from multiple conflicting rules that are applicable in certain situations.

Paraconsistent reasoning (Section 3.8) While most of the approaches above aim at resolving conflicts, e.g., by performing a repair, or revision of the knowledge base, or by deciding which arguments should be upheld and which should be rejected, paraconsistent reasoning studies logics which are able to derive meaningful conclusions also from inconsistent theories and data sets, e.g., by ignoring the inconsistent premisses and drawing conclusions only from the consistent part of the knowledge.

2 Ambient Intelligence, Knowledge and Conflicts

2.1 Ambient Intelligence

The advent and penetration of interconnected mobile devices into our everyday life has triggered a shift in computing towards sensor-rich environments with pervasive technologies, often referred to as smart spaces. Driven by the *ubiquitous computing* paradigm, a term coined by Mark Weiser’s 1991 vision of a new generation of computer systems (Weiser, 1991), the new research area of AmI has emerged. AmI places the human user at the center of attention aiming at creating intelligent environments with the ability to adapt to human preferences, serve their needs and goals, and communicate with their inhabitants utilizing novel means. This paradigm implies a seamless medium of interaction, advanced networking technology, and efficient knowledge management, in order to deploy an environment that is aware of the characteristics of human presence and the diversities of personalities, and is also capable of responding intelligently and proactively to the users’ needs.

AmI systems are intended to be (Zelkha, 1998; Aarts et al., 2001):

1. embedded within the environment: users do not need to be concerned with their operation,
2. context aware: they are able to recognize the user and the situation,
3. personalized: they can serve different users according to their own needs,
4. adaptive: they can change in response to the environment and users’ actions,
5. anticipatory: they can understand users’ needs and act upon them pro-actively, as opposed to only responding user generated requests.

Agent-based paradigms, most commonly Multi-Agent Systems (MAS) are often used to design and develop AmI environments. These contain embedded software entities, called agents, able to perceive and reason upon the current context, exploit the functionality of devices installed within the environment, and pursue specific goals while exhibiting autonomous behaviour.

However there are a few features that distinguish AmI applications from MAS in general. The agents are assumed to be physically placed within a real world environment, also inhabited by human users –

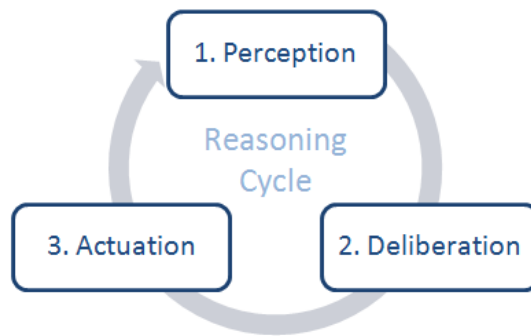


Figure 1 The reasoning cycle of autonomous devices in smart spaces

i.e., the smart space. As a variety of elements and devices serving diverse purposes are typically installed in smart spaces, it is unreasonable to assume that all agents will share a unique agent architecture – in fact, the agents may be rather heterogeneous in their implementation. Particularly, their cognitive skills may range from simple reactive agents whose behaviour is based on the most recent sensor readings, to complex knowledge-based and deliberative agents that perform elaborate reasoning in order to infer relevant context, make estimates over the users’ intentions, and communicate and negotiate with the other agents in collaborative manner.

Given the complex tasks that AmI applications are to carry out as a whole, we typically assume that a smart space hosts at least a small number of the latter type of agents. While often such rational agents are modelled using the BDI architecture (Rao and Georgeff, 1991, 1995; Cohen and Levesque, 1990; Bratman, 1987), comprising beliefs (i.e., some knowledge), desires (i.e., certain set of goals), and intentions (i.e., commitment to execute actions in order to meet a chosen set of the goals). Aiming to provide a suitable abstraction of agents for the need of this survey, we simplify the agent architecture and assume that the agents comprise especially the following components:

- A *knowledge base* (KB) of certain sort, comprising as a distinguished part the *context model* of the current situation respective to the agent, and possibly some additional background and domain knowledge used by the agent. There may be different kinds of beliefs that we may need to distinguish. Each agent may keep track of different aspects of the world and represent them differently than the other agents.
- A set of *goals* the agent is able to follow in all possible situations to serve its purpose, from which the agent selects some, depending on the current perceived context.
- Either some predefined plans of *actions* to execute to achieve each goal, or the ability to plan the actions accordingly when needed.
- Some way to *communicate* with other agents with the aim to exchange knowledge and cooperate the next actions (e.g., messages, queries, bridge rules, etc.).

It should be remarked that in AmI systems the general aim of an agent is to perceive and accommodate the goals of the users and to help them in carrying out actions to achieve these goals. For this reasons, agents may likewise model users simply as agents having goals and executing actions. This abstraction is indeed useful when studying AmI environments as a whole, however we must keep in mind that there is a distinction between the goals of an agent and that of a user, which are not always easily specifiable.

An abstract loop that can characterize the basic internal reasoning phases carried out by an agent is shown in Fig. 1 and involves the phases of perception, deliberation and actuation. This cycle is triggered by specific sensory inputs that the agent is monitoring (or the lack of them) and captures the ability to both deliberate about how best to *interpret changes* that occur in their dynamically changing world, as well as to *make decisions* about the most appropriate course of actions that needs to be taken to support the human

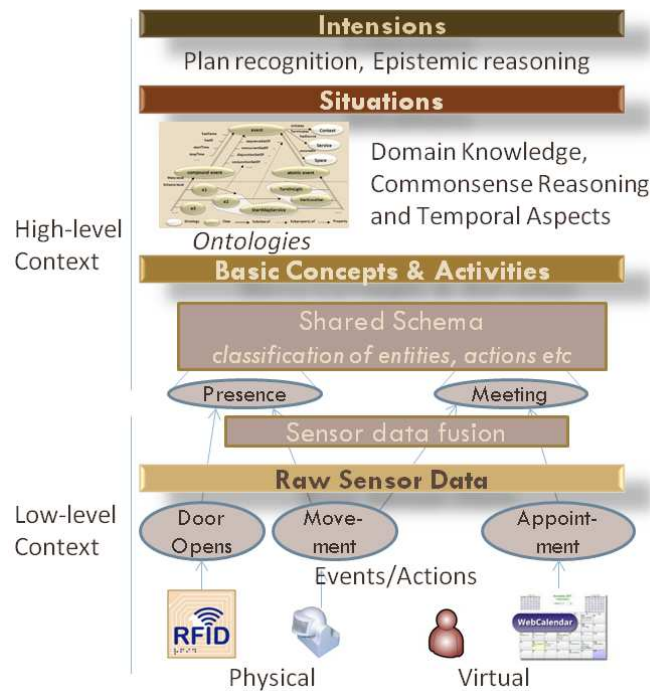


Figure 2 Conceptual layers of contextual knowledge

users' activities. While many approaches have been proposed to study each phase alone, recent studies (e.g., Pecora et al., 2012; Chen and Khalil, 2011) argue about the need for a seamless integration of the tasks of perception, recognition and acting in a coherent loop, in order to synthesize support services in smart environments with proper and verifiable behaviour.

In addition to its dynamic nature, the aspect of heterogeneity is an equally challenging factor for developing AmI services. Agents operating in smart spaces may possess different reasoning skills, obtain access to distinct knowledge repositories, local or shared, and evaluate incoming information based on different trust criteria. A real-world smart system needs to respect the fact that the way high-level context is inferred by each involved agent is not an objective process. Being highly distributed, these environments produce information that can be interpreted in a totally different manner by the various intelligent agents; as such, it is not uncommon for the latter to end up having incoherent and conflicting views of the current context. Devising intelligent automated mechanisms for identifying, preventing or resolving conflicts is of utmost importance, in order to appropriately balance between the two main design principles that have been set for the success of smart spaces: being as less intrusive as possible minimizing the need for user input, while still letting humans feel confident that they have control over their space.

2.2 Relevance of Knowledge Representation and Reasoning in Ambient Intelligence

A typical AmI application, as described in the previous section, needs to deal with a certain amount of knowledge, in order to evaluate the situation and to carry out the required tasks to serve its users. This knowledge must be represented, and processed within the system. In this section, we take a closer look into the types of knowledge that need to be represented and the kind of processing (i.e., reasoning) that is needed. Many of the problems involved have been thoroughly studied for years in the area of knowledge representation and reasoning.

In Fig. 2, we observe the different conceptual layers of knowledge within an AmI system. As a first type of knowledge, the bottom layers are concerned with identifying the current *context* in which the AmI system is placed. There are typically two layers of contextual knowledge – the lower level containing

raw sensor readings, and the higher level in which these readings are interpreted on a more abstract level, using a set of concepts based on a commonly agreed schema.

Climbing up the levels of abstraction, the need to represent rich knowledge structures by means of expressive models becomes more apparent. User activities, such as the process of making coffee, are highly goal-driven, typically follow specific patterns, and pre-assume a significant extent of background and domain knowledge with respect to their causal effects and ramifications. In addition, their compositions, often referred to as *situations*, such as the preparation of breakfast, have rich structural and temporal aspects, as for instance location, duration, frequency, causality, and action. In order for AmI systems to fulfil complex tasks, they may need to consider also specific domain knowledge, and data from external data sources. Expressive symbolic modelling not only allows to combine all necessary reasoning tasks, but in addition significantly enhances the reasoning capacity of smart applications by enabling developers to hide the complexities and noise of sensor readings (Bikakis et al., 2008; Bikakis and Antoniou, 2010b), while exploiting the implicit structure of the activities being observed and data that needs to be processed (Ye et al., 2012; Loke, 2004). Languages with expressive and formal semantics, relying on ontologies (Staab and Studer, 2004; Sowa, 2000) or dealing with commonsense reasoning (Mueller, 2014; Kuipers, 1984), spatio-temporal issues (Cohn and Hazarika, 2001; Gabbay et al., 2003), action and planning (Ginsberg and Smith, 1988; Lifschitz, 1999; Eiter et al., 2003b) have long been studied in knowledge representation.

AmI systems comprise autonomous entities, which need to act in synchronized fashion and collaborate, in order to meet the users' goals. This inherently imposes the need of distributed processing of the knowledge involved in the overall application. Different agents may hold different viewpoints on the context, and they may have access to different knowledge resources. Therefore, ambiguity and conflicts frequently arise and must be resolved at execution time, in order to assure smooth operation of the system (Henricksen and Indulska, 2004). While current AmI implementations often take a simplistic and centralized approach to conflict resolution (Resendes et al., 2014), within KR a multitude of approaches have been devised with the aim to deal with distributed knowledge sources modelled from the perspective of distinct viewpoints (Giunchiglia, 1993), possibly inconsistent knowledge (Bertossi et al., 2005), and reaching agreement (Ossowski, 2013).

While many of the KR techniques presented above may not have achieved yet the requirement of scalability in an extent suitable for immediate application to AmI systems, they certainly focus on a number of issues that are central to AmI. Their further development will give significant potential for improving the capabilities of AmI systems. In this survey, we concentrate our focus to the problem of resolving conflicts once they arise, and specifically consider the relevant KR approaches that address this problem in different settings and with different goals. Before presenting the surveyed areas in Section 3, we first take a look at the different types of conflicts that are met in AmI systems.

2.3 Conflicting Knowledge

One of the problems that has long been studied in KR is how to deal with conflicting knowledge. This problem is particularly relevant in scenarios where multiple distributed knowledge sources that have to be combined for reasoning come into play. As we argued above, AmI systems and architectures often fall into this case, especially if they incorporate multiple autonomous agents that need to cooperate, in order to achieve common goals. Indeed, this has been noted by other researchers working in the field (Resendes et al., 2014; Henricksen and Indulska, 2004; Muñoz Ortega et al., 2010).

Resendes et al. (2014) analyze different types of conflicts that may arise in AmI systems and organize them into a taxonomy, as listed in Table 1. They identify four basic broad categories of conflicts, which are dubbed *dimensions* in order to stress their orthogonality, i.e., the fact that one conflict can be independently classified with respect to each of them.

The *source* dimension indicates where/how each conflict originates – it may be the case that users (or applications) are conflicting over some resource allocation, or it is not possible to execute some action due to policy, or there are conflicting user profiles. The *interveniens* dimension considers the different parties between which the conflict arises: there might be conflicting intentions within a single user, between

Table 1 Taxonomy of conflicts (Resendes et al., 2014)

Dimension	Possible types
Source	resource application policy role
Intervenients	single user user vs. user user vs. space
Time of detection	a priori when it occurs a posteriori
Solvability	conflict avoidance conflict resolution acknowledge inability acknowledge occurrence

multiple users, or between user and the space. The *time of detection* dimension sorts conflicts into those that are (can be) detected a priori, at the time they occur, or only a posteriori. Finally, the *solvability* dimension indicates at which level can conflicts be resolved – before they happen (i.e., to avoid them), or immediately when they happen, or, possibly, some conflicts cannot be resolved in sensible time, and these are further split into those which cannot be resolved at all, and those which cannot be resolved due to being detected too late.

For an example, let us consider that there are two conflicting instructions for the temperature to be set in the living room, say one 20 °C and the other one 23 °C. Using the taxonomy of Resendes et al., we classify such a conflict as resource conflict w.r.t. the source dimension. In the intervenients dimension there are two likely cases: either this is a conflict between the space and a user, one instruction coming from each, or this might be a conflict between two distinct users. The case of a single user providing both instructions is not ruled out, but indeed less likely. Moving to the time of detection dimension, such conflicts are typically detected when they occur, if the instructions are immediate. They can also be detected a priori, if they concern a setting relevant for some future period. According to the solvability dimension, the conflict is likely solvable by conflict resolution (e.g., based on some policy, preference between the user and the space, between the two users, or alike). Alternatively, conflict avoidance may be achieved if the users are warned and they can agree on a common setting. This may however decrease unobtrusiveness, praised in AmI systems, if not executed with caution – for example a warning at the time when the user is providing the setting may be considered sufficiently unobtrusive.

Homola and Patkos (2015) propose an additional dimension to be added to the taxonomy of Resendes et al., namely that of the *knowledge type*. As each type of knowledge is processed differently, and in a different point of the agents reasoning cycle (cf. Fig. 1), conflicts in distinct types of knowledge may need to be processed differently in AmI systems.

Table 2 Taxonomy of conflicts (Homola and Patkos, 2015)

Dimension	Possible types
Knowledge type	sensory input context domain/background goal action

The possible values of the knowledge type dimension are listed in Table 2 and in more detail they are described as follows:

Sensory input conflict: if a conflicting reading of some sensors appears. Either multiple readings of the same sensor, or similar sensors may be conflicting. Or the reading may be of different sensors, however the outputs are mutually exclusive (the agents know that these outputs cannot occur at the same time). The conflict may arise within a single agent, but it may also be distributed between more than one agent (each containing part of the conflicting readings). The latter option may subsequently possibly cause a contextual conflict.

Contextual conflict: if two (or more) agents are part of the same situation, their models of the world are conflicting, implying, e.g., a different location, or perceived activity of the user, etc. This type of conflict may likely be caused by a previous unnoticed sensory input conflict. But it may also be caused by different evaluation of the situation.

Domain and background knowledge conflict: domain and background knowledge refer to the information the agent possesses and uses in order to fulfil its purpose. For instance, a calendar scheduling agent associated with a user records information about months in a year, days in a week, working days, holidays, etc. This is the knowledge respective to the domain of the agent's tasks. To contrast this with contextual knowledge, the fact that Monday follows Sunday is part of unchanging domain knowledge, respective to the calendar domain, while the current date and time, first day of week are in reality contextual knowledge, which changes from situation to situation. It is apparent, that conflict in domain and background knowledge should occur less frequently in AmI systems, in comparison to the remaining four kinds of conflicts discussed here. Also, if they occur, they may require a different kind of solution, due to domain and background knowledge being most typically considered unchanging and fully specified (to the extent required by the application). Hence redesign of the agent's knowledge base by its creator may be required, in contrast to automatic dealing with the conflict.

Goal conflict: if two (or more) agents are part of the same situation, their models of the world are compatible, but they have mutually conflicting goals. Note that we do not consider it a goal conflict if agents have conflicting goals in different models of the world, because it is natural to have different goals in different situations.

Action conflict: if two (or more) agents share a compatible model of the world, and a compatible set of goals, yet decide to follow a contradictory course of actions to carry out their goals.

According to the knowledge type dimension, the conflict between two desired temperature settings, as discussed above, is a goal conflict rather than an action conflict, as there are really two conflicting goals for setting the temperature to a different value. This is different than the case when there is a common goal and the conflict is merely between the actions chosen to be executed by distinct parties.

As a case of an action conflict, consider an ambient assisted living scenario in which the agents monitoring a user detect that the user has left the stove on while she went to take a shower. With the common goal to secure the stove an action, conflict arises when one agent decides to alert the user, while another decides to call external assistance. The actions are not directly contradictory, however one of them is redundant if executed at the same time. It is more preferred if the agents avoid such a conflict by reaching an agreement on the best action to execute immediately.

As further observed by Homola and Patkos (2015), these five knowledge types can be sorted on the scale from lower to higher level of knowledge: (a) sensory input, (b) contextual, domain and background knowledge, (c) goals, and (d) actions, in the respective order. Distinguishing between these five types is important also due to the following conjecture: solving conflicts in a lower level knowledge may possibly prevent occurrence of further conflicts in the higher levels of knowledge. Consider an example in which two agents have a conflict in the contextual knowledge, that is, their interpretation of the situation in which they both participate is not compatible (e.g., they may have conflicting information about location). If the conflict is resolved at this level, it is less likely that the agents will come up with conflicting goals and consequently action plans.

In the remainder of this survey, we overview and compare different KR formalisms and tools that are suitable to resolve conflicts. Each formalism can be suitable for different type of conflicts, and the knowledge type dimension is often important to consider. As noted above, conflicts in domain and background knowledge most likely require a manual solution, and hence they are not in our main focus. There are, however, a few formalisms capable of addressing these conflicts, as we note below in the survey.

3 Conflict Resolution Approaches

In the previous section we have learned of the different types of conflicts that may arise in AmI scenarios. In this section, we will present different ways in which agents could resolve such conflicts. In particular, we will focus on the research areas of *context modelling*, *multi-context systems*, *belief change*, *ontology evolution*, *ontology debugging*, *argumentation*, and *preferential reasoning*. Each of these areas is relevant to the problem of resolving conflicts, however, each of them uses a different approach to resolve the conflict. Even though these fields were developed and motivated in different contexts, we feel that the ideas and approaches used there can be easily applied for AmI-related problems which we also highlight in this section.

3.1 Context Modelling and Recognition

There is a plethora of methodologies that investigate issues related to the recognition of context. These methodologies are distinct in the way in which they model, represent or reason over the involved information (relevant surveys include Ye et al., 2012; Chen and Khalil, 2011; Sadri, 2011; Yang, 2009). Among the different classifications that can be made, a commonly adopted one is related to how the information is being processed, which leads to the very broad distinction between *data-driven* and *knowledge-based* approaches. The former rely on a numerical characterization of the uncertainty in inferring context, while the latter apply symbolic reasoning techniques from the field of KR. In the sequel, we investigate the main advantages and weaknesses of methodologies in both categories, as well as recent approaches that attempt to combine prominent features in hybrid models. The main objective of these approaches is to process information from the lower levels of context and produce inferences on the higher levels, with the big majority of approaches focusing on activity recognition (Figure 2). Although the topic of automatic conflict detection and recognition is inherent in the construction of smart spaces, addressing it as part of the context recognition research has only produced partial solutions, as we will see later on.

3.1.1 Data-Driven Approaches

Table 3 Characteristic features of data-driven context recognition fields

Advantages	Weak Points
<ul style="list-style-type: none"> ● Effective handling of uncertainty and conflicts at the sensor level ● Learning process ● Under conditions, can deal with noisy sensor data ● Easily extract patterns and complicated associations 	<ul style="list-style-type: none"> ● Poor portability, scalability and reusability of the models ● Require large amount of training data ● Data annotation is a laborious task ● Lack of formal semantics ● Prone to domain-dependent performance

Data-driven approaches adopt primarily a probabilistic and statistical view of information and widely rely on the enormous impact of machine learning techniques in real-world applications. Although further classification can be made, e.g., based on whether supervised or unsupervised methods are being used,

one distinctive characteristic of data-driven activity recognition algorithms is their capacity to model uncertainty. They apply quantitative measures to evaluate plausibility of observed data, which renders them a popular solution for deciding how best to resolve contradictory sensor inputs. For instance, the problem of domestic activity recognition for a single user was approached by training multiple naive Bayesian models enhanced with ranking features and reliability factors to detect interleaved activities and unexpected sensor malfunction (Lu and Fu, 2009). The same topic for multiple users was investigated with the application of Hidden Markov Models (HMM) that can benefit the process of recognition taking into consideration temporal patterns of data (Singla et al., 2010).

The ability to learn from datasets is a big leverage for data-driven models for tackling conflicts at the sensor level, but often becomes their main point of weakness when attempting to address problems related to the recognition of high-level context. The performance of data-driven approaches is largely dependent on the availability of big amounts of - labelled or unlabelled - training data, thus compromising their capacity to offer scalable, reusable and portable solutions. Due to the pragmatic difficulty to monitor the behaviour of different humans for a long period of time while they perform everyday activities, the models produced exclusively from data-driven techniques are often prone to domain-dependent performance, limiting their reusability and portability (Ye et al., 2012).

Moreover, for abnormal or exceptional behaviours, such as for recognizing a heart attack, specialized equipment is needed (e.g., heartbeat/pulse sensors), which is often considered too intrusive for the daily activities of ordinary users. Without such equipment though it is difficult to train a typical AmI system properly, which is why for instance the applicability of certain approaches, like the one proposed by Jakkula et al. (2009), is limited to frequent and predictable behaviours only. For less common situations, there is also the problem of overfitting, i.e., when the training of a system is based on a small set of annotated data, which cannot be disregarded, as understood in the work of Lester et al. (2005). Table 3 summarises the main points of these approaches.

3.1.2 Knowledge-Based Approaches

Table 4 Characteristic features of knowledge-based context recognition fields

Advantages	Weak Points
<ul style="list-style-type: none"> • Semantically clear • Enhanced interoperability, sharing and portability • Verifiable and intelligible behaviour • Consistency checking • Portability due to the incorporation of domain knowledge • Flexibly extensive with new context types 	<ul style="list-style-type: none"> • Treatment of fuzziness and uncertainty • Learning capacity • Quantified confidence weights of inferred models • Scalability can be an issue in some approaches

With knowledge-based approaches the rules of inference are modelled from first principles, rather than learned from raw data, and typically rely on formal specifications of their syntax and semantics, exploiting symbolic modelling and logic-based reasoning techniques. The expressive power, along with the capacity to verify the properties of their axiomatizations, are key advantages of these methodologies. Among knowledge-based approaches for context recognition, ontology-based models are arguably the most popular ones.

Ontology languages have rich and formal semantics that enables them to express complex knowledge using a wide set of primitives. These languages are utilized in modelling high-level contextual information, due to their ability to incorporate rich domain knowledge and heuristics in a machine processable way, thus offering a number of advantages in terms of expressiveness and quality of

representation, automation and interoperability, composition and level of formality (Bettini et al., 2010; Strang and Linnhoff-Popien, 2004). The most elaborate recent studies in the field of ontology-based context recognition are probably (Riboni and Bettini, 2011b) and (Springer and Turhan, 2009). Based on expressive languages (i.e., OWL 2 Description Logic (DL) and OWL DL, respectively) and decidable reasoners, they enable a context-aware system to detect inconsistencies, infer out activities occurring in the environment, and reproduce knowledge to model activities that share similar characteristics.

Although a multitude of pervasive computing systems have applied ontologies in modelling and reasoning on context knowledge (e.g., (Preuveneers et al., 2004; Patkos et al., 2010; Ye et al., 2007)), most of them try to avoid the generation of conflicts relying on centralized solutions, whereas only few try to explicitly incorporate a solution for resolving conflicts about context within a distributed environment. An early example is Semantic Space (Wang et al., 2004; Tan et al., 2005), a context infrastructure for building smart spaces that investigates a variety of issues, such as context modelling, storage, inference, querying and dynamic discovery of available context providers (wrappers). Context wrappers that obtain raw context information from various software and hardware sources transform them into semantic knowledge (markups) based on the system's context model and store this knowledge in the KB. As an attempt to avoid potential conflicts generated by application-specific inferences, the higher-level context inferred from markups using general purpose reasoners is not explicitly stored in the KB; instead, when needed, specific rulesets are applied to obtain the required knowledge on-the-fly. While this approach may be sufficient for restricted domains, it can be problematic for the general AmI setting, where the different entities often need to obtain a ubiquitous and commonly agreed view of the current situation, in order to decide the best actions to perform and support humans accordingly.

Similar approaches that perform rule-based reasoning are also proposed by Fuchs et al. (2005), as well as in the SOCAM system (Gu et al., 2005), where first-order logic (FOL) rules are applied to reason about the context data and resolve possible conflicts between data coming from different sources. Sets of rules are defined for the classification and evaluation of the quality of the observed context data.

One serious limitation and a key reason for the superiority of data-driven approaches over ontology-based ones, is the limited support for temporal reasoning by ontology-based languages, according to Riboni et al. (2011). Nevertheless, temporal extensions of Semantic Web languages start to become popular (Gutierrez et al., 2007; Batsakis et al., 2011). More importantly, the inherent uncertainty of the information that exists in ubiquitous domains is difficult to handle at the pure symbolic level. While context recognition with respect to coherent incoming information is where most of the aforementioned studies are focusing on, the resolution of information that is conflicting has not been extensively considered so far. Indeed, Semantic Web-based approaches mostly deal with the problem of context disambiguation up to the point of acknowledging that certain parameters can be regarded as unknown.

The problem becomes more pronounced when the recognition task involves high-level complex situations, which ultimately leads in having to resolve the two other types of conflicts defined in Section 2.3, namely goal and action conflicts. For instance, Sadri (2010) proposes an approach to recognize the intentions of a user by means of identifying plans using action graphs. Human intentions are often unclear, cannot be directly measured with physical devices and may be the result of controversial desires. Committing to specific potential human intentions typically means for an agent to decide which of the conflicting knowledge to keep and which to drop. Considering the fact that often multiple and heterogeneous entities are employed in a smart space to perform such reasoning tasks, it becomes evident that dealing with the resolution of the conflicting viewpoints adopted will be inevitable in the next generation of smart systems. This aspect is starting to become an important research topic by considering the integration of techniques from other fields, such as argumentation, as it often requires extra-logical information, as we will see in the following sections. Table 4 gives an overview of the main features of the approaches belonging to this category.

3.1.3 Hybrid and Other Approaches

Although data-driven methods are currently the mainstream choice to activity recognition, with most effective being the supervised learning methods, numerous recent studies justify the attention that

knowledge-based approaches have attracted over the last years. Yet, experience showed that both lines of investigation suffer from limitations that restrict the former to the lower levels of data abstraction and the latter to high-level knowledge. A seamless integration of methodologies for all levels is essential for the materialization of AmI objectives. Much of current research is working towards this end. The COSAR system (Riboni and Bettini, 2011a) for example loosely-couples ontological OWL DL reasoning with statistical inferencing, where the latter predicts the set of possible activities without considering context parameters, in order to make the task manageable, while the former is applied to refine the results. In a similar style, the approach by Roy et al. (2011) applies possibility theory to model qualitatively incomplete knowledge, coupled with a DL representation of context to characterize the subsumption relation of actions, whereas the one by Helaoui et al. (2012) presents a first attempt to use probabilistic DL for activity recognition. Still, these frameworks provide only limited or no support for temporal reasoning, inheriting some of the deficiencies of ontological reasoning in expressiveness, as discussed before.

A coupling of uncertainty with rich temporal relations is presented by Helaoui et al. (2011) that uses Markov Logic Networks (MLNs), a statistical relational framework, to introduce uncertainty measures in logical statements to recognize simultaneous and nested activities. Skarlatidis et al. (2011) go even further to combine MLNs with the Event Calculus, a theory for reasoning about action and change, in order to exploit certain properties of the latter, such as the persistence of activities, and soften them as appropriate. Closely related to MLNs is the approach presented by Augusto et al. (2008) that integrates confidence values to Event-Condition-Action rules. The authors define a new syntax for expressing temporal relations among events, with Dempster-Shafer theory being used to assign confidence values to both antecedents and conclusions of rules. The belief rules they define can be used both to monitor a user's interactions and to recognize exceptional situations.

3.1.4 Summarizing

Although the domain of AmI environments demands the generation of collective context-aware applications for group of users, where individual agents with personal goals seek collaborative execution of tasks (see for instance (Braga Silva et al., 2011; Muñoz Ortega et al., 2010; Muñoz et al., 2011)), relevant literature on collective context conflict resolution is rather scarce. A large portion of contemporary approaches, summarized in Table 5, rely largely on centralized architectures, where a single reasoning entity handles conflict resolution. Those that deploy distributed settings, on the other hand, either try to avoid conflict occurrences (Muñoz Ortega et al., 2010) or focus primarily on policy conflict detection, i.e., establishing behaviour schemes that guarantee acceptable system states (Resendes et al., 2014).

The heterogeneity of the reasoning entities inhabiting smart spaces, as well as the processing load required to reason about context renders unreasonable the assumption that full perception is owed by all agents in a smart environment for describing the world state and the situation the users are involved in. Similarly, the need to combine context inference with actuation cannot be overlooked. As evidenced by Pecora et al. (2012), inference, sensing and actuation must operate in close cooperation, in order to manage an effective integration of the cognitive capabilities of an intelligent system to be both context-aware and proactive. As the various smart entities need to interact and negotiate with one another, individually or collectively, in order to make decisions and synchronize their actions, conflicts inevitably emerge at the actuation level as well, a topic that only recently started to attract attention by current research in smart spaces, as we see in later sections and especially in Section 3.5.

3.2 Multi-Context Systems and Distributed Knowledge Representation

AmI systems are particularly peculiar in that they are inherently distributed and decentralized, and their components are supposed to act autonomously. The agents should be able to carry out their tasks in cooperation with other agents, but also independently, e.g., if communication, or perhaps other parts of the system are broken. Such assumptions pose increased requirements on knowledge processing, particularly reasoning, which was traditionally investigated especially for the single KB/single agent case. Although reasoning agents (e.g., based on BDI architecture (Rao and Georgeff, 1991, 1995)) were investigated in the context of distributed multi-agent systems (Wooldridge and Jennings, 1995; Jennings et al., 1998), the

Table 5 Context modelling approaches

Methodology	Uncertainty	Formal Semantics	Learning Capacity	Interoperability
Naive Bayes and Bayesian Networks (Lu and Fu, 2009; Wu et al., 2007) HMM (Singla et al., 2010; Jakkula et al., 2009) Case-based (Knox et al., 2010) Others (Neural Networks, Support Vector Machines, Suffix trees)	Yes, especially for low-level context	No	Yes	Limited
Ontology and rule-based (Riboni and Bettini, 2011b; Springer and Turhan, 2009) Classic Logic-based (Artikis et al., 2010; Mastrogiovanni et al., 2011; Rugnone et al., 2007; Sadri, 2010) Defeasible Logic (Pajares Ferrando and Onaindia, 2012; Bikakis and Antoniou, 2010a)	Limited, usually in the form of non-determinism	Yes	Very limited	Yes
Hybrid (Riboni and Bettini, 2011a; Roy et al., 2011; Augusto et al., 2008) Markov Logic Networks (Helaoui et al., 2011; Skarlatidis et al., 2011) Evidence Theory (McKeever et al., 2010; Sebbak et al., 2012; Hong et al., 2009) Constraint-based reasoning (Pecora et al., 2012; Cirillo et al., 2009)	Can handle quantitative and qualitative uncertainty	Some cases	Yes	Still limited

concern was usually about how should the agents process the newly acquired knowledge (e.g., possibly by revising their KB, see Section 3.3) and what should the resulting knowledge state of the agent be, upon which they would then act. However, it was not traditionally investigated what should the resulting knowledge state of the whole system be, and how the knowledge of one agent can influence the others, etc.

3.2.1 Multi-Context Systems

This interesting problem was pursued by Giunchiglia (1993), and Giunchiglia and Serafini (1994), who proposed Multi-Contexts Systems (MCS). In MCS, we deal with a collection of knowledge bases $\mathcal{K}_1, \dots, \mathcal{K}_n$. Each of the knowledge bases \mathcal{K}_i is a collection of formulae in its own language \mathcal{L}_i . The

knowledge bases of an MCS, commonly called *contexts*, represent different pieces of knowledge that are to be combined in a unified reasoning system. The contexts may represent different knowledge sources, knowledge bases of communicating and cooperating agents, etc. The issue of multi-contextuality is captured by MCS in various ways. Not only with different languages that the contexts may possibly use, but also with the fact that each context may be respective to a different situation and, therefore, may contain diverse facts. Each context may even represent similar information differently. The combination of these assumptions renders MCS very flexible in modelling scenarios of diverse levels of knowledge heterogeneity, from completely homogeneous, up to ones involving largely heterogeneous knowledge sources/agents.

Logical semantics of the contexts is assumed, in the sense that we have either entailment or a proof system by which we can derive when a formula $\phi \in \mathcal{L}_i$ is true in \mathcal{K}_i . The knowledge from different contexts is combined with *bridge rules* of the form $i : \phi \leftarrow j : \psi$, meaning that if the formula $\psi \in \mathcal{L}_j$ is true in \mathcal{K}_j then also the formula $\phi \in \mathcal{L}_i$ must be true in \mathcal{K}_i . That is, bridge rules allow to derive consequences in one context (target context) based on premises previously proven in some other context (source context). Hence, bridge rules allow to characterize knowledge transfer between contexts, but also to translate from the language used in one context to that of another, which may be necessary given their possible heterogeneity. That is, given the bridge rule $i : \phi \leftarrow j : \psi$, the recipient agent \mathcal{K}_i upon receipt of information ψ from the sender agent \mathcal{K}_j concludes ϕ , where ϕ represents the recipient's own representation and interpretation of the senders statement ψ . A more general form of bridge rules allows more assumptions coming from different contexts on the right hand side (e.g., $i : \phi \leftarrow j_1 : \psi_1, \dots, j_n : \psi_n$). Such rules fire if ψ_k is derived in \mathcal{K}_{j_k} for all j_k .

Inference in MCS was first characterized by a proof theory, where bridge rules are used as calculus rules, that are combined with the local calculus of each context (Giunchiglia, 1993; Giunchiglia and Serafini, 1994). A model theoretical semantics for MCS, called *local model semantics*, was introduced by Giunchiglia and Ghidini (1998; 2001). In this semantics, the model of a whole MCS is a collection of local models over which additional semantic constraints are asserted that are derived from the bridge rules. Local model semantics was particularly influential and formed the base for further research. We will now survey the main areas of research associated with MCS.

3.2.2 Distributed Logics and Distributed Ontologies

Due to their capacity to combine logical reasoning over multiple knowledge sources, MCS have been used to formalize Distributed First Order Logic (DFOL) (Ghidini and Serafini, 1998), and later Distributed Description Logic (DDL) (Borgida and Serafini, 2003). The latter approach, in particular, proved to be influential and sparked considerable interest in distributed ontologies, where the power of MCS is used either to make alignments between heterogeneous and possibly ambiguous ontologies that are to be combined in reasoning (Ghidini et al., 2008; Ghidini and Serafini, 2008), or to facilitate truly distributed inference (Serafini and Tamin, 2004; Serafini et al., 2005; Homola and Serafini, 2010). Both directions may be useful in AmI systems, in cases when more than one ontology is employed within a system, possibly governed by independent agents. Other related approaches to distributed ontologies include \mathcal{E} -connections (Kutz et al., 2002, 2003; Cuenca Grau et al., 2004), Context OWL (Bouquet et al., 2004), Integrated Distributed Description Logics, and Package-based Description Logics (Bao et al., 2009). For a comparison of their expressive power, we refer the interested reader to Homola (2010, chap. 6).

3.2.3 Non-Monotonic Multi-Context Systems

Logic-based multi-agent systems often rely on non-monotonic logics, in which the agents are able to reason with assumption, and derive new consequences from assumptions as long as they are not disproven. In order to plug non-monotonic contexts into MCS, it was desirable to enable also non-monotonic bridge rules. Such rules are of the form $i : \phi \leftarrow j_1 : \psi_1, \dots, j_k : \psi_k, \mathbf{not} \ k+1 : \psi_{k+1}, \dots, \mathbf{not} \ l : \psi_l$, and they allow to assert consequences in some context also based on the fact that some evidence is *not* proven in a source context of the bridge rule. For example, consider a situation in which the control agent is instructed to switch lights on during the night if a person is present and switch them off if a person is not present,

using knowledge from the calendar and detector agents. While the former case is easily captured by a monotonic bridge rule (1), the latter is not; we need a non-monotonic bridge rule (2) for that:

$$ctrl : lights_on \leftarrow cal : night, det : preson_present \quad (1)$$

$$ctrl : lights_off \leftarrow cal : night, \mathbf{not} det : preson_present \quad (2)$$

Local model semantics was not sufficient to handle non-monotonic bridge rules. First steps towards such extensions were taken by Roelofsen and Serafini (2005) and Brewka et al. (2007), but the semantics which later became generally accepted as a de-facto standard for non-monotonic MCSs is the equilibrium semantics given by Brewka and Eiter (2007).

The expressive power of MCS is further increased in Managed Multi-Context Systems (mMCS) Brewka et al. (2011). While so far we dealt with bridge rules, which always result in addition of a formula into the target context, mMCS introduce new operations, such as deletion of a formula, revision by a formula (in the sense of belief revision, cf. Section 3.3).

3.2.4 Contextual Knowledge Representation

While bridge rules allow to change and accommodate information transferred between contexts in arbitrary way, they do not suggest how this should be done. The reason is that MCS enable to connect knowledge bases pertaining to different contexts, but do not provide any apparent means to capture the characteristics of these differing contexts. Already in his early work on contextual reasoning, McCarthy (1993) described transfer of information between contexts as knowledge lifting. This operation, which in MCS is implemented with bridge rules, was also studied under other names, such as knowledge push and pop (Benerecetti et al., 2000). It is understood that the knowledge is changed or adjusted during the transfer, in order to fit into the target context. What is more, these changes are influenced by the characteristics of the source and the target context, sometimes also called contextual meta knowledge. This meta knowledge may refer, e.g., to a particular location, period of time, topic, etc., associated with a context. Once such meta information is assigned to contexts, contextual relations between them are studied, e.g., one context preceding another in time, or, is associated with a broader topic, and so on. Thanks to these relations, contexts can be organized into a contextual space (Lenat, 1998; Benerecetti et al., 2001). Thus, for instance, the statement $President(Bill_Clinton)$ associated with the context of US in year 2000 may be changed to $ExPresident(Bill_Clinton)$ when lifted into another context associated with some future period of time.

Contextualized Knowledge Repositories (CKR) (Serafini and Homola, 2012) can be seen as extension of MCS that addresses this issue. In CKR, users may initialize a number of contextual dimensions, with respective values and their relations. Such dimensional values are then assigned to contexts as a form of meta knowledge. Thus, we can have a context associated with, e.g., US politics 2000, similarly as illustrated above. CKR relies on the mechanism of knowledge importing, which enables to reuse the knowledge of a context in another one. For example, in any context we can access the predicate $President_{US,politics,2000}()$, which will import relevant instances respective to $President()$ from the context of US politics in 2000. Or to give a very simple example from the AmI domain, consider that we divide a smart home into separate rooms and we use a single dimension $location$ which takes values, such as $kitchen$, $livingroom$, etc. In the context of the living room, we can access the relevant users present in the kitchen as instances of $UserPresent_{kitchen}$. This way we transfer the knowledge from one context to another. This way, the user does not deal directly with bridge rules. Further versions of CKR extend the framework with more expressive meta theories than just simple dimensional properties (Bozzato and Serafini, 2013) and also with non-monotonic defeasible reasoning capabilities (Bozzato et al., 2014). Similarly, Description Logics of Context (Klarman and Gutiérrez-Basulto, To appear) allow to model a set of contexts and a meta theory that describes their relations, and information between contexts is then accessed using dedicated modal operators.

CKR and similar formalisms may particularly be useful to AmI applications to develop agents which need to combine numerous amounts of knowledge imported from various sources. This can be information from sensors and other agents, or external knowledge available in the form of linked data datasets from the

Web. Each piece of information can be associated with respective contextual metadata and then seamlessly combined in reasoning.

3.2.5 Conflict Resolution and Argumentation in MCS

Apart from resolving sensory input conflicts, MCS can potentially be applied to resolve any types of conflicts that may arise between the agents in AmI environments. However, their main efficacy lies with resolving or avoiding contextual conflicts between the agents, as documented by the studies of Serafini and Homola (2012); Bikakis and Antoniou (2010a); Benerecetti et al. (2000); Ghidini et al. (2008). It is also apparent from the foundational works that MCS were built upon (Lenat, 1998).

MCS immediately allow for localized conflict resolution, i.e., if a context \mathcal{K}_i imports mutually conflicting information from some other contexts $\mathcal{K}_j, \mathcal{K}_l$, this can be resolved within \mathcal{K}_i . For instance, we may choose to prefer the information from \mathcal{K}_j and neglect the one from \mathcal{K}_l , or vice versa, or we may decide to ignore it entirely, or to react in some other appropriate way.

A global view on inconsistency handling in MCS was studied by Eiter et al. (2010b) who looked at MCS systems which have no equilibrium and propose possible explanations why this happens. Confining local inconsistencies and preventing them from polluting the entire system was also one of the design goals of DDL (Serafini et al., 2005) and CKR (Serafini and Homola, 2012).

The problem with localized conflict resolution is that two separate, autonomous agents may face the same conflict differently, choosing two different resolutions and act upon them. This may possibly disturb the overall coordination of agents in the system. Negotiating about conflicting knowledge between autonomous entities has been studied in the argumentation theory (see Section 3.5). Combining MCS with argumentation therefore seems to be a particularly promising direction in this respect. Already Parsons et al. (1998) propose to use argumentation within an MCS-based agent architecture in order to resolve conflicts that arise between agents. More recently, Bikakis and Antoniou (2010a) study an application of MCS and argumentation in the context of AmI systems. They proposed so called Contextual Defeasible Logic, built as an MCS with defeasible logic used inside contexts that uses argumentation to resolve conflicts in each context. Such an approach, however, is still localized: each context resolves the conflicts locally, based on its local preferences. Brewka and Eiter (2009) introduce Argumentation Context Systems. This approach takes further steps towards reaching a certain level of agreement between the agents, in order to resolve mutual conflicts: MCS are enriched with so called mediators, which import relevant information from other contexts using bridge rules that resolve any conflicts relying on an argumentation semantics.

3.2.6 Towards Applicability of MCS in AmI

Some of the systems described above were also developed into working prototypes. Support for \mathcal{E} -connections is implemented in the Pellet reasoner (Sirin et al., 2007). A distributed reasoner prototype for DDL was released under the name DRAGO (Serafini and Tamin, 2005). It enables to combine and reason with OWL ontologies with expressive power up to \mathcal{SHIQ} DL (Horrocks et al., 2000). It was developed as an extension of the Pellet reasoner (Sirin et al., 2007). The DDL framework and the DRAGO system were consecutively employed in tasks such as debugging of automatically generated ontology mappings (Meilicke et al., 2007) and also providing reasoning support for manual ontology mapping revisions (Meilicke et al., 2009).

An implementation of an RDF-based CKR was showed by Joseph and Serafini (2011). It is an extension of the OWLIM semantic data store (Bishop et al., 2011). A use-case based exploration of practical applicability of contextualized modelling with CKR was published by Bozzato et al. (2013).

A working prototype of an MCS system was developed by Bögl et al. (2010), in order to demonstrate the method for finding explanations for inconsistency in MCS by Eiter et al. (2010b). This implementation is based on the tool named dlhex (Eiter et al., 2006).

Regarding the question of what role should MCS actually play in AmI applications, they were proposed as basis of agent architectures. Parsons et al. (1998), Casali et al. (2005), and Sabater et al. (2002) use MCS to develop the internal architecture of an agent. While Parsons et al. and Casali et al. essentially

Table 6 MCS: Comparison of approaches

	Local Language	Interconnection	Non-monotonic	Practicality
MCS	various	bridge rules	no	–
DFOL	FOL	bridge rules	no	–
Distributed ontologies	DL/OWL	bridge rules links imports	no	reasoners available
Contextual representations	DL	imports modalities	some	reasoners available
MCS/equilibrium semantics	various	bridge rules	yes	implemented and evaluated
mMCS	various	bridge rules with operations	yes	–
argMCS	various	mediators	yes	–
Contextual Defeasible Logic	defeasible logic	bridge rules	yes	implemented and evaluated
poss-MCS	possibilistic LP	bridge rules	yes	–

design an MCS implementation of the BDI architecture, Sabater et al. propose a more elaborate, modular architecture which extends BDI. Other works (Cimatti and Serafini, 1995; Bikakis and Antoniou, 2010a) use the notion of context to encapsulate an agent and exploit bridge-rules to implement inter-agent exchange of knowledge. This enables to reason about the knowledge distributed in the whole MAS. While such an approach can be used within each agent to model its perception of the other agents' knowledge, Parsons et al. (1998) envisaged that bridge rules could potentially be exploited also for capturing the actual inter-agent communication and thus the MCS would take the role of a fully fledged multi-agent architecture.

Above, we highlighted the application of MCS in resolving conflicts, especially when referring to contextual conflicts. Recently, work has been done towards integrating also normative reasoning into MCS (Knorr et al., 2014), which can further enhance the capabilities of AmI systems relying upon the MCS architecture.

As mentioned above, Bikakis and Antoniou (2010a) studied Contextual Defeasible Logic, which is capable of conflict resolution using argumentation. Basing the approach on MCS, they yield a fully distributed contextual platform suitable for AmI applications. The authors compared their solution with a centralized one (Bikakis and Antoniou, 2010b). This approach was also experimentally evaluated by an implementation (Bikakis et al., 2011), in which various conflict resolution strategies were studied. The communication complexity of these strategies was evaluated, shown to be ranging from polynomial to exponential, depending on their expressiveness. Further implementation details were given by Moawad et al. (2013).

Yet another extension to MCS that may be potentially useful in AmI applications as it allows to deal with uncertainty is the possibilistic extension of MCS (dubbed poss-MCS, Jin et al. (2012)). This work is still preliminary, offering a basic framework with an equilibria-based semantics and an initial decision algorithm. The authors envisage work on more efficient algorithms in the future.

The works from the area of multi-context systems and distributed knowledge representation are summarized in Table 6, where we compare them with respect to the local language of the contexts, type of interconnection between contexts, whether they are non-monotonic or not, and whether a practical implementation is known.

3.3 Belief Change

3.3.1 Classical Belief Change

Belief change is important in the context of conflict resolution, because it prevents inconsistencies (and therefore conflicts) from arising, by taking appropriate actions during changes, i.e., during the process of adding into the KB the new knowledge that came from sensors or other agents. Thus, following the classification of Resendes et al. (2014), belief change falls under the “conflict avoidance” type. A recent survey of the field (Fermé and Hansson, 2011) shows that the problem is very challenging, both from the theoretical and the practical perspective.

In terms of the classification of conflict types appearing in Section 2.3, one could say that belief change can, in principle, deal with all types. However, belief change approaches are most suitable for contextual conflicts, as they were designed to deal with conflicts in agents’ models and beliefs. Some types of conflicts, namely goal and action conflicts, would require special variations or applications of belief change theories into languages that they were not in their original focus, e.g., BDI models or action languages, whereas others (e.g., sensory input and domain and background knowledge conflicts) may require assumptions that are usually considered in subfields of belief change (rather than the mainstream), such as the assumptions considered in non-prioritized belief change. Further issues that restrict the applicability of belief change theories in AmI environments are provided in the analysis below.

As far as the theoretical perspective is concerned, the main challenges arise from the fact that it is often difficult, even in toy examples, to identify the appropriate result of a change operation, and several philosophical considerations need to be taken into account (such as the debate related to the adoption of coherence or the foundational viewpoint (Gärdenfors, 1990), the use of belief sets or belief bases (Hansson, 1991), the differences between static-world and dynamic-world changes (Katsuno and Mendelzon, 1991), the semantics of minimal change (Fuhrmann, 1991; Hansson, 1996; Makinson, 1987), and others). As already mentioned, most of the belief change field is focusing on understanding and resolving those challenges, i.e., identifying the optimal way to resolve conflicts (logical inconsistencies) arising due to the introduction of new, conflicting information.

From the practical perspective, the main challenges are related to the fact that belief change algorithms are usually intractable. This is partly attributed to the types of languages considered (propositional and first-order logic) whose reasoning problems are intractable (at best). However, intractability should also be attributed to the inherent difficulty of the problem itself and the difficult challenges that it poses. Unfortunately, most of the works in belief change are not dealing with the practical aspects of the problem.

All the works related to belief change can be broadly classified in three different categories: *postulation approaches*, *construction approaches* and *explicit definition of concrete operators*. These are defined in detail below:

- *Postulation* amounts to defining a set of formal requirements (postulates) that determine whether any given operator behaves in a “rational” manner. Depending on the application at hand, the engineer can choose the exact semantics of the operator to use; as long as the postulates are satisfied, the operator is guaranteed to produce rational results and have certain desirable properties. Not surprisingly, there is no single set of postulates that works well for all cases (Flouris et al., 2004; Ribeiro et al., 2013), even though there are some proposals that are more widely accepted than others (Alchourrón et al., 1985; Hansson, 1991).
- A set of postulates tells us the desired properties of a rational change operator, but gives us no clue as to how to construct one. The latter is the role of *construction approaches*, which essentially define a construction methodology for change operators satisfying a particular set of postulates. At the heart of such approaches is usually a representation theorem that proves that the constructed family coincides with the family of operators that satisfy the intended postulates.
- The *explicit definition of concrete operators* is a more direct approach to belief change, where a specific change operator is provided, often for use in a certain application or context. Such operators usually employ some kind of application-specific hard-coded or parameterized methodology to

define the result, as this usually involves extra-logical considerations. In addition, such operators are often shown to satisfy certain postulates or are based on some construction method.

The most seminal work on belief change is that of Alchourrón et al. (1985), a postulation attempt for the operators of contraction and revision. These postulates are often referred to as *the AGM postulates* by the initials of the authors. The AGM postulates formed the cornerstone of later approaches on belief change. Alchourrón et al. considered three operators: *expansion*, which is the trivial addition of information without regards to possible inconsistencies that could arise; *revision*, which deals with adding information consistently; and *contraction*, which deals with removing information.

These operators deal with the assumption of a static world, i.e., in cases where a new observation, experiment or other information forces us to change our conceptualization of the world; the world itself does not change, but our modelling of the world does. On the other hand, under the dynamic world assumption a belief change operation is caused by a change in the world itself; in this case, there is nothing wrong with our original conceptualization, but the world itself is evolving and we need to keep our conceptualization up-to-date. These two settings have different semantics, so another pair of operations (*update* and *erasure*), along with a set of postulates, were defined by Katsuno and Mendelzon (1991). These are the dynamic counterparts of revision and contraction respectively. Note that the dynamic setting is more relevant for the AmI domain.

The presence of postulates allowed to formally show a number of interesting results for these operators. In particular, the operations of contraction and revision, were shown to be interdefinable in the presence of the postulates (Alchourrón et al., 1985). Further results showed that update and erasure are also interdefinable, and revealed connections between the static-world operations and their dynamic-world counterparts (Katsuno and Mendelzon, 1991). In most contexts, including the AmI context, revision and update are the most relevant operators; however, due to contraction/erasure being simpler, and in the light of the above results, most works in the literature deal with contraction.

Obviously, the intuition behind the AGM postulates is not valid for all settings. The most controversial postulate in the AGM set was the postulate of recovery, which captures the informal *principle of minimal change* for contraction; this principle states that change operators should have the minimal possible effect (or “impact”) on the original KB. The recovery postulate was criticized as non-intuitive by several authors (Fuhrmann, 1991; Hansson, 1996), and its status was the subject of several debates (e.g., Makinson, 1987). Alternative postulates were proposed, the most prominent one being the postulate of relevance (Hansson, 1991), which captured the intuition of minimality in a different way. Surprisingly, relevance, despite being proposed as a more intuitive alternative to recovery, was formally shown to be equivalent to recovery in the presence of the other postulates under the assumptions considered by the AGM work (Hansson, 1991).

All the above works (and most of the works related to belief change in general) are dealing with *prioritized belief change*, i.e., they assume that the new information is unconditionally accepted (an assumption known as the *principle of primacy of new information* (Dalal, 1988) or the *principle of success* (Alchourrón et al., 1985)), and this is also captured in one of the AGM postulates. The effects of dropping this assumption were studied in the subfield of *non-prioritized belief change* (Hansson, 1997; Hansson et al., 2001; Shakarian et al., 2014). Non-prioritized belief change is important for the AmI setting, where the cause of a conflict may be found in the input, e.g., a faulty sensor reading (sensory input conflict), and not in the agent’s KB.

Most construction approaches are based on the AGM postulates, in the sense that they show that the resulting family of operators coincides with the family of operators satisfying the AGM postulates. One such construction was provided by Alchourrón et al. themselves in their original paper (Alchourrón et al., 1985), but others were proposed as well (Gärdenfors and Makinson, 1988; Grove, 1988; Alchourrón and Makinson, 1985; Hansson, 1994; Rott, 1992).

Works that explicitly propose a concrete belief change operator are rather scarce in the belief change literature. Unfortunately, this makes the application of belief revision methods to practical domains (like AmI) more difficult. Works that propose an explicit belief change operator are based on the idea of “closeness” between (sets of) models: they view a KB as a set of models (i.e., those that the KB satisfies),

and the result of a change application (e.g., contraction or revision) is the KB (i.e., set of models) that satisfies the required postulates, while being the “closest” to the set of models satisfied by the original KB. The difference in these works stems from the different definition of “closeness” between models. Two of the most important concrete operators that have been proposed are those of Dalal (1988) and Chou and Winslett (1994).

Another family of works, known as Truth Maintenance Systems, provide explicit operators via a step-wise, formula-based approach, where a set of rules determines the facts to be added/removed from the KB in each step, towards reaching a KB satisfying a set of conditions (i.e., requirements that correspond to the considered postulates). Eventually, a state is reached where no more facts need to be added/removed to achieve the required properties, at which point the result is returned (Doyle, 1979).

Note that most of the above works are not touching practical issues such as efficient implementation of the related algorithms. This is an inherent problem of belief change methods, as they are dealing with logics which are intractable at best. However, most of the employed techniques are also problematic when used in less expressive logics, as they involve identifying minimal sets of formulae that cause conflicts and selecting one of them in an optimal manner, a process that is also (usually) intractable.

An additional drawback of belief change techniques, as related to their applicability in the AmI setting, is the fact that they do not deal with distributed, multi-agent settings, but consider scenarios where a single agent autonomously collects information from its environment and incorporates some information in its own KB, without regards to the existence (or not) of other cooperating (or competing) agents.

For these reasons, belief change techniques were only rarely considered in AmI settings so far, e.g., by Bosse and Sharpanskykh (2010). Nevertheless, we argue that belief change techniques should be viewed (and used) for what they offer, namely a robust understanding of the process of change and evolution (which includes conflict resolution as an integral process) and a rich set of theoretical results that describe this process. Under this light, belief change literature could be re-used to understand and describe the intricacies of the conflict resolution process in AmI settings, but this would require revisiting existing belief change approaches under this prism.

3.3.2 Belief Change in Semantic Web and Other Non-Classical Logics

The AGM approach, as well as most belief change approaches, are based on some relatively strong assumptions regarding the underlying language; this essentially limits their applicability to KBs represented using the so-called *classical logics*, which basically amount to propositional and first-order logic. For a complete list of these assumptions, as well as their effects on the supported languages, see the works of Alchourrón et al. (1985), Ribeiro et al. (2013), Creignou et al. (2014), Cuenca Grau et al. (2012).

However, changes also happen in different settings, where other knowledge representation languages are used. For example, we could mention handling of changes in logic programming (Lloyd, 1984), multi-context systems (Giunchiglia, 1993), horn logic (Horn, 1951), or datasets based on Semantic Web languages (such as Description Logics (Baader et al., 2003) or OWL (OWL Working Group, 2009)).

For addressing dynamicity in logic programming, different approaches have been considered, some of which consider a non-standard set of postulates that is more suitable for the characteristics of logic programming and its different variants (Leite and Pereira, 1998; Alferes et al., 2000; Leite, 2002; Shakarian et al., 2014). In multi-context systems, different variants have been proposed addressing changes in the knowledge itself (Gonçalves et al., 2014b), or in the corresponding bridge rules (Gonçalves et al., 2014a). The field of belief change for horn KBs has been addressed in various papers (Delgrande and Wassermann, 2010; Langlois et al., 2008; Booth et al., 2011, 2009; Delgrande, 2008; Zhuang and Pagnucco, 2010, 2012; Adaricheva et al., 2011), where most approaches are again trying to adapt belief change ideas to apply for the reduced expressiveness of horn KBs. A more general approach appears in Creignou et al. (2014); in that work, the authors assume any given fragment of propositional logic where the models of the formulas are closed under a Boolean function (such as horn logic), and propose a general approach for adapting standard operators (e.g., the revision operators presented in (Dalal, 1988), (Chou and Winslett, 1994)) in a way that the result of the revision remains in the fragment under consideration.

The case of changes in Semantic Web datasets and ontologies is much more relevant for this survey, as there is an increasing volume of works that employ Semantic Web languages to address AmI-related problems and/or exploit data in the Linked Data cloud for AmI applications (Celino et al., 2012; Emaldi et al., 2012; Lécué et al., 2012; Daly et al., 2013). For this reason, this subsection is mostly focusing on the dynamics of Semantic Web datasets. In that context, the problem has been addressed in the field of ontology evolution, where it has been argued that the adaptation of belief change methods and ideas in ontology evolution would provide several benefits (Flouris and Plexousakis, 2006).

The idea of applying belief change theories in the ontological setting was introduced in a series of works that studied the feasibility and consequences of applying the AGM postulates in the ontological setting (Flouris, 2006b,a; Flouris et al., 2006a; Flouris and Plexousakis, 2006; Flouris et al., 2004, 2005). Even though the AGM postulates can be easily reformulated to apply for ontological languages, it so happens that most DLs are not closed with respect to updates, in the sense that one can find examples where none of the “expected results” (per the postulates) is expressible in the underlying DL.

Subsequent work by the same authors proposed new postulates, like *optimal recovery* (Flouris et al., 2006b) or a generalized form of relevance (Ribeiro et al., 2013), which share most of the good properties of the standard AGM postulates, while being more widely applicable. The latter (generalized relevance) was shown to be applicable for a large class of logics, which includes all compact logics (Ribeiro et al., 2013). Given that most Semantic Web languages are compact, this work is very relevant for the AmI domain (and non-classical logics in general). Other works provided further insights on why certain Semantic Web languages cannot comply with belief change methods (that were developed for classical logics), resulting in the so-called *inexpressibility results* (Cuenca Grau et al., 2012). Other similar negative results appear in (De Giacomo et al., 2007; Liu et al., 2006; Calvanese et al., 2010).

These results motivated the search for ways to circumvent this problem. One approach used approximation techniques, i.e., evolution approaches resulting to an ontology whose set of models is as close as possible to the desired one (De Giacomo et al., 2007; Wang et al., 2010; Qi et al., 2015). Others chose to develop new DLs which provably avoid such problems (De Giacomo et al., 2009; Liu et al., 2006).

Some works adopt a more direct approach by proposing specific operators (inspired by belief change ideas), which are applicable for certain DLs. For example, Lee and Meyer (2004) deal with ontologies represented in the \mathcal{ALU} DL fragment; OWL ontologies are handled by Halaschek-Wiener and Katz (2006); Qi and Du (2009) propose three different revision operators for DLs; and Ribeiro et al. (2007) deal in general with knowledge representation formalisms that do not support negation (making it applicable to RDF/S ontologies, as well as ontologies represented using certain DL fragments).

The maxi-adjustment algorithm (Benferhat et al., 2004), is an approach for repairing inconsistencies in stratified propositional KBs in a minimal manner; the works by Qi et al. (2006b,a), based on this approach, develop evolution algorithms that guarantee the validity of the result in the context of stratified ontologies. Note however that this line of work assumes that ontologies are expressed using disjunctive DLs (Meyer et al., 2005), an extension of standard DLs that supports disjunction of axioms.

Gutierrez et al. (2006) consider the operator of erasure for RDF/S ontologies. Due to the simplicity of the underlying language, the main problem considered by Gutierrez et al. (2006) is how to prevent the removed triple from reappearing in the ontology as the result of RDFS entailment. The approach of Gutierrez et al. (2006) addresses this problem using a technique inspired by belief revision.

The resolution of conflicts in the Semantic Web languages is a critical task for the AmI setting, as more and more works are employing such languages to address AmI-related problems. In addition, the wealth of information existing in the Semantic Web (as Linked Open Data), is increasingly being exploited in various AmI applications, especially in the context of Smart Cities (Celino et al., 2012; Emaldi et al., 2012; Lécué et al., 2012; Daly et al., 2013). As a result, resolving the conflicts that appear in the underlying data, represented using Semantic Web languages, will become increasingly important, and one possible approach in this direction is the application of belief change technologies in such languages (as advocated by the works presented in this subsection).

The works related to the generalization of belief change approaches to Semantic Web (or other) languages, are mainly focusing on the feasibility of such an application, and are thus not concerned with

the tractability properties of the corresponding approach. More work is needed in this respect to verify that these approaches can scale when applied to practical situations. Therefore, as with classical belief change approaches, one should view the works presented here as a means to understand the process of change and conflict resolution in representation languages that are useful in AmI settings.

3.4 Ontologies and Belief Change

3.4.1 Ontology Evolution

Ontology evolution deals with the process of modifying an ontology in response to a certain change in the domain or its conceptualization (Flouris et al., 2008), and its main objective is to prevent conflicts from appearing in the ontology during the evolution process (where the term “ontology” refers to both the data and the schema). Thereby ontology evolution falls under conflict avoidance with respect to the classification of conflicts given by Resendes et al. (2014). Given the popularity of ontology-based methods in the AmI field, ontology evolution is highly relevant for this survey, as it could be directly adapted for conflict resolution in smart spaces. Recent surveys on ontology evolution were done by Flouris et al. (2008) and Zablith et al. (2015).

As with belief change, all the conflict types appearing in Section 2.3 are relevant for ontology evolution. However, since many of the approaches are dealing with the data part of the ontology, they are only applicable to sensory input and contextual conflicts, but some of the more recent works are also dealing with the schema part, making them applicable for the other types of knowledge conflicts as well.

In ontology evolution, two different types of conflicts are considered, namely *incoherency* and *inconsistency*. Incoherency appears when a class is unsatisfiable (Flouris et al., 2006a). Inconsistency is closer to the notion of logical inconsistency and appears when an ontology has no models (Flouris et al., 2006a).

In early approaches to ontology evolution, the application of changes upon ontologies was performed manually by the editor/curator using ontology editors (e.g., Protégé (Noy et al., 2006, 2000), OilEd (Bechhofer et al., 2001)) and reasoners used to pinpoint conflicts. Later on, more specialized tools appeared, which can identify the changes to be performed to guarantee validity, possibly with some user interaction. User interaction may be direct, through an intuitive interface (e.g., Lam et al., 2005b), or indirect through parameters, like evolution strategies (e.g., Stojanovic et al., 2002). Examples of such tools are KAON (Gabel et al., 2004), OntoStudio (formerly OntoEdit, Sure et al., 2003), and ReTax++ (Lam et al., 2005a,b). It is obvious that such approaches cannot be applied in the AmI setting, because it is assumed that agents should resolve conflicts (and inconsistencies/incoherencies) in an automated manner.

RUL (Magiridou et al., 2005) is a declarative language for data updating in RDF/S ontologies, which takes into account RDFS semantics, as well as a fixed set of constraints on the resulting RDF/S ontology. For every change requested by the user, the language automatically checks whether the application of said change would cause any problems related to the above constraints (taking into account RDFS semantics), and, if so, it automatically adds further changes (side-effects) to guarantee that the end result will have no conflicts.

In EvoPat (Rieß et al., 2010), the identification of conflicts is performed using SPARQL queries; each conflict is associated with one or more SPARQL Update statements that resolve it. The same idea of identifying conflicting patterns (in various ways) and resolving them (in a user-defined way, or using some hard-coded, predetermined process) was employed in various works (e.g., Djedidi and Aufaure, 2009, 2010; Liu et al., 2006; Roger et al., 2002).

A formal method for applying changes in the presence of custom validity rules was proposed by Konstantinidis et al. (2008a,b) and Flouris et al. (2013), where the incorporation of changes is performed automatically, taking care that the validity rules are not violated at the end of the process (see also the discussion on invalidity in Section 3.4.2).

Ontology evolution approaches are, by conception, meant to be applied in real settings where ontologies are used, and often the intended application area is the Semantic Web. As a result, scalability and tractability is generally an objective for these approaches, and sometimes applicability and formal rigour are sacrificed to achieve good tractability properties. The main drawback for many evolution

approaches is the fact that they rely on manual or semi-automatic processes, which makes them unsuitable for the AmI setting. Therefore, further research efforts towards a fully automated ontology evolution process would be highly relevant for AmI.

Table 7 shows the works related to the evolution of ontologies and includes works presented in this section, as well as related works presented in Section 3.3.2. The referenced works have been grouped according to their properties. The second column in the table mentions the language(s) that the corresponding works consider (OWL, RDF/S, various types of DLs, or arbitrary logics). The third column specifies the properties that the corresponding method respects during evolution: some works support the definition of custom (provided by the user) constraints to be respected during evolution, whereas others focus on preserving coherence and/or consistency. Finally, the fourth column determines the resolution method, which can be manual, semi-automatic or automatic; approximate; or inspired by/reusing other methodologies/fields, such as belief change approaches or the maxi-adjustment algorithm.

Table 7 Summary of ontology evolution approaches

Referenced Work(s)	Supported Language	Properties Considered	Resolution Method
Protégé (Noy et al., 2006, 2000) OilEd (Bechhofer et al., 2001)	OWL	Custom	Manual (Editors)
KAON (Gabel et al., 2004) OntoStudio (Sure et al., 2003) ReTax++ (Lam et al., 2005a,b)	OWL	Coherence Consistency	Semi-automatic
EvoPat (Rieß et al., 2010) (Konstantinidis et al., 2008a,b; Flouris et al., 2013) (Djedidi and Aufaure, 2009, 2010)	RDF/S	Custom	Automatic
RUL (Magiridou et al., 2005)	RDF/S (Data Only)	Custom	Automatic
(Liu et al., 2006; Roger et al., 2002)	DL	Coherence Consistency	Automatic
(Lee and Meyer, 2004)	<i>ACU</i> DL	Consistency	Belief change
(Halaschek-Wiener and Katz, 2006)	OWL	Consistency	Belief change
(Ribeiro et al., 2007)	Logics Without Negation	Consistency	Belief change
(Gutierrez et al., 2006)	RDF/S	Consistency	Belief change
(Flouris, 2006b,a; Flouris et al., 2006a; Flouris and Plexousakis, 2006; Flouris et al., 2004, 2005, 2006b) (Ribeiro et al., 2013) (Cuenca Grau et al., 2012)	General (Arbitrary Logic)	Consistency	Belief change
(De Giacomo et al., 2007) (Wang et al., 2010) (De Giacomo et al., 2009)	DL	Consistency	Approximate
(Qi et al., 2006b,a)	Disjunctive DL (Stratified)	Consistency	Maxi-adjustment
(Qi and Du, 2009)	DL	Consistency	Belief change

3.4.2 Ontology Debugging

The field of ontology debugging addresses conflicts after they have already appeared in the KB (cf. Section 2.3). In contrast to ontology evolution, the reason that caused the conflict is unknown (or

irrelevant) in this field, i.e., it is not considered during the resolution of the conflict. Works in ontology debugging are not only dealing with inconsistencies and incoherencies, but also with *invalidities*, which are violations of one or more custom validity rules that express application- or domain-specific requirements on the underlying ontology (Roussakis et al., 2011). For a related survey see the one by Flouris et al. (2008).

Several recent works have acknowledged the need for imposing custom, application-specific or user-defined requirements (in the form of validity rules) upon ontologies (Lausen et al., 2008; Motik et al., 2009; Serfiotis et al., 2005; Tao et al., 2010). Thus, identifying and resolving cases where an ontology violates the imposed requirements, either after a reckless change or for other reasons, is paramount for the seamless functionality of the associated applications. Such validity rules are also important for smart spaces, where agents can employ commonsense background knowledge (in the form of rules) to improve their performance in supporting the user in the smart space, for context recognition, or to overrule unreasonable data (e.g., sensor readings); such rules should be respected by the agents' KB at all times.

Validity rules are often encoded as part of the ontological schema (e.g., as OWL rules (Horrocks et al., 2005)); however, in some cases the ontological language is not rich enough to encode the necessary rules, so another "rule layer" is considered on top of the ontology, encoded in some more expressive logical language. In both cases, an important invariant in ontology debugging is that the rules are considered fixed and do not change. Thus, in the former case (rules in the schema) ontology debugging only applies changes in the data part (instance level) of the ontology to resolve a conflict, whereas in the latter (rules in an external layer) it may affect both the schema and the data. Due to this invariant, ontology debugging is not suitable for resolving conflicts related to the rule level, which typically encodes (parts of) the domain and background knowledge.

Regardless of the type of conflicts considered (coherency, consistency or custom validity rules), ontology debugging comprises of two different stages. The first, called *diagnosis* refers to the identification of the conflicts, as well as the possible causes behind such conflicts; the second, called *repair*, refers to the determination of the best way to resolve the identified conflicts. Essentially, diagnosis provides the necessary information for repair to resolve the conflicts.

Standard reasoners are of little help for the task of diagnosis, because, even though they can identify the existence of a contradiction, they provide little support for resolving and eliminating it (cf. Section 2.3, the solvability dimension (Resendes et al., 2014)). On the other hand, manual identification of the sources of a conflict (contradiction) is not feasible, especially in a smart space setting. Therefore, a more powerful approach is required in order to identify the part(s) of the ontology that led to the contradiction (Flouris et al., 2008).

Repairing is even more difficult, because, in addition to identifying the causes of a conflict (diagnosis), one must determine the "optimal" (under some measure of optimality) way to resolve such a conflict. This process is very similar to the process of identifying the "minimal change" in the belief change/ontology evolution context, and often requires some kind of user feedback, as the choice involves non-logical considerations. Due to this fact, most of the works related to the field of ontology debugging actually deal with the problem of diagnosis only, leaving the problem of repairing to human experts. However, this is not enough for most AmI applications.

Many approaches use some tableau-based algorithm for diagnosis. One of the most influential approaches was given by Schlobach and Cornet (2003), where a tableau-based algorithm for identifying the causes of an incoherency for a specific DL was presented. Similar tableaux-based algorithms for diagnosis were also proposed (Plessers and Troyer, 2006; Meyer et al., 2006; Wang et al., 2005). In all these techniques, diagnosis reports the axioms responsible for a conflict; a more fine-grained approach would be to identify the parts of the axioms that are responsible for the conflict (see the works of Kalyanpur et al. (2006) and Lam et al. (2008) for such approaches).

As already mentioned, the process of repairing usually employs some kind of user interaction. In some cases (e.g., in ontology editors such as Protégé (Noy et al., 2006, 2000), or in the case of reasoners, such as Stardog <http://stardog.com/>, or QUONTO (Acciarri et al., 2005)), this interaction is direct, i.e., the user is presented with the conflicts (possibly with some support regarding the results of the diagnosis)

and asked to manually resolve them. In the ORE tool (Lehmann and Bühmann, 2010) a similar interactive approach is used, where the system presents the user with a set of suggestions for resolving the conflict. Such approaches are not useful in the AmI context, where agents need to decide how to resolve conflicts in an automated manner.

Automated approaches for repairing either employ ad-hoc solutions for resolving conflicts (e.g., (Masotti et al., 2011; Chortis and Flouris, 2015), or take advantage of some kind of implicit user feedback (e.g., (Roussakis et al., 2011; Qi and Pan, 2007; Meyer et al., 2005)). For example, Qi and Pan (2007); Meyer et al. (2005) take into account external information related to the stratification of knowledge to identify the optimal resolution option, whereas Roussakis et al. (2011) relies on user feedback that is provided at input time via a set of user-defined “preferences”. These preferences act as high-level declarative specifications for the “ideal” repair, based on which the system is able to automatically determine the optimal resolution of conflicts in order to produce a repair that is as close as possible to the “ideal” one, as specified by the preferences. The same technique, using preferences based on metadata (such as trust or provenance) was applied in a real setting by Flouris et al. (2012).

In an effort to achieve scalability, some works consider ontological languages with a limited expressiveness, such as $DL-Lite_{\mathcal{A}}$ (a tractable DL). Examples include the works of Masotti et al. (2011); Chortis and Flouris (2015) both of which provide algorithms for automated repairing of $DL-Lite_{\mathcal{A}}$ KBs.

A rather original approach for repairing, proposed by Moguillansky et al. (2008), employs ideas from argumentation frameworks to identify and resolve conflicts. In particular, a conflict is defined as an “attack” between arguments (which can be easily identified using logical reasoning), whereas repairing consists in determining accepted/rejected ontological axioms based on the standard acceptability semantics of argumentation frameworks. This approach can be used both for ontology evolution and for ontology debugging.

An alternative approach to repairing, originally proposed for relational databases, but also applied in the ontological setting, is *consistent query answering* (Arenas et al., 1999). Consistent query answering advocates the use of query rewriting techniques that allow “hiding” the conflict from the user by providing consistent answers over a conflicting dataset. This is an indirect approach to repairing, as the conflicts are allowed to persist in the dataset, rather than being resolved.

In the context of relational databases, consistent query answering approaches have been proposed in various works dealing with different classes of constraints, mainly key constraints (e.g., (Grieco et al., 2005; Wijssen, 2009; Chomicki and Marcinkowski, 2005; Haase et al., 2005)). Recently, there has also been some relevant research for knowledge bases expressed in DLs, such as the works of Lembo et al. (2010, 2011), which deal with different variants of inconsistency-tolerant semantics to reach a good compromise between expressive power of the semantics and computational complexity of inconsistency-tolerant query answering. In (Martinez et al., 2014) the $Datalog^{\pm}$ language (Cali et al., 2009) is extended with defeasible semantics. This essentially allows the incorporation of inconsistency tolerant semantics for $Datalog^{\pm}$ ontologies.

Ontology debugging approaches are also concerned with the scalability properties of the proposed algorithms. In most of the works presented (both for repairing and for consistent query answering), one can find experimental results, as well as theoretical analyses of their computational complexity. Of course, the scalability of approaches for ontology debugging is greatly depending on the expressiveness of the considered underlying language and integrity constraints; when considering expressive DLs or expressive integrity constraints, the problem of diagnosis/repair is inherently intractable.

Ontology debugging is very relevant to the AmI setting where agents should make sure that their KBs satisfy the imposed validity rules at all times; this is especially relevant for recognizing the context and for reacting appropriately to its changes. For the same reason, ontology debugging is mostly useful for contextual conflicts, but also for the other types of conflicts which are identifiable through a set of rationality constraints (validity rules) based on the background knowledge about the domain.

Table 8 shows the works related to the problem of ontology debugging, grouped according to their properties. The table shows the supported language for the input KB per case, the type of problem

considered (diagnosis, diagnosis and repair, or consistent query answering), as well as some details on the methodology used (manual, semi-automatic, automatic etc.).

Table 8 Summary of ontology debugging approaches

Referenced Work(s)	Supported Language	Problem Considered	Methodology
Protégé (Noy et al., 2006, 2000) Stardog (http://stardog.com/)	OWL	Diagnosis Repair	Manual (Editors and reasoners)
QUONTO (Acciarri et al., 2005)	<i>DL-Lite_A</i>	Diagnosis Repair	Manual (Editors and reasoners)
(Lehmann and Bühmann, 2010)	OWL	Diagnosis Repair	Semi-automatic
(Plessers and Troyer, 2006; Meyer et al., 2006; Wang et al., 2005) (Kalyanpur et al., 2006; Lam et al., 2008)	DL	Diagnosis	Tableaux-based
(Qi and Pan, 2007; Meyer et al., 2005)	DL	Diagnosis Repair	Automatic (Stratification)
(Roussakis et al., 2011) (Flouris et al., 2012)	RDF/S	Diagnosis Repair	Automatic (Preferences)
(Masotti et al., 2011) (Chortis and Flouris, 2015)	<i>DL-Lite_A</i>	Diagnosis Repair	Automatic
(Moguillansky et al., 2008)	<i>ALC DL</i>	Diagnosis Repair	Automatic (Argumentation)
(Grieco et al., 2005; Wijzen, 2009; Chomicki and Marcinkowski, 2005; Haase et al., 2005)	Relational	Consistent query answering	Consistent query answering
(Lembo et al., 2011, 2010)	DL	Consistent query answering	Consistent query answering
(Martinez et al., 2014)	Datalog [±]	Consistent query answering	Consistent query answering

3.5 Argumentation

Argumentation is nowadays a very popular conflict resolution approach. Its semantics can be adapted to both centralized and decentralized distributed settings and some solvers have already been implemented, making it relevant also for the AmI domain.

The research on argumentation covers a wide range of disciplines: from psychology, philosophy and social sciences in general, to cognitive science and AI (Rahwan and Simari, 2009; Besnard and Hunter, 2008; Prakken and Vreeswijk, 2002). In the latter in particular, the focus of relevant research is devoted to formal models of argumentation. One of the main challenges is to design a formal system that enjoys desirable semantic properties and tractable computational complexity, while being theoretically easy to understand.

In this section, we overview existing approaches and discuss how argumentation can be suitable for dealing with conflicting information in AmI environment.

Formal models of argumentation can be divided according to whether they focus on a specific logical language and the structure of arguments, or not. Thus, we usually distinguish between *abstract* and *structured* argumentation (Prakken and Vreeswijk, 2002).

3.5.1 Abstract Argumentation

The most influential work on abstract argumentation is by Dung (1995), where *abstract argumentation frameworks* (AFs) have been introduced. Abstract argumentation does not consider any structure of arguments or conditions defining the conflict (attack) between arguments.

An *argument* in abstract argumentation is an atomic term that is understood as anything that a rational agent can argue with/about and the *attack* between arguments is an arbitrary binary relation representing any inconsistency between these arguments. The advantages of this abstract approach are simple elegant semantics and generality. The main issue in abstract argumentation is to determine which arguments are accepted, which are rejected, and which are left undecided. Generally, if an argument a attacks an argument b then arguments a and b cannot be accepted together, and b is rejected whenever a is accepted.




Figure 3 Nodes denote arguments and arrows denote the attack relation (e.g., a attacks b)

Intuitively, an argumentation semantics prescribes a set of sets of arguments, called extensions, to each argumentation framework. Different semantics have been proposed by Dung (1995), based on the notion of *admissibility*, and several of them have been defined with different motivations in mind. An argument a is *defended* by a set of arguments S if S attacks all arguments attacking a . An extension is said to be *admissible* if it is *conflict-free* (i.e., it does not contain any pair of attacking arguments) and defends all its members. For example, given the AF shown in Figure 3, the admissible extensions are \emptyset , $\{a\}$, $\{c\}$, $\{d\}$, $\{a, c\}$, and $\{a, d\}$.

In our example, not all admissible sets are relevant as an intended meaning of the argumentation framework. Various semantics based on admissibility have been proposed by Dung (1995) and others, with the aim to capture the notion of relevant extensions. For example, admissible extensions $\{c\}$, $\{d\}$ do not provide intuitive meaning, since they do not contain all the arguments they defend (argument a namely). The semantics requiring such property is called *complete*. There are three complete extensions $\{a\}$, $\{a, c\}$, $\{a, d\}$ for the given AF. As it can be seen, complete extensions can be included in one another resulting in two different semantics: *skeptical* and *credulous*. The *skeptical* semantics is the most careful one, where arguments cannot be defended by themselves. The extension corresponding to the *skeptical* semantics is called *grounded* and we can see that $\{a\}$ is the grounded extension. On the other hand, there are two credulously accepted extensions $\{a, c\}$, $\{a, d\}$ called *preferred*. The last semantics proposed by Dung (1995) is called *stable* and requires all arguments outside of an extension to be attacked by an argument in the extension. The only stable extension of our AF is $\{a, d\}$. Note, however, that stable semantics is not always defined – there are AFs with no stable extensions, such as the AF consisting of only one self-attacking argument. As sometimes this is unsatisfactory, a weaker semantics called *semi-stable* have been defined (Caminada et al., 2012) which is universally defined and coincides with the stable semantics if stable extensions exist.

Given an arbitrary complete extension E , a justification status *in*, *out*, *undec* can be assigned to arguments (Caminada and Gabbay, 2009). Intuitively, an argument is:

- *in* if it is defended by E ,
- *out* if it is attacked by E ,
- *undec* otherwise.

Arguments that are *in* are included in the extension and are always accepted. Rejected arguments (status *out*) are attacked by the extension and, consequently, are never accepted. Finally, the acceptance of undecided arguments (status *undec*) is not determined, since they are attacked by undecided arguments only.

Argumentation can be described as a discussion between two players: a proponent trying to justify his claim and an opponent who counterargues. The inclusion of an argument in some of the semantics

described above can be decided procedurally by creating and evaluating such discussions, called *argument games*. Intuitively, a player wins a discussion if she has the last word.

Dung (1995) also studied the relationship of argumentation with other non-monotonic formalisms, such as default logic (Reiter, 1980), inductive defeasible system (Pollock, 1995), stable (Gelfond and Lifschitz, 1988) and well-founded (Van Gelder et al., 1988) semantics for logic programming. Furthermore, significant research has been conducted on studying proof theories (Modgil and Caminada, 2009), and complexity analysis (Dunne and Wooldridge, 2009) for the varying semantics. In Table 9 we summarize the complexity results of Dunne and Wooldridge (2009), for different abstract argumentation semantics and for both the credulous and the skeptical case (the trailing “-c” means that the given problem is complete for the given class).

Various extensions of abstract AFs have been studied. For instance, Dung’s AF only features the single relation attack between arguments – the attack relation. But when a_1 attacks a_2 , and a_2 attacks a_3 , then in a sense a_1 indirectly and implicitly supports a_3 . However, as argued by Cayrol and Lagasquie-Schiex (2009), in certain application domains also explicit support relations between arguments may be present. To capture this, they present Bipolar AFs featuring both attack and support.

Abstract Dialectical Frameworks (ADF), a generalization of abstract AFs, were introduced by Brewka and Woltran (2010). Acceptance conditions are added to arguments in ADF. They allow to specify different types of interactions between arguments. Both attack and support are expressible in ADF. ADFs are characterized by Brewka and Woltran (2010) in terms of dependency graphs together with acceptance conditions.

Assumption-Based Frameworks (ABF) were introduced in (Bondarenko et al., 1997). ABFs understand non-monotonic reasoning as a deduction from assumptions. The structure of ABF is an abstraction of well-known non-monotonic formalisms and it was shown that many of them are instances of ABF. Argumentation semantics were applied to (sub)sets of assumptions. It was shown that traditional semantic characterizations of non-monotonic theories, e.g., stable models of logic programs or extensions of default theories, correspond to argumentation semantic characterizations. Assumption-based argumentation (Dung et al., 2009) has been later developed over the basis of ABF.

Other extensions of abstract argumentation frameworks include the addition of preferences (Amgoud and Vesic, 2011), or weighted attacks (Dunne et al., 2011; Coste-Marquis et al., 2012) into the context of abstract argumentation. This is often required by various application domains, and therefore it considerably extends the flexibility and applicability of abstract argumentation.

Table 9 Complexity of abstract argumentation (Dunne and Wooldridge, 2009)

	Admissible	Grounded	Complete	Preferred	Stable
Credulous	NP-c	P	NP-c	NP-c	NP-c
Skeptical	trivial	P	P	Π_2^P -c	coNP-c

3.5.2 Structured Argumentation

For certain application areas Dung’s approach may be too abstract to be directly used in practice. Take for example the legal domain, where problems can be naturally formalized in a logical language comprising facts, implications, or rules, etc. In abstract argumentation, the arguments are just atomic “letters” but here the formulae are more complex and the notion of attack will depend on their structure and the semantics of the language.

The usual methodology in order to employ argumentation in such cases is to instantiate an abstract argumentation framework (e.g., Dung’s AF). This is often done by defining the following four notions: a logical language, the structure of an argument, the attack relation, and the status of an argument (Prakken and Vreeswijk, 2002). Each one of these notions can be expressed by means of the previous. The status of an argument depends on the notion of arguments and the attack relation. The attack relation depends on

the notion of arguments and the underlying logical language. The structure of arguments is defined with respect to the underlying logical language.

Once the instantiation is done, extensions may be computed treating the result as abstract AF (e.g., using one of Dung’s semantics). The semantics of the structured argumentation framework is then derived from the arguments present in the extensions.

As for the logical language, classical propositional language may be used (Besnard and Hunter, 2001). More often the language of Defeasible Logic Programming (DeLP) (García and Simari, 2004) or other logic programming or defeasible logic languages (Prakken, 2010; Governatori et al., 2004; Prakken and Sartor, 1997) are considered. If the system uses a language of classical logic (Besnard and Hunter, 2001) then we speak of *deductive argumentation*. On the other hand, the languages that feature both strict and defeasible rules are also called *defeasible theories* (Caminada and Amgoud, 2007).

The language of DeLP is particularly interesting, as its semantics is usually derived from argumentation frameworks. Two kinds of rules are distinguished in the language of DeLP: strict (\rightarrow) and defeasible (\Rightarrow). While strict rules are used to represent some kind of deductive reasoning (i.e., whenever the preconditions hold, we accept the conclusion), defeasible rules formalize tentative, uncertain knowledge, where the validity of the precondition of a rule usually (but not necessary) implies the validity of the head of the rule. Thus, defeasible rules can be defeated, in light of some conflicting evidence.

Let us demonstrate this on a very simple example from the intelligent household domain. We want to encode the following rules: the heating is generally turned on before the user usually arrives (for simplicity, assume that *arrival_time* becomes true at the right moment). However, if based on the events scheduled in the user’s calendar we know that the user is having a meeting in a remote location (for simplicity, *in_meeting* is true in such case) then the first rule should be overridden and the heating is not turned on). We may encode such rules in a DeLP program, specifically the first two rules below:

$$arrival_time \Rightarrow heating_on \quad (3)$$

$$in_meeting \rightarrow \neg heating_on \quad (4)$$

$$\rightarrow arrival_time \quad (5)$$

$$\rightarrow in_meeting \quad (6)$$

Rule (3) is defeasible, it represents knowledge that is usually valid. Rule (4) is strict, it represents knowledge that is always valid. There are also two (strict) facts (5–6) representing that both *arrival_time* and *in_meeting* are true in the current instance. In DeLP, arguments are actually constructed through the chaining of the rules. For example, in the program above, we may construct the following arguments: $A_1 = [\rightarrow arrival_time]$, $A_2 = [\rightarrow in_meeting]$, $A_3 = [A_1 \Rightarrow heating_on]$, and $A_4 = [A_1, A_2 \rightarrow \neg heating_on]$. We can see that arguments have an inherently recursive structure. Based on conflicting conclusions of the arguments we can construct the attack relation: there will be a single attack, A_4 attacking A_3 , intuitively due their conflicting conclusions and due to the fact that the strict rule overrides the defeasible one. Thus we have constructed an abstract AF and we can use any of the Dung’s semantics to compute the extensions and consequently also the meaning of the given DeLP. Taking, e.g., the grounded semantics, there is a single extension $E = \{A_1, A_2, A_4\}$ (in fact, in this simple case the complete, preferred, and stable semantics give also the very same result). Finally, we derive the semantics of the DeLP program as the conclusions of the arguments present in the extension, i.e., the set of literals $\{arrival_time, in_meeting, \neg heating_on\}$.

García and Simari (2004) deal with the language of defeasible logic. The semantics is determined by the set of literals, which is computed procedurally via an argument game. However, the argument game is not admissibility-based and therefore it departs from Dung’s semantics. It is interesting to note though that an online solver has been implemented.¹

On the other hand, Prakken and Sartor (1997) directly instantiate Dung’s AF. They define an argument as a sequence of derivations and apply argumentation games to compute the grounded semantics. According to the authors, several non-standard design decisions were motivated by the legal domain.

¹Online solver for DeLP: http://lidia.cs.uns.edu.ar/delp_client/

Table 10 Structured argumentation-based formalisms

	Language	Postulates	Complexity	Implementation
(Prakken and Sartor, 1997)	defeasible	unsatisfied	P	–
(Besnard and Hunter, 2001)	classical	irrelevant	PSPACE-c	(Efstathiou and Hunter, 2010)
(García and Simari, 2004)	defeasible	unsatisfied	P	Yes ¹
(Governatori et al., 2004)	defeasible	unsatisfied	P	(Rock, 2009; Aceto, 2010)
(Prakken, 2010)	defeasible	satisfied	P/Π_2^P -c	(Snaith and Reed, 2012)
(Baláž et al., 2014a,b, 2015)	defeasible	satisfied	P/Π_2^P -c	Transformation to GLP

A more recent work by Prakken (2010) introduces an argumentation framework with structured arguments, called ASPIC⁺. It is basically a framework for structured argumentation rather than a particular system, where a particular language and argument ordering (preferences on arguments used in conflict resolutions) is left to be determined by the user. The online Java system TOAST that implements ASPIC⁺ has been developed by Snaith and Reed (2012).

A different methodology was applied by Baláž et al. (2014a), where the status of an argument does not depend on attacks between arguments, but instead on attacks between conflict resolutions. Intuitively, a conflict resolution is a recipe describing how a conflict is resolved. Within the language containing defeasible rules and default literals a conflict may be resolved by either attacking a default literal or a defeasible rule.

Some approaches working with defeasible theories (García and Simari, 2004; Governatori et al., 2004; Prakken and Sartor, 1997) fail to meet the so called *rationality postulates*, as Caminada and Amgoud (2007) pointed out. Specifically, properties known as *indirect consistency* (i.e., conclusions of extensions and conclusions of the intersection of all extensions must be consistent) and *closure under strict rules* (i.e., conclusions of extensions are closed under strict rules) are violated. Violation of these postulates can result into justification of absurdities or incomplete results, where some conclusions are missing.

While ASPIC⁺ (Prakken, 2010) satisfies the aforementioned postulates, it relies on transposition of rules. As pointed out by and Baláž et al. (2014a), this is not in line with the intuitions of some non-monotonic rule-based formalisms, like logic programming, where the rules are treated as unidirectional. Hence, using such a framework with the aim to derive a semantics for DeLP may lead to some undesirable consequences. Baláž et al. (2014a, 2015) propose alternate approaches which avoid transposition and the associated problems while still satisfying the desired postulates. Both works extend ABF. Baláž et al. (2015) extend ABF with a formal notion of conflict resolution strategy (CRS), with the idea that this component should not be fixed or limited to a small number of strategies, but rather an input parameter that can change from one application to another. It is showed that the postulates are satisfied for any given input CRS.

In a related work, Baláž et al. (2014b) study the relationship between the argumentation semantics of DeLP (i.e., defeasible reasoning) and more traditional Generalized Logic Programs (i.e., default reasoning). They show a transformation by which it is possible to compute a number of the argumentation semantics of DeLP in terms of a transformed GLP program, relying on existing GLP semantics and solvers.

Table 10 summarizes several approaches and their properties with the focus on the underlying language (classical theory or some defeasible theory), the satisfaction of all relevant postulates, theoretical complexity, and whether there is a practical implementation of the given approach. If the reader is interested in further details regarding structured argumentation, we suggest to refer to the special issue of Argument & Computation dedicated to this topic (A&C 5, 1).

3.5.3 Applications in Ambient Intelligence

Argumentation has already been applied in AmI as a KR paradigm for dealing with both incomplete (partial) and inconsistent (contradictory) knowledge (Pajares Ferrando and Onaindia, 2012; Bikakis and Antoniou, 2010a; Moraitis and Spanoudakis, 2007; Muñoz et al., 2011; Muñoz Ortega et al., 2010; Muñoz et al., 2010). The approach presented by Muñoz et al. (2010), for instance, uses argumentation techniques, in order to tailor services to the preferences of multiple users that share the same resources (i.e., a TV set). An internal dialogue is structured whenever conflicting preferences arise.

Notable are also the studies of Pajares Ferrando and Onaindia (2012); Bikakis and Antoniou (2010a); Moraitis and Spanoudakis (2007) that respect the distributed nature of contextual information, where different entities possess locally a partial and tentative view of the actual world state. In these studies, defeasible rules are defined to represent uncertainty during context recognition, while techniques from the argumentation theory are applied in an attempt to resolve conflicts and reach a consensus about the actual context. Argumentation techniques are well tailored to resolving such types of *contextual* conflicts. In particular, Pajares Ferrando and Onaindia (2012) implemented and experimentally evaluated a DeLP Multi-Agent Partial Order Planning framework. Their objective is to choose a plan respecting both the desire to minimize the computational overhead and to maximize the quality of the solution plan. Bikakis and Antoniou (2010a) aimed primarily at representation and reasoning issues. They defined a contextual defeasible logic (MCS instantiated with defeasible logic in each of the contexts), providing a decentralized platform and a set of distributed algorithms for query evaluation. To resolve all possible conflicts, a total preference ordering on the system contexts is assumed. Further, Bikakis and Antoniou (2011) study also partial preferences as well as different conflict resolution techniques (ambiguity blocking, ambiguity propagation, team defeat).

Although not directly applied in an AmI setting, an interesting application of argumentation techniques was proposed by Leite and Martins (2011) in the Social Web area where the social voting determine arguments strength and consequently also the semantics of the system (valuation of all arguments).

As we may observe, applications in AmI will often require to consider certain levels of uncertainty, preference or weights on arguments and attacks. Therefore, the studies extending argumentation frameworks with these capabilities (as discussed above) are quite relevant to this domain. But also other relevant directions, e.g., investigating applications of argumentation for decision making (Amgoud and Prade, 2009; Amgoud and Vesic, 2012), may be relevant.

To conclude, argumentation is a well-investigated field and one of its advantages over the other KR formalisms is its user-friendliness; argumentation-based semantics can be intuitively explained for both researchers in the KR domain and people not familiar with formal logic.

3.6 Belief Change and Argumentation

In the previous sections we noted that the area of belief change may amplify AmI systems with the ability to deal with newly acquired information and to accommodate changes in the current situation. Moreover, the field of argumentation has contributed works that enable agents with contrary beliefs or goals, to achieve agreement.

In this section, we survey some works proposed recently, that investigate the interrelations between these two fields. This issue is gaining an increasing research attention: three recent events, Madeira Workshop on Belief Revision and Argumentation (Fermé et al., 2013), Luxembourg Workshop on the Dynamics of Argumentation, Rules and Conditionals (DARC 2012), and Dagstuhl Seminar on Belief Change and Argumentation in Multi-Agent Scenarios (Dix et al., 2013) focused on the relations of these fields.

Being a relatively novel research direction, it has mainly developed theoretical models so far. Even so, it can be considered relevant to the current survey as it may extend the state of the art in both fields of belief change and argumentation. In realistic multi-agent settings as is AmI, these two important capabilities of agents need to be combined to ensure that the changes in beliefs and goals are executed in a mutually compatible way.

3.6.1 General Considerations on Interrelation of Argumentation and Belief Change

Standard argumentation theory deals with a set of arguments and an attack relation. However, when argumentation becomes dynamic by adding or removing new arguments/attacks, interesting problems arise. Belief change methods are helpful in such situations. On the other hand, a variety of argumentation semantics exists, representing different views on compatible sets of arguments, which set the basis for a more flexible approach to belief change.

In works by Baroni et al. (2013) and Fermé et al. (2013), argumentation and belief change are considered as reasoning process and are thoroughly compared for commonalities and differences. They conclude that both fields capture partially distinct but overlapping research problems, and identify a number of interesting open research questions, posed by their comparison (e.g., they note the rising importance of postulates in argumentation which were long considered central in belief change, and call for proposals of “reasoning benchmarks” which could be used to evaluate different approaches in both fields).

According to Falappa et al. (2011), some argumentation formalisms can be used to define belief change operators, and belief change techniques have been used for modeling the dynamics of beliefs in such argumentation formalisms. Complementary roles of belief change and argumentation in understanding and modeling complex reasoning processes are stressed. The analysis of connections between argumentation and belief change within a complex reasoning process is based mainly on the ideas of Falappa et al. (2009). A complex reasoning process consists usually of the following basic reasoning steps: (1) reception of new information, (2) evaluation of it, (3) change of beliefs, and (4) inference. Basically, argumentation can be used mainly in step 2 and belief change can be used in step 3. However, a more detailed analysis shows that there are complex interrelations between argumentation and belief change within the different reasoning steps.

Rotstein et al. (2010) propose Dynamic Argumentation Framework (DAF), in which a new feature, *evidence*, is introduced. Evidence enables to distinguish valid arguments. Change is represented at different levels: change on evidence, on arguments, on conflicts, or on a preference relation.

The works surveyed so far, study the general connections between belief change and argumentation. Next, more technical approaches follow, categorized in two directions: those, applying methods of belief change to argumentation and those, that use the argumentation viewpoint to introduce some new features about belief change.

3.6.2 Belief Change Applied to Argumentation

Computational aspects of argumentation frameworks updating are studied by Liao et al. (2011). A modular approach to updates is implemented as follows: if an update operation is specified for a given argumentation framework, the updated argumentation framework is divided into three parts: arguments affected by the update, arguments unaffected, and conditioning arguments. The latter are unaffected arguments, which attack affected arguments.

The role of conditioning arguments is essential from the computational point of view. Computation of the status of arguments can be divided in two parts. The status of unaffected arguments is the same as in the original argumentation framework; it is not changed by the update. The status of affected arguments is computed in a Conditioned Argumentation Framework (CAF), where attacks of conditioning arguments against affected arguments influence the status of the second. An algorithm implementing this method is described.

The next works focus on elementary change operations in abstract argumentation frameworks. They provide a view on the change of basic components of argumentation frameworks. Different basic change operations are considered.

Cayrol et al. (2010) defined four basic change operations on argumentation frameworks: adding an attack between arguments, removing an attack, adding an argument together with attacks involving it and removing an argument together with the involved attacks. The case of adding one argument is studied in details. The main focus is on the impact that the changes cause on the structure of extensions and on the status of some particular arguments.

Boella et al. (2009b) are focused on impact from adding attack relations to the semantics of an abstract argumentation framework. The work is continued in Boella et al. (2009a), where the removal of attacks and arguments is studied. Only the case of a semantics with precisely one extension is considered. The main focus is on principles for the argumentation dynamics.

Coste-Marquis et al. (2014) and Maily (2013) study revision of attack relations in argumentation as minimal change of the arguments status. The principle of minimal change plays an important role in the belief change research: it states that it is appropriate to preserve as much from the given knowledge set as possible. It is further shown how AGM belief revision postulates (Alchourrón et al., 1985) can be translated to the case of argumentation systems.

A different approach was undertaken by Baumann and Brewka (2010; 2012; 2013). So called *enforcing problem* was described and solved by Baumann and Brewka (2010). It poses the question whether it is possible, given a specific set of allowed operations, to modify a given argumentation framework \mathcal{A} into \mathcal{A}' such that a desired set of arguments E is contained in some extension of the modified AF. Some conditions, under which enforcements are possible, were identified.

An important special case of the enforcing problem – how to reach that goal by a minimal change – was studied by Baumann (2012). Given an argumentation framework A and a set of arguments E , the minimal number of additions or removals of attacks needed to reach an enforcement of E is called the *characteristic of E* . This number depends on the underlying semantics and the type of allowed modifications. It was shown that in certain cases there are local criteria, allowing to determine the characteristic. Local in the sense that the criteria are based on properties of the underlying argumentation framework and they enable to determine the characteristic in a finite number of steps.

The *spectrum problem* is studied by (Baumann and Brewka, 2013). Given a set of semantics and a modification type, the task is to determine for the pairs (σ, Φ) , where σ is a semantics and Φ is a modification type, the set of all natural numbers which are characteristics of arbitrarily argumentation framework A and a set of arguments E . Surprisingly, this rather abstract problem yields interesting insights into relations of stable, semi-stable and preferred semantics: it may be arbitrarily more difficult to enforce arguments using stable rather than semi-stable semantics, and also using semi-stable rather than preferred semantics.

Some researchers addressed problems connected to belief change in instantiated argumentation systems (structured argumentation frameworks with a subargument relation). Moguillansky et al. (2013) studied argumentation within Defeasible Logic Programming (DeLP). They defined prioritized argument revision operators for a given DeLP. The newly inserted argument becomes undefeated after the revision, hence its conclusion becomes warranted. In order to ensure this warrant, the program has to be changed in accordance with a minimal change principle.

3.6.3 Argumentation Applied to Belief Change

A relatively smaller part of research is devoted to this aspect of interrelations between belief change and argumentation. As already mentioned, Moguillansky et al. (2008) employed argumentation to belief change in ontologies (particularly, ontology debugging). They build an argumentation framework, in which mutually inconsistent ontological axioms attack each other, and argumentation semantics thus determines possible repairs.

Liao (2013) constructed a layered (abstract) argumentation framework with subargument relation (AFwS). The semantics of AFwS provide a basis for the study of updating a layered AFwS and its properties. Among motivations for this research is a scenario as follows. An argumentation component is put before a belief revision component. Suppose that a new piece of (updating) information is given. The argumentation component serves as a filter: a set of accepted arguments (and therefore, their conclusions) is obtained by means of argumentation. As a consequence, only new information justified by the argumentation process inputs into the process of belief change. This can serve as a contribution to standard implementations of belief change.

Krümpelmann et al. (2012) proposed a way how to distinguish whether new information should be accepted. Deductive argumentation (Besnard and Hunter, 2001) is used to assess the value of new

Table 11 Belief change and argumentation: Comparison of works

	BC of AF	BC by AF	Studied Problem
Baroni et al. (2013)	Yes	Yes	general comparison
(Falappa et al., 2009),(Falappa et al., 2011)	Yes	Yes	general comparison, inter-applicability
(Rotstein et al., 2010)	Yes		dynamic evidence-based argumentation
Conditioned AFs (Liao et al., 2011)	Yes		update of AF
Cayrol et al. (2010), Boella et al. (2009b,a)	Yes		addition/removal of attacks/arguments
Coste-Marquis et al. (2014), Maily (2013)	Yes		revision in AF, minimal change
Baumann and Brewka (2010; 2012; 2013)	Yes		enforcing and related problems
Moguillansky et al. (2013)	Yes		prioritized revision of arguments
Moguillansky et al. (2008)		Yes	ontology debugging
Liao (2013)		Yes	belief revision
Krumpelmann et al. (2012)		Yes	non-prioritized belief revision

information. Hereby is obtained a revision operator which accepts new information only if the new information is justifiable.

3.6.4 Summary

Understanding the mutual interrelations between belief change and argumentation, presents a contribution to both these fields, which we previously found relevant for AmI applications. In particular, it may contribute to develop more flexible argumentation frameworks, capable of updating, and thus reflecting the necessary changes that may be required by a change of a situation faced by a particular AmI application. On the other hand, it may contribute to the development of more flexible and more effective belief change operators implemented with argumentation procedures.

Given that this research direction is fairly new, most of the proposals are yet at the theoretical level, lacking reasoning support and implementations. We see this as a notable research challenge, especially for the KR community. Out of these theoretical works we would like to particularly highlight the notion of complex reasoning processes (Falappa et al., 2009), which highlights interactions between related reasoning tasks of a rational agent and which is fairly in line with the reasoning cycle of agents in AmI systems (cf. Figure 1). The only work with more practical implications in this area is that of Liao et al. (2011). In Table 11 the overviewed approaches are classified with respect to whether they are an application of belief change in argumentation, or an application of argumentation to belief change (columns “BC of AF” and “BC by AF,” respectively). Some of the approaches fall into both categories. The main focus of each work is also summarized.

3.7 Preferential Reasoning

The notion of preferences is part of the everyday human reasoning. For example, in case of two laws giving conflicting instructions, the instructions given by the law with more “power” precede. Another example is that doctors usually prefer non invasive procedures over invasive ones.

Almost every KR formalism was extended to support preferences. However, the term “preferences” is too abstract, and means slightly different things in different approaches. A common underlying intuition is that preferences select between multiple options. In this section, we focus our attention to logic programming, a widely used non-monotonic formalism. Logic programming uses if-then rules to express the knowledge about a domain. In contrast, e.g., with production systems, which also use if-then rules, logic programming is purely declarative. Logic programming is very relevant for AmI, as it is a generic knowledge representation formalism, which was applied in many areas, e.g., agent programming (Köster

et al., 2009), assisted living (Mileo et al., 2008a,b) decision support (Nogueira et al., 2001), diagnosis (Balduccini and Gelfond, 2003), multi-agent planning (Son et al., 2009), planning (Dimopoulos et al., 1997), policies (Son and Lobo, 2001). For additional references to applications we refer the reader to Schaub (2011). For the survey of preference handling approaches in other non-monotonic formalisms we refer the reader to Delgrande et al. (2004b). In the context of logic programming, preferences are used in the following ways:

- Preferences on rules are used to control the applicability of rules. Having two conflicting rules that are both applicable, we use the preferred one.
- Preferences on literals are used to prefer answer sets containing some literals over answer sets containing others.

3.7.1 Preferences on Rules

Consider we have a rule encoding that an agent should execute an action A , and a second rule encoding that the agent should not execute the action A . If the “if” parts of both rules are satisfied, we have a conflict. This conflict can be resolved using preferences and preferential reasoning. Assuming that we prefer the second rule, e.g., because it is based on more specific information, preferential reasoning makes the first rule inapplicable whenever the second one is applicable.

One of the standard semantics for logic programming is the answer set semantics (Gelfond and Lifschitz, 1991). It assigns to a logic program a collection of answer sets, each corresponding to an alternative set of beliefs an agent can accept. The so called *selective* preference handling approaches, studied, e.g., by Brewka and Eiter (1999), Wang et al. (2000), Delgrande et al. (2003), Zhang and Foo (1997), Sakama and Inoue (2000), Šefranek (2008), and Šimko (2013), select a subset of standard answer sets as *preferred*. They do so in order to stay compatible with the answer set semantics, instead of inventing a completely new semantics.

The approaches studied by Delgrande et al. (2003), Wang et al. (2000), and Brewka and Eiter (1999) can be characterized as *prescriptive* (Delgrande et al., 2004b). Preferences on rules are interpreted as the order, in which rules are applied. As a consequence, a less preferred rule cannot defeat a preferred rule. Each of the approaches puts slightly different conditions on the order in which rules have to be applied. Schaub and Wang (2003) showed that the approaches form a hierarchy. The biggest difference between the approaches is that the one by Brewka and Eiter (1999) handles only direct conflicts, while those by Delgrande et al. (2003) and Wang et al. (2000) handle indirect conflicts. Delgrande et al. (2003), Eiter et al. (2003a), Grell et al. (2005), and Asuncion and Zhang (2009) deal with the issue of computing the semantics.

On the opposite side of prescriptive approaches lie *descriptive* (Delgrande et al., 2004b) approaches. They do not see preferences as the order of application of rules, instead preferences are handled in a more declarative fashion. Zhang and Foo (1997) view preference handling as a removal of less preferred rules. Sakama and Inoue (2000) define preference handling as a comparison of the rules that generate answer sets. Šefranek (2008), Šefranek and Šimko (2011), and Šefranek and Šimko (2013) look at preference handling as a form of argumentation. Šimko (2013) uses preferences to transform conflicting rules into rules defining exceptions: a preferred rule defines exception to a less preferred rule, but not the other way around.

The computational complexity of various decision problems for the aforementioned semantics has been studied. While for some semantics the decision problems are polynomially decidable, the decision problems for the remaining approaches are NP-complete, or lie on a higher level of the polynomial hierarchy. This is expressed by the values “P” and “NP-c/worse” in Table 12.

Besides the answer set semantics, the well-founded semantics is the second standard semantics for logic programs. It can be computed in polynomial time. Brewka (1996), Schaub and Wang (2002), and Wang et al. (2000) defined preferred well-founded semantics for logic programs with preferences, which can also be computed in polynomial time.

Table 12 Preferential reasoning: Preferences on rules

	Type	Underlying semantics	Complexity	Practicality
(Brewka and Eiter, 1999)	prescriptive	answer set	NP-c/worse	reduction
(Delgrande et al., 2003)	prescriptive	answer set	NP-c/worse	native
(Wang et al., 2000)	prescriptive	answer set	NP-c/worse	reduction
(Zhang and Foo, 1997)	descriptive	answer set	NP-c/worse	–
(Sakama and Inoue, 2000)	descriptive	answer set	NP-c/worse	algorithm
(Šefránek, 2008; Šefránek and Šimko, 2011, 2013)	descriptive	answer set	NP-c/worse	–
Šimko (2013)	descriptive	answer set	NP-c/worse	reduction
(Brewka, 1996)	prescriptive	well-founded	P	–
(Schaub and Wang, 2002)	prescriptive	well-founded	P	–
(Wang et al., 2000)	prescriptive	well-founded	P	–

Delgrande et al. (2003) and Zhang and Foo (1997) consider *dynamic preferences*. Preferences not only change the semantics of logic programs, they are also subject of reasoning.

The main shortcoming of the literature in the field is that it provides little or no insight into which semantics to use in a domain at hand. The study of principles for preferential reasoning can help to fill this gap. Some development was done by Brewka and Eiter (1999), Šefránek (2008), and Šefránek and Šimko (2013). Ideally, given an application domain, a suitable semantics is selected based on a subset of relevant principles. However, work in this direction is still needed. So far, existing principles do not sufficiently differentiate between the approaches and are unable to guide in selecting a right semantics for a task at hand.

Table 12 summarises the approaches for reasoning with preferences on rules. The values in the column “practicality” have the following meaning: (i) native – a solver with an algorithm specifically tailored for the approach is implemented, (ii) algorithm – there is an algorithm specifically tailored for the approach, but no implementation is available, (iii) reduction – reduction of an approach to logic programming without preferences exist.

3.7.2 Preferences on Literals

If the answer set programming methodology is used, a problem is encoded into a program in a way that the answer sets of the program correspond to the solutions of the problem, e.g., using answer set programming for planning, each answer set of the program corresponds to a valid plan. In some situations, having multiple solutions means a conflict, e.g., as we can execute only one plan of actions, we need to select one of the valid plans. While there is always an option of selecting a random plan, we usually have additional preferences that need to be taken into account. For example, if it is possible we want to select a plan containing non-destructive actions. Preferences on literals and preferential reasoning allow us to express and reason with exactly this kind of preferences. Usually, an order on answer sets is computed based on the preferences, and maximal answer sets with respect to the order are selected.

Sakama and Inoue (2000) extend logic programs with a preference relation on literals. The relation is then transferred to a relation on answer sets and maximal answer sets are selected as preferred. In this way, preferred answer sets contain preferred literal. Sakama and Inoue also showed how preferences on literals can be applied to various forms of commonsense reasoning: minimal abduction, prioritized abduction, default reasoning, prioritized default reasoning, circumscription, and prioritized circumscription.

Brewka (2002) introduces logic programs with ordered disjunction, in which preference of a literal is given by its position in a disjunction. The intuition behind the rule with ordered disjunction $A \times B \leftarrow C$ is as follows. If C is contained in an answer set S , then A is in S if possible. But if it is not possible, then

Table 13 Preferential reasoning: Preferences on literals

	Preferences as	Complexity	Practicality
Sakama and Inoue (2000)	a relation on literals	NP-c/worse	algorithm
Brewka (2002)	rules with ordered disjunction in the head	NP-c/worse	algorithm
Brewka et al. (2003)	a preference program	NP-c/worse	algorithm

(at least) B is in S (Brewka, 2002). Brewka also shows how programs with ordered disjunction can be used in the configuration domain.

Brewka et al. (2003) consider answer set optimization programs consisting of two parts. The generating program produces answer sets representing solutions of a problem. The preference program expresses user preferences using the rules of the form $A_1 > \dots > A_n \leftarrow B_1, \dots, B_m$. A preference relation on the answer sets of the generating program is based on the degree in which the rules of the preference program are satisfied.

Complexity of various decision problems for the aforementioned semantics have been studied. The decision problems for all the reviewed approaches are NP-complete, or lie on a higher level of the polynomial hierarchy. This is expressed by the value “NP-c/worse” in Table 13.

Table 13 summarises the approaches for reasoning with preferences on rules. The column “practicality” contains a unique value “algorithm” – there is an algorithm specifically tailored for the approach.

Related to the aforementioned approaches are CP-nets (Boutilier et al., 1999, 2004). They deal with the problem of inferring preferences on vectors of feature values given conditional preferences on feature values. This problem can be very naturally expressed by all the reviewed approaches for preferences on literals. The benefits of these logic-programming-based approaches are: (i) the availability of default negation allows to represent defeasible information, and (ii) while CP-nets are concerned only with one fixed space of vectors, the answer set optimization programs have a substantial flexibility in defining the space (Brewka et al., 2003). The same holds also for the Sakama and Inoue’s preferences on literals, and logic programs with ordered disjunction. Brewka et al. (2003) also show that (i) the semantics of the CP-nets and the answer set optimization programs are different, mainly due to the different interpretations of the phrase “other things being equal” which lies at the heart of both semantics, and (ii) preferences inferred by CP-nets can be approximated by preferences inferred by answer set optimization programs.

Closely related to logic programming are relational databases, which are also relevant for AmI. It has been explored how to extend relational databases with preferences on database rows, which are preferences on literals in logic-programming terminology. For the survey of the approaches adding preferences to relational databases we refer the reader to Stefanidis et al. (2011). Another interesting approach is Lukasiewicz et al. (2013). It extends Datalog+/-, a family of ontology languages, with preferences. Syntactically the preferences are on literals, similar to that of Brewka et al.’s answer set optimization programs, but the semantics follows the database approaches.

3.7.3 Summary

We have given an overview of the approaches for preference handling in logic programming. In this section we give some pointers how the approaches can be used in AmI.

In Section 2.3, five types of conflicts with respect to the knowledge type dimension were introduced. With a suitable encoding, all conflicts seem to be solvable using preference handling. However, some might feel little unnatural, and probably other approach would be more appropriate.

Preferences on literals are especially suitable for handling conflicting goals and plans. If answer sets correspond to different goals/plans, we can select preferred ones. Preferences on literals can, e.g., prefer literals representing non-destructive actions over destructive ones, or actions using less expensive resources. For the issue of incorporating domain-specific preferences in planning systems we refer the reader to Delgrande et al. (2004a).

Preferences on rules are usable for handling both conflicts inside an agent and conflicts between agents. One way at looking at preferences on rules is as a handy way of encoding exceptions between rules. When writing logic programs, we usually use general rules with exceptions. Preferences on rules allow us to express exceptions in a more easily and change tolerant way. Preferences on rules were already used in AmI. Bikakis and Antoniou (2010a) use trust level of agents to determine preferences on conflicting rules, although they use a different formalism than we discuss in this section.

Regarding the applicability of the approaches in real environments, algorithms and prototypical solvers exists. However, additional work needs to be done as no production ready solver for preferences exists.

3.8 Paraconsistent Reasoning

One of the key features of AmI systems is the ability of agents operating within them to handle knowledge that originates from multiple sources, that may be incomplete, ambiguous, or even inconsistent. AmI systems are not supposed to halt or report errors when they face such problematic situations, instead they should be able to react to such situations appropriately, reconstructing and reusing the consistent and reliable parts of the knowledge at hand. In this section, we give an overview of *paraconsistent reasoning*, the area of KR that comprehensively addresses the problem of reasoning with inconsistent knowledge.

Inconsistency is a challenge for classical logic-based systems, as in classical logic meaningful reasoning is not possible once inconsistency arises. Once an inconsistency occurs, the *ex falso quodlibet* principle is applied, and then all formulae are derived as consequences. Therefore, studies in the area of paraconsistent reasoning focus on identifying the sources of inconsistency in the knowledge and on developing methods to constrain and isolate inconsistent knowledge, as well as consequences that are only supported by inconsistent knowledge and there is no way to derive them from consistent premises. Consequently, also the goal of proposing repairs and removing the inconsistencies from the affected knowledge sources is often considered (Bertossi et al., 2005).

Related goals are indeed addressed by a number of approaches that we already reviewed: belief revision, ontology repair, and argumentation. All these approaches aim at conflict resolution, or at least avoidance. In this section, we will focus on *paraconsistent logics* which, instead, focus on derivation of sound conclusions from knowledge that may possibly contain inconsistencies, without necessarily attempting to repair the knowledge base. Such approaches may be beneficial if an agent is acting upon the currently available knowledge, without the need to store it for later reuse.

3.8.1 Propositional Case

Paraconsistent logic can be split into several types (Bertossi et al., 2005): *signed systems*, which involve the renaming of literals and then restoring the non-conflicting part of the original theory by adding equivalences with their renamings; *weakly-negative logics* and *quasi-classical logic*, which use a restricted subset of classical proof theory, or rely on natural deduction to apply the proof rules more carefully, in order not to avoid the explosive derivation of all conclusions; and *multi-valued logics*, which employ a dedicated semantics, in which the truth and the falsity of each statement are considered independently.

Besnard and Schaub (1998) define a paraconsistent semantics for propositional theories by a signed system, in which they represent each positive literal a as a^+ and each negative literal $\neg a$ as a^- . They employ default logic (Reiter, 1980) which allows to interpret a^+ as a and a^- as $\neg a$, but only as long as the equivalence between a^+ and $\neg a^-$ can be assumed, that is as long as it does not cause the conflict between a^+ and a^- to occur. Thus, a consistent part of the original theory is effectively reconstructed. Unlike some other paraconsistent logics, in the consistent case the entailment within such a system coincides with classical entailment. Decision procedures for default logics (e.g., Junker and Konolige, 1990; Ben-Eliyahu and Dechter, 1991; Niemelä, 1995) can be used, although even in propositional case default entailment is known to be complete with respect to the second level of the polynomial hierarchy (Gottlob, 1992). This stream of development has largely gave way to (answer-set) logic programming, with relevant paraconsistent extensions being discussed in Section 3.8.2.

Besnard and Hunter (1995) and Hunter (2000b) proposed and developed quasi-classical logic, which they show to possess useful properties to reason with inconsistent knowledge. A number of interesting

applications have been studied (e.g., Hunter, 2000a; Byrne and Hunter, 2004), and the logic is known to be decidable Hunter (2000b), however no implementations are known.

The four-valued propositional logic was developed by Belnap (1977). This logic works independently with truth and falsity, yielding two new truth values of a statement (apart from the classical *true* and *false*), namely *unknown* and *inconsistent*. This approach was later generalized for more than four truth values (Ginsberg, 1988; Fitting, 1991a). Translations of this logic into first-order logic are known (Rodrigues and Russo, 1998), which enable the use of classical first-order provers, e.g., VAMPIRE (Riazanov and Voronkov, 2002), for reasoning.

Arieli and Denecker (2003) also investigate on paraconsistent propositional logic with multi-valued semantics and its translation to classical logic. In addition, they employ preferential reasoning (cf. Section 3.7), with the help of which they accept only models in which the inconsistent part of the knowledge base is minimized. They provide a polynomial translation, yielding a first-order theory with an addition of circumscriptive second-order formulae (McCarthy, 1980) representing the preferential criteria. This allows to use known translations of circumscriptive formulae (Ohlbach, 1996; Doherty et al., 1997; Gustafsson, 1996), finally yielding a first-order theory and again enabling to resort to classical first-order provers.

Besnard et al. (2005) encode various paraconsistent systems (maximal-consistent subsets, signed systems, and multi-valued approaches) into quantified propositional logic (QBF). This enables their comparison, but also the use of QBF solvers (e.g., Feldmann et al., 2000; Giunchiglia et al., 2001; Letz, 2002) for reasoning. This paves the way towards more practical applications, given the recent increased interest and developments in the area of QBF solvers.

3.8.2 Paraconsistent Logic Programs

The logic programming paradigm has been successfully applied in agent-based applications. Development of paraconsistent semantics for logic programs has been driven by the ability of agents (and other systems) to deal with situations, in which inconsistent information simply cannot be ruled out. Real-time applications, and distributed systems with autonomous entities and decentralized information sources fall under this category.

Paraconsistent semantics for logic programming was largely built on top of multi-valued logics (Belnap, 1977; Ginsberg, 1988; Fitting, 1991a). First, such semantics for logic programs was developed by Blair and Subrahmanian (1987, 1989); other early studies in this area include those of Fitting (1991b) and Kifer and Lozinskii (1992). All use four-valued semantics and only work with classical negation.

Sakama (1992) concentrates on extended logic programs (ELP, Gelfond and Lifschitz, 1991), which feature both classical and default negation, and possibly negation in the head. Sakama proposes a paraconsistent version of the well founded semantics (Van Gelder et al., 1991) for this class of logic programs, where in order to distinguish between the classical and the default negation he resorts to the seven-valued semantics of Ginsberg (1988).

This work is further extended by Sakama and Inoue (1995) who developed a paraconsistent stable-model semantics for extended disjunctive logic programs (EDLP). This semantics was also implemented on top of the MGTP reasoner (Inoue et al., 1992). Further evaluation of this semantics was done by Alcântara et al. (2004) and Odintsov and Pearce (2005).

An alternative paraconsistent version of the well-founded semantics for extended logic programs, dubbed $WFSX_p$, was proposed by Alferes et al. (1995). This semantics is based on the principles of coherence and introspection (Damásio and Pereira, 1995), ensuring, e.g., that the default negation of an atom (i.e., the weaker one) is always entailed from the classical negation of the same atom. A dedicated decision procedure called SLX, which uses a procedure similar to the standard PROLOG SLDNF procedure (Lloyd, 1984) was described and implemented (Alferes et al., 1995). Further extensions of $WFSX_p$ towards other common logic programming semantics are proposed by (Damásio and Pereira, 1995).

In the same paper discussed above, Sakama and Inoue (1995) proposed also the so called semi-stable semantics, which has the feature that it is able to derive consequences in cases when the classical stable-model semantics has no model, but coincides with it in cases when it has models. Such an approach, dubbed *paracoherent* is not intended to draw conclusions from truly inconsistent knowledge bases, but merely to overcome non-existence of models in some cases due to some rather technical reasons (e.g., cyclic dependencies). This line of work was further extended by Eiter et al. (2010a) who provided a model-theoretic characterization, and proposed several improvements. They show the complexity of reasoning, which is one level up in the polynomial hierarchy when compared to the classical stable-model semantics. Finally, they briefly describe a prototype implementation.

3.8.3 Other Paraconsistent Logics

Given the popularity of Semantic Web ontology languages such as RDF and OWL 2 in AmI applications, their paraconsistent variants may contribute to the ability of AmI systems to deal with inconsistent information sources. Four-valued paraconsistent extensions of description logics (i.e. the family of logics which provides the formal semantics for OWL 2) were already investigated by Patel-Schneider (1989) and Straccia (1997). The approach of Ma et al. (2007) who propose $\mathcal{ALC}4$, a four-valued extension of \mathcal{ALC} (cf. Baader et al., 2003), is particularly interesting to us as paraconsistent reasoning is obtained with no additional computational cost by reduction to classical description logic. In the follow-up work, Ma and Hitzler (2009) extended this approach towards the \mathcal{SROIQ} DL (Horrocks et al., 2006), reaching the full expressivity of OWL 2. They also investigate the tractable fragments of OWL 2. Their approach has been implemented into the RaDON plug-in (Ji et al., 2009) of the NeOn ontology engineering toolkit (Erdmann and Waterfeld, 2012).

A quasi-classical variant of description logic was developed by Zhang and Lin (2012). A tableau algorithm was also described by Zhang et al. (2009). These studies carry over the quasi-classical approach into the area of ontologies.

Hybrid KBs feature both ontologies and rules, a combination which is often required by practical applications, as ontologies carry the terminological knowledge, while rules are a more powerful tool for implementing reasoning. Paraconsistent reasoning has been investigated for each component separately, but recently also for hybrid KBs by Kaminski et al. (2015) who have investigated two multi-valued semantics in this context.

Another approach to inconsistency is based on estimating the amount of contradicting knowledge within a given information source or between different information sources with suitable inconsistency measures (Hunter and Konieczny, 2005). Such an approach may provide agents with more fine grained assessment of diverse information sources with respect to their degree of inconsistency in comparison to the current agent's belief and/or their degree of self-contradiction. Some sources may be rejected straight ahead, while some, more valuable sources, may be integrated with the agent's beliefs directly or by suitable conflict resolution methods. See the work of Hunter and Konieczny (2005) for an overview of this area.

3.8.4 Applicability of Paraconsistent Reasoning

Paraconsistent reasoning has found applications in databases (Arieli et al., 2004), including medical knowledge bases (da Costa and Subrahmanian, 1989), and its applications in spatial databases were also conceived (Rodríguez, 2005).

Paraconsistent reasoning has been further applied for inconsistency management in areas, such as software engineering, and on problems, such as combining inconsistent specifications (Hunter and Nuseibeh, 1998) and requirements (Ernst et al., 2012).

Quasi-classical logic has been applied in dealing with inconsistency in structured text excerpts (Hunter, 2000a), e.g., structured news reports (Byrne and Hunter, 2004). A similar approach may be valuable in AmI applications with the need to process textual inputs from multiple users. An interesting observation in these works is that the presence of inconsistency need not necessarily be an indicator of error; for instance, if multiple news reports are found inconsistent with the background knowledge, then this may be

Table 14 Paraconsistent reasoning: Comparison of approaches

	Language	Type	Practicality
Besnard and Schaub (1998)	propositional	signed	reduction to default logic
Belnap (1977); Arieli and Denecker (2003)	propositional	multi-valued	reduction to FOL
Arieli and Denecker (2003)	propositional	quasi-classical	decidable
Besnard et al. (2005)	propositional	multi-valued/signed	reduction to QBF
Blair and Subrahmanian (1987, 1989); Fitting (1991b); Kifer and Lozinskii (1992)	LP	multi-valued	–
Sakama (1992); Sakama and Inoue (1995)	EDLP	multi-valued	implemented
Alferes et al. (1995)	ELP	multi-valued	implemented
Sakama and Inoue (1995); Eiter et al. (2010a)	EDLP	multi-valued/paracoherent	implemented
Paraconsistent OWL (Ma et al., 2007)	DL	multi-valued	implemented
Paraconsistent hybrid KBs (Kaminski et al., 2015)	MKNF	multi-valued	–
Quasi-classical DL (Zhang and Lin, 2012)	DL	quasi-classical	reasoning algorithm

an indicator of interesting new developments in the domain, that have to be accommodated and processed (Byrne and Hunter, 2004).

We are not aware of any direct application or case study of paraconsistent reasoning in AmI. Nevertheless, these approaches can be applied to any type of conflict resolution, most notably sensory input conflicts, and also for resolving conflicts between new observations and background knowledge (Byrne and Hunter, 2004; Ernst et al., 2012).

As already noted above, current AmI applications already make use of ontologies. If inconsistency handling of ontological data becomes needed, paraconsistent OWL (Ma et al., 2007; Ma and Hitzler, 2009) may become handy (or alternatively some of the ontology debugging approaches surveyed in Section 3.4.2 may be applied).

While all the approaches surveyed in previous sections mostly aim at removing conflicts and repairing the knowledge base, paraconsistent reasoning shifts the point of view into simply being able to reason also with inconsistent knowledge, without necessarily requiring some kind of repair. This point of view may be useful in certain AmI scenarios, where agents need to react based on their inputs while maintaining the knowledge base is of secondary interest. Many of the logics which we surveyed have favourable (polynomial) reasoning complexity (Coste-Marquis and Marquis, 2005). Many of the paraconsistent logic programming extensions have developed into experimental implementations, and some of them are also known to be tractable (e.g., $WFSX_p$ (Alferes et al., 2003)). Further research may be however needed in order to make them effective enough for realtime AmI applications.

3.8.5 Summary

In this section we have discussed approaches that allow to reason also with inconsistent knowledge bases, evaluate them and draw conclusions from sound parts of the knowledge. Most of them are summarized and compared in Table 14 where we distinguish them based on the underlying representation language and different type of paraconsistent reasoning. We also point out practical results achieved in this area. Measures of inconsistency (Hunter and Konieczny, 2005, cf.) are not included as the main focus in this area is different from the rest of the surveyed works, as explained above.

4 Summary

In this work, we surveyed a number of research areas in KR that we believe to be relevant to the problem of conflict resolution, that, as we noted, is crucial in AmI, and in many other knowledge-intensive application scenarios.

The body of research concentrating on modelling context within AmI, also addresses conflict resolution to a certain extent. The main attention is given to resolving conflicts within context, originating, for instance, from faulty or incompatible sensor readings, or as a result of situation change. This part is well elaborated in the literature, and it is also efficiently handled, for example by resorting to hybrid techniques integrating KR and machine learning methods. The field, however, relies mostly on centralized context models, which may not be sufficient in real world AmI environments, as we discussed above. Therefore, this branch may largely benefit from cross-fertilization with the other KR branches discussed in the paper.

Multi-context systems (MCS) and related approaches allow to represent heterogeneous and interconnected systems composed of knowledge bases, each of which may be modelled in different language and from a different contextual perspective (therefore they are called contexts). As such, MCS allow to resolve conflicts that may arise between the contexts, although mostly in a localized fashion, i.e., each context is capable to resolve the conflicts locally and independently from the other contexts. Some recent efforts may help to make conflict resolution more shared (Brewka and Eiter, 2009) and its understanding global (Eiter et al., 2010b). In this respect, cross-fertilization between MCS and argumentation seems to be a promising approach (Brewka and Eiter, 2009; Bikakis and Antoniou, 2010a). Some studies related to MCS also considered resolving, or, at least isolating conflicts inside contexts; this can be seen as some limited way how to handle background or domain knowledge conflicts. The research in MCS has been advanced to the point, where multiple implementations and evaluations are known and even for experimental applications in AmI settings (Bikakis and Antoniou, 2010a).

Conflict resolution has long been studied in the area of belief revision, where knowledge bases are combined with newer, or more important knowledge, and conflicts need to be resolved in order to yield a consistent result. As such, belief change methodologies mostly fall under conflict avoidance, as they prevent the conflict from creeping into the knowledge base. The main body of research in this field focuses mostly on foundational research, trying to characterize suitable conflict resolution strategies with postulates and devising revision operators, that behave accordingly. The area may contribute to AmI as a foundation for suitable conflict resolution strategies to be applied in AmI systems. More practical approach is undertaken in ontology evolution and ontology debugging fields, where real algorithmic support, feasibility and effectiveness are considered relevant. Nevertheless, further research will be needed before the results can be applicable in real time, as most of the current methodologies, especially in ontology evolution, are semi-automatic and require human intervention. In ontology debugging several approaches are known, which go in the fully-automated direction (Qi and Pan, 2007; Meyer et al., 2005; Roussakis et al., 2011).

Argumentation is a representation technique that formalizes notions of argument, attack, and support, and enables to resolve conflicts by identifying sets of arguments that are consistent, and defending the set from external attacks. The notion of argument is rather abstract, which enables to resolve conflict between beliefs, goals, actions, etc. The increasing popularity of argumentation is given by its rich and flexible formalism with a family of semantics with well established theoretical properties, some of which also enjoy feasible complexity results. It was applied on a number of practical problems, and notably, also in AmI (Pajares Ferrando and Onaindia, 2012; Bikakis and Antoniou, 2010a; Moraitis and Spanoudakis, 2007), especially as a decision making technique for autonomous agents in the presence of conflicting information.

The position of argumentation as essential and effective conflict resolution technique is further assured by the fact that researchers from the other fields try to integrate argumentation into their approaches when conflicts need to be resolved. We noted such attempts in multi-context systems (Brewka and Eiter, 2009) and ontology debugging (Moguillansky et al., 2008), but the interchange between argumentation and belief change seems to be especially fruitful, as we documented in Section 3.6. This line of research

is only very recent, and mostly theoretical results were yet achieved, though we believe its possible future impact on practical applications, including AmI, is quite likely.

Preferences are sometimes combined with different reasoning formalisms, in order to select a rule to be applied in a given situation from a set of possibly conflicting rules, or in order to distinguish between multiple possible derivations. The former case can be seen as conflict avoiding, while the latter case is a more delicate indirect conflict resolution method, as it allows to choose from the set of all possible solutions, some of them possibly conflicting, based on predefined preferences. As we have argued, such a strategy may be useful in resolving conflicts in goals and actions, if preferences are paired with a formalism that can capture planning such as, for instance, logic programming. Algorithms for reasoning with preferences were devised, and prototypical reasoners were implemented. Preferential reasoning was also already applied in AmI (Bikakis and Antoniou, 2010a).

Table 15 Summary of conflict resolution fields

Field	Conflict types	Resolution method	Theoretical focus	Tractable variants (complexity)	Applications in AmI (and elsewhere)
Current Context Modelling Approaches	Contextual Sensory Bckg./Domain Goal Action	Mostly Prevention	–	–	Yes
Multi-Context Systems	Contextual (Sensory) (Bckg./Domain)	Resolution, Isolation	Yes	–	Yes
Belief Change	Contextual Bckg./Domain (Sensory) (Goal) (Action)	Prevention	Yes	–	–
Ontology Evolution	Contextual Bckg./Domain Sensory	Prevention	–	Yes, depending on underlying Description Logic	–
Ontology Debugging	Contextual Bckg./Domain Sensory	Repair	–	Yes, depending on the expressive power of the constraint language	–
Argumentation	Contextual Bckg./Domain Goal Action	Resolution	Yes	Grounded semantics (P)	Yes
Preferential Reasoning	Goal Action	Prevention, Indirect	Yes	Preferred well-founded semantics (P)	(in planning, configuration)
Paraconsistent Reasoning	Sensory Context Bckg./Domain	Isolation	Yes	Multiple (P)	(in data integra- tion)

While most of the above approaches work by removing or avoiding conflicts, paraconsistent reasoning takes a slightly different direction and concentrates on identifying and isolating the inconsistent parts of

the knowledge base and carefully drawing conclusions only from the consistent knowledge. In this area, different approaches were theoretically studied, however, especially paraconsistent logic programming and paraconsistent ontologies were developed also in practice and prototypical reasoners have been implemented. We believe that they can be potentially useful in AmI, especially when the AmI environments are to react to the current situation (which may feature inconsistencies) without the need to necessarily store all the current information for future processing.

A summary of our observations is presented in Table 15, where we compare the surveyed areas with respect to the different types of conflicts they are typically suited to resolve depending on which type of knowledge they occur, and how the conflicts are most typically resolved in each given area. Then, in the next two columns we highlight whether the area has a strong theoretical focus, but also whether the development has led into computationally feasible algorithms or implementations. In the final column, we indicate whether some applications in AmI (or at least some other relevant areas) are known. The details can be found in each respective subsection of Section 3.

We conclude that indeed KR has been fruitful in addressing the problem of conflict resolution from many different points of view and with diverse applications in different use cases. The different approaches to conflict resolution are very well theoretically developed, in the sense that the semantics is established and its properties are investigated. But, as we saw, for number of the approaches also effective (in the sense of polynomial) algorithms were developed, and some of them were implemented and experimentally evaluated. Finally, we have also pointed out a number of works that are already trying to apply KR methods in AmI (Moraitis and Spanoudakis, 2007; Mileo et al., 2008b; Muñoz et al., 2010; Bikakis and Antoniou, 2010a; Muñoz et al., 2010; Muñoz Ortega et al., 2010; Pajares Ferrando and Onaindia, 2012).

Acknowledgements

We would like to thank to Antonis Bikakis, Luciano Serafini, and Stefan Woltran who provided some very useful feedback on the preliminary version of this survey. We would also like to thank to Júlia Pukancová and Daniel Skalický for help with proofreading. This work resulted from the Slovak–Greek bilateral project “Multi-context Reasoning in Heterogeneous environments”, registered on the Slovak side under no. SK-GR-0070-11 with the APVV agency and co-financed by the Greek General Secretariat of Science and Technology and the European Union. It was further supported from the Slovak national VEGA project no. 1/1333/12.

References

- Emile Aarts, Rick Harwig, and Martin Schuurmans. Ambient intelligence. In Peter J. Denning, editor, *The Invisible Future: The Seamless Integration of Technology into Everyday Life*. McGraw-Hill Companies, New York, 2001.
- A&C 5(1), 2014. *Argument & Computation, Special Issue: Tutorials on Structured Argumentation Original Articles, 5(1)*, 2014.
- Andrea Acciarri, Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, Mattia Palmieri, and Riccardo Rosati. QuOnto: Querying ontologies. In *Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference, July 9–13, 2005, Pittsburgh, Pennsylvania, USA*, pages 1670–1671. AAAI Press, 2005.
- Giacomo Aceto. *Implementation of a non-ground meta-interpreter for defeasible logic*. PhD thesis, Università di Bologna, Italy, 2010.
- Kira V. Adaricheva, Robert H. Sloan, Balázs Szörényi, and György Turán. Horn belief contraction: Remainders, envelopes and complexity. In *Logical Formalizations of Commonsense Reasoning, Papers from the 2011 AAAI Spring Symposium, Technical Report SS-11-06, Stanford, California, USA, March 21–23, 2011*. AAAI, 2011.

- João Alcântara, Carlos Viegas Damásio, and Luís Moniz Pereira. A declarative characterization of disjunctive paraconsistent answer sets. In *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI'2004, including Prestigious Applicants of Intelligent Systems, PAIS 2004, Valencia, Spain, August 22–27, 2004*, pages 951–952. IOS Press, 2004.
- Carlos E. Alchourrón and David Makinson. On the logic of theory change: Safe contraction. *Studia Logica*, 44(4):405–422, 1985.
- Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- José Júlio Alferes, Carlos Viegas Damásio, and Luís Moniz Pereira. A logic programming system for nonmonotonic reasoning. *Journal of Automated Reasoning*, 14(1):93–147, 1995.
- José Júlio Alferes, João Alexandre Leite, Luís Moniz Pereira, Halina Przymusinska, and Teodor C. Przymusinski. Dynamic updates of non-monotonic knowledge bases. *The Journal of Logic Programming*, 45(1–3):43–70, 2000.
- José Júlio Alferes, Carlos Viegas Damásio, and Luís Moniz Pereira. Semantic web logic programming tools. In *Principles and Practice of Semantic Web Reasoning, International Workshop, PPSWR 2003, Mumbai, India, December 8, 2003, Proceedings*, volume 2901 of *LNCS*, pages 16–32. Springer, 2003. ISBN 3-540-20582-9.
- Leila Amgoud and Henri Prade. Using arguments for making and explaining decisions. *Artificial Intelligence*, 173(3–4):413–436, 2009.
- Leila Amgoud and Srdjan Vesic. A new approach for preference-based argumentation frameworks. *Annals of Mathematics and Artificial Intelligence*, 63(2):149–183, 2011.
- Leila Amgoud and Srdjan Vesic. On the use of argumentation for multiple criteria decision making. In *Advances in Computational Intelligence – 14th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, IPMU 2012, Catania, Italy, July 9–13, 2012, Proceedings, Part IV*, volume 300 of *CCIS*, pages 480–489. Springer, 2012.
- Marcelo Arenas, Leopoldo E. Bertossi, and Jan Chomicki. Consistent query answers in inconsistent databases. In *Proceedings of the Eighteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, May 31–June 2, 1999, Philadelphia, Pennsylvania, USA*, pages 68–79. ACM Press, 1999.
- Ofer Arieli and Marc Denecker. Reducing preferential paraconsistent reasoning to classical entailment. *Journal of Logic and Computation*, 13(4):557–580, 2003.
- Ofer Arieli, Marc Denecker, Bert Van Nuffelen, and Maurice Bruynooghe. Database repair by signed formulae. In *Foundations of Information and Knowledge Systems, Third International Symposium, FoIKS 2004, Wilhelminenberg Castle, Austria, February 17–20, 2004, Proceedings*, volume 2942 of *LNCS*, pages 14–30. Springer, 2004.
- Alexander Artikis, Marek Sergot, and Georgios Paliouras. A logic programming approach to activity recognition. In *Proceedings of the 2nd ACM international workshop on Events in multimedia, EiMM '10*, pages 3–8, 2010.
- Vernon Asuncion and Yan Zhang. Translating preferred answer set programs to propositional logic. In *Logic Programming and Nonmonotonic Reasoning, 10th International Conference, LPNMR 2009, Potsdam, Germany, September 14–18, 2009. Proceedings*, volume 5753 of *LNCS*, pages 396–401. Springer, 2009.

- Juan Carlos Augusto, Jun Liu, Paul McCullagh, Hui Wang, and Yang Jian-Bo. Management of uncertainty and spatio-temporal aspects for monitoring and diagnosis in a smart home. *International Journal of Computational Intelligence Systems*, 1(4):361–378, 2008.
- Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, New York, NY, USA, 2003. ISBN 0-521-78176-0.
- Martin Baláž, Jozef Frtús, and Martin Homola. Conflict resolution in structured argumentation. In *LPAR 2013, 19th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, December 12–17, 2013, Stellenbosch, South Africa, Short papers proceedings*, volume 26 of *EPiC Series*, pages 23–34. EasyChair, 2014a.
- Martin Baláž, Jozef Frtús, Martin Homola, Ján Sefránek, and Giorgos Flouris. Embedding defeasible logic programs into generalized logic programs. In *Proceedings of the 28th Workshop on (Constraint) Logic Programming (WLP 2014) Proceedings of the 23rd International Workshop on Functional and (Constraint) Logic Programming, Wittenberg, Germany, September 15–17, 2014.*, volume 1335 of *CEUR Workshop Proceedings*, pages 11–25, 2014b.
- Martin Baláž, Jozef Frtús, Giorgos Flouris, Martin Homola, and Ján Sefránek. Conflict resolution in assumption-based frameworks. In *Multi-Agent Systems – 12th European Conference, EUMAS 2014, Prague, Czech Republic, December 18–19, 2014, Revised Selected Papers*, volume 8953 of *LNCS*, pages 360–369. Springer, 2015.
- Marcello Balduccini and Michael Gelfond. Diagnostic reasoning with A-Prolog. *Theory and Practice of Logic Programming*, 3(4):425–461, 2003.
- Jie Bao, George Voutsadakis, Giora Slutzki, and Vasant Honavar. Package-based description logics. In Heiner Stuckenschmidt, Christine Parent, and Stefano Spaccapietra, editors, *Modular Ontologies: Concepts, Theories and Techniques for Knowledge Modularization*, volume 5445 of *LNCS*, pages 349–371. Springer, 2009.
- Pietro Baroni, Massimiliano Giacomin, and Guillermo R. Simari. Belief revision and argumentation: a reasoning process view. In Eduardo L. Fermé, Dov M. Gabbay, and Guillermo R. Simari, editors, *Trends in Belief Revision and Argumentation Dynamics*, volume 48 of *Logic and Cognitive Systems*. College Publications, 2013.
- Sotiris Batsakis, Kostas Stravoskoufos, and Euripides G. M. Petrakis. Temporal reasoning for supporting temporal queries in OWL 2.0. In *Knowledge-Based and Intelligent Information and Engineering Systems – 15th International Conference, KES 2011, Kaiserslautern, Germany, September 12–14, 2011, Proceedings, Part I*, volume 6881 of *LNCS*, pages 558–567. Springer, 2011.
- Ringo Baumann. What does it take to enforce an argument? minimal change in abstract argumentation. In *ECAI 2012 – 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27–31, 2012*, volume 242 of *FAIA*, pages 127–132. IOS Press, 2012.
- Ringo Baumann and Gerhard Brewka. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Computational Models of Argument: Proceedings of COMMA 2010, Desenzano del Garda, Italy, September 8–10, 2010.*, volume 216 of *FAIA*, pages 75–86. IOS Press, 2010.
- Ringo Baumann and Gerhard Brewka. Spectra in abstract argumentation: An analysis of minimal change. In *Logic Programming and Nonmonotonic Reasoning, 12th International Conference, LPNMR 2013, Corunna, Spain, September 15–19, 2013. Proceedings*, volume 8148 of *LNCS*, pages 174–186. Springer, 2013.

- Sean Bechhofer, Ian Horrocks, Carole A. Goble, and Robert Stevens. OilEd: A reason-able ontology editor for the semantic web. In *KI 2001: Advances in Artificial Intelligence, Joint German/Austrian Conference on AI, Vienna, Austria, September 19–21, 2001, Proceedings*, volume 2174 of *LNCS*, pages 396–408. Springer, 2001.
- Nuel D. Belnap. A Useful Four-Valued Logic. In J. Michael Dunn and George Epstein, editors, *Modern Uses of Multiple-Valued Logic*, volume 2 of *Episteme*, pages 5–37. Springer, 1977.
- Rachel Ben-Eliyahu and Rina Dechter. Default logic, propositional logic, and constraints. In *Proceedings of the 9th National Conference on Artificial Intelligence, Anaheim, CA, USA, July 14–19, 1991, Volume 1*, pages 379–385. AAAI Press, 1991.
- Massimo Benerecetti, Paolo Bouquet, and Chiara Ghidini. Contextual reasoning distilled. *Journal of Experimental & Theoretical Artificial Intelligence*, 12(3):279–305, 2000.
- Massimo Benerecetti, Paolo Bouquet, and Chiara Ghidini. On the dimensions of context dependence: Partiality, approximation, and perspective. In *Modeling and Using Context, Third International and Interdisciplinary Conference, CONTEXT, 2001, Dundee, UK, July 27–30, 2001, Proceedings*, volume 2116 of *LNCS*, pages 59–72. Springer, 2001.
- Salem Benferhat, Souhila Kaci, Daniel Le Berre, and Mary-Anne Williams. Weakening conflicting information for iterated revision and knowledge integration. *Artificial Intelligence*, 153(1–2):339–371, 2004.
- Leopoldo E. Bertossi, Anthony Hunter, and Torsten Schaub. Introduction to inconsistency tolerance. In Leopoldo E. Bertossi, Anthony Hunter, and Torsten Schaub, editors, *Inconsistency Tolerance*, volume 3300 of *LNCS*, pages 1–14. Springer, 2005.
- Philippe Besnard and Anthony Hunter. Quasi-classical logic: Non-trivializable classical reasoning from inconsistent information. In *Symbolic and Quantitative Approaches to Reasoning and Uncertainty, European Conference, ECSQARU’95, Fribourg, Switzerland, July 3–5, 1995, Proceedings*, volume 946 of *LNCS*, pages 44–51. Springer, 1995.
- Philippe Besnard and Anthony Hunter. A logic-based theory of deductive arguments. *Artificial Intelligence*, 128(12):203–235, 2001.
- Philippe Besnard and Anthony Hunter. *Elements of Argumentation*. MIT Press, 2008.
- Philippe Besnard and Torsten Schaub. Signed systems for paraconsistent reasoning. *Journal of Automated Reasoning*, 20(1):191–213, 1998.
- Philippe Besnard, Torsten Schaub, Hans Tompits, and Stefan Woltran. Representing paraconsistent reasoning via quantified propositional logic. In Leopoldo E. Bertossi, Anthony Hunter, and Torsten Schaub, editors, *Inconsistency Tolerance*, volume 3300 of *LNCS*, pages 84–118. Springer, 2005.
- Claudio Bettini, Oliver Brdiczka, Karen Henriksen, Jadwiga Indulska, Daniela Nicklas, Anand Ranganathan, and Daniele Riboni. A survey of context modelling and reasoning techniques. *Pervasive and Mobile Computing*, 6(2):161–180, 2010.
- Antonis Bikakis and Grigoris Antoniou. Defeasible contextual reasoning with arguments in ambient intelligence. *IEEE Transactions on Knowledge and Data Engineering*, 22(11):1492–1506, 2010a.
- Antonis Bikakis and Grigoris Antoniou. Rule-based contextual reasoning in ambient intelligence. In *Semantic Web Rules – International Symposium, RuleML 2010, Washington, DC, USA, October 21–23, 2010. Proceedings*, volume 6403 of *LNCS*, pages 74–88. Springer, 2010b.

- Antonis Bikakis and Grigoris Antoniou. Partial preferences and ambiguity resolution in contextual defeasible logic. In *Logic Programming and Nonmonotonic Reasoning – 11th International Conference, LPNMR 2011, Vancouver, Canada, May 16–19, 2011. Proceedings*, volume 6645 of LNCS, pages 193–198. Springer, 2011.
- Antonis Bikakis, Theodore Patkos, Grigoris Antoniou, and Dimitris Plexousakis. A survey of semantics-based approaches for context reasoning in ambient intelligence. In *Constructing Ambient Intelligence – Aml 2007 Workshops Darmstadt, Germany, November 7–10, 2007 Revised Papers*, volume 11 of CCIS, pages 14–23. Springer, 2008.
- Antonis Bikakis, Grigoris Antoniou, and Panayiotis Hassapis. Strategies for contextual reasoning with conflicts in ambient intelligence. *Knowledge and Information Systems*, 27(1):45–84, 2011.
- Barry Bishop, Atanas Kiryakov, Damyan Ognyanoff, Ivan Peikov, Zdravko Tashev, and Ruslan Velkov. OWLIM: A family of scalable semantic repositories. *Semantic Web*, 2(1):33–42, 2011.
- Howard A. Blair and V. S. Subrahmanian. Paraconsistent logic programming. In *Foundations of Software Technology and Theoretical Computer Science, Seventh Conference, Pune, India, December 17–19, 1987, Proceedings*, volume 287 of LNCS, pages 340–360. Springer, 1987.
- Howard A. Blair and V. S. Subrahmanian. Paraconsistent logic programming. *Theoretical Computer Science*, 68(2):135–154, 1989.
- Guido Boella, Souhila Kaci, and Leendert van der Torre. Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 10th European Conference, ECSQARU 2009, Verona, Italy, July 1–3, 2009. Proceedings*, volume 5590 of LNCS, pages 107–118. Springer, 2009a.
- Guido Boella, Souhila Kaci, and Leendert W. N. van der Torre. Dynamics in argumentation with single extensions: Attack refinement and the grounded extension (extended version). In *Argumentation in Multi-Agent Systems, 6th International Workshop, ArgMAS 2009, Budapest, Hungary, May 12, 2009. Revised Selected and Invited Papers*, volume 6057 of LNCS, pages 150–159. Springer, 2009b.
- Markus Bögl, Thomas Eiter, Michael Fink, and Peter Schüller. The MCS-IE system for explaining inconsistency in multi-context systems. In *Logics in Artificial Intelligence – 12th European Conference, JELIA 2010, Helsinki, Finland, September 13–15, 2010. Proceedings*, volume 6341 of LNCS, pages 356–359. Springer, 2010.
- Andrei Bondarenko, Phan Minh Dung, Robert A. Kowalski, and Francesca Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93:63–101, 1997.
- Richard Booth, Thomas Meyer, and Ivan José Varzinczak. Next steps in propositional Horn contraction. In *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11–17, 2009*, pages 702–707, 2009.
- Richard Booth, Thomas Meyer, Ivan Varzinczak, and Renata Wassermann. On the link between partial meet, kernel, and infra contraction and its application to Horn logic. *Journal of Artificial Intelligence Research*, 42:31–53, 2011.
- Alexander Borgida and Luciano Serafini. Distributed description logics: Assimilating information from peer sources. *Journal on Data Semantics*, 1:153–184, 2003.
- Tibor Bosse and Alexei Sharpanskykh. A framework for modeling and analysis of ambient agent systems: Application to an emergency case. In *Ambient Intelligence and Future Trends – International Symposium on Ambient Intelligence (ISAMI 2010), Guimarães, Portugal, 16–18 June 2010*, volume 72 of AISC, pages 121–129. Springer, 2010.

- Paolo Bouquet, Fausto Giunchiglia, Frank van Harmelen, Luciano Serafini, and Heiner Stuckenschmidt. Contextualizing ontologies. *Journal of Web Semantics*, 1(4):325–343, 2004.
- Craig Boutilier, Ronen I. Brafman, Holger H. Hoos, and David Poole. Reasoning with conditional ceteris paribus preference statements. In *UAI '99: Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence, Stockholm, Sweden, July 30–August 1, 1999*, pages 71–80. Morgan Kaufmann, 1999.
- Craig Boutilier, Ronen I Brafman, Carmel Domshlak, Holger H Hoos, and David Poole. CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.
- Loris Bozzato and Luciano Serafini. Materialization calculus for contexts in the semantic web. In *Informal Proceedings of the 26th International Workshop on Description Logics, Ulm, Germany, July 23–26, 2013*, volume 1014 of *CEUR Workshop Proceedings*, pages 552–572, 2013.
- Loris Bozzato, Chiara Ghidini, and Luciano Serafini. Comparing contextual and flat representations of knowledge: a concrete case about football data. In V. Richard Benjamins, Mathieu d’Aquin, and Andrew Gordon, editors, *Proceedings of the 7th International Conference on Knowledge Capture, K-CAP 2013, Banff, Canada, June 23–26, 2013*, pages 9–16. ACM, 2013.
- Loris Bozzato, Thomas Eiter, and Luciano Serafini. Contextualized knowledge repositories with justifiable exceptions. In Meghyn Bienvenu, Magdalena Ortiz, Riccardo Rosati, and Mantas Simkus, editors, *Informal Proceedings of the 27th International Workshop on Description Logics, Vienna, Austria, July 17–20, 2014.*, volume 1193 of *CEUR Workshop Proceedings*, pages 112–123. CEUR-WS.org, 2014.
- Thais R. M. Braga Silva, Linnyer B. Ruiz, and Antonio A. F. Loureiro. Conflicts treatment for ubiquitous collective and context-aware applications. *Journal of Applied Computing Research*, 1(1):33–47, 2011.
- Michael E. Bratman. *Intention, plans, and practical reason*. Harvard University Press, Cambridge, MA, United States, 1987.
- Gerhard Brewka. Well-Founded Semantics for Extended Logic Programs with Dynamic Preferences. *Journal of Artificial Intelligence Research*, 4:19–36, 1996.
- Gerhard Brewka. Logic programming with ordered disjunction. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence and Fourteenth Conference on Innovative Applications of Artificial Intelligence, July 28–August 1, 2002, Edmonton, Alberta, Canada.*, pages 100–105. AAAI Press, 2002.
- Gerhard Brewka and Thomas Eiter. Preferred answer sets for extended logic programs. *Artificial Intelligence*, 109(1–2):297–356, 1999.
- Gerhard Brewka and Thomas Eiter. Equilibria in heterogeneous nonmonotonic multi-context systems. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, July 22–26, 2007, Vancouver, British Columbia, Canada*, pages 385–390. AAAI Press, 2007.
- Gerhard Brewka and Thomas Eiter. Argumentation context systems: A framework for abstract group argumentation. In *Logic Programming and Nonmonotonic Reasoning, 10th International Conference, LPNMR 2009, Potsdam, Germany, September 14–18, 2009. Proceedings*, volume 5753 of *LNCS*, pages 44–57. Springer, 2009.
- Gerhard Brewka and Stefan Woltran. Abstract dialectical frameworks. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010, Toronto, Ontario, Canada, May 9–13, 2010*, 2010.

- Gerhard Brewka, Ilkka Niemelä, and Mirosław Truszczyński. Answer set optimization. In *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 9–15, 2003*, pages 867–872. Morgan Kaufmann, 2003.
- Gerhard Brewka, Floris Roelofsen, and Luciano Serafini. Contextual default reasoning. In *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6–12, 2007*, pages 268–273, 2007.
- Gerhard Brewka, Thomas Eiter, Michael Fink, and Antonius Weinzierl. Managed multi-context systems. In *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16–22, 2011*, pages 786–791, 2011.
- Emma Byrne and Anthony Hunter. Man bites dog: looking for interesting inconsistencies in structured news reports. *Data & Knowledge Engineering*, 48(3):265–295, 2004.
- Andrea Cali, Georg Gottlob, and Thomas Lukasiewicz. Datalog[±]: a unified approach to ontologies and integrity constraints. In *Proceedings of the 12th International Conference on Database Theory (ICDT-09)*, 2009.
- Diego Calvanese, Evgeny Kharlamov, Werner Nutt, and Dmitriy Zheleznyakov. Evolution of DL-Lite knowledge bases. In *The Semantic Web – ISWC 2010, 9th International Semantic Web Conference, ISWC 2010, Shanghai, China, November 7–11, 2010, Revised Selected Papers, Part I*, volume 6496 of *LNCS*, pages 112–128. Springer, 2010.
- Martin Caminada and Leila Amgoud. On the evaluation of argumentation formalisms. *Artificial Intelligence*, 171(5–6):286–310, 2007.
- Martin W. A. Caminada and Dov M. Gabbay. A logical account of formal argumentation. *Studia Logica*, 93(2–3):109–145, 2009.
- Martin W. A. Caminada, Walter A. Carnielli, and Paul E. Dunne. Semi-stable semantics. *Journal of Logic and Computation*, 22(5):1207–1254, 2012.
- Ana Casali, Lluís Godo, and Carles Sierra. Graded BDI models for agent architectures. In *Computational Logic in Multi-Agent Systems, 5th International Workshop, CLIMA V, Lisbon, Portugal, September 29–30, 2004, Revised Selected and Invited Papers*, volume 3487 of *LNCS*. Springer, 2005.
- Claudette Cayrol and Marie-Christine Lagasque-Schiex. Bipolar abstract argumentation systems. In Iyad Rahwan and Guillermo R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 65–84. Springer, 2009.
- Claudette Cayrol, Florence Dupin de Saint-Cyr, and Marie-Christine Lagasque-Schiex. Change in abstract argumentation frameworks: Adding an argument. *Journal of Artificial Intelligence Research*, 38:49–84, 2010.
- Irene Celino, Simone Contessa, Marta Corubolo, Daniele Dell’Aglia, Emanuele Della Valle, Stefano Fumeo, and Thorsten Krüger. Linking smart cities datasets with human computation – the case of UrbanMatch. In *The Semantic Web – ISWC 2012, 11th International Semantic Web Conference, Boston, MA, USA, November 11–15, 2012, Proceedings, Part II*, volume 7650 of *LNCS*, pages 34–49. Springer, 2012.
- Marie Chan, Daniel Estève, Christophe Escriba, and Eric Campo. A review of smart homes-present state and future challenges. *Computer Methods and Programs Biomedicine*, 91(1):55–81, 2008.
- Liming Chen and Ismail Khalil. Activity recognition: Approaches, practices and trends. In Liming Chen, Chris Nugent, Jit Biswas, and Jesse Hoey, editors, *Activity Recognition in Pervasive Intelligent Environments*, volume 4 of *Atlantis Ambient and Pervasive Intelligence*, pages 1–31. Atlantis Press, 2011.

- Jan Chomicki and Jerzy Marcinkowski. On the computational complexity of minimal-change integrity maintenance in relational databases. In Leopoldo E. Bertossi, Anthony Hunter, and Torsten Schaub, editors, *Inconsistency Tolerance*, volume 3300 of *LNCS*. Springer, 2005.
- Michalis Chortis and Giorgos Flouris. A diagnosis and repair framework for *DL-Lite_A* KBs. In *Proceedings of the First DIACHRON Workshop on Managing the Evolution and Preservation of the Data Web co-located with 12th European Semantic Web Conference (ESWC 2015), Portorož, Slovenia, May 31, 2015.*, volume 1377 of *CEUR Workshop Proceedings*, pages 1–15, 2015.
- Seng-cho Timothy Chou and Marianne Winslett. A model-based belief revision system. *Journal of Automated Reasoning*, 12(2):157–208, 1994.
- Alessandro Cimatti and Luciano Serafini. Multi-agent reasoning with belief contexts: The approach and a case study. In *Intelligent Agents, ECAI-94 Workshop on Agent Theories, Architectures, and Languages, Amsterdam, The Netherlands, August 8–9, 1994, Proceedings*, volume 890 of *LNCS*. Springer, 1995.
- Marcello Cirillo, F. Lanzelotto, Federico Pecora, and Alessandro Saffiotti. Monitoring domestic activities with temporal constraints and components. In *Intelligent Environments 2009 – Proceedings of the 5th International Conference on Intelligent Environments, Barcelona, Spain 2009*, volume 2 of *AISE*, pages 117–124. IOS Press, 2009.
- Philip R. Cohen and Hector J. Levesque. Intention is choice with commitment. *Artificial intelligence*, 42(2):213–261, 1990.
- Anthony G. Cohn and Shyamanta M. Hazarika. Qualitative spatial representation and reasoning: An overview. *Fundamenta informaticae*, 46(1):1–29, 2001.
- Diane J. Cook, Juan C. Augusto, and Vikramaditya R. Jakkula. Ambient intelligence: Technologies, applications, and opportunities. *Pervasive and Mobile Computing*, 5(4):277–298, 2009.
- Sylvie Coste-Marquis and Pierre Marquis. On the complexity of paraconsistent inference relations. In Leopoldo E. Bertossi, Anthony Hunter, and Torsten Schaub, editors, *Inconsistency Tolerance*, volume 3300 of *LNCS*, pages 151–190. Springer, 2005.
- Sylvie Coste-Marquis, Sébastien Konieczny, Pierre Marquis, and Mohand Akli Ouali. Weighted attacks in argumentation frameworks. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012, Rome, Italy, June 10–14, 2012*. AAAI Press, 2012.
- Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. On the revision of argumentation systems: Minimal change of arguments statuses. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference, KR 2014, Vienna, Austria, July 20–24, 2014*. AAAI Press, 2014.
- Nadia Creignou, Odile Papini, Reinhard Pichler, and Stefan Woltran. Belief revision within fragments of propositional logic. *Journal of Computer and Systems Sciences*, 80(2):427–449, 2014.
- Bernardo Cuenca Grau, Bijan Parsia, and Evren Sirin. Working with multiple ontologies on the semantic web. In *The Semantic Web – ISWC 2004: Third International Semantic Web Conference, Hiroshima, Japan, November 7–11, 2004. Proceedings*, volume 3298 of *LNCS*, pages 620–634. Springer, 2004.
- Bernardo Cuenca Grau, Evgeny Kharlamov, and Dmitriy Zheleznyakov. Ontology contraction: Beyond the propositional paradise. In *Proceedings of the 6th Alberto Mendelzon International Workshop on Foundations of Data Management, Ouro Preto, Brazil, June 27–30, 2012*, volume 866 of *CEUR Workshop Proceedings*, pages 62–74, 2012.
- Newton C. A. da Costa and V. S. Subrahmanian. Paraconsistent logics as a formalism for reasoning about inconsistent knowledge bases. *Artificial Intelligence in Medicine*, 1(4):167–174, 1989.

- Mukesh Dalal. Investigations into a theory of knowledge base revision: Preliminary report. In *Proceedings of the 7th National Conference on Artificial Intelligence*. St. Paul, MN, August 21–26, 1988., pages 475–479. AAAI Press, 1988.
- Elizabeth M. Daly, Freddy Lécué, and Veli Bicer. Westland row why so slow? Fusing social media and linked data sources for understanding real-time traffic conditions. In *18th International Conference on Intelligent User Interfaces, IUI '13, Santa Monica, CA, USA, March 19–22, 2013*, pages 203–212. ACM, 2013.
- Carlos Viegas Damásio and Luís Moniz Pereira. A model theory for paraconsistent logic programming. In *Progress in Artificial Intelligence, 7th Portuguese Conference on Artificial Intelligence, EPIA '95, Funchal, Madeira Island, Portugal, October 3–6, 1995, Proceedings*, volume 990 of LNCS, pages 377–386. Springer, 1995.
- DARC 2012. Workshop on the Dynamics of Argumentation, Rules and Conditionals. University of Luxembourg, 2012. See online: <http://icr.uni.lu/darc/DARC/>, 2012.
- Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, and Riccardo Rosati. On the approximation of instance level update and erasure in description logics. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, July 22–26, 2007, Vancouver, British Columbia, Canada*, pages 403–408. AAAI Press, 2007.
- Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, and Ricardo Rosati. On instance-level update and erasure in description logic ontologies. *Journal of Logic and Computation*, 19(5):745–770, 2009.
- James P. Delgrande. Horn clause belief change: Contraction functions. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16–19, 2008*, pages 156–165. AAAI Press, 2008.
- James P. Delgrande and Renata Wassermann. Horn clause contraction functions: Belief set and belief base approaches. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010, Toronto, Ontario, Canada, May 9–13, 2010*. AAAI Press, 2010.
- James P. Delgrande, Torsten Schaub, and Hans Tompits. A framework for compiling preferences in logic programs. *Theory and Practice of Logic Programming*, 3(2):129–187, 2003.
- James P. Delgrande, Torsten Schaub, and Hans Tompits. Domain-specific preferences for causal reasoning and planning. In *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004), June 3–7 2004, Whistler, British Columbia, Canada*, pages 63–72. AAAI, 2004a.
- James P. Delgrande, Torsten Schaub, Hans Tompits, and Kewen Wang. A classification and survey of preference handling approaches in nonmonotonic reasoning. *Computational Intelligence*, 20(2):308–334, 2004b.
- Yannis Dimopoulos, Bernhard Nebel, and Jana Koehler. Encoding planning problems in nonmonotonic logic programs. In *Recent Advances in AI Planning, 4th European Conference on Planning, ECP'97, Toulouse, France, September 24–26, 1997, Proceedings*, volume 1348 of LNCS, pages 169–181. Springer, 1997.
- Jürgen Dix, Sven Ove Hansson, Gabriele Kern-Isberner, and Guillermo Ricardo Simari. Belief change and argumentation in multi-agent scenarios (dagstuhl seminar 13231). *Dagstuhl Reports*, 3(6):1–21, 2013.

- Rim Djedidi and Marie-Aude Aufaure. Change management patterns (CMP) for ontology evolution process. In *Proceedings of the 3rd International Workshop on Ontology Dynamics, (IWOD 2009), collocated with the 8th International Semantic Web Conference (ISWC-2009), Washington DC, USA, October 26, 2009*, volume 519 of *CEUR Workshop Proceedings*, 2009.
- Rim Djedidi and Marie-Aude Aufaure. *ONTO-EVO^AL* an ontology evolution approach guided by pattern modeling and quality evaluation. In *Foundations of Information and Knowledge Systems, 6th International Symposium, FoIKS 2010, Sofia, Bulgaria, February 15–19, 2010. Proceedings*, volume 5956 of *LNCS*, pages 286–305. Springer, 2010.
- Patrick Doherty, Witold Lukaszewicz, and Andrzej Szalas. Computing circumscription revisited: A reduction algorithm. *Journal of Automated Reasoning*, 18(3):297–336, 1997.
- Jon Doyle. A truth maintenance system. *Artificial Intelligence*, 12(3):231–272, 1979.
- Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.
- Phan Minh Dung, Robert A Kowalski, and Francesca Toni. Assumption-based argumentation. In Iyad Rahwan and Guillermo R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 199–218. Springer, 2009.
- Paul E. Dunne and Michael Wooldridge. Complexity of abstract argumentation. In Iyad Rahwan and Guillermo R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 85–104. Springer, 2009.
- Paul E. Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, and Michael Wooldridge. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artificial Intelligence*, 175(2): 457–486, 2011.
- Vasiliki Efstathiou and Anthony Hunter. Jargue: An implemented argumentation system for classical propositional logic. In *COMMA 2010 Third International Conference on Computational Models of Argument Desenzano del Garda (Italy), 8th–10th September 2010*, 2010. Demo paper, available online: http://www.ing.unibs.it/~comma2010/demos/Efstathiou_etal.pdf.
- Thomas Eiter, Wolfgang Faber, Nicola Leone, and Gerald Pfeifer. Computing preferred answer sets by meta-interpretation in answer set programming. *Theory and Practice of Logic Programming*, 3(4–5): 463–498, 2003a.
- Thomas Eiter, Wolfgang Faber, Nicola Leone, Gerald Pfeifer, and Axel Polleres. A logic programming approach to knowledge-state planning, II: The DLV^K system. *Artificial Intelligence*, 144(1–2):157–211, 2003b.
- Thomas Eiter, Giovambattista Ianni, Roman Schindlauer, and Hans Tompits. dlhex: A system for integrating multiple semantics in an answer-set programming framework. In *20th Workshop on Logic Programming, Vienna, Austria, February 22–24, 2006*, volume 1843-06-02 of *INFSYS Research Report*, pages 206–210. Technische Universität Wien, Austria, 2006.
- Thomas Eiter, Michael Fink, and João Moura. Paracoherent answer set programming. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010, Toronto, Ontario, Canada, May 9–13, 2010*. AAAI Press, 2010a.
- Thomas Eiter, Michael Fink, Peter Schüller, and Antonius Weinzierl. Finding explanations of inconsistency in multi-context systems. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010, Toronto, Ontario, Canada, May 9–13, 2010*. AAAI Press, 2010b.

- Mikel Emaldi, Jon Lazaro, Xabier Laiseca, and Diego Lopez-de Ipina. LinkedQR: Improving tourism experience through linked data and QR codes. In *Ubiquitous Computing and Ambient Intelligence*, volume 7656 of LNCS, pages 371–378. Springer, 2012.
- Michael Erdmann and Walter Waterfeld. Overview of the NeOn toolkit. In María del Carmen Suárez-Figueroa, Asunción Gómez-Pérez, Enrico Motta, and Aldo Gangemi, editors, *Ontology Engineering in a Networked World*, pages 281–301. Springer, 2012.
- Neil A. Ernst, Alexander Borgida, John Mylopoulos, and Ivan Jureta. Agile requirements evolution via paraconsistent reasoning. In *Advanced Information Systems Engineering – 24th International Conference, CAiSE 2012, Gdansk, Poland, June 25–29, 2012. Proceedings*, volume 7328 of LNCS, pages 382–397. Springer, 2012.
- Marcello A. Falappa, Gabriele Kern-Isberner, and Guillermo R. Simari. Belief revision and argumentation theory. In Iyad Rahwan and Guillermo R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 341–360. Springer, 2009.
- Marcelo A. Falappa, Alejandro Javier García, Gabriele Kern-Isberner, and Guillermo Ricardo Simari. On the evolving relation between belief revision and argumentation. *The Knowledge Engineering Review*, 26(1):35–43, 2011.
- Rainer Feldmann, Burkhard Monien, and Stefan Schamberger. A distributed algorithm to evaluate quantified boolean formulae. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence, July 30–August 3, 2000, Austin, Texas, USA*, pages 285–290. AAAI Press, 2000.
- Eduardo L. Fermé and Sven Ove Hansson. AGM 25 years: Twenty-five years of research in belief change. *Journal of Philosophical Logic*, 40(2):295–331, 2011.
- Eduardo L. Fermé, Dov M. Gabbay, and Guillermo R. Simari, editors. *Trends in Belief Revision and Argumentation*, volume 48 of *Logic and Cognitive Systems*. College Publications, 2013. Procs. of the 2012 Workshop on Belief Revision and Argumentation, Funchal, Madeira.
- Melvin Fitting. Kleene’s logic, generalized. *Journal of Logic and Computation*, 1(6):797–810, 1991a.
- Melvin Fitting. Bilattices and the semantics of logic programming. *Journal of Logic Programming*, 11(1&2):91–116, 1991b.
- Giorgos Flouris. On belief change in ontology evolution. *AI Communications Journal (AI-Com)*, 19(4): 395–397, 2006a. PhD Thesis Summary.
- Giorgos Flouris. *On Belief Change and Ontology Evolution*. PhD thesis, University of Crete, 2006b.
- Giorgos Flouris and Dimitris Plexousakis. Bridging ontology evolution and belief change. In *Advances in Artificial Intelligence, 4th Hellenic Conference on AI, SETN 2006, Heraklion, Crete, Greece, May 18–20, 2006, Proceedings*, volume 3955 of LNCS, pages 486–489. Springer, 2006.
- Giorgos Flouris, Dimitris Plexousakis, and Grigoris Antoniou. Generalizing the AGM postulates: Preliminary results and applications. In *10th International Workshop on Non-Monotonic Reasoning (NMR 2004), Whistler, Canada, June 6–8, 2004, Proceedings*, pages 171–179, 2004.
- Giorgos Flouris, Dimitris Plexousakis, and Grigoris Antoniou. On applying the AGM theory to DLs and OWL. In *The Semantic Web – ISWC 2005, 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6–10, 2005, Proceedings*, volume 3729 of LNCS, pages 216–231. Springer, 2005.

- Giorgos Flouris, Zhisheng Huang, Jeff Z. Pan, Dimitris Plexousakis, and Holger Wache. Inconsistencies, negations and changes in ontologies. In *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, July 16–20, 2006, Boston, Massachusetts, USA*, pages 1295–1300. AAAI Press, 2006a.
- Giorgos Flouris, Dimitris Plexousakis, and Grigoris Antoniou. On generalizing the AGM postulates. In *STAIRS 2006 – Proceedings of the Third Starting AI Researchers’ Symposium, Riva del Garda, Trentino, Italy, 2006*, volume 142 of *FAIA*, pages 132–143. IOS Press, 2006b.
- Giorgos Flouris, Dimitris Manakanatas, Haridimos Kondylakis, Dimitris Plexousakis, and Grigoris Antoniou. Ontology change: Classification and survey. *The Knowledge Engineering Review*, 26(2): 117–152, 2008.
- Giorgos Flouris, Yannis Roussakis, Maria Poveda-Villalon, Pablo N. Mendes, and Irimi Fundulaki. Using provenance for quality assessment and repair in linked open data. In *Proceedings of the 2nd Joint Workshop on Knowledge Evolution and Ontology Dynamics, in conjunction with the 11th International Semantic Web Conference (ISWC 2012), Boston, MA, USA, November 12, 2012*, volume 890 of *CEUR Workshop Proceedings*, 2012.
- Giorgos Flouris, George Konstantinidis, Grigoris Antoniou, and Vassilis Christophides. Formal foundations for RDF/S KB evolution. *Knowledge and Information Systems*, 35(1):153–191, 2013.
- Florian Fuchs, Iris Hochstatter, Michael Krause, and Michael Berger. A Metamodel Approach to Context Information. In *3rd IEEE Conference on Pervasive Computing and Communications Workshops (PerCom 2005 Workshops), 8–12 March 2005, Kauai Island, HI, USA*, pages 8–14. IEEE Computer Society, 2005.
- A. Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20(2):175–203, 1991.
- Dov M Gabbay, Agi Kurucz, Frank Wolter, and Michael Zakharyashev. *Many-Dimensional Modal Logics: Theory and Applications*, volume 148 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, 2003.
- Thomas Gabel, York Sure, and Johanna Voelker. KAON – ontology management infrastructure. SEKT informal deliverable D3.1.1.a, 2004. Available online: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.1.5228&rep=rep1&type=pdf>.
- Alejandro J. García and Guillermo R. Simari. Defeasible logic programming: an argumentative approach. *Theory and Practice of Logic Programming*, 4(1–2):95–138, 2004.
- Peter Gärdenfors. The dynamics of belief systems: Foundations versus coherence theories. *Revue Internationale de Philosophie*, 44:24–46, 1990.
- Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, March 1988*, pages 83–95. Morgan Kaufmann, 1988.
- Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. In *Logic Programming, Proceedings of the Fifth International Conference and Symposium, Seattle, Washington, August 15–19, 1988*, pages 1070–1080. MIT Press, 1988.
- Michael Gelfond and Vladimir Lifschitz. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing*, 9(3/4):365–386, 1991.
- Chiara Ghidini and Fausto Giunchiglia. Local models semantics, or contextual reasoning = locality + compatibility. *Artificial Intelligence*, 127(2):221–259, 2001.

- Chiara Ghidini and Luciano Serafini. Distributed first order logics. In *Frontiers of Combining Systems, FroCoS '98, Second International Workshop, Amsterdam, The Netherlands*, pages 121–140. Research Studies Press, 1998.
- Chiara Ghidini and Luciano Serafini. Mapping properties of heterogeneous ontologies. In *Artificial Intelligence: Methodology, Systems, and Applications, 13th International Conference, AIMSA 2008, Varna, Bulgaria, September 4–6, 2008. Proceedings*, volume 5253 of *LNCS*, pages 181–193. Springer, 2008.
- Chiara Ghidini, Luciano Serafini, and Sergio Tessaris. Bridging heterogeneous representations of binary relations: First results. In *Proceedings of the 21st International Workshop on Description Logics (DL2008), Dresden, Germany, May 13–16, 2008*, volume 353 of *CEUR Workshop Proceedings*, 2008.
- Matthew L. Ginsberg. Multivalued logics: a uniform approach to reasoning in artificial intelligence. *Computational Intelligence*, 4(3):265–316, 1988.
- Matthew L. Ginsberg and David E. Smith. Reasoning about action i: A possible worlds approach. *Artificial intelligence*, 35(2):165–195, 1988.
- Enrico Giunchiglia, Massimo Narizzano, and Armando Tacchella. QuBE: A system for deciding quantified boolean formulas satisfiability. In *Automated Reasoning, First International Joint Conference, IJCAR 2001, Siena, Italy, June 18–23, 2001, Proceedings*, volume 2083 of *LNCS*, pages 364–369. Springer, 2001.
- Fausto Giunchiglia. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine*, 16:345–364, 1993.
- Fausto Giunchiglia and Chiara Ghidini. Local models semantics, or contextual reasoning = locality + compatibility. In *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98), Trento, Italy, June 2–5, 1998.*, pages 282–291. Morgan Kaufmann, 1998.
- Fausto Giunchiglia and Luciano Serafini. Multilanguage hierarchical logics or: How we can do without modal logics. *Artificial Intelligence*, 65(1):29–70, 1994.
- Ricardo Gonçalves, Matthias Knorr, and João Leite. Evolving bridge rules in evolving multi-context systems. In *Computational Logic in Multi-Agent Systems – 15th International Workshop, CLIMA XV, Prague, Czech Republic, August 18–19, 2014. Proceedings*, volume 8624 of *LNCS*, pages 52–69. Springer, 2014a.
- Ricardo Gonçalves, Matthias Knorr, and João Leite. Evolving multi-context systems. In *ECAI 2014 – 21st European Conference on Artificial Intelligence, 18–22 August 2014, Prague, Czech Republic, Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, volume 263 of *FAIA*, pages 375–380. IOS Press, 2014b.
- Georg Gottlob. Complexity results for nonmonotonic logics. *Journal of Logic and Computation*, 2(3): 397–425, 1992.
- Guido Governatori, Michael J. Maher, Grigoris Antoniou, and David Billington. Argumentation semantics for defeasible logic. *Journal of Logic and Computation*, 14(5):675–702, 2004.
- Susanne Grell, Kathrin Konczak, and Torsten Schaub. $\text{nomore}^<$: A System for Computing Preferred Answer Sets. In *Logic Programming and Nonmonotonic Reasoning, 8th International Conference, LPNMR 2005, Diamante, Italy, September 5–8, 2005, Proceedings*, volume 3662 of *LNCS*, pages 394–398. Springer, 2005.

- Luca Grieco, Domenico Lembo, Riccardo Rosati, and Marco Ruzzi. Consistent query answering under key and exclusion dependencies: algorithms and experiments. In *Proceedings of the 2005 ACM CIKM International Conference on Information and Knowledge Management, Bremen, Germany, October 31–November 5, 2005*, pages 792–799. ACM, 2005.
- Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17(2):157–170, 1988.
- Tao Gu, Hung Keng Pung, and Da Qing Zhang. A service-oriented middleware for building context-aware services. *Journal of Network and Computer Applications*, 28(1):1–18, 2005.
- Joakim Gustafsson. An Implementation and Optimization of an Algorithm for Reducing Formulas in Second-Order Logic. Technical report, Department of Mathematics, Linköping University, Sweden, 1996.
- Claudio Gutierrez, Carlos A. Hurtado, and Alejandro A. Vaisman. The meaning of erasing in RDF under the Katsuno-Mendelzon approach. In *Ninth International Workshop on the Web and Databases, WebDB 2006, Chicago, Illinois, USA, June 30, 2006*, 2006.
- Claudio Gutierrez, Carlos A. Hurtado, and Alejandro Vaisman. Introducing time into RDF. *IEEE Transactions on Knowledge and Data Engineering*, 19(2):207–218, 2007.
- Peter Haase, Frank van Harmelen, Zhisheng Huang, Heiner Stuckenschmidt, and York Sure. A framework for handling inconsistency in changing ontologies. In *The Semantic Web – ISWC 2005, 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6–10, 2005, Proceedings*, volume 3729 of *LNCS*, pages 353–367. Springer, 2005.
- Christian Halaschek-Wiener and Yarden Katz. Belief base revision for expressive description logics. In *Proceedings of the OWLED*06 Workshop on OWL: Experiences and Directions, Athens, Georgia, USA, November 10–11, 2006*, volume 216 of *CEUR Workshop Proceedings*, 2006.
- Sven Ove Hansson. Belief contraction without recovery. *Studia Logica*, 50(2):251–260, 1991.
- Sven Ove Hansson. Kernel contraction. *Journal of Symbolic Logic*, 59(3):845–859, 1994.
- Sven Ove Hansson. Knowledge-level analysis of belief base operations. *Artificial Intelligence*, 82(1–2): 215–235, 1996.
- Sven Ove Hansson. What’s new isn’t always best. *Theoria: Special Issue on Non-Prioritized Belief Revision*, 63(1–2):1–13, 1997.
- Sven Ove Hansson, Eduardo L. Fermé, John Cantwell, and Marcelo A. Falappa. Credibility limited revision. *Journal of Symbolic Logic*, 66(4):1581–1596, 2001.
- Rim Helaoui, Mathias Niepert, and Heiner Stuckenschmidt. Recognizing interleaved and concurrent activities using qualitative and quantitative temporal relationships. *Pervasive and Mobile Computing*, 7(6):660–670, 2011.
- Rim Helaoui, Daniele Riboni, Mathias Niepert, Claudio Bettini, and Heiner Stuckenschmidt. Towards activity recognition using probabilistic description logics. In *Activity Context Representation: Techniques and Languages, Workshops at the Twenty-Sixth AAAI Conference on Artificial Intelligence*, volume WS-12-05 of *AAAI Technical Report*. AAAI, 2012.
- Karen Henriksen and Jadwiga Indulska. Modelling and using imperfect context information. In *2nd IEEE Conference on Pervasive Computing and Communications Workshops (PerCom 2004 Workshops), 14–17 March 2004, Orlando, FL, USA*, pages 33–37. IEEE Computer Society, 2004.
- Martin Homola. *Semantic Investigations in Distributed Ontologies*. PhD thesis, Comenius University in Bratislava, Slovakia, 2010.

- Martin Homola and Theodore Patkos. Different types of conflicting knowledge in ami environments. In *Knowledge Engineering and Knowledge Management – EKAW 2014 Satellite Events, VISUAL, EKMI, and ARCOE-Logic, Linköping, Sweden, November 24–28, 2014. Revised Selected Papers.*, volume 8982 of *LNCS*, pages 52–63. Springer, 2015.
- Martin Homola and Luciano Serafini. Augmenting subsumption propagation in distributed description logics. *Applied Artificial Intelligence*, 24(1–2):39–76, 2010.
- Xin Hong, Chris Nugent, Maurice Mulvenna, Sally McClean, Bryan Scotney, and Steven Devlin. Evidential fusion of sensor data for activity recognition in smart homes. *Pervasive and Mobile Computing*, 5(3):236–252, June 2009.
- Alfred Horn. On sentences which are true of direct unions of algebras. *Journal of Symbolic Logic*, 16(1): 14–21, 1951.
- Ian Horrocks, Ulrike Sattler, and Stephan Tobies. Practical reasoning for very expressive description logics. *Logic Journal of the IGPL*, 8(3):239–263, 2000.
- Ian Horrocks, Peter F Patel-Schneider, Sean Bechhofer, and Dmitry Tsarkov. OWL rules: A proposal and prototype implementation. *Journal of Web Semantics*, 3(1):23–40, 2005.
- Ian Horrocks, Oliver Kutz, and Ulrike Sattler. The even more irresistible SROIQ. In *Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning, Lake District of the United Kingdom, June 2–5, 2006*, pages 57–67. AAAI Press, 2006.
- Anthony Hunter. Merging potentially inconsistent items of structured text. *Data & Knowledge Engineering*, 34(3):305–332, 2000a.
- Anthony Hunter. Reasoning with contradictory information using quasi-classical logic. *Journal of Logic and Computation*, 10(5):677–703, 2000b.
- Anthony Hunter and Sébastien Konieczny. Approaches to measuring inconsistent information. In Leopoldo Bertossi, Anthony Hunter, and Torsten Schaub, editors, *Inconsistency tolerance*, volume 3300 of *LNCS*, pages 191–236. Springer, 2005.
- Anthony Hunter and Bashar Nuseibeh. Managing inconsistent specifications: Reasoning, analysis, and action. *ACM Transactions on Software Engineering and Methodology*, 7(4):335–367, 1998.
- Information Society Technologies Advisory Group (ISTAG). Ambient intelligence: from vision to reality. ISTAG Report, available online: ftp://ftp.cordis.europa.eu/pub/ist/docs/istag-ist2003_consolidated_report.pdf, 2003.
- Katsumi Inoue, Miyuki Koshimura, and Ryuzo Hasegawa. Embedding negation as failure into a model generation theorem prover. In *Automated Deduction – CADE-11, 11th International Conference on Automated Deduction, Saratoga Springs, NY, USA, June 15–18, 1992, Proceedings*, volume 607 of *LNCS*, pages 400–415. Springer, 1992.
- Vikramaditya R. Jakkula, Aaron S. Crandall, and Diane J. Cook. Enhancing anomaly detection using temporal pattern discovery. In Achilles D. Kameas, Victor Callagan, Hani Hagraas, Michael Weber, and Wolfgang Minker, editors, *Advanced Intelligent Environments*, pages 175–194. Springer, 2009.
- Nicholas R. Jennings, Katia P. Sycara, and Michael Wooldridge. A roadmap of agent research and development. *Autonomous Agents and Multi-Agent Systems*, 1(1):7–38, 1998.
- Qiu Ji, Peter Haase, Guilin Qi, Pascal Hitzler, and Steffen Stadtmüller. RaDON – repair and diagnosis in ontology networks. In *The Semantic Web: Research and Applications, 6th European Semantic Web Conference, ESWC 2009, Heraklion, Crete, Greece, May 31–June 4, 2009, Proceedings*, volume 5554 of *LNCS*, pages 863–867. Springer, 2009.

- Yifan Jin, Kewen Wang, and Lian Wen. Possibilistic reasoning in multi-context systems: Preliminary report. In *PRICAI 2012: Trends in Artificial Intelligence – 12th Pacific Rim International Conference on Artificial Intelligence, Kuching, Malaysia, September 3–7, 2012. Proceedings*, volume 7458 of *LNCS*, pages 180–193. Springer, 2012.
- Mathew Joseph and Luciano Serafini. Simple reasoning for contextualized RDF knowledge. In *Modular Ontologies – Proceedings of the Fifth International Workshop, WoMO 2011, Ljubljana, Slovenia, August 2011*, volume 230 of *FAIA*, pages 79–93. IOS Press, 2011.
- Ulrich Junker and Kurt Konolige. Computing the extensions of autoepistemic and default logics with a truth maintenance system. In *Proceedings of the 8th National Conference on Artificial Intelligence. Boston, Massachusetts, July 29–August 3, 1990*, pages 278–283. AAAI Press, 1990.
- Aditya Kalyanpur, Bijan Parsia, Evren Sirin, and Bernardo Cuenca Grau. Repairing unsatisfiable concepts in OWL ontologies. In *The Semantic Web: Research and Applications, 3rd European Semantic Web Conference, ESWC 2006, Budva, Montenegro, June 11–14, 2006, Proceedings*, volume 4011 of *LNCS*, pages 170–184. Springer, 2006.
- Tobias Kaminski, Matthias Knorr, and João Leite. Efficient paraconsistent reasoning with ontologies and rules. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25–31, 2015*, pages 3098–3105. AAAI Press, 2015.
- Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91). Cambridge, MA, USA, April 22–25, 1991.*, pages 387–394. Morgan Kaufmann, 1991.
- Michael Kifer and Eliezer L. Lozinskii. A logic for reasoning with inconsistency. *Journal of Automated Reasoning*, 9(2):179–215, 1992.
- Szymon Klarman and Víctor Gutiérrez-Basulto. Description logics of context. *Journal of Logic and Computation*, To appear. doi: 10.1093/logcom/ext011.
- Matthias Knorr, Martin Slota, João Leite, and Martin Homola. What if no hybrid reasoner is available? hybrid MKNF in multi-context systems. *Journal of Logic and Computation*, 24(6):1279–1311, 2014.
- Stephen Knox, Lorcan Coyle, and Simon Dobson. Using ontologies in case-based activity recognition. In *Proceedings of the Twenty-Third International Florida Artificial Intelligence Research Society Conference, May 19–21, 2010, Daytona Beach, Florida*. AAAI Press, 2010.
- George Konstantinidis, Giorgos Flouris, Grigoris Antoniou, and Vassilis Christophides. A formal approach for RDF/S ontology evolution. In *ECAI 2008 – 18th European Conference on Artificial Intelligence, Patras, Greece, July 21–25, 2008, Proceedings*, volume 178 of *FAIA*, pages 70–74. IOS Press, 2008a.
- George Konstantinidis, Giorgos Flouris, Grigoris Antoniou, and Vassilis Christophides. On RDF/S ontology evolution. In *Semantic Web, Ontologies and Databases, VLDB Workshop, SWDB-ODSIS 2007, Vienna, Austria, September 24, 2007, Revised Selected Papers*, volume 5005 of *LNCS*, pages 21–42. Springer, 2008b.
- Michael Köster, Peter Novák, David Mainzer, and Bernd Fuhrmann. Two case studies for Jazzyk BSM. In *Agents for Games and Simulations, Trends in Techniques, Concepts and Design [AGS 2009, The First International Workshop on Agents for Games and Simulations, May 11, 2009, Budapest, Hungary]*, volume 5920 of *LNCS*, pages 33–47. Springer, 2009.

- Patrick Krümpelmann, Matthias Thimm, Marcelo A. Falappa, Alejandro Javier García, Gabriele Kern-Isberner, and Guillermo Ricardo Simari. Selective revision by deductive argumentation. In *Theory and Applications of Formal Argumentation – First International Workshop, TFAA 2011, Barcelona, Spain, July 16–17, 2011, Revised Selected Papers*, volume 7132 of *LNCS*, pages 147–162. Springer, 2012.
- Benjamin Kuipers. Commonsense reasoning about causality: deriving behavior from structure. *Artificial intelligence*, 24(1):169–203, 1984.
- Oliver Kutz, Frank Wolter, and Michael Zakharyashev. Connecting abstract description systems. In *Proceedings of the Eighth International Conference on Principles and Knowledge Representation and Reasoning (KR-02), Toulouse, France, April 22–25, 2002*, pages 215–226. Morgan Kaufmann, 2002.
- Oliver Kutz, Carsten Lutz, Frank Wolter, and Michael Zakharyashev. \mathcal{E} -connections of description logics. In *Proceedings of the 2003 International Workshop on Description Logics (DL2003), Rome, Italy September 5–7, 2003*, volume 81 of *CEUR Workshop Proceedings*, 2003.
- Joey Sik Chun Lam, Derek H. Sleeman, Jeff Z. Pan, and Wamberto Weber Vasconcelos. A fine-grained approach to resolving unsatisfiable ontologies. *Journal on Data Semantics*, 10:62–95, 2008.
- Sik Chun Lam, Derek. Sleeman, and Wamberto Vasconcelos. ReTAX++: A tool for browsing and revising ontologies. In *ISWC 2005, 4th International Semantic Web Conference, Galway, Ireland, November 6–10, 2005, Posters Track, 2005a*. Available online: <http://homepages.abdn.ac.uk/d.sleeman/pages/published-papers/publications-2005/p155-ISWC-poster-joeypdf.pdf>.
- Sik Chun Joey Lam, Derek Sleeman, and Wamberto Vasconcelos. Retax+: A cooperative taxonomy revision tool. In *Applications and Innovations in Intelligent Systems XII Proceedings of AI-2004, the Twenty-fourth SGA International Conference on Innovative Techniques and Applications of Artificial Intelligence*, pages 64–77. Springer, 2005b.
- Marina Langlois, Robert H. Sloan, Balázs Szörényi, and György Turán. Horn complements: Towards Horn-to-Horn belief revision. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, Chicago, Illinois, USA, July 13–17, 2008*, pages 466–471. AAAI Press, 2008.
- Georg Lausen, Michael Meier, and Michael Schmidt. SPARQLing constraints for RDF. In *EDBT 2008, 11th International Conference on Extending Database Technology, Nantes, France, March 25–29, 2008, Proceedings*, volume 261 of *ACM ICPS*, pages 499–509. ACM, 2008.
- Freddy Lécué, Anika Schumann, and Marco Luca Sbodio. Applying semantic web technologies for diagnosing road traffic congestions. In *The Semantic Web – ISWC 2012, 11th International Semantic Web Conference, Boston, MA, USA, November 11–15, 2012, Proceedings, Part II*, volume 7650 of *LNCS*, pages 114–130. Springer, 2012.
- Kevin Lee and Thomas Meyer. A classification of ontology modification. In *AI 2004: Advances in Artificial Intelligence, 17th Australian Joint Conference on Artificial Intelligence, Cairns, Australia, December 4–6, 2004, Proceedings*, volume 3339 of *LNCS*, pages 248–258. Springer, 2004.
- Jens Lehmann and Lorenz Bühmann. ORE – A tool for repairing and enriching knowledge bases. In *The Semantic Web – ISWC 2010, 9th International Semantic Web Conference, ISWC 2010, Shanghai, China, November 7–11, 2010, Revised Selected Papers, Part II*, volume 6497 of *LNCS*, pages 177–193. Springer, 2010.
- João Leite and João Martins. Social abstract argumentation. In Toby Walsh, editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16–22, 2011*, pages 2287–2292. IJCAI/AAAI, 2011.
- João Alexandre Leite. *Evolving Knowledge Bases*, volume 81 of *FAIA*. IOS Press, 2002.

- João Alexandre Leite and Luís Moniz Pereira. Generalizing updates: From models to programs. In *Logic Programming and Knowledge Representation, Third International Workshop, LPKR '97, Port Jefferson, New York, USA, October 17, 1997, Selected Papers*, volume 1471 of *LNCS*, pages 224–246. Springer, 1998.
- Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati, Marco Ruzzi, and Domenico Fabio Savo. Inconsistency-tolerant semantics for description logics. In *Web Reasoning and Rule Systems – Fourth International Conference, RR 2010, Bressanone/Brixen, Italy, September 22–24, 2010. Proceedings*, volume 6333 of *LNCS*, pages 103–117. Springer, 2010.
- Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati, Marco Ruzzi, and Domenico Fabio Savo. Query rewriting for inconsistent DL-Lite ontologies. In *Web Reasoning and Rule Systems – 5th International Conference, RR 2011, Galway, Ireland, August 29–30, 2011. Proceedings*, volume 6902 of *LNCS*, pages 155–169. Springer, 2011.
- Doug Lenat. The Dimensions of Context-Space. Technical report, CYCorp, 1998. Published online <http://www.cyc.com/doc/context-space.pdf>.
- Jonathan Lester, Tanzeem Choudhury, Nicky Kern, Gaetano Borriello, and Blake Hannaford. A hybrid discriminative/generative approach for modeling human activities. In Leslie Pack Kaelbling and Alessandro Saffiotti, editors, *Proceedings of the 19th international joint conference on Artificial intelligence*, pages 766–772. Professional Book Center, 2005.
- Reinhold Letz. Lemma and model caching in decision procedures for quantified boolean formulas. In *Automated Reasoning with Analytic Tableaux and Related Methods, International Conference, TABLEAUX 2002, Copenhagen, Denmark, July 30–August 1, 2002, Proceedings*, volume 2381 of *LNCS*, pages 160–175. Springer, 2002.
- Beishui Liao. Layered argumentation frameworks with subargument relation and their dynamics. In Eduardo L. Fermé, Dov M. Gabbay, and Guillermo R. Simari, editors, *Trends in Belief Revision and Argumentation Dynamics*, volume 48 of *Logic and Cognitive Systems*. College Publications, 2013.
- Beishui Liao, Li Jin, and Robert C. Koons. Dynamics of argumentation systems: A division-based method. *Artificial Intelligence*, 175(11):1790–1814, 2011.
- Vladimir Lifschitz. Action languages, answer sets, and planning. In Krzysztof R. Apt, Victor W. Marek, Mirek Truszczynski, and David S. Warren, editors, *The Logic Programming Paradigm: A 25 Years Perspective*, pages 357–373. Springer, 1999.
- Hongkai Liu, Carsten Lutz, Maja Miličić, and Frank Wolter. Updating description logic aboxes. In *Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning, Lake District of the United Kingdom, June 2–5, 2006*, pages 46–56. AAAI Press, 2006.
- John W. Lloyd. *Foundations of Logic Programming*. Springer, 1984.
- Seng W. Loke. Representing and reasoning with situations for context-aware pervasive computing: a logic programming perspective. *The Knowledge Engineering Review*, 19(3):213–233, 2004.
- Ching-Hu Lu and Li-Chen Fu. Robust location-aware activity recognition using wireless sensor network in an attentive home. *IEEE Transactions on Automation Science and Engineering*, 6(4):598–609, 2009.
- Thomas Lukasiewicz, Maria Vanina Martinez, and Gerardo Ignacio Simari. Preference-based query answering in datalog+/- ontologies. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3–9, 2013*. AAAI, 2013.
- Yue Ma and Pascal Hitzler. Paraconsistent reasoning for OWL 2. In *Web Reasoning and Rule Systems, Third International Conference, RR 2009, Chantilly, VA, USA, October 25–26, 2009, Proceedings*, volume 5837 of *LNCS*, pages 197–211. Springer, 2009.

- Yue Ma, Pascal Hitzler, and Zuoquan Lin. Algorithms for paraconsistent reasoning with OWL. In *The Semantic Web: Research and Applications, 4th European Semantic Web Conference, ESWC 2007, Innsbruck, Austria, June 3–7, 2007, Proceedings*, volume 4519 of *LNCS*, pages 399–413. Springer, 2007.
- Matoula Magiridou, S. Sahtouris, Vassilis Christophides, and Manolis Koubarakis. RUL: A declarative update language for RDF. In *The Semantic Web – ISWC 2005, 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6–10, 2005, Proceedings*, volume 3729 of *LNCS*, pages 506–521. Springer, 2005.
- Jean-Guy Mailly. Dynamic of argumentation frameworks. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3–9, 2013*. AAAI, 2013.
- David Makinson. On the status of the postulate of recovery in the logic of theory change. *Journal of Philosophical Logic*, 16(4):383–394, 1987.
- Maria Vanina Martinez, Cristhian Ariel David Deagustini, Marcelo A. Falappa, and Guillermo Ricardo Simari. Inconsistency-tolerant reasoning in datalog[±] ontologies via an argumentative semantics. In *Advances in Artificial Intelligence – IBERAMIA 2014, 14th Ibero-American Conference on AI, Santiago de Chile, Chile, November 24–27, 2014, Proceedings*, volume 8864 of *LNCS*, pages 15–27. Springer, 2014.
- Giulia Masotti, Riccardo Rosati, and Marco Ruzzi. Practical abox cleaning in dl-lite (progress report). In *Proceedings of the 24th International Workshop on Description Logics (DL 2011), Barcelona, Spain, July 13–16, 2011*, volume 745 of *CEUR Workshop Proceedings*, 2011.
- Fulvio Mastrogiovanni, Antonello Scalmato, Antonio Sgorbissa, and Renato Zaccaria. Smart environments and activity recognition: A logic-based approach. In Liming Chen, Chris Nugent, Jit Biswas, and Jesse Hoey, editors, *Activity Recognition in Pervasive Intelligent Environments*, volume 4 of *Atlantis Ambient and Pervasive Intelligence*, pages 83–109. Atlantis Press, 2011.
- John McCarthy. Circumscription – A form of non-monotonic reasoning. *Artificial Intelligence*, 13(1–2): 27–39, April 1980.
- John McCarthy. Notes on formalizing context. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence. Chambéry, France, August 28–September 3, 1993*, pages 555–562. Morgan Kaufmann, 1993.
- Susan McKeever, Juan Ye, Lorcan Coyle, Chris Bleakley, and Simon Dobson. Activity recognition using temporal evidence theory. *Journal of Ambient Intelligence and Smart Environments*, 2(3):253–269, 2010.
- Christian Meilicke, Heiner Stuckenschmidt, and Andrei Tamilin. Repairing ontology mappings. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, July 22–26, 2007, Vancouver, British Columbia, Canada*, pages 1408–1413. AAAI Press, 2007.
- Christian Meilicke, Heiner Stuckenschmidt, and Andrei Tamilin. Reasoning support for mapping revision. *Journal of Logic and Computation*, 19(5):807–829, 2009.
- Thomas Meyer, Kevin Lee, and Richard Booth. Knowledge integration for description logics. In *Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference, July 9–13, 2005, Pittsburgh, Pennsylvania, USA*, pages 645–650. AAAI Press, 2005.

- Thomas Meyer, Kevin Lee, Richard Booth, and Jeff Z. Pan. Finding maximally satisfiable terminologies for the description logic \mathcal{ALC} . In *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, July 16–20, 2006, Boston, Massachusetts, USA*, pages 269–274. AAAI Press, 2006.
- Alessandra Mileo, Davide Merico, and Roberto Bisiani. Wireless sensor networks supporting context-aware reasoning in assisted living. In *Proceedings of the 1st ACM International Conference on Pervasive Technologies Related to Assistive Environments, PETRA 2008, Athens, Greece, July 16–18, 2008*, volume 282 of *ACM ICPS*, page 54. ACM, 2008a.
- Alessandra Mileo, Davide Merico, and Roberto Bisiani. A logic programming approach to home monitoring for risk prevention in assisted living. In *Logic Programming, 24th International Conference, ICLP 2008, Udine, Italy, December 9–13 2008, Proceedings*, volume 5366 of *LNCS*, pages 145–159. Springer, 2008b.
- Assaad Moawad, Antonis Bikakis, Patrice Caire, Grégory Nain, and Yves Le Traon. A rule-based contextual reasoning platform for ambient intelligence environments. In *Theory, Practice, and Applications of Rules on the Web – 7th International Symposium, RuleML 2013, Seattle, WA, USA, July 11–13, 2013. Proceedings*, volume 8035 of *LNCS*, pages 158–172. Springer, 2013.
- Sanjay Modgil and Martin Caminada. Proof theories and algorithms for abstract argumentation frameworks. In Iyad Rahwan and Guillermo R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 105–129. Springer, 2009.
- Martin O. Moguillansky, Nicolás D. Rotstein, and Marcelo A. Falappa. A theoretical model to handle ontology debugging & change through argumentation. In *Proceedings of the 2nd International Workshop on Ontology Dynamics (IWOD 2008), collocated with the 7th International Semantic Web Conference (ISWC 2008), October 26–30, Karlsruhe, Germany, 2008*. Available online: <http://users.ics.forth.gr/~fgeo/files/IWOD08Proc.pdf>.
- Martín O. Moguillansky, Nicolás D. Rotstein, Marcelo A. Falappa, Alejandro Javier García, and Guillermo Ricardo Simari. Dynamics of knowledge in *DeLP* through argument theory change. *Theory and Practice of Logic Programming*, 13(6):893–957, 2013.
- Pavlos Moraitis and Nikolaos Spanoudakis. Argumentation-based agent interaction in an ambient-intelligence context. *IEEE Intelligent Systems*, 22(6):84–93, 2007.
- Boris Motik, Ian Horrocks, and Ulrike Sattler. Bridging the gap between OWL and relational databases. *Journal of Web Semantics*, 7(2):74–89, 2009.
- Andrés Muñoz, Juan A. Botía, and Juan Carlos Augusto. Intelligent decision-making for a smart home environment with multiple occupants. In Da Ruan, editor, *Computational Intelligence in Complex Decision Systems*, volume 2 of *Atlantis Computational Intelligence Systems*, pages 325–371. Atlantis Press, 2010.
- Andrés Muñoz, Juan Carlos Augusto, Ana Villa, and Juan Antonio Botía. Design and evaluation of an ambient assisted living system based on an argumentative multi-agent system. *Personal Ubiquitous Computing*, 15(4):377–387, April 2011.
- Andrés Muñoz Ortega, Juan A. Botía Blaya, Félix J. García Clemente, Gregorio Martínez Pérez, and Antonio F. Gómez Skarmeta. Solving conflicts in agent-based ubiquitous computing systems: A proposal based on argumentation. In Eleni Mangina, Javier Carbo, and Jose M. Molina, editors, *Agent-Based Ubiquitous Computing*, volume 1 of *Atlantis Ambient and Pervasive Intelligence*, pages 1–12. Atlantis Press, 2010.
- Erik T. Mueller. *Commonsense Reasoning: An Event Calculus Based Approach*. Morgan Kaufmann, 2nd edition, 2014.

- Ilkka Niemelä. Towards efficient default reasoning. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, IJCAI 95, Montréal Québec, Canada, August 20–25 1995*, 2 Volumes, pages 312–318. Morgan Kaufmann, 1995.
- Monica Nogueira, Marcello Balduccini, Michael Gelfond, Richard Watson, and Matthew Barry. An a-prolog decision support system for the space shuttle. In *Practical Aspects of Declarative Languages, Third International Symposium, PADL 2001, Las Vegas, Nevada, March 11–12, 2001, Proceedings*, volume 1990 of LNCS, pages 169–183. Springer, 2001.
- Natalya F. Noy, Ray W. Ferguson, and Mark A. Musen. The knowledge model of Protégé-2000: Combining interoperability and flexibility. In *Knowledge Acquisition, Modeling and Management, 12th International Conference, EKAW 2000, Juan-les-Pins, France, October 2–6, 2000, Proceedings*, volume 1937 of LNCS, pages 17–32. Springer, 2000.
- Natalya F. Noy, Abhita Chugh, William Liu, and Mark A. Musen. A framework for ontology evolution in collaborative environments. In *The Semantic Web – ISWC 2006, 5th International Semantic Web Conference, ISWC 2006, Athens, GA, USA, November 5–9, 2006, Proceedings*, volume 4273 of LNCS, pages 544–558. Springer, 2006.
- Sergei P. Odintsov and David Pearce. Routley semantics for answer sets. In *Logic Programming and Nonmonotonic Reasoning, 8th International Conference, LPNMR 2005, Diamante, Italy, September 5–8, 2005, Proceedings*, volume 3662 of LNCS, pages 343–355. Springer, 2005.
- Hans Jürgen Ohlbach. SCAN – Elimination of predicate quantifiers. In *Automated Deduction – CADE-13, 13th International Conference on Automated Deduction, New Brunswick, NJ, USA, July 30–August 3, 1996, Proceedings*, volume 1104 of LNCS, pages 161–165. Springer, 1996.
- Sascha Ossowski, editor. *Agreement Technologies*. Springer, 2013.
- OWL Working Group, editor. *OWL 2 Web Ontology Language Document Overview*. W3C Recommendation. 27 October 2009.
- Sergio Pajares Ferrando and Eva Onaindia. Defeasible argumentation for multi-agent planning in ambient intelligence applications. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4–8, 2012*, pages 509–516. IFAAMAS, 2012.
- Simon Parsons, Carles Sierra, and Nicholas R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.
- Peter F. Patel-Schneider. A four-valued semantics for terminological logics. *Artificial Intelligence*, 38(3): 319–351, 1989.
- Theodore Patkos, Ioannis Chrysakis, Antonis Bikakis, Dimitris Plexousakis, and Grigoris Antoniou. A reasoning framework for ambient intelligence. In *Artificial Intelligence: Theories, Models and Applications, 6th Hellenic Conference on AI, SETN 2010, Athens, Greece, May 4–7, 2010, Proceedings*, volume 6040 of LNCS, pages 213–222. Springer, 2010.
- Federico Pecora, Marcello Cirillo, Francesca Dell’Osa, Jonas Ullberg, and Alessandro Saffiotti. A constraint-based approach for proactive, context-aware human support. *Journal of Ambient Intelligence and Smart Environments*, 4(4):347–367, 2012.
- Peter Plessers and Olga De Troyer. Resolving inconsistencies in evolving ontologies. In *The Semantic Web: Research and Applications, 3rd European Semantic Web Conference, ESWC 2006, Budva, Montenegro, June 11–14, 2006, Proceedings*, volume 4011 of LNCS, pages 200–214. Springer, 2006.
- John L. Pollock. *Cognitive Carpentry: A Blueprint for How to Build a Person*. MIT Press, 1995.

- Henry Prakken. An abstract framework for argumentation with structured arguments. *Argument & Computation*, 1(2):93–124, 2010.
- Henry Prakken and Giovanni Sartor. Argument-based logic programming with defeasible priorities. *Journal of Applied Non-classical Logics*, 7(1):25–75, 1997.
- Henry Prakken and Gerard Vreeswijk. Logics for defeasible argumentation. In Dow M. Gabbay and Guenther Franz, editors, *Handbook of Philosophical Logic*, volume 4, pages 219–318. Kluwer Academic Publishers, 2002.
- Davy Preuveneers, Jan Van den Bergh, Dennis Wagelaar, Andy Georges, Peter Rigole, Tim Clerckx, Yolande Berbers, Karin Coninx, Viviane Jonckers, and Koen De Bosschere. Towards an extensible context ontology for ambient intelligence. In *Ambient Intelligence: Second European Symposium, EUSAI 2004, Eindhoven, The Netherlands, November 8–11, 2004. Proceedings*, volume 3295 of LNCS, pages 148–159. Springer, 2004.
- Guide Qi and Jeff Z. Pan. A stratification-based approach for inconsistency handling in description logics. In *Proceedings, International Workshop on Ontology Dynamics (IWOD-07), Held as part of the 4th European Semantic Web Conference (ESWC-07), Innsbruck, Austria, June 7, 2007*, pages 83–96, 2007. Available online: <http://users.ics.forth.gr/~fgeo/files/IWOD07Proc.pdf>.
- Guilin Qi and Jianfeng Du. Model-based revision operators for terminologies in description logics. In *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11–17, 2009*, pages 891–897, 2009.
- Guilin Qi, Weiru Liu, and David A. Bell. A revision-based approach to handling inconsistency in description logics. *Artificial Intelligence Review*, 26(1-2):115–128, 2006a.
- Guilin Qi, Weiru Liu, and David A. Bell. Knowledge base revision in description logics. In *Logics in Artificial Intelligence, 10th European Conference, JELIA 2006, Liverpool, UK, September 13–15, 2006, Proceedings*, volume 4160 of LNCS, pages 386–398. Springer, 2006b.
- Guilin Qi, Zhe Wang, Kewen Wang, Xuefeng Fu, and Zhiqiang Zhuang. Approximating model-based abox revision in dl-lite: Theory and practice. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25–30, 2015, Austin, Texas, USA.*, pages 254–260. AAAI Press, 2015.
- Iyad Rahwan and Guillermo R. Simari, editors. *Argumentation in Artificial Intelligence*. Springer, 2009.
- Anand S. Rao and Michael P. Georgeff. Modeling rational agents within a BDI-architecture. In *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91). Cambridge, MA, USA, April 22–25, 1991*, pages 473–484. Morgan Kaufmann, 1991.
- Anand S. Rao and Michael P. Georgeff. BDI agents: From theory to practice. In *Proceedings of the First International Conference on Multiagent Systems, June 12–14, 1995, San Francisco, California, USA*, pages 312–319. The MIT Press, 1995.
- Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1–2):81–132, 1980.
- Sílvia Resendes, Paulo Carreira, and André C. Santos. Conflict detection and resolution in home and building automation systems: a literature review. *Journal of Ambient Intelligence and Humanized Computing*, 5(5):699–715, 2014.
- Alexandre Riazanov and Andrei Voronkov. The design and implementation of VAMPIRE. *AI Communications*, 15(2–3):91–110, 2002.

- Márcio Moretto Ribeiro, , and Renata Wassermann. Base revision in Description Logics – preliminary results. In *Proceedings, International Workshop on Ontology Dynamics (IWOD-07), Held as part of the 4th European Semantic Web Conference (ESWC-07), Innsbruck, Austria, June 7, 2007*, pages 69–82, 2007. Available online: <http://users.ics.forth.gr/~fgeo/files/IWOD07Proc.pdf>.
- Márcio Moretto Ribeiro, Renata Wassermann, Giorgos Flouris, and Grigoris Antoniou. Minimal change: Relevance and recovery revisited. *Artificial Intelligence*, 201:59–80, 2013.
- Daniele Riboni and Claudio Bettini. COSAR: hybrid reasoning for context-aware activity recognition. *Personal and Ubiquitous Computing*, 15(3):271–289, 2011a.
- Daniele Riboni and Claudio Bettini. OWL 2 modeling and reasoning with complex human activities. *Pervasive and Mobile Computing*, 7(3):379–395, 2011b.
- Daniele Riboni, Linda Pareschi, Laura Radaelli, and Claudio Bettini. Is ontology-based activity recognition really effective? In *Ninth Annual IEEE International Conference on Pervasive Computing and Communications, PerCom 2011, 21–25 March 2011, Seattle, WA, USA, Workshop Proceedings*, pages 427–431. IEEE, 2011.
- Christoph Rieß, Norman Heino, Sebastian Tramp, and Sören Auer. EvoPat – Pattern-based evolution and refactoring of RDF knowledge bases. In *The Semantic Web – ISWC 2010, 9th International Semantic Web Conference, ISWC 2010, Shanghai, China, November 7–11, 2010, Revised Selected Papers, Part I*, volume 6496 of *LNCS*, pages 647–662. Springer, 2010.
- Andrew Rock. Deimos: Query answering defeasible logic system. Technical report, Griffith University, School of Computing and Information Technology, 2009. Available online: <http://www.ict.griffith.edu.au/arock/defeasible/doc/Deimos-long.pdf>, see also the online system: <http://www.ict.griffith.edu.au/arock/defeasible/Defeasible.cgi>.
- Odinaldo Rodrigues and Alessandra Russo. A translation method for belnap logic. Research Report Doc 98/7, Imperial College Longon, 1998. Available online: <http://www.doc.ic.ac.uk/research/technicalreports/1998/DTR98-7.pdf>.
- Andrea Rodríguez. Inconsistency issues in spatial databases. In Leopoldo E. Bertossi, Anthony Hunter, and Torsten Schaub, editors, *Inconsistency Tolerance*, volume 3300 of *LNCS*, pages 237–269. Springer, 2005.
- Floris Roelofsen and Luciano Serafini. Minimal and absent information in contexts. In *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30–August 5, 2005*, pages 558–563. Professional Book Center, 2005.
- Mathieu Roger, Ana Simonet, and Michel Simonet. Toward updates in description logics. In *Proceedings of the 9th International Workshop on Knowledge Representation meets Databases (KRDB 2002), Toulouse France, April 21, 2002*, volume 54 of *CEUR Workshop Proceedings*, 2002.
- Nicolás D. Rotstein, Martín O. Moguillansky, Alejandro J. García, and Guillermo R. Simari. A dynamic argumentation framework. In *Computational Models of Argument: Proceedings of COMMA 2010, Desenzano del Garda, Italy, September 8–10, 2010.*, volume 216 of *FAIA*, pages 427–438. IOS Press, 2010.
- Hans Rott. Preferential belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information*, 1(1):45–78, 1992.
- Yannis Roussakis, Giorgos Flouris, and Vassilis Christophides. Declarative repairing policies for curated KBs. In *In Proceedings of the 10th Hellenic Data Management Symposium (HDMS-11)*, 2011.

- Patrice C. Roy, Sylvain Giroux, Bruno Bouchard, Abdenour Bouzouane, Clifton Phua, Andrei Tolstikov, and Jit Biswas. A possibilistic approach for activity recognition in smart homes for cognitive assistance to Alzheimer's patients. In Liming Chen, Chris D. Nugent, Jit Biswas, and Jesse Hoey, editors, *Activity Recognition in Pervasive Intelligent Environments*, volume 4 of *Atlantis Ambient and Pervasive Intelligence*, pages 33–58. Atlantis Press, 2011.
- Paul Rubel, Jocelyne Fayn, Lucas Simon-Chautemps, Hussein Atoui, Mattias Ohlsson, David Telisson, Stefano Adami, Sébastien Arod, Marie Claire Forlini, Cesare Malossi, Joël Placide, Gian Luca Ziliani, Deodato Assanelli, and Philippe Chevalier. New paradigms in telemedicine: Ambient intelligence, wearable, pervasive and personalized. In Andreas Lymberis and Danilo de Rossi, editors, *Wearable eHealth Systems for Personalised Health Management*, volume 108 of *Studies in Health Technology and Informatics*, pages 123–132. IOS Press, 2004.
- Alberto Rugnone, Enrico Vicario, Chris D. Nugent, Mark P. Donnelly, David Craig, Cristiano Paggetti, and Elena Tamburini. HomeTL: A visual formalism, based on temporal logic, for the design of home based care. In *IEEE Conference on Automation Science and Engineering, CASE 2007, September 22–25, 2007, Scottsdale, Arizona, USA*, pages 747–752. IEEE, 2007.
- Jordi Sabater, Carles Sierra, Simon Parsons, and Nicholas R. Jennings. Engineering executable agents using multi-context systems. *Journal of Logic and Computation*, 12(3):413–442, 2002.
- Fariba Sadri. Intention Recognition with Event Calculus Graphs. In *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence and International Conference on Intelligent Agent Technology – Workshops, Toronto, Canada, August 31–September 3, 2010*, pages 386–391. IEEE Computer Society, 2010.
- Fariba Sadri. Ambient intelligence: A survey. *ACM Computing Surveys*, 43(4):36:1–36:66, 2011.
- Chiaki Sakama. Extended Well-Founded Semantics for Paraconsistent Logic Programs. In *Fifth Generation Computer Systems '92: Proceedings of the International Conference on Fifth Generation Computer Systems. 1992, June 1–5, Tokyo, Japan*, pages 592–599. IOS Press, 1992.
- Chiaki Sakama and Katsumi Inoue. Paraconsistent Stable Semantics for Extended Disjunctive Programs. *Journal of Logic and Computation*, 5(3):265–285, 1995.
- Chiaki Sakama and Katsumi Inoue. Prioritized logic programming and its application to commonsense reasoning. *Artificial Intelligence*, 2000.
- Torsten Schaub. Collection on answer set programming (asp) and more, 2011. Available online: <http://www.cs.uni-potsdam.de/~torsten/asp/>.
- Torsten Schaub and Kewen Wang. Preferred well-founded semantics for logic programming by alternating fixpoints: preliminary report. In *9th International Workshop on Non-Monotonic Reasoning (NMR 2002), April 19–21, Toulouse, France, Proceedings*, pages 238–246, 2002.
- Torsten Schaub and Kewen Wang. A semantic framework for preference handling in answer set programming. *Theory and Practice of Logic Programming*, 3(4–5):569–607, 2003.
- Stefan Schlobach and Ronald Cornet. Non-standard reasoning services for the debugging of description logic terminologies. In *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 9–15, 2003*, pages 355–362. Morgan Kaufmann, 2003.
- Faouzi Sebbak, Abdelghani Chibani, Yacine Amirat, Farid Benhammadi, and Aïcha Mokhtari. An evidential fusion approach for activity recognition under uncertainty in ambient intelligence environments. In *The 2012 ACM Conference on Ubiquitous Computing, Ubicomp'12, Pittsburgh, PA, USA, September 5–8, 2012*, pages 834–840. ACM, 2012.

- Ján Šefránek. Preferred answer sets supported by arguments. In *Proceedings of the Twelfth International Workshop on Non-Monotonic Reasoning*, volume 819 of *UNSW-CSE-TR*, pages 232–240. The University of New South Wales, School of Computer Science and Engineering, 2008.
- Ján Šefránek and Alexander Šimko. Warranted Derivations of Preferred Answer Sets. In *19th International Conference on Applications of Declarative Programming and Knowledge Management (INAP 2011) and 25th Workshop on Logic Programming (WLP 2011), Vienna, September 28–30, 2011 proceedings*, volume 1842-11-06 of *INFSYS Research Report*, pages 195–207. Technische Universität Wien, Austria, 2011.
- Ján Šefránek and Alexander Šimko. A descriptive approach to preferred answer sets. In *Applications of Declarative Programming and Knowledge Management – 19th International Conference, INAP 2011, and 25th Workshop on Logic Programming, WLP 2011, Vienna, Austria, September 28–30, 2011, Revised Selected Papers*, volume 7773 of *LNCS*, pages 195–214. Springer, 2013.
- Luciano Serafini and Martin Homola. Contextualized knowledge repositories for the Semantic Web. *Journal of Web Semantics, Special Issue: Reasoning with context in the Semantic Web*, 12, 2012.
- Luciano Serafini and Andrei Taminin. Local tableaux for reasoning in distributed description logics. In *Proceedings of the 2004 International Workshop on Description Logics (DL2004), Whistler, British Columbia, Canada, June 6–8, 2004*, volume 104 of *CEUR Workshop Proceedings*, 2004.
- Luciano Serafini and Andrei Taminin. Drago: Distributed reasoning architecture for the semantic web. In *The Semantic Web: Research and Applications, Second European Semantic Web Conference, ESWC 2005, Heraklion, Crete, Greece, May 29–June 1, 2005, Proceedings*, volume 3532 of *LNCS*, pages 361–376. Springer, 2005.
- Luciano Serafini, Alexander Borgida, and Andrei Taminin. Aspects of distributed and modular ontology reasoning. In *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30–August 5, 2005*, pages 570–575, 2005.
- Giorgos Serfiotis, Ioanna Koffina, Vassilis Christophides, and Val Tannen. Containment and minimization of RDF/S query patterns. In *The Semantic Web – ISWC 2005, 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6–10, 2005, Proceedings*, volume 3729 of *LNCS*, pages 607–623. Springer, 2005.
- Paulo Shakarian, Gerardo I. Simari, and Marcelo A. Falappa. Belief revision in structured probabilistic argumentation. In *Foundations of Information and Knowledge Systems – 8th International Symposium, FoIKS 2014, Bordeaux, France, March 3–7, 2014. Proceedings*, volume 8367 of *LNCS*, pages 324–343. Springer, 2014.
- Alexander Šimko. Extension of Gelfond-Lifschitz Reduction for Preferred Answer Sets : Preliminary Report. In *Kiel Declarative Programming Days 2013: 20th International Conference on Applications of Declarative Programming and Knowledge Management (INAP 2013), 22nd International Workshop on Functional and (Constraint) Logic Programming (WFLP 2013), 27th Workshop on Logic Programming (WLP 2013)*, volume 1306 of *Technical Report/Bericht*, pages 2–16. Institut für Informatik der Christian-Albrechts-Universität zu Kiel, 2013.
- Geetika Singla, Diane J. Cook, and Maureen Schmitter-Edgecombe. Recognizing independent and joint activities among multiple residents in smart environments. *Journal of Ambient Intelligence and Humanized Computing*, 1(1):57–63, 2010.
- Evren Sirin, Bijan Parsia, Bernardo Cuenca Grau, Aditya Kalyanpur, and Yarden Katz. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics*, 5(2):51–53, 2007.

- Anastasios Skarlatidis, Georgios Paliouras, George A. Vouros, and Alexander Artikis. Probabilistic Event Calculus Based on Markov Logic Networks. In *Rule-Based Modeling and Computing on the Semantic Web, 5th International Symposium, RuleML 2011-America, Ft. Lauderdale, FL, Florida, USA, November 3–5, 2011. Proceedings*, volume 7018 of *LNCS*, pages 155–170. Springer, 2011.
- Mark Snaith and Chris Reed. TOAST: online ASPIC⁺ implementation. In *Computational Models of Argument, Proceedings of COMMA 2012, Vienna, Austria, September 10–12, 2012*, volume 245 of *FAIA*, pages 509–510. IOS Press, 2012.
- Tran Cao Son and Jorge Lobo. Reasoning about policies using logic programs. In *Answer Set Programming, Towards Efficient and Scalable Knowledge Representation and Reasoning, Proceedings of the 1st Intl. ASP'01 Workshop, Stanford, March 26–28, 2001*. AAAI, 2001.
- Tran Cao Son, Enrico Pontelli, and Chiaki Sakama. Logic programming for multiagent planning with negotiation. In *Logic Programming, 25th International Conference, ICLP 2009, Pasadena, CA, USA, July 14–17, 2009. Proceedings*, volume 5649 of *LNCS*, pages 99–114. Springer, 2009.
- John F. Sowa. *Knowledge representation: logical, philosophical, and computational foundations*. Brooks/Cole Publishing, Pacific Grove, CA, USA, 2000.
- Thomas Springer and Anni-Yasmin Turhan. Employing description logics in ambient intelligence for modeling and reasoning about complex situations. *Journal of Ambient Intelligence and Smart Environments*, 1(3):235–259, 2009.
- Steffen Staab and Rudi Studer, editors. *Handbook on Ontologies*. Springer, 2004.
- Kostas Stefanidis, Georgia Koutrika, and Evaggelia Pitoura. A survey on representation, composition and application of preferences in database systems. *ACM Transactions on Database Systems*, 36(3):19, 2011.
- Ljiljana Stojanovic, Alexander Maedche, Boris Motik, and Nenad Stojanovic. User-driven ontology evolution management. In *Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web, 13th International Conference, EKAW 2002, Sigüenza, Spain, October 1–4, 2002, Proceedings*, volume 2473 of *LNCS*, pages 285–300. Springer, 2002.
- Umberto Straccia. A sequent calculus for reasoning in four-valued description logics. In *Automated Reasoning with Analytic Tableaux and Related Methods, International Conference, TABLEAUX '97, Pont-à-Mousson, France, May 13–16, 1997, Proceedings*, volume 1227 of *LNCS*, pages 343–357. Springer, 1997.
- Thomas Strang and Claudia Linnhoff-Popien. A context modeling survey. In *Workshop on Advanced Context Modelling, Reasoning and Management, UbiComp 2004 – The Sixth International Conference on Ubiquitous Computing, Nottingham, England, September 7, 2004*, 2004.
- York Sure, Jürgen Angele, and Steffen Staab. OntoEdit: Multifaceted inferencing for ontology engineering. *Journal on Data Semantics*, 1(1):128–152, 2003.
- Joo Geok Tan, Daqing Zhang, Xiaohang Wang, and Heng Seng Cheng. Enhancing semantic spaces with event-driven context interpretation. In Hans-Werner Gellersen, Roy Want, and Albrecht Schmidt, editors, *Proceedings of the Third international conference on Pervasive Computing*, pages 80–97, 2005.
- Jiao Tao, Evren Sirin, Jie Bao, and Deborah L. McGuinness. Extending OWL with integrity constraints. In *Proceedings of the 23rd International Workshop on Description Logics (DL 2010), Waterloo, Ontario, Canada, May 4–7, 2010*, volume 573 of *CEUR Workshop Proceedings*, 2010.

- Allen Van Gelder, Kenneth A. Ross, and John S. Schlipf. Unfounded sets and well-founded semantics for general logic programs. In *Proceedings of the Seventh ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, March 21–23, 1988, Austin, Texas, USA*, pages 221–230. ACM, 1988.
- Allen Van Gelder, Kenneth A. Ross, and John S. Schlipf. The well-founded semantics for general logic programs. *Journal of the ACM*, 38(3):620–650, 1991.
- Hai Wang, Matthew Horridge, Alan L. Rector, Nick Drummond, and Julian Seidenberg. Debugging OWL-DL ontologies: A heuristic approach. In *The Semantic Web – ISWC 2005, 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6–10, 2005, Proceedings*, volume 3729 of LNCS, pages 745–757. Springer, 2005.
- Kewen Wang, Lizhu Zhou, and Fangzhen Lin. Alternating fixpoint theory for logic programs with priority. In *Computational Logic – CL 2000, First International Conference, London, UK, 24–28 July, 2000, Proceedings*, volume 1861 of LNCS, pages 164–178. Springer, 2000.
- Xiaohang Wang, Jin Song Dong, Chung-Yau Chin, Sanka Ravipriya Hettiarachchi, and Daqing Zhang. Semantic space: An infrastructure for smart spaces. *IEEE Pervasive Computing*, 3(3):32–39, 2004.
- Zhe Wang, Kewen Wang, and Rodney Topor. A new approach to knowledge base revision in DL-Lite. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11–15, 2010*. AAAI Press, 2010.
- Mark Weiser. The Computer for the 21st Century. *Scientific American*, 265(3):94–104, 1991.
- Jef Wijsen. Consistent query answering under primary keys: a characterization of tractable queries. In *Database Theory – ICDT 2009, 12th International Conference, St. Petersburg, Russia, March 23–25, 2009, Proceedings*, volume 361 of ACM ICPS, pages 42–52. ACM, 2009.
- Michael Wooldridge and Nicholas R Jennings. Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.
- Jianxin Wu, Adebola Osuntogun, Tanzeem Choudhury, Matthai Philipose, and James M. Rehg. A scalable approach to activity recognition based on object use. In *IEEE 11th International Conference on Computer Vision, ICCV 2007, Rio de Janeiro, Brazil, October 14–20, 2007*, pages 1–8. IEEE, 2007.
- Qiang Yang. Activity recognition: Linking low-level sensors to high-level intelligence. In *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11–17, 2009*, pages 20–25, 2009.
- Juan Ye, Lorcan Coyle, Simon Dobson, and Paddy Nixon. Ontology-based models in pervasive computing systems. *The Knowledge Engineering Review*, 22(4):315–347, 2007.
- Juan Ye, Simon Dobson, and Susan McKeever. Situation identification techniques in pervasive computing: A review. *Pervasive and Mobile Computing*, 8(1):36–66, 2012.
- Fouad Zablith, Grigoris Antoniou, Mathieu d’Aquin, Giorgos Flouris, Haridimos Kondylakis, Enrico Motta, Dimitris Plexousakis, and Marta Sabou. Ontology evolution: a process-centric survey. *The Knowledge Engineering Review*, 30(1):45–75, 2015.
- Eli Zelkha. The future of information appliances and consumer devices. In *Palo Alto Ventures, Palo Alto, California, (unpublished document)*, 1998.
- Xiaowang Zhang and Zuoquan Lin. Quasi-classical description logic. *Multiple-Valued Logic and Soft Computing*, 18(3–4):291–327, 2012.

- Xiaowang Zhang, Guohui Xiao, and Zuoquan Lin. A tableau algorithm for handling inconsistency in OWL. In *The Semantic Web: Research and Applications, 6th European Semantic Web Conference, ESWC 2009, Heraklion, Crete, Greece, May 31–June 4, 2009, Proceedings*, volume 5554 of *LNCS*, pages 399–413. Springer, 2009.
- Yan Zhang and Norman Y. Foo. Answer sets for prioritized logic programs. In *Logic Programming, Proceedings of the 1997 International Symposium, Port Jefferson, Long Island, NY, USA, October 13–16, 1997*, pages 69–83. MIT Press, 1997.
- Zhi Qiang Zhuang and Maurice Pagnucco. Horn contraction via epistemic entrenchment. In *Logics in Artificial Intelligence – 12th European Conference, JELIA 2010, Helsinki, Finland, September 13–15, 2010. Proceedings*, volume 6341 of *LNCS*, pages 339–351. Springer, 2010.
- Zhi Qiang Zhuang and Maurice Pagnucco. Model based Horn contraction. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012, Rome, Italy, June 10–14, 2012*. AAAI Press, 2012.