



Merging Existential Rules Programs in Multi-Agent Contexts through Credibility Accrual

Cristhian A.D. Deagustini, Juan Carlos L. Teze, M. Vanina Martinez, Marcelo A. Falappa, Guillermo R. Simari

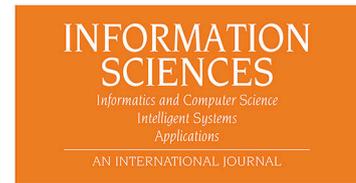
PII: S0020-0255(20)31042-2
DOI: <https://doi.org/10.1016/j.ins.2020.10.050>
Reference: INS 15964

To appear in: *Information Sciences*

Received Date: 10 March 2020
Accepted Date: 16 October 2020

Please cite this article as: C.A.D. Deagustini, J.C. L. Teze, M. Vanina Martinez, M.A. Falappa, G.R. Simari, Merging Existential Rules Programs in Multi-Agent Contexts through Credibility Accrual, *Information Sciences* (2020), doi: <https://doi.org/10.1016/j.ins.2020.10.050>

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



Merging Existential Rules Programs in Multi-Agent Contexts through Credibility Accrual

Cristhian A. D. Deagustini^{a,b,*}, Juan Carlos L. Teze^{a,b}, M. Vanina Martinez^c,
Marcelo A. Falappa^{a,b}, Guillermo R. Simari^b

^a*Grupo de Investigación en Agentes y Sistemas Inteligentes (GINAySI),
Laboratorio de Informática y Sistemas, Facultad de Ciencias de la Administración,
Universidad Nacional de Entre Ríos, (E3202BHR) Concordia, Entre Ríos, Argentina*

^b*Institute for Computer Science and Engineering (UNS-CONICET),*

Department of Computer Science and Engineering,

Universidad Nacional del Sur, Alem 1253, (B8000CPB) Bahía Blanca, Argentina.

^c*Instituto de Investigación en Ciencias de la Computación, Universidad Nacional de
Buenos Aires, Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET),
Pabellón I - Ciudad Universitaria, (C1428EGA) CABA, Argentina.*

Abstract

Merging operators represent a significant tool to extract a consistent and informative view from a set of agents. The consideration of practical scenarios where some agents can be more credible than others has contributed to substantially increase the interest in developing systems working with trust models. In this context, we propose an approach to the problem of merging knowledge in a multiagent scenario where every agent assigns to other agents a value reflecting its perception on how credible each agent is. The focus of this paper is the introduction of an operator for merging Datalog[±] ontologies considering agents' credibility. We present a procedure to enhance a conflict resolution strategy by exploiting the credibility attached to a set of formulas; the approach is based on accrual functions that calculate the value of formulas according to the credibility of the agents that inform them. We show how our new operator can obtain the best-valued knowledge base among consistent bases available, according to the credibilities attached to the sources.

Keywords: Belief Revision, Ontologies Merging, Belief Accrual, Multi-agent Systems, Trust

1. Introduction

Information sharing in practical applications is nowadays a common aspect that may come in several different flavors, such as the creation of collective knowledge in social environments (knowledge introduced and maintained by

*Corresponding author

Email address: caddeagustini@gmail.com (Cristhian A. D. Deagustini)

different members of the community), data repositories shared across organizations, news repositories receiving information from a multitude of sources, and many other possible scenarios. While this design decision has several advantages, like greater consensus and bias removal, the decentralization of knowledge maintenance could also have disadvantages, in particular regarding conflicting pieces of information leading to inconsistencies in the knowledge. Through the years many proposals dealing with such issues were introduced, ranging from the field of AI [6, 7, 11] and Belief Revision [1, 13, 24] to that of Databases [5, 33, 38], to cite a few of an extensive literature.

Seeking generality, we will focus on multi-agent application domains where agents of a community communicate with each other sharing their knowledge. Each agent can act as the source of some piece of information (or as an informant, in [48]'s terminology) to other agents, and some agents can be more credible than others. In particular, we present a formalization where every agent will store the credibility degree that it associates with other agents in the community, and this value will reflect its perception of that agent's credibility.

Our proposal introduces an approach to deal with the problem of merging potentially conflicting knowledge in a multi-agent scenario, exploiting the credibility associated with the different agents in the community. Despite its direct application to real-world problems, the mechanisms to deal with this problem have not been extensively studied in the literature. For instance, some works propose a formalization of merging operators [30], but they are not based on trust models. The use of stratified belief bases [18] is another path taken by researchers when dealing with these scenarios. Such approaches to multi-agent merging propose to construct a stratified knowledge base from a set of input knowledge bases; as an example, in [39], the authors introduce methods to assign a particular value to formulas during a merging process. Nevertheless, such approaches differ from ours in the importance we give to agents' trust to guide the merging of information.

The usefulness of defining a trust model has been widely accepted in the knowledge representation and reasoning literature. As stated in [42], the evolution of the multi-agent systems paradigm and e-commerce have contributed substantially to an increasing interest in representing trust; many applications in information and communication technologies have arisen as the result of works on trust models [22]. Clearly, some form of trust requires to be modeled in any multi-agent problem where critical decisions are made based on credibility among those agents. On the other hand, the integration of different systems, and the interaction resulting from this integration, may led to different practical problems. In this context, the management of conflicting information is an important and challenging issue that has to be faced [6, 7]. Several significant results aiming at dealing with such conflicts have been developed in the area of Belief Revision, where the central contribution is the definition of merging operators conducted by a set of rationality postulates.

In agent societies, consensus and unification of credibility criteria among its members is an important characteristic that should be taken into account when integrating local knowledge bases of the society. Following this direction, dif-

ferent proposals from Belief Revision theory have been developed to define and construct revision operators over an agent's knowledge base [29, 36, 47, 48]. In particular, these proposals focus on modeling the overall belief revision behavior of a multi-agent society, where agents need to communicate, cooperate, coordinate, and negotiate with each other to achieve a given goal. Even when our proposal does not address multi-agent negotiation directly, it presents mechanisms that contribute to the strengthening of capabilities of general multi-agent systems. Liu and Williams in [36] introduced an analysis of Belief Revision in this particular type of system.

The work in [30], which presented a formalization for merging propositional knowledge bases, has inspired and established the foundations to numerous works in Belief Revision. Several other works have also focused on defining merging operators, such as [14, 27]. The formal account of merging operators allows not only the representation of features intrinsically associated with pieces of information but also provides a natural and easy way to represent other types of domain information such as the identification of the sources from where information comes, as well as properties of such informants. [29].

Our novel approach represents a significant advance on the formalization of the process of merging knowledge bases where formulas have values attached to them measuring the credibility assigned to these formulas; then, these values are employed to decide which formulas to keep and which ones should be removed to solve conflicts. The procedure is based on the accrual of beliefs, that is, several formulas collaboratively aggregate their respective credibility values attempting to prevail in the resolution of conflicts; we will focus our research on exploring the theoretical aspects of the proposed approach. Next, we will show a running example that will also motivate our proposal's main ideas.

Example 1. *Consider a scenario where a group of colleagues is trying to decide on which of the following two hotel construction projects they will invest in: 1) the construction of a Hotel & Office complex in the downtown area of the city, or 2) the construction of a complex near the beach zone of the city. Each one of the members of the group will put forward its personal preference over 1) and 2); the goal is to find a final consensual decision on which project they will invest as a group. This group decision will be directly influenced by the credibility (or trust) that each group member has about the others. Suppose that the group is formed by five colleagues $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4,$ and \mathcal{A}_5 . We will assume that each member of the group can assign a degree of credibility to other agents in the group, e.g., for \mathcal{A}_1 , \mathcal{A}_2 is less credible than \mathcal{A}_3 and \mathcal{A}_4 , while it does not say anything about \mathcal{A}_5 . As we are trying to obtain the most credible option as a group, we can use the individual credibility orderings to evaluate each (conflicting) option's credibility. For example, assume option 1) is reported as the preferred option for $\mathcal{A}_1, \mathcal{A}_2,$ and \mathcal{A}_4 , whereas \mathcal{A}_3 and \mathcal{A}_5 prefer option 2). From this perspective, the credibility for option 1) can be obtained establishing how credible are $\mathcal{A}_1, \mathcal{A}_2,$ and \mathcal{A}_4 , and the credibility for option 2) establishing how credible $\mathcal{A}_1, \mathcal{A}_2$ are. As the example shows, it seems very natural that the best group decision is determined by comparing conflicting options' credibility values. In the rest of*

the paper, we will show different alternatives to accomplish this type of task.

In this article, we are particularly interested in multi-agent settings where agents need to reach a consensus on how conflicts among information pieces will be handled. This consensus is expected to be in line with the credibility associated with agents at the community level.

The applications in real-world scenarios for this particular context are numerous, such as in medical or health care systems, industrial robots, or social network platforms. A particular example could be the Twitter social network, which can be seen as a complex network. In such a case, the interactions among agents (users) such as retweets, replies, and mentions are part of the network's implicit dynamics and can be used to extract useful information. Nevertheless, it is hard to evaluate the trustworthiness of such information effectively, so the trustworthiness of shared contents and the reputation of agents play an essential part in such a platform. Another challenging application domain is that of forecasting the stock market. Recent developments consider the monitoring of Twitter content as a possible tool for enhancement, and several advances can be considered in this regard [41]. For instance, one possible development is to consider reciprocal retweets to calculate the degree of trust agents assign to other agents; this credibility value could be used to give a final community prediction based on the tweets reflecting stock market movements. The potential exploitation of the integration of many data sources is vast and can be employed in many application domains besides the few mentioned above.

Finally, knowledge bases in the form of ontologies are becoming a useful device that provides a convenient way to deal with the massive amounts of data we can encounter in real-world environments. Moreover, the expressive power of ontologies allows us to perform essential tasks on data integration [33]. In this context, Datalog^{\pm} is an attractive language to represent the knowledge bases of different agents, offering a significant tradeoff between expressibility and decidability, which favors the applicability to various application domains.

Given these considerations, we will focus on merging knowledge bases represented as Datalog^{\pm} ontologies to obtain a new ontology that incorporates, as much as possible, the knowledge of the original bases, focusing on enforcing consistency. Therefore, we present a multi-agent scenario where the agents' knowledge is expressed as a Datalog^{\pm} ontology; besides that ontology, each agent may associate with other community agents a particular credibility value. Our proposal is based on credibility accrual functions; these functions will assign values to formulas based on the agents' credibility value that inform them. Then, to resolve information conflicts, we define incision functions based on accrual functions [24, 13] that select which formulas should be removed from the conflicting sets. The ontology resulting from this process is the set obtained by integrating all formulas that are not removed by these incision functions.

The main contributions of this work are the following:

- We formalize a multi-agent setting where agents have their knowledge expressed as Datalog^{\pm} ontologies and assign values to other agents in the

community, considering its particular opinion on how credible they are. We also introduce conflicts in such a scenario.

- We formalize a framework for merging Datalog[±] ontology in the multi-agent context presented, where conflicts arise in the process. In order to achieve some consensus about how conflicts among pieces of information should be solved, we devise a general way to use strategies based on the credibility values that agents ascribe to each other.
- In this setting, credibilities are assigned to formulas through functions on the credibilities associated with agents that inform them; for this calculation, we propose a family of distinct functions called *credibility accrual functions*.
- We define merging operators based on accrual functions to integrate the Datalog[±] ontologies (corresponding to members in a community) into a new consistency preserving ontology; we also provide insights on such operators' behavior.
- Finally, we present particular examples of credibility accrual functions and show for each the corresponding merging process.

The rest of the paper is structured as follows. Section 2 presents some preliminary concepts. In Section 3, we introduce both credibility accrual functions and merging operators based on them. Section 4 contains specializations of the general credibility accrual functions previously introduced. Later, in Section 5, we formally compare our approach with a related well-known Belief Revision approach. Finally, Sections 6 and 7 further relate our work with the literature and offer some conclusions.

2. Background

We begin by briefly recalling the research context of belief change theory, focusing on particular in merging operators. Later in this section, we introduce some basic notation that will be used throughout the paper.

2.1. Belief Change and Knowledge Merging

As it is usual in the development of research, how far one should go back to find the origins of belief change theory is a matter of possibly heated discussion. Nonetheless, it can be argued, that some of the first steps were given in Isaac Levi's work in [34], where the fundamental problems concerning this field of research were discussed, and in the work of William Harper's in [26], where a rational way to interrelate belief change operators were proposed. Later, essential developments in the field came from the work of Carlos Alchourrón and David Makinson [2], and also by Peter Gärdenfors [21]. These works eventually came together with the three authors laying the main building blocks of the

AGM model in their seminal work [1], setting the foundations from where belief change theory eventually evolved.

The work in Belief Revision may be divided into two main lines. On the one hand, in the line known as the *coherence model*, we have those works focused on changes in sets of formulas that are closed under some consequence relation, called *belief sets* [1]. On the other hand, another line is referred to as the *foundational model*, which involves those works formalizing changes on non-closed sets, called *belief bases* [24, 25].

In the *foundational model*, kernel contraction [24] is based on the selection among the minimal subsets of a belief base K that contribute to make K imply α (α -kernels); kernel contraction is known to be more general than partial meet contraction for belief bases [25, p. 88-92], and hence more general than the AGM approach to contraction [24, 25]. The approach is based on the use of incision functions, which define the formulas that are to be deleted from each α -kernels. In this work, we will exploit this approach to attack a problem that is inherently different from knowledge contraction/revision: knowledge merging.

The problem of merging knowledge is one that arises in several situations. In particular, in fields outside of Computer Science, it is often needed to combine several sources of information. Moreover, such a combination needs to be done in a way that, when found, conflicts are taken care of. As discussed in [19], merging multiple sources of information can be applied in many different scenarios like distributed databases, multi-agent systems, data warehousing, *etc.*, where it is necessary to integrate multiple databases into a single and preferably consistent database. One of the most relevant applications of merging is in the integration of expert systems [50], where each one has a belief base representing its belief state, which can be either knowledge expressed in a propositional language; or (a subset of) a first-order logic.

As opposed to most revision approaches, merging opens the possibility that new evidence is partially or even completely ignored if old information is suitably well-entrenched. The merge operation joins old and new information in a consistent set without giving undue precedence to one or the other. One of the most prominent works on merging is that of [30], further research on the topic has emerged from this foundational proposal [19].

2.2. Preliminaries

Here, we will deal with application environments where knowledge is expressed in an ontological language; several such languages were developed to achieve such knowledge representation. In particular, we focus on a representation consisting of a database instance together with a set of existential rules, also known as Datalog[±] [10]. Existential rules extend Datalog rules by allowing existential quantification in the rule heads. We will now recall the basic notions of Datalog[±] ontologies from [10], which further extends the language with other types of rules and constraints (both syntactic and semantic) that limit the expressive power in favor of computational tractability.

We will assume that the domain of discourse consists of a countable set of *data constants* Δ , a countable set of nulls Δ_N (as placeholders for unknown

values), and a countable set of variables \mathcal{V} . We also assume that different constants represent different values (Unique Names Assumption). To distinguish constants from variables, we adopt the standard notation from Logic Programming, where variable names begin with uppercase letters, while constants and predicate symbols begin with lowercase letters.

We assume a *relational schema* \mathcal{R} , which is a finite set of *predicate symbols* (or simply *predicates*). A *term* t is a constant, a null, or a variable. An *atom* a has the form $p(t_1, \dots, t_n)$, where p is an n -ary predicate and t_1, \dots, t_n are terms; an atom is ground iff all terms in it are constants. Let \mathcal{L} be a first-order language such that $\mathcal{R} \subset \mathcal{L}$; then $\mathcal{L}_{\mathcal{R}}$ denotes the sublanguage generated by \mathcal{R} . A *database (instance)* of \mathcal{R} is a finite set of atoms with predicates in \mathcal{R} and terms in $\Delta \cup \Delta_N$. A *homomorphism* on constants, nulls and variables is a mapping $h : \Delta \cup \Delta_N \cup \mathcal{V} \rightarrow \Delta \cup \Delta_N \cup \mathcal{V}$ such that (i) $c \in \Delta$ implies $h(c) = c$, (ii) $c \in \Delta_N$ implies $h(c) \in \Delta \cup \Delta_N$, and (iii) h is naturally extended to atoms, sets of atoms, and conjunctions of atoms.

Given a relational schema \mathcal{R} , a *tuple-generating dependency (TGD)* σ is a first-order formula of the form $\forall \mathbf{X} \forall \mathbf{Y} \Phi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} \Psi(\mathbf{X}, \mathbf{Z})$ where $\Phi(\mathbf{X}, \mathbf{Y})$ and $\Psi(\mathbf{X}, \mathbf{Z})$ are conjunctions of atoms over \mathcal{R} called the *body* (denoted $body(\sigma)$) and the *head* (denoted $head(\sigma)$), respectively. Consider a database D for a relational schema \mathcal{R} , and a TGD σ on \mathcal{R} of the form $\Phi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} \Psi(\mathbf{X}, \mathbf{Z})$. Then, σ is *applicable* to D if there exists a homomorphism h that maps the atoms of $\Phi(\mathbf{X}, \mathbf{Y})$ to atoms in D . Let σ be applicable to D , and h' be a homomorphism that extends h as follows: for each $X_i \in \mathbf{X}$, $h'(X_i) = h(X_i)$; for each $Z_j \in \mathbf{Z}$, $h'(Z_j) = z_j$, where z_j is a “fresh” null, *i.e.*, $z_j \in \Delta_N$, z_j does not occur in D , and z_j lexicographically follows all other nulls already introduced. The *application* of σ on D adds to D the atom $h'(\Psi(\mathbf{X}, \mathbf{Z}))$ if it is not already in D . After the application we say that σ is satisfied by D . The *Chase* for a database D and a set of TGDs Σ_T , denoted $chase(D, \Sigma_T)$, is the exhaustive application of the TGDs [10] in a breadth-first (level-saturating) fashion, which leads to a (possibly infinite) chase for D and Σ . It is important to remark that BCQs Q over D and Σ_T can be evaluated on the chase for D and Σ_T , *i.e.*, $D \cup \Sigma_T \models Q$ is equivalent to $chase(D, \Sigma_T) \models Q$ [10]. We will also use $D \cup \Sigma_T \vdash Q$ when $chase(D, \Sigma_T) \models Q$.

Negative constraints (NCs) are first-order formulas of the form $\forall \mathbf{X} \Phi(\mathbf{X}) \rightarrow \perp$, where $\Phi(\mathbf{X})$ is a conjunction of atoms (without nulls) and the head is the truth constant *false*, denoted \perp . A NC τ is satisfied by a database D under a set of TGDs Σ_T iff there does not exist a homomorphism h that maps the atoms of $\Phi(\mathbf{X})$ to D , where D is such that every TGD in Σ_T is satisfied, *i.e.*, the atoms in the body cannot all be true together.

Equality-generating dependencies (EGDs) are first-order formulas of the form $\forall \mathbf{X} \Phi(\mathbf{X}) \rightarrow X_i = X_j$, where $\Phi(\mathbf{X})$ is a conjunction of atoms, and X_i and X_j are variables from \mathbf{X} . An EGD σ is satisfied in a database D for \mathcal{R} iff, whenever there exists a homomorphism h such that $h(\Phi(\mathbf{X})) \subseteq D$, it holds that $h(X_i) = h(X_j)$.

NCs and EGDs play an important role in the matter of conflicts in Datalog[±] ontologies. In fact, the approach that we present in this work ensures that

neither NCs nor EGDs are violated in the resulting ontology.

As is the usual case in the literature, in general, the universal quantifiers in TGDs, NCs, and EGDs are omitted, and both the sets of dependencies and the set of constraints are assumed to be finite. Now that we have presented the different ways of expressing knowledge in Datalog[±], we are ready to formally define Datalog[±] ontologies, which will serve as a means to represent both the database instances and existential rules programs our merging operators will deal with.

Definition 1 (Datalog[±] Ontology). *A Datalog[±] ontology $KB = (D, \Sigma)$, where $\Sigma = \Sigma_T \cup \Sigma_E \cup \Sigma_{NC}$, consists of a database instance D that is a finite set of ground atoms (without nulls), a set of TGDs Σ_T , a set of separable EGDs Σ_E , and a set of NCs Σ_{NC} .*

Unless explicitly said, through the paper, when it is clear from the context, we will refer to the component Σ in KB as the set of constraints in the ontology, without distinguishing between dependencies and constraints. Given a database D for \mathcal{R} and a set of constraints $\Sigma = \Sigma_T \cup \Sigma_E \cup \Sigma_{NC}$, the set of *models* of D and Σ , denoted $\text{mods}(D, \Sigma)$, is the set of all databases B such that $D \subseteq B$ and every formula in Σ is satisfied. The following example shows a simple Datalog[±] ontology; the ontology describes knowledge about the scenario introduced in the running example.

Example 2 (Datalog[±] Ontology).

$$KB_2 = \left\{ \begin{array}{l} D : \{a_1 : \text{expST}(p1), a_2 : \text{unemployment}(p1), \\ a_3 : \text{incSBZ}(p1), a_4 : \text{incOTB}(p1), \\ a_5 : \text{fund}(c1, p1), a_6 : \text{fund}(c1, p2)\} \\ \Sigma_{NC} : \{\tau_1 : \text{investDA}(P) \wedge \text{investBZ}(P) \rightarrow \perp\} \\ \Sigma_E : \{\nu_1 : \text{fund}(E, P) \wedge \text{fund}(E, P') \rightarrow P = P'\} \\ \Sigma_T : \{\sigma_1 : \text{expST}(P) \wedge \text{incOTB}(P) \rightarrow \text{investBZ}(P), \\ \sigma_2 : \text{expST}(P) \rightarrow \text{investDA}(P), \\ \sigma_3 : \text{unemployment}(P) \wedge \text{incSBZ}(P) \rightarrow \text{investBZ}(P)\} \end{array} \right\}$$

Observe that the set KB has six ground atoms reflecting different things: that p_1 intends to exploit summer tourism ($\text{expST}(p1)$), its goal of increasing occupancy rate of hotels ($\text{incOTB}(p1)$), that it aims to increase sales in beach zones ($\text{incSBZ}(p1)$), p_1 's objective of reducing unemployment ($\text{unemployment}(p1)$), that project p_1 is funded by company c_1 ($\text{fund}(c1, p1)$), and that project p_2 is funded by company c_1 ($\text{fund}(c1, p2)$).

The set Σ_T of TGDs expresses the following dependencies: TGDs σ_1 and σ_3 state reasons to invest in beach zones, while TGD σ_2 gives reasons to invest in a Hotel & Office complex into downtown area. The only NC τ_1 states that a project can not promote investments into the downtown area and beach zones at

the same time, and EGD ν_1 states that every company can not fund more than one project.

Following the classical notion of consistency, we say that a consistent Datalog[±] ontology has a non-empty set of models.

Definition 2 (Consistency). *A Datalog[±] ontology $KB = (D, \Sigma)$ is consistent iff $\text{mods}(D, \Sigma) \neq \emptyset$. Otherwise, we will say that KB is inconsistent.*

Example 3 (Consistency). *Consider the Datalog[±] ontology from the example above; this ontology is clearly inconsistent. Database instance D is clearly not a model in itself since at least the TGD σ_2 is applicable to D , but there is no superset of D such that it satisfies all TGDs and constraints in Σ at the same time. For instance, the TGDs σ_2 is applicable in D resulting in the new atom $\text{investDA}(p1)$, which together with $\text{investZB}(p1)$, which is obtained from σ_3 , violates the NC τ_1 .*

For the rest of the paper, unless explicitly stated, $KB = (D, \Sigma)$ will denote a Datalog[±] ontology with $\Sigma = \Sigma_T \cup \Sigma_E \cup \Sigma_{NC}$, where D is a database instance, Σ_T is the set of all TGDs, Σ_E the set of all EGDs, and Σ_{NC} is the set of all NCs in Σ .

3. Merging ontologies in multi-agent settings

Merging processes aim to integrate several knowledge sources (in this work, in the form of Datalog[±] ontologies) into a new knowledge base that represents the original sources as much as possible (according to a certain measure) while preserving certain characteristics, in particular consistency [30, 31].

Here, we define merging processes in a multi-agent scenario, where individual agents' knowledge will be merged. In such scenario, it is possible to define conflict resolution strategies that are influenced by the degrees of credibility that agents associate to each other in the community. In this section we introduce merging operators that account for the credibility of agents when dealing with conflicts. We begin by establishing the multi-agent context in which the operators work. Then, we move on to show how conflicts in the community of agents may arise; and finally we look into how such conflicts are solved and the properties of this general merging approach.

3.1. Context

As briefly explained before, in this work we aim to merge knowledge coming from several different sources. In particular, we will express such knowledge as Datalog[±] ontologies, and obtain as the merging result a new Datalog[±] ontology reflecting the original ones as much as possible without violating the constraints included in the result.

To do this we will define a multi-agent scenario, where the final merged ontology obtained can be seen as the agreed knowledge of the community of

agents. These agents will then be constituted by two different components. On the one hand, we have the agent's knowledge expressed by the Datalog[±] ontology. On the other hand, each agent will establish a credibility degree with respect to other agents in the community, reflecting how much it is willing to accept information coming from those other agents. We will then exploit the different credibilities that agents in the community associate to each other in order to define how conflicts should be solved. To do this, we will link formulas to certain values calculated from the credibilities associated in the community to the agents that are the informants of such formulas. This calculation will be made by exploiting particular accrual functions [13] tailored to our multiagent scenario.

As said, agents will attach to other agents in a community a measure of how much they are willing to accept a formula coming from that other agent. We call this measure a *credibility degree*, which is a pair of an agent (the one that is being associated the credibility) and a value. Formally this definition is as follows.

Definition 3 (Credibility degree). *Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a set of agents, a credibility degree associated to $\mathcal{A} \in \mathbb{A}$ is a real number $c \in [0, 1]$. We denote this as the pair (\mathcal{A}, c) .*

Each agent will include a set of credibility degrees, that will reflect how much trust the agent grants to agents in the community. This set of credibility degrees are known as the agent's credibility vector.

Definition 4 (Agent's credibility vector). *Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a set of agents, $n > 1$, the credibility vector of agent $\mathcal{A} \in \mathbb{A}$, denoted $\mathbb{V}_{\mathcal{A}}$, is the set $\{(\mathcal{A}_1, c_1), (\mathcal{A}_2, c_2), \dots, (\mathcal{A}_m, c_m)\}$, $m \leq n - 1$, of credibility degrees such that:*

1. *there is no $(\mathcal{A}_j, c_j) \in \mathbb{V}_{\mathcal{A}}$ such that $\mathcal{A}_j = \mathcal{A}$, i.e., there is no credibility degree about the agent \mathcal{A} itself, and*
2. *there is no $(\mathcal{A}_j, c), (\mathcal{A}_k, c') \in \mathbb{V}_{\mathcal{A}}$ such that $j = k$, i.e., the credibility degree associated to an agent is unique in $\mathbb{V}_{\mathcal{A}}$.*

The set $\mathbb{V}_{\mathcal{A}}$ could be empty when the agent \mathcal{A} has no credibility degree to any other agent.

As can be seen in Definition 4, there are some considerations regarding the credibility degrees that an agent can store in its credibility vector. These considerations are made because we plan to use credibilities as the basis of a conflict resolution strategy. The strategy will obtain a credibility degree for each agent, and this degree can be seen as the credibility degree assigned to every agent by the community as a whole. Thus, we want to prevent some improper influences in the calculation of such an agreed degree. To begin with, an agent should not influence how the community regards it by giving itself any credibility degree, that in extreme cases could be maximal or minimal; therefore, we impose the restriction that the credibility vector of an agent cannot include a credibility

degree for the agent itself. Additionally, another form of unfair influence on the credibility degree associated with an agent by the community as a whole may arise if some other agent assigns to it more than one credibility degree, that is, an agent \mathcal{A} could intentionally assign to agent \mathcal{B} several credibility degrees with values close to 1 or 0 to promote or demote the credibility of agent \mathcal{B} . Finally, it may be the case that agent \mathcal{A} has no reason to assign a credibility degree to agent \mathcal{B} ; for instance, if they have had no interactions in the past. Thus, we define credibility vectors as partial assignments of credibility to agents in the community, instead of total ones, although the assignment may be total in particular cases. Nevertheless, we will assume throughout the paper that every agent \mathcal{A} has at least one credibility assigned to it by some other agent \mathcal{B} in the community.

Example 4 (Credibility Vector). *Consider the application domain introduced in Example 1. Suppose that we have a set of agents $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$, and that agent \mathcal{A}_1 assigns to agent \mathcal{A}_2 a credibility of 0.3 and both \mathcal{A}_3 and \mathcal{A}_4 a credibility of 0.7, while it does not assign a credibility degree for agent \mathcal{A}_5 . Then, for that agent we have that*

$$\mathbb{V}_{\mathcal{A}_1} = \{(\mathcal{A}_2, 0.3), (\mathcal{A}_3, 0.7), (\mathcal{A}_4, 0.7)\}.$$

We are now ready to refine the definition of agents and to introduce the notion of communities in our framework. As explained, agents have their knowledge expressed as a Datalog $^\pm$ ontology, and they maintain a credibility vector where they state how much they trust other agents in the community. A community of agents, or just a community, is a collection of agents such as the ones just described.

Definition 5 (Agents and Communities). *An agent \mathcal{A} is a pair $(KB_{\mathcal{A}}, \mathbb{V}_{\mathcal{A}})$, where $KB_{\mathcal{A}}$ is a Datalog $^\pm$ ontology and $\mathbb{V}_{\mathcal{A}}$ is the agent's credibility vector. A community of agents is a set of agents $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$. We will refer to every $\mathcal{A} \in \mathbb{A}$ as an agent of the community.*

Without loss of generality, we will assume that for all $\mathcal{A} \in \mathbb{A}$ the Datalog $^\pm$ knowledge base $KB_{\mathcal{A}}$ is defined over a common relational schema \mathcal{R} . Although this may not be the case in practice, there exist effective tools for ontology alignment that can integrate heterogeneous schemas [3, 4].

Example 5 (Agents and Communities). *Consider the example illustrated in Figure 1. In the figure, we can see a community of agents with five members $\mathbb{A}_5 = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$, where*

- $\mathcal{A}_1 = (KB_{\mathcal{A}_1}, \mathbb{V}_{\mathcal{A}_1})$ such that
 $KB_{\mathcal{A}_1} = \{a_4, a_5, a_6, \tau_1, \nu_1\}$ and $\mathbb{V}_{\mathcal{A}_1} = \{(\mathcal{A}_2, 0.3), (\mathcal{A}_3, 0.7), (\mathcal{A}_4, 0.7)\}$.
- $\mathcal{A}_2 = (KB_{\mathcal{A}_2}, \mathbb{V}_{\mathcal{A}_2})$ such that
 $KB_{\mathcal{A}_2} = \{a_1, a_4, \sigma_1\}$ and $\mathbb{V}_{\mathcal{A}_2} = \{(\mathcal{A}_1, 0.5), (\mathcal{A}_3, 0.4), (\mathcal{A}_4, 0.4), (\mathcal{A}_5, 0.6)\}$.

- $\mathcal{A}_3 = (KB_{\mathcal{A}_3}, \mathbb{V}_{\mathcal{A}_3})$ such that
 $KB_{\mathcal{A}_3} = \{a_2, a_3, \sigma_3\}$ and $\mathbb{V}_{\mathcal{A}_1} = \{(\mathcal{A}_1, 0.6), (\mathcal{A}_2, 0.2), (\mathcal{A}_4, 0.7), (\mathcal{A}_5, 0.8)\}$.
- $\mathcal{A}_4 = (KB_{\mathcal{A}_4}, \mathbb{V}_{\mathcal{A}_4})$ such that
 $KB_{\mathcal{A}_4} = \{a_1, a_3, a_5, \sigma_2\}$ and $\mathbb{V}_{\mathcal{A}_1} = \{(\mathcal{A}_1, 0.6), (\mathcal{A}_2, 0.9), (\mathcal{A}_5, 0.8)\}$.
- $\mathcal{A}_5 = (KB_{\mathcal{A}_5}, \mathbb{V}_{\mathcal{A}_5})$ such that
 $KB_{\mathcal{A}_5} = \{a_2, a_3, a_6, \sigma_3\}$ and $\mathbb{V}_{\mathcal{A}_1} = \{(\mathcal{A}_1, 0.8), (\mathcal{A}_2, 0.3), (\mathcal{A}, 0.7), (\mathcal{A}_4, 0.3)\}$.

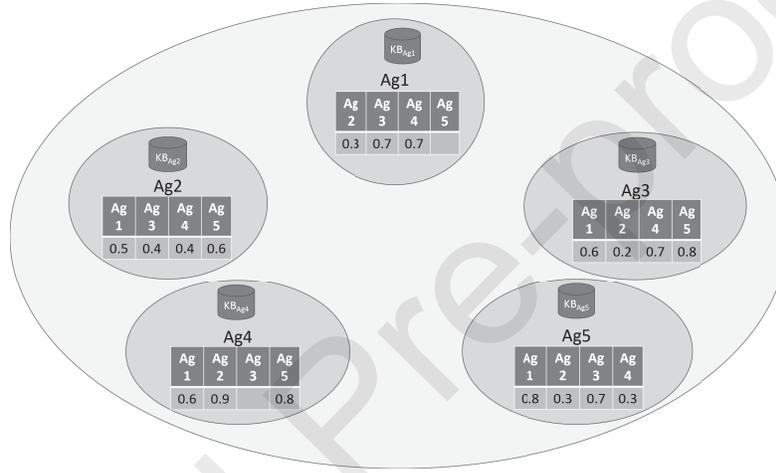


Figure 1: Agents and Communities

3.2. Conflicts in the community

Knowledge dynamics is a problem widely studied in the literature of Belief Revision. Several constructions for revision, contraction, or merging operators have been developed over the years. A particularly interesting perspective regarding our approach in this work is [24], which shows how a contraction operation on belief bases can be modeled employing *incision functions*. The authors in [24] utilize incision functions to perform contractions over belief bases by some formula α considering minimal sets, or α -kernels, entailing that formula and then performing “incisions” on these sets. Since the α -kernels are minimal sets, the sets resulting from incisions no longer entail α . Finally, the resulting belief base is settled as the union of all formulas that are not removed by the incisions. This approach is known as *kernel contraction*; the task of restoring consistency is also known in the belief revision literature as contraction by *falsum* [23].

Instead of defining contraction processes, in this work we aim to perform merging processes over several Datalog[±] ontologies. Nevertheless, Kernel Contraction does provide the foundations for how our approach deals with conflicts,

since we define the merging process as the application of incision functions on the union of Datalog[±] ontologies while still considering the different sources of the formulas by means of accrual functions. In particular, instead of directly considering minimal inconsistent subsets of formulas in the ontology (which are equivalent to \perp -kernels), in this work we consider incisions over structures called *clusters* [37, 12, 13] that groups together related kernels.

Thus, we begin by defining kernels, the fundamental conflicts that can arise in the community. As explained, kernels are conflictive sets of formulas that are minimal from a set inclusion perspective; that is, sets that lead to a violation of some of the restrictions and, because of their minimality, every subset of them does not produce a violation. In [12], a kernel, and thus also a cluster, is defined as either a set of atoms or a set of TGDs, depending on whether conflicts in the database instance or the set of TGDs are to be solved, respectively. We will not follow such an approach here; instead, we will consider kernels formed together by both atoms from the database instance and TGDs. Additionally, we will examine conflicts at a community level; therefore, we will look for conflicts in the agents' cumulative knowledge at once (instead of looking into each particular ontology).

Definition 6 (Kernels). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents on the form $\mathcal{A} = (KB_{\mathcal{A}}, \mathbb{V}_{\mathcal{A}})$ such that $KB_{\mathcal{A}} = (D_{\mathcal{A}}, \Sigma_{\mathcal{A}})$ is the knowledge base and $\mathbb{V}_{\mathcal{A}}$ is the credibility vector for agent \mathcal{A} . Also, for some $S \subseteq \cup_{1 \leq i \leq n} KB_{\mathcal{A}_i}$, let $S_D = S \cap \cup_{1 \leq i \leq n} D_{\mathcal{A}_i}$ and $S_{\Sigma} = S \cap \cup_{1 \leq i \leq n} \Sigma_{\mathcal{A}_i}$ be the set of atoms and the set of rules, respectively.

The set of kernels for \mathbb{A} , denoted with $\coprod_{\mathbb{A}}$, is the set of all X such that

- $X \subseteq \bigcup_{1 \leq i \leq n} KB_{\mathcal{A}_i}$,
- $\text{mods}(X_D, X_{\Sigma}) = \emptyset$, and
- for every $X' \subsetneq X$ it holds that $\text{mods}(X'_D, X'_{\Sigma}) \neq \emptyset$.

Example 6 (Kernels). Let $\mathbb{A}_5 = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4, \mathcal{A}_5\}$ be community of agents introduced in Example 5, and $\coprod_{\mathbb{A}_5}$ the set of kernels for \mathbb{A}_5 . Then, we can obtain from $\coprod_{\mathbb{A}_5}$ the following kernels:

- $X_1 = \{a_1, a_2, a_3, \sigma_2, \sigma_3, \tau_1\}$
- $X_2 = \{a_1, a_4, \sigma_1, \sigma_2, \tau_1\}$
- $X_3 = \{a_5, a_6, \nu_1\}$

We will identify relations among minimal conflicts in the knowledge of the community if any such relations exist. To do this, we group related kernels together in a new structure called a cluster, by means of an overlapping relation. We now recall the definition of overlapping from [12].

Definition 7 (Overlapping, Equivalence [12]). Let \mathcal{L} be a first order language, $\mathcal{R} \subset \mathcal{L}$ be a relational schema, and $\mathcal{L}_{\mathcal{R}}$ the sublanguage generated by \mathcal{R} . Given $A \subset \mathcal{L}_{\mathcal{R}}$ and $B \subset \mathcal{L}_{\mathcal{R}}$, we say they overlap, denoted $A \theta B$, iff $A \cap B \neq \emptyset$. Furthermore, given a multi-set of first order formulas $\mathcal{M} \subset 2^{\mathcal{L}_{\mathcal{R}}}$ we denote as $\theta_{\mathcal{M}}^*$ the equivalence relation obtained over \mathcal{M} through the transitive closure of θ .

Based on the overlapping relation, we transitively obtain clusters as the union of those kernels that do overlap. We now adapt the definition of clusters in [12] to our setting where kernels are obtained based on the community knowledge as a whole.

Definition 8 (Clusters). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents, and $\mathbb{K}_{\mathbb{A}}$ be the set of kernels for \mathbb{A} . Let θ be the overlapping relation, and $\mathcal{K} = \mathbb{K}_{\mathbb{A}} / \theta_{\mathbb{K}_{\mathbb{A}}}^*$ the quotient set for the equivalence relation obtained over $\mathbb{K}_{\mathbb{A}}$.

A cluster is a set $\varsigma = \bigcup_{\kappa \in [\kappa]} \kappa$, where $[\kappa] \in \mathcal{K}$. We denote by $\mathbb{C}_{\mathbb{A}}$ the set of all clusters for \mathbb{A} .

Example 7 (Clusters). Consider for instance the set of kernels $\mathbb{K}_{\mathbb{A}} = \{\{a_1, a_2, a_3, \sigma_2, \sigma_3, \tau_1\}, \{a_1, a_4, \sigma_1, \sigma_2, \tau_1\}, \{a_5, a_6, \nu_1\}\}$ for our running example introduced above in Example 6. Then, we see that there are two clusters based in these kernels, $\mathbb{C}_{\mathbb{A}} = \{C_1, C_2\}$ such that $C_1 = \{a_1, a_2, a_3, a_4, \sigma_1, \sigma_2, \sigma_3, \tau_1\}$ and $C_2 = \{a_5, a_6, \nu_1\}$.

3.3. Conflict Resolution

Up to this point, we have looked into the nature of the conflicts that can arise in a merging scenario where beliefs come from different agents with different epistemic states. Once we have found the conflicts, the next logical step is to solve them somehow. To do so, as explained in Section 2 in this work, we will exploit the findings in [13], defining accrual functions as a way to measure the credibility associated with a particular formula, based on the credibilities attached to those agents that claim said formula.

In [13], accrual functions are defined for revision processes in a general way, considering a particular context; that is, a particular knowledge base under revision (the epistemic state), and a particular formula that triggers the revision (the epistemic input). Also, they consider that we can have different *features* that represent different dimensions we want to consider when choosing how to deal with some conflict. However, in this work, we will directly consider accrual functions that are particular for our goal. As such, we will define them directly under the light of the conflicts found in the knowledge of the community as a whole and will consider only the particular feature we want to exploit (agent's credibility). However, it should be noted that should we ever want to include other aspects in our calculation (*e.g.*, authority degree, cost or benefits associated to formulas) then to enhance our model to consider several different aspects (*i.e.*, features) is pretty straightforward, assuming that we can establish an order of importance among those features. In such a scenario, we can simply

use accrual functions such as the ones we will introduce here for the different features and then use orders *à la* [13] to establish our conflict resolution strategy.

Next, we will adapt the definition of accrual functions introduced in [13] to our multi-agent knowledge merging scenario, where the selection of which formulas will be deleted from the community knowledge when solving conflicts depends on the credibility attached to those agents that have the formulas mentioned earlier in their particular knowledge.

Definition 9 (Credibility Accrual Function). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents where each $\mathcal{A} \in \mathbb{A}$ is described as $\mathcal{A} = (KB_{\mathcal{A}}, \mathbb{V}_{\mathcal{A}})$ such that $KB_{\mathcal{A}}$ is the knowledge base and $\mathbb{V}_{\mathcal{A}}$ is the credibility vector for agent \mathcal{A} , $KB = \cup_{1 \leq i \leq n} KB_{\mathcal{A}_i}$, $\prod_{\mathbb{A}}$ be the set of Kernels for \mathbb{A} , and $\prod_{\mathbb{A}}$ be the set of clusters obtained from $\prod_{\mathbb{A}}$. Let $\mathbf{S} = \{A, B, C, \dots, Z\}$ be such that for all $X \in \mathbf{S}$ it holds that $X \subseteq KB$.

A function $\mathcal{V} : 2^{KB} \mapsto \mathbb{R}$ is a credibility accrual function iff all of the following conditions hold:

- **Transitivity:** if $\mathcal{V}_{[\text{cred}]}(\{A\}) \leq \mathcal{V}_{[\text{cred}]}(\{B\})$ and $\mathcal{V}_{[\text{cred}]}(\{B\}) \leq \mathcal{V}_{[\text{cred}]}(\{C\})$, then $\mathcal{V}_{[\text{cred}]}(\{A\}) \leq \mathcal{V}_{[\text{cred}]}(\{C\})$.
- **Domination:** Let $\mathcal{P}(A) = \{A_1, A_2, \dots, A_n\}$ be a partition of A , if for all A_i ($1 \leq i \leq n$) it holds that $\mathcal{V}_{[\text{cred}]}(\{A_i \setminus B\}) \leq \mathcal{V}_{[\text{cred}]}(\{A_i \setminus C\})$ then we have that $\mathcal{V}_{[\text{cred}]}(\{A \setminus B\}) \leq \mathcal{V}_{[\text{cred}]}(\{A \setminus C\})$.
- **Pertinence:** for all $A \subseteq \bigcup_{1 \leq i \leq n} KB_{\mathcal{A}_i}$ and $\beta \in A$ it holds that if $\beta \notin \prod_{\mathbb{A}}$ then $\mathcal{V}_{[\text{cred}]}(\{A\}) = \mathcal{V}_{[\text{cred}]}(\{A \cup \{\beta\}\})$.

When originally introduced in [13], accrual functions were devised to work on multi-dimensionally valued knowledge bases, and thus they were able to account for several distinct features. In this work, we only consider a single feature, that is, the credibility of agents in the community; therefore, we often use \mathcal{V} (omitting the index) instead of $\mathcal{V}_{[\text{cred}]}$.

As can be seen, Definition 9 does not introduce a particular, constructive way of obtaining the credibility associated with some set of formulas, but rather define a family of such functions. The introduction of particular accrual functions is reserved for future work, but it should be noted that for any function considered, we aim to define a behavior that is consistent no matter the particular accrual function used. Thus, when dealing with theoretical developments we will consider an arbitrary \mathcal{V} defined as in definition 9.

An additional property of credibility accrual functions follows from Definition 9, that of **Neutrality**. This property states that, when evaluating sets of formulas that do not include any conflicting formulas, then accrual functions gracefully degrade to the neutral case; that is, they evaluate such sets as equals to the empty one; therefore, they do not influence the conflict resolution process in any way.

Lemma 1. Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents where for $\mathcal{A} \in \mathbb{A}$, $\mathcal{A} = (KB_{\mathcal{A}}, \mathbb{V}_{\mathcal{A}})$ such that $KB_{\mathcal{A}}$ is the knowledge base and $\mathbb{V}_{\mathcal{A}}$ is the credibility vector for agent \mathcal{A} , $\mathbb{K}_{\mathbb{A}}$ be the set of kernels for \mathbb{A} , and $\mathbb{C}_{\mathbb{A}}$ be the set of clusters obtained from $\mathbb{K}_{\mathbb{A}}$. Let $A \subseteq \cup_{1 \leq i \leq n} KB_{\mathcal{A}_i}$ and let \mathcal{V} be a credibility accrual function. Then, \mathcal{V} satisfies

- **Neutrality:** if A is such that $A \cap \mathbb{K}_{\mathbb{A}} = \emptyset$, then it holds that $\mathcal{V}(A) = \mathcal{V}(\emptyset)$.

Proof Neutrality follows from **Pertinence**. Let $A \subseteq \cup_{1 \leq i \leq n} KB_{\mathcal{A}_i}$ be such that $A \cap \mathbb{K}_{\mathbb{A}} = \emptyset$. Then, for all $\alpha \in A$ it holds that $\alpha \notin \mathbb{K}_{\mathbb{A}}$, and thus by **Pertinence** $\mathcal{V}(\{A\}) = \mathcal{V}(\{A \cup \{\alpha\}\})$. Since $\alpha \notin \mathbb{K}_{\mathbb{A}}$ holds for any arbitrary $\alpha \in A$, then we have that $\mathcal{V}(\emptyset) = \mathcal{V}(\emptyset \cup A)$. This last equality is equivalent to $\mathcal{V}(\emptyset) = \mathcal{V}(A)$. \square

The **Neutrality** property will help to simplify the formalization of the particular accrual functions that we shall introduce later in the paper. Accrual functions then will be used to establish the value of a set of formulas in order to define which formulas should be deleted to solve conflicts. Based on accrual functions we will define a variety of incision functions.

Definition 10 (Cluster Incision functions). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents where for $\mathcal{A} \in \mathbb{A}$, $\mathcal{A} = (KB_{\mathcal{A}}, \mathbb{V}_{\mathcal{A}})$ such that $KB_{\mathcal{A}}$ is the knowledge base and $\mathbb{V}_{\mathcal{A}}$ is the credibility vector for agent \mathcal{A} , $KB = \cup_{1 \leq i \leq n} KB_{\mathcal{A}_i}$, $\mathbb{K}_{\mathbb{A}}$ be the set of Kernels for \mathbb{A} , $\mathbb{C}_{\mathbb{A}}$ be the set of clusters obtained from $\mathbb{K}_{\mathbb{A}}$, and \mathcal{V} be a credibility accrual function for \mathbb{A} .

Then, if \mathcal{R} is a relational schema for KB , \mathcal{L} is a first-order language such that $\mathcal{R} \subset \mathcal{L}$, and $\mathcal{L}_{\mathcal{R}}$ is the sublanguage generated by \mathcal{R} , an Incision Function for \mathbb{A} is a function $\varrho : 2^{\mathcal{L}_{\mathcal{R}}} \rightarrow \mathcal{L}_{\mathcal{R}}$ such that all of the following conditions hold:

- $\varrho(\mathbb{C}_{\mathbb{A}}) \subseteq \cup(\mathbb{K}_{\mathbb{A}})$.
- For all $\varsigma \in \mathbb{C}_{\mathbb{A}}$ and $\kappa \in \mathbb{K}_{\mathbb{A}}$ such that $\kappa \subseteq \varsigma$ it holds that $\kappa \cap \varrho(\mathbb{C}_{\mathbb{A}}) \neq \emptyset$.
- For all $\varsigma \in \mathbb{C}_{\mathbb{A}}$ it holds that if $\varsigma \cap \varrho(\mathbb{C}_{\mathbb{A}}) = I$, then for all $F \subset \varsigma$ where F satisfies the two previous conditions it holds that $\mathcal{V}(\varsigma \setminus F) \leq \mathcal{V}(\varsigma \setminus I)$.

Example 8 (Incision functions). Assume that we are using \mathcal{V} as an accrual function already defined, and ϱ takes clusters and selects sets of formulas in them that has the smallest credibility value to be deleted. Consider the clusters $C_1 = \{a_1, a_2, a_3, a_4, \sigma_1, \sigma_2, \sigma_3, \tau_1\}$ and $C_2 = \{a_5, a_6, \nu_1\}$ from Example 7. For every cluster we could use \mathcal{V} to compute credibility values of each set of formulas in them. Following this idea, for cluster C_1 we assume $\{a_1, a_3, a_4\}$ has the smallest value, and $\{\nu_1\}$ for cluster C_2 . Thus, $\{a_1, a_3, a_4, \nu_1\}$ is the set of formulas to be deleted in order to restore consistency. Note that we are assuming that all conditions of ϱ hold. Later in the paper we will introduce some particular accrual functions achieving the behavior specified by such conditions.

Finally, we define the merging operator that obtains a conflict-free knowledge base by solving all conflicts that arise when considering the community knowledge as a whole.

Definition 11 (ck-Merging Operator). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents, where for each $\mathcal{A} \in \mathbb{A}$, $\mathcal{A} = (KB_{\mathcal{A}}, \mathbb{V}_{\mathcal{A}})$ such that $KB_{\mathcal{A}}$ is the knowledge base and $\mathbb{V}_{\mathcal{A}}$ is the credibility vector for agent \mathcal{A} , $\mathbb{K}_{\mathbb{A}}$ be the set of Kernels for \mathbb{A} , $\mathbb{C}_{\mathbb{A}}$ be the set of clusters obtained from $\mathbb{K}_{\mathbb{A}}$, and ϱ be an incision function.

A community-knowledge merging operator Υ is defined as

$$\Upsilon(\mathbb{A}) = \bigcup_{(KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i}) \in \mathbb{A}} KB_{\mathcal{A}_i} \setminus \varrho(\mathbb{C}_{\mathbb{A}})$$

We will use ck-merging operator for this operator and ck-merging when no confusion my arise.

Example 9 (ck-Merging). Consider the Datalog[±] ontology KB_2 presented in Example 2 which is the union of the particular knowledge bases of the agents in the community \mathbb{A}_5 presented in Example 5. Consider also the set $\{a_1, a_3, a_4, \nu_1\}$ obtained from incision function ϱ in Example 8. If we remove this set from KB_2 , then the following consistent Datalog[±] ontology can be obtained:

$$KB_2 = \left\{ \begin{array}{l} D : \{a_2 : \text{unemployment}(p1), a_5 : \text{fund}(c1, p1), \\ \quad a_6 : \text{fund}(c1, p2)\} \\ \Sigma_{NC} : \{\tau_1 : \text{investDA}(P) \wedge \text{investBZ}(P) \rightarrow \perp\} \\ \Sigma_E : \{\} \\ \Sigma_T : \{\sigma_1 : \text{expTS}(P) \wedge \text{incOTB}(P) \rightarrow \text{investBZ}(P), \\ \quad \sigma_2 : \text{expTS}(P) \rightarrow \text{investDA}(P), \\ \quad \sigma_3 : \text{unemployment}(P) \wedge \text{incSBZ}(P) \rightarrow \text{investBZ}(P)\} \end{array} \right\}$$

3.4. Properties of the operator

We will now look into some properties of our ck-merging operator. One of the most important properties we expect from this operator is that the final merged belief base always reaches a consensus regarding conflicts among the pieces of knowledge provided by the different agents, *i.e.*, that the obtained knowledge base is always consistent.

Proposition 1 (Consistency). Let Υ be ck-merging operator. Then, Υ satisfies

- **Consistency:** $\text{mods}(\Upsilon(\mathbb{A})_D, \Upsilon(\mathbb{A})_\Sigma) \neq \emptyset$.

Proof Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents, where for each $\mathcal{A} \in \mathbb{A}$, $\mathcal{A} = (KB_{\mathcal{A}}, \mathbb{V}_{\mathcal{A}})$ such that $KB_{\mathcal{A}}$ is the knowledge base and $\mathbb{V}_{\mathcal{A}}$ is the credibility vector for agent \mathcal{A} , $\coprod_{\mathbb{A}}$ be the set of kernels for \mathbb{A} , $\coprod_{\mathbb{A}}$ be the set of clusters obtained from $\coprod_{\mathbb{A}}$, and ϱ be an incision function. Let Υ be defined as in Definition 11.

Suppose by *reductio ad absurdum* that $\text{mods}(\Upsilon(\mathbb{A})_D, \Upsilon(\mathbb{A})_\Sigma) = \emptyset$. Then, there must exist $X \subseteq \Upsilon(\mathbb{A})$ such that $\text{mods}(X_D, X_\Sigma) = \emptyset$, and for every possible $X' \subsetneq X$ it holds that $\text{mods}(X'_D, X'_\Sigma) \neq \emptyset$ (since $\text{mods}(\emptyset) \neq \emptyset$).

Let $KB = (\cup_{\mathcal{A}_i \in \mathbb{A}} KB_{\mathcal{A}_i})$. Since $\Upsilon(\mathbb{A}) = KB \setminus \varrho(\coprod_{\mathbb{A}})$ and $X \subseteq \Upsilon(\mathbb{A})$, then it holds that $X \cap \varrho(\coprod_{\mathbb{A}}) = \emptyset$. However, since $\Upsilon(\mathbb{A}) = KB \setminus \varrho(\coprod_{\mathbb{A}})$, then it holds that $X \subseteq KB$. Thus, we have that $X \subseteq KB$, $\text{mods}(X_D, X_\Sigma) = \emptyset$, and for every $X' \subsetneq X$ it holds that $\text{mods}(X'_D, X'_\Sigma) \neq \emptyset$. By Definition 6 this results in $X \in \coprod_{\mathbb{A}}$. Since $X \in \coprod_{\mathbb{A}}$, then from Definition 10 it follows that $X \cap \varrho(\coprod_{\mathbb{A}}) \neq \emptyset$.

Therefore, we have that $X \cap \varrho(\coprod_{\mathbb{A}}) = \emptyset$ and that $X \cap \varrho(\coprod_{\mathbb{A}}) \neq \emptyset$, a contradiction coming from our initial supposition that $\text{mods}(\Upsilon(\mathbb{A})_D, \Upsilon(\mathbb{A})_\Sigma) = \emptyset$, and it holds that $\text{mods}(\Upsilon(\mathbb{A})_D, \Upsilon(\mathbb{A})_\Sigma) \neq \emptyset$. \square

Now that we have proven that the knowledge base obtained by the cmerging operator is conflict-free, it remains to be seen how the operator behaves regarding the accrual function used to perform the incisions. It can be shown that the operator is indeed optimal in this aspect; *i.e.*, it obtains the best-valued knowledge base among the consistent ones.

Proposition 2 (Optimality). *Let Υ be defined as in Definition 11. Then, Υ satisfies*

- **Optimality:** *if $KB = \cup_{\mathcal{A}_i \in \mathbb{A}} KB_{\mathcal{A}_i}$ is inconsistent and $KB' \subset K$ is consistent then it holds that $\mathcal{V}(KB') \leq \mathcal{V}(\Upsilon(\mathbb{A}))$.*

Proof Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents, where for each $\mathcal{A} \in \mathbb{A}$, $\mathcal{A} = (KB_{\mathcal{A}}, \mathbb{V}_{\mathcal{A}})$ such that $KB_{\mathcal{A}}$ is the knowledge base and $\mathbb{V}_{\mathcal{A}}$ is the credibility vector for agent \mathcal{A} , $\coprod_{\mathbb{A}}$ be the set of kernels for \mathbb{A} , $\coprod_{\mathbb{A}} = \{\varsigma_1, \varsigma_2, \dots, \varsigma_n\}$ be the set of clusters obtained from $\coprod_{\mathbb{A}}$, and ϱ be an incision function. Let Υ be defined as in Definition 11, and let \mathbb{A} be such that $KB = \cup_{\mathcal{A}_i \in \mathbb{A}} KB_{\mathcal{A}_i}$ is inconsistent.

This proof is analogous to the one provided for contraction operators in [13]. It is based on the fact that the choice in every cluster is optimal, and that such result can be expanded to the entire knowledge base by means of **Domination**.

Let $F = KB \setminus \coprod_{\mathbb{A}}$ be the set of formulas that are not included in any conflict. We have then that $\{\varsigma_1, \varsigma_2, \dots, \varsigma_n, F\}$ is a partition of KB .

Consider an arbitrary $K' \subseteq KB$ such that KB' is consistent. From **Neutrality** in Lemma 1 it follows that set F does not influence in the credibility value of the final knowledge base obtained, and thus there is no need to consider it here. We do need to show, however, that the removal in each cluster retains a more credible subset of the original ontologies than the resolution used to obtain KB' .

To do this consider an arbitrary $\varsigma_i \in \mathbb{III}_{\mathbb{A}}$. Let $R_{\varsigma_i} = (KB \cap \varsigma_i) \setminus KB'$ be the sets of formulas removed from ς_i to build KB' . R_{ς_i} satisfies the first condition on Definition 10: since $R_{\varsigma_i} \subseteq KB \cap \varsigma_i$ and $\varsigma_i \in \mathbb{III}_{\mathbb{A}}$, then it holds that $R_{\varsigma_i} \subseteq \bigcup(\mathbb{III}_{\mathbb{A}})$. Moreover, R_{ς_i} also satisfies the second condition: since KB' is consistent then for every kernel $\kappa \in \mathbb{II}_{\mathbb{A}}$ such that $\kappa \subseteq \varsigma$ it holds that $\kappa \cap R_{\varsigma_i} \neq \emptyset$ (otherwise KB' will be rendered inconsistent, see proof of Proposition 1). Since R_{ς_i} satisfies those conditions then by Definition 10 we have that $I_{\varsigma_i} = \varsigma_i \cap \varrho(\mathbb{III}_{\mathbb{A}})$ (that is, the incision made over ς_i by the function ϱ) is such that $\mathcal{V}(\varsigma_i \setminus R_{\varsigma_i}) \leq \mathcal{V}(\varsigma_i \setminus I_{\varsigma_i})$.

This holds for any arbitrary $\varsigma_i \in \mathbb{III}_{\mathbb{A}}$. Now, we have that $\{\varsigma_1, \varsigma_2, \dots, \varsigma_n, F\}$ is a partition of K , and that for any part ς_i ($1 \leq i \leq n$) it holds that $\mathcal{V}(\varsigma_i \setminus R_{\varsigma_i}) \leq \mathcal{V}(\varsigma_i \setminus I_{\varsigma_i})$, and also by **Neutrality** $\mathcal{V}(KB) = \mathcal{V}(KB \setminus F)$. Finally, from **Domination** it holds that

$$\mathcal{V}(KB \setminus (\bigcup_{\varsigma_i \in \mathbb{III}_{\mathbb{A}}} R_{\varsigma_i} \cup F)) \leq \mathcal{V}(KB \setminus (\bigcup_{\varsigma_i \in \mathbb{III}_{\mathbb{A}}} I_{\varsigma_i} \cup F)).$$

This is equivalent to $\mathcal{V}(KB') \leq \mathcal{V}(\Upsilon(\mathbb{A}))$. □

Note that the optimality of the ck-merging operator is inherently connected to the credibility of the knowledge bases, rather than following an approach based on the number of removed formulas. That is, it can be the case that the operator removes more formulas than required from a quantitative point of view, thus making some “superfluous” removals. For instance, we can have a kernel $\{a_1, a_2, a_3, \tau_1\}$ without overlapping with any other kernel (so it is also a cluster), and the incision function choose to remove the set $\{a_1, a_2\}$, while the removal of any singleton subset of the kernel is sufficient to solve the conflict. Nevertheless, this is our purpose since we aim to maximize credibility instead of minimizing the number of formulas removed.

4. Specialized semantics for credibility-based merging

So far, in our work, we have introduced a general credibility-based approach to merge a community’s knowledge. To do this, we rely on a formal definition of credibility accrual functions. This definition does not entail a particular way to solve the conflicts but rather an entire family of functions satisfying properties that ensure that the knowledge base obtained by performing incisions based on them result to be the most credible ones, taking into consideration the community’s consensus.

Since Definition 9 is a general characterization of these functions, it is possible to consider many options to describe particular subsets of the mentioned family of functions. Therefore, we can define different functions that are tailored to particular application environments, thus entailing diverse semantics that can impact on the merging process. Nevertheless, to discuss different functions is out of the scope of this work. However, in this section, we will introduce some particular credibility accrual functions, along with a complete example that illustrates how they work.

Next, we begin by establishing how we measure the credibility of a particular formula. Then, we introduce a proper credibility accrual function that will calculate the credibility of a set of formulas.

4.1. Credibility of a single formula

We will introduce some definitions concerning the way the credibility of particular agents in a community is measured and how that affects the credibility associated with particular pieces of knowledge.

To do that, we introduce some auxiliary definitions that will be used on the rest of the paper. The set of sources for a certain formula is the set of those agents that have said formula in their knowledge base. Formally,

Definition 12 (Source set of a formula). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents on the form $\mathcal{A}_i = (KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i})$ such that $KB_{\mathcal{A}_i}$ is the $Datalog^\pm$ ontology and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i . Let $\alpha \in K_{\mathcal{L}}$ be a formula. The source set of α over \mathbb{A} is $\mathcal{S}_{\mathbb{A}}(\alpha) = \{\mathcal{A}_i \mid \mathcal{A}_i \in \mathbb{A} \text{ and } \alpha \in KB_{\mathcal{A}_i}\}$.

Example 10 (Source set of a formula). Consider the formula $\text{fund}(c1, p1)$ and our community of agents \mathbb{A}_5 introduced in Example 5. We can see that two agents have $\text{fund}(c1, p1)$ in their ontologies. Then, they both are a source of this formula, $\mathcal{S}_{\mathbb{A}_5}(\text{fund}(c1, p1)) = \{\mathcal{A}_1, \mathcal{A}_4\}$.

The credibility set of a particular agent in the community is the set of those credibilities associated to that agent in the credibility vectors of other agents in the community.

Definition 13 (Credibility set of an agent). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents on the form $\mathcal{A}_i = (KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i})$ such that $KB_{\mathcal{A}_i}$ is the $Datalog^\pm$ ontology and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i . The set of credibilities for agent \mathcal{A}_i over \mathbb{A} is

$$\widehat{C}(\mathcal{A}_i) = \{(\mathcal{A}_j, c_i) \mid (\mathcal{A}_i, c_i) \in \mathbb{V}_{\mathcal{A}_j} \text{ for some } \mathcal{A}_j \in \mathbb{A}\}.$$

In what follows, when stating the credibility set of an agent \mathcal{A}_i we will often omit the Agent \mathcal{A}_j that associate the credibility to \mathcal{A}_i , unless it is deemed as necessary for clarity reasons.

Note that by definition the credibility set of an agent may be empty, that is, it may not have any credibility associated to it. Nevertheless, as stated before in the rest of the paper given a set of agents \mathbb{A} we assume that there is no $\mathcal{A}_i \in \mathbb{A}$ such that $\widehat{C}(\mathcal{A}_i) = \emptyset$.

Example 11 (Credibility set of an agent). Consider the community of agents \mathbb{A}_5 depicted in Figure 1. We have that the credibilities for every agent in \mathbb{A}_5 are defined as follows,

- $\widehat{C}(\mathcal{A}_1) = \{(\mathcal{A}_2, 0.5), (\mathcal{A}_3, 0.6), (\mathcal{A}_4, 0.6), (\mathcal{A}_5, 0.8)\}$
- $\widehat{C}(\mathcal{A}_2) = \{(\mathcal{A}_1, 0.3), (\mathcal{A}_3, 0.2), (\mathcal{A}_4, 0.9), (\mathcal{A}_5, 0.3)\}$

- $\widehat{C}(\mathcal{A}_3) = \{(\mathcal{A}_1, 0.7), (\mathcal{A}_2, 0.4), (\mathcal{A}_5, 0.7)\}$
- $\widehat{C}(\mathcal{A}_4) = \{(\mathcal{A}_1, 0.7), (\mathcal{A}_2, 0.4), (\mathcal{A}_3, 0.7), (\mathcal{A}_5, 0.3)\}$
- $\widehat{C}(\mathcal{A}_5) = \{(\mathcal{A}_2, 0.6), (\mathcal{A}_3, 0.8), (\mathcal{A}_4, 0.8)\}$

When assessing the credibility associated with formulas in the community's knowledge, we will need to consider the credibility of several agents (those that claim such formula) as a whole.

We begin by establishing that the credibility associated to a particular agent in the community will be the minimal among those credibilities that other agents in the community associate to it. Formally,

Definition 14 (Credibility of an agent). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents on the form $\mathcal{A}_i = (KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i})$ such that $KB_{\mathcal{A}_i}$ is the $Datalog^\pm$ ontology and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i , and let $\mathcal{A}_i \in \mathbb{A}$. The agent credibility function $\widetilde{C} : \mathbb{A} \mapsto [0, 1]$ is defined as

$$\widetilde{C}(\mathcal{A}_i) = \min(\{c_i \mid (\mathcal{A}_j, c_i) \in \widehat{C}(\mathcal{A}_i)\})$$

Then, in this work the credibility of a set of agents is established as the credibility of the most credible of the agents in it.

Definition 15 (Credibility of a set of agents). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents on the form $\mathcal{A}_i = (KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i})$ such that $KB_{\mathcal{A}_i}$ is the $Datalog^\pm$ ontology and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i , and let $\mathbb{A}' \subseteq \mathbb{A}$. The agent set credibility function $\mathcal{C} : 2^{\mathbb{A}} \mapsto [0, 1]$ is defined as

$$\mathcal{C}(\mathbb{A}') = \max_{\mathcal{A}_i \in \mathbb{A}'}(\{c_i \mid c_i = \widetilde{C}(\mathcal{A}_i)\})$$

Example 12 (Credibility of a set of agents). For instance, consider the formula $\text{fund}(c1, p1)$, and its sources $\mathcal{S}_{\mathbb{A}_5}(\text{fund}(c1, p1)) = \{\mathcal{A}_1, \mathcal{A}_4\}$ introduced above in Example 10. The credibility of this set is computed as follows. First, we identify the minimal value in the credibility set for agents \mathcal{A}_1 and \mathcal{A}_4 , that is 0.5 and 0.3, respectively. Then, we obtain the maximal value among these values and the resulting credibility is $\mathcal{C}(\mathcal{S}_{\mathbb{A}_5}(\text{fund}(c1, p1))) = 0.5$.

Based on the credibility of a set of agents, we can define the credibility attached to a particular formula in the community knowledge. The value of a particular formula is established as the credibility associated with the set of agents that are the source of the formula.

Definition 16 (Credibility of a formula). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents on the form $\mathcal{A}_i = (KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i})$ such that $KB_{\mathcal{A}_i}$ is the $Datalog^\pm$ ontology and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i , and $\alpha \in K_{\mathcal{L}}$. The formula credibility function $\nu : K_{\mathcal{L}} \mapsto [0, 1]$ is defined as

$$\nu(\alpha) = \mathcal{C}(\mathcal{S}_{\mathbb{A}}(\alpha))$$

Example 13 (Credibility of a formula). Consider the community of agents introduced in Example 5, and also cluster $\{a_5, a_6, \nu_1\}$ from Example 7. Then, for the formulas in the cluster we have

$$\begin{aligned}\nu(a_5) &= \max(\min(0.5, 0.6, 0.6, 0.8), \min(0.7, 0.4, 0.7, 0.3)) = \max(0.5, 0.3) = 0.5 \\ \nu(a_6) &= \max(\min(0.5, 0.6, 0.6, 0.8), \min(0.6, 0.8, 0.8)) = \max(0.5, 0.6) = 0.6 \\ \nu(\nu_1) &= \max(\min(0.5, 0.6, 0.6, 0.8)) = \max(0.5) = 0.5\end{aligned}$$

4.2. Enhancing the credibility of a formula to sets of formulas: accumulation and absorption semantics

The previous definition gives us a way to measure the credibility of a particular formula; however, to accurately define our conflict resolution strategies, we need to measure the credibility attached to a set of formulas. As we previously observed, to present several examples of such functions is out of the scope of this work; nevertheless, we will now introduce some specialized credibility accrual functions that will serve as particular examples of the general credibility accrual functions defined earlier.

When **Neutrality** was introduced, we noticed that this property implies that sets of formulas that do not have any conflicting formulas in it are evaluated as the empty set. For our particular semantics, we will impose the additional restriction that the value assigned to the empty set is zero; that is, in what follows, we assume that $\mathcal{V}(\emptyset) = 0$.

Different options arise when it becomes necessary to define a particular credibility accrual function. An approach that it is reasonable is to define particularizations of Definition 9 that are based on the accumulation of the different values for formulas in a set.

Definition 17 (Accumulation Credibility Accrual Functions). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents on the form $\mathcal{A}_i = (KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i})$ such that $KB_{\mathcal{A}_i}$ is the Datalog[±] ontology and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i , and $\coprod_{\mathbb{A}}$ be the set of Kernels for \mathbb{A} . Also, let $\mathbb{K} = \bigcup_{K \in \coprod_{\mathbb{A}}} K$ the set of every formula in some kernel in $\coprod_{\mathbb{A}}$. The Accumulation credibility accrual function $\mathcal{V}^{\uparrow} : 2^{K_{\mathcal{C}}} \mapsto \mathbb{R}$ is defined as

$$\mathcal{V}^{\uparrow}(A) = \sum_{\alpha \in A \cap \mathbb{K}} (\nu(\alpha))$$

Example 14 (Accumulation Accrual Function). Consider the set of clusters $\coprod_{\mathbb{A}} = \{C_1, C_2\}$ where

$$C_1 = \{a_1, a_2, a_3, a_4, \sigma_1, \sigma_2, \sigma_3, \tau_1\} \text{ and } C_2 = \{a_5, a_6, \nu_1\}.$$

For instance, take cluster $C_2 = \{a_5, a_6, \nu_1\}$, and the credibility formulas in it shown in Example 13. The accumulation accrual function returns the following results for its subsets:

- $\mathcal{V}^{\uparrow}(\{a_5, a_6, \nu_1\}) = \mathcal{V}(\{a_5\}) + \mathcal{V}(\{a_6\}) + \mathcal{V}(\{\nu_1\}) = 0.5 + 0.6 + 0.5 = 1.6$

- $\mathcal{V}^\uparrow(\{a_5\}) = \mathcal{V}(\{a_5\}) = 0.5$
- $\mathcal{V}^\uparrow(\{a_6\}) = \mathcal{V}(\{a_6\}) = 0.6$
- $\mathcal{V}^\uparrow(\{\nu_1\}) = \mathcal{V}(\{\nu_1\}) = 0.5$
- $\mathcal{V}^\uparrow(\{a_5, a_6\}) = \mathcal{V}(\{a_5\}) + \mathcal{V}(\{a_6\}) = 0.5 + 0.6 = 1.1$
- $\mathcal{V}^\uparrow(\{a_5, \nu_1\}) = \mathcal{V}(\{a_5\}) + \mathcal{V}(\{\nu_1\}) = 0.5 + 0.5 = 1$
- $\mathcal{V}^\uparrow(\{a_6, \nu_1\}) = \mathcal{V}(\{a_6\}) + \mathcal{V}(\{\nu_1\}) = 0.6 + 0.5 = 1.1$

For space reasons we omit the calculations for the subsets of cluster C_1 , which will proceed in similar manner.

A different option could be, for instance, to associate the credibility of a set of formulas to some specific formula in it (with most/least credible formulas being the natural choices). We call such semantics as absorption semantics. Next, we formally define such a semantics, tailored to our framework.

Definition 18 (Maximality Credibility Accrual Functions). *Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents on the form $\mathcal{A}_i = (KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i})$ such that $KB_{\mathcal{A}_i}$ is the Datalog[±] ontology and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i , and $\coprod_{\mathbb{A}}$ be the set of Kernels for \mathbb{A} . Also, let $\mathbb{K} = \bigcup_{K \in \coprod_{\mathbb{A}}} K$ the set of*

every formula in some kernel in $\coprod_{\mathbb{A}}$.

The Maximality credibility accrual function $\mathcal{V}^{\max} : 2^{K\mathcal{L}} \mapsto [0, 1]$ is such that $\mathcal{V}^{\max}(A) = \nu(\alpha)$ iff $\alpha \in A \cap \mathbb{K}$ and there not exists $\beta \in A \cap \mathbb{K}$ such that $\nu(\alpha) < \nu(\beta)$

For space reasons we will not define and present results for Minimality credibility accrual functions, but it is fairly easy to extrapolate them from the ones for Maximality credibility accrual functions.

Example 15 (Maximality Credibility Accrual Functions). *Consider again the set of clusters $\coprod_{\mathbb{A}} = \{C_1, C_2\}$ where $C_1 = \{a_1, a_2, a_3, a_4, \sigma_1, \sigma_2, \sigma_3, \tau_1\}$ and $C_2 = \{a_5, a_6, \nu_1\}$.*

- $\mathcal{V}^{\max}(\{a_5, a_6, \nu_1\}) = \max(\{\mathcal{V}(\{a_5\}), \mathcal{V}(\{a_6\}), \mathcal{V}(\{\nu_1\})\}) = 0.6$
- $\mathcal{V}^{\max}(\{a_5\}) = \max(\{\mathcal{V}(\{a_5\})\}) = 0.5$
- $\mathcal{V}^{\max}(\{a_6\}) = \max(\{\mathcal{V}(\{a_6\})\}) = 0.6$
- $\mathcal{V}^{\max}(\{\nu_1\}) = \max(\{\mathcal{V}(\{\nu_1\})\}) = 0.5$
- $\mathcal{V}^{\max}(\{a_5, a_6\}) = \max(\{\mathcal{V}(\{a_5\}), \mathcal{V}(\{a_6\})\}) = 0.6$
- $\mathcal{V}^{\max}(\{a_5, \nu_1\}) = \max(\{\mathcal{V}(\{a_5\}), \mathcal{V}(\{\nu_1\})\}) = 0.5$
- $\mathcal{V}^{\max}(\{a_6, \nu_1\}) = \max(\{\mathcal{V}(\{a_6\}), \mathcal{V}(\{\nu_1\})\}) = 0.6$

For space reasons we omit the calculations for the subsets of cluster C_1 , which will proceed in similar manner.

The following result shows that the functions introduced in Definition 17 and Definition 18 are indeed proper accrual functions with respect to Definition 9.

Proposition 3. *Let \mathcal{V}^\uparrow be an accumulation accrual function. Then, \mathcal{V}^\uparrow is a credibility accrual function; that is to say, \mathcal{V}^\uparrow satisfies*

- **Transitivity:** if $\mathcal{V}^\uparrow(\{A\}) \leq \mathcal{V}^\uparrow(\{B\})$ and $\mathcal{V}^\uparrow(\{B\}) \leq \mathcal{V}^\uparrow(\{C\})$, then $\mathcal{V}^\uparrow(\{A\}) \leq \mathcal{V}^\uparrow(\{C\})$.
- **Domination:** if $\mathcal{P}(A) = \{A_1, A_2, \dots, A_n\}$ is a partition of A and for all A_i , ($1 \leq i \leq n$), it holds that $\mathcal{V}^\uparrow(\{A_i \setminus B\}) \leq \mathcal{V}^\uparrow(\{A_i \setminus C\})$, then we have that $\mathcal{V}^\uparrow(\{A \setminus B\}) \leq \mathcal{V}^\uparrow(\{A \setminus C\})$.
- **Pertinence:** for all $A \subseteq \bigcup_{1 \leq i \leq n} KB_{\mathcal{A}_i}$ and $\beta \in A$ it holds that if $\beta \notin \prod_{\mathbb{A}}$, then $\mathcal{V}^\uparrow(\{A\}) = \mathcal{V}^\uparrow(\{A \cup \{\beta\}\})$.

Proof

Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents of the form $\mathcal{A}_i = (KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i})$ such that $KB_{\mathcal{A}_i}$ is the Datalog[±] ontology and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i , and $\prod_{\mathbb{A}}$ be the set of Kernels for \mathbb{A} . Also, let $\mathbb{K} = \bigcup_{K \in \prod_{\mathbb{A}}} K$ the set of every formula in some kernel in $\prod_{\mathbb{A}}$.

- **Transitivity:** if $\mathcal{V}^\uparrow(\{A\}) \leq \mathcal{V}^\uparrow(\{B\})$ and $\mathcal{V}^\uparrow(\{B\}) \leq \mathcal{V}^\uparrow(\{C\})$, then $\mathcal{V}^\uparrow(\{A\}) \leq \mathcal{V}^\uparrow(\{C\})$.

Follows directly from the fact that the function is defined over \mathbb{R} .

- **Domination:** if $\mathcal{P}(A) = \{A_1, A_2, \dots, A_n\}$ is a partition of A and for all A_i ($1 \leq i \leq n$) it holds that $\mathcal{V}^\uparrow(\{A_i \setminus B\}) \leq \mathcal{V}^\uparrow(\{A_i \setminus C\})$ then we have that $\mathcal{V}^\uparrow(\{A \setminus B\}) \leq \mathcal{V}^\uparrow(\{A \setminus C\})$.

Let A be such that $\mathcal{P}(A) = \{A_1, A_2, \dots, A_n\}$ is a partition of A and for all A_i ($1 \leq i \leq n$) it holds that $\mathcal{V}^\uparrow(\{A_i \setminus B\}) \leq \mathcal{V}^\uparrow(\{A_i \setminus C\})$.

From Definition 17 we have that this is

$$(1) \quad \sum_{\alpha \in \{A_i \setminus B\} \cap \mathbb{K}} (\nu(\alpha)) < \sum_{\alpha \in \{A_i \setminus C\} \cap \mathbb{K}} (\nu(\alpha))$$

for every $A_i \in \mathcal{P}(A)$.

Since $\mathcal{P}(A)$ is a partition of A , then for every $A_i, A_j \in \mathcal{P}(A)$ we have that $A_i \cap A_j = \emptyset$. Then, since no formula belongs to more than one subset no interaction between parts is possible. From this and (1) we have that

$$(2) \quad \sum_{\alpha \in \{A_1 \setminus B\} \cap K} (\nu(\alpha)) + \sum_{\alpha \in \{A_2 \setminus B\} \cap K} (\nu(\alpha)) + \cdots + \sum_{\alpha \in \{A_n \setminus B\} \cap K} (\nu(\alpha))$$

$$<$$

$$\sum_{\alpha \in \{A_1 \setminus C\} \cap K} (\nu(\alpha)) + \sum_{\alpha \in \{A_2 \setminus C\} \cap K} (\nu(\alpha)) + \cdots + \sum_{\alpha \in \{A_n \setminus C\} \cap K} (\nu(\alpha)).$$

Also, since $\mathcal{P}(A)$ is a partition of A , then $A_1 \cup A_2 \cup \cdots \cup A_n = A$. Then, it holds that $\{A_1 \setminus B \cap K\} \cup \{A_2 \setminus B \cap K\} \cup \cdots \cup \{A_n \setminus B \cap K\} = \{A \setminus B \cap K\}$. Since every factor in the first member of equation (2) is smaller than its correspondent in the second member, $A_i \cap A_j = \emptyset$, and $A_1 \cup A_2 \cup \cdots \cup A_n = A$ then from (1) and (2) we have that

$$\sum_{\alpha \in \{A \setminus B\} \cap K} (\nu(\alpha)) < \sum_{\alpha \in \{A \setminus C\} \cap K} (\nu(\alpha)).$$

From Definition 17 this is $\mathcal{V}^\uparrow(\{A \setminus B\}) \leq \mathcal{V}^\uparrow(\{A \setminus C\})$.

- **Pertinence:** for all $A \subseteq \bigcup_{1 \leq i \leq n} KB_{A_i}$ and $\beta \in A$ it holds that if $\beta \notin \coprod_A$ then $\mathcal{V}^\uparrow(\{A\}) = \mathcal{V}^\uparrow(\{A \cup \{\beta\}\})$.

From Definition 17 we have that

$$\mathcal{V}^\uparrow(\{A \cup \{\beta\}\}) = \sum_{\alpha \in \{A \cup \{\beta\}\} \cap K} (\nu(\alpha)).$$

Since $\beta \notin \coprod_A$, then $\beta \notin K$, and by extension $\beta \notin \{A \cup \{\beta\}\} \cap K$. Thus, it holds that $\{A \cup \{\beta\}\} \cap K = \{A\} \cap K$, and

$$\sum_{\alpha \in \{A \cup \{\beta\}\} \cap K} (\nu(\alpha)) = \sum_{\alpha \in \{A\} \cap K} (\nu(\alpha)).$$

From Definition 17 this is $\mathcal{V}^\uparrow(\{A\}) = \mathcal{V}^\uparrow(\{A \cup \{\beta\}\})$.

□

Proposition 4. Let \mathcal{V}^{\max} be a maximality credibility accrual function. Then, \mathcal{V}^{\max} is a credibility accrual function, that is, it satisfies

- **Transitivity:** if $\mathcal{V}^{\max}(\{A\}) \leq \mathcal{V}^{\max}(\{B\})$ and $\mathcal{V}^{\max}(\{B\}) \leq \mathcal{V}^{\max}(\{C\})$, then $\mathcal{V}^{\max}(\{A\}) \leq \mathcal{V}^{\max}(\{C\})$.
- **Domination:** if $\mathcal{P}(A) = \{A_1, A_2, \dots, A_n\}$ is a partition of A and for all A_i ($1 \leq i \leq n$) it holds that $\mathcal{V}^{\max}(\{A_i \setminus B\}) \leq \mathcal{V}^{\max}(\{A_i \setminus C\})$ then we have that $\mathcal{V}^{\max}(\{A \setminus B\}) \leq \mathcal{V}^{\max}(\{A \setminus C\})$.

- **Pertinence:** for all $A \subseteq \bigcup_{1 \leq i \leq n} KB_{\mathcal{A}_i}$ and $\beta \in A$ it holds that if $\beta \notin \prod_{\mathbb{A}}$ then $\mathcal{V}^{\max}(\{A\}) = \mathcal{V}^{\max}(\{A \cup \{\beta\}\})$.

Proof Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ be a community of agents on the form $\mathcal{A}_i = (KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i})$ such that $KB_{\mathcal{A}_i}$ is the Datalog[±] ontology and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i , and $\prod_{\mathbb{A}}$ be the set of Kernels for \mathbb{A} . Also, let $\mathbf{K} = \bigcup_{K \in \prod_{\mathbb{A}}} K$ the set of every formula in some kernel in $\prod_{\mathbb{A}}$.

- **Transitivity:** if $\mathcal{V}^{\max}(\{A\}) \leq \mathcal{V}^{\max}(\{B\})$ and $\mathcal{V}^{\max}(\{B\}) \leq \mathcal{V}^{\max}(\{C\})$, then $\mathcal{V}^{\max}(\{A\}) \leq \mathcal{V}^{\max}(\{C\})$.

Follows directly from the fact that the function is defined over \mathbb{R} .

- **Domination:** if $\mathcal{P}(A) = \{A_1, A_2, \dots, A_n\}$ is a partition of A and for all A_i ($1 \leq i \leq n$) it holds that $\mathcal{V}^{\max}(\{A_i \setminus B\}) \leq \mathcal{V}^{\max}(\{A_i \setminus C\})$ then we have that $\mathcal{V}^{\max}(\{A \setminus B\}) \leq \mathcal{V}^{\max}(\{A \setminus C\})$.

Let A be such that $\mathcal{P}(A) = \{A_1, A_2, \dots, A_n\}$ is a partition of A and for all A_i ($1 \leq i \leq n$) it holds that $\mathcal{V}^{\max}(\{A_i \setminus B\}) \leq \mathcal{V}^{\max}(\{A_i \setminus C\})$.

By Definition 18 we have that if for every A_i it holds that $\mathcal{V}^{\max}(\{A_i \setminus B\}) \leq \mathcal{V}^{\max}(\{A_i \setminus C\})$ then there exists $\gamma \in \{\{A_i \setminus B\} \cap \mathbf{K}\}$ such that $(\nu(\gamma)) \leq (\nu(\beta))$ for every $\beta \in \{A_i \setminus C\} \cap \mathbf{K}$. This holds for every $A_i \in \mathcal{P}(A)$, and thus there exists $\gamma' \in \{\bigcup_{1 \leq i \leq n} (A_i \setminus B) \cap \mathbf{K}\}$ such that $(\nu(\gamma')) \leq (\nu(\beta'))$ for every $\beta' \in \{\bigcup_{1 \leq i \leq n} (A_i \setminus C) \cap \mathbf{K}\}$.

Since $\bigcup_{1 \leq i \leq n} (A_i) = A$ because $\mathcal{P}(A)$ is a partition of A , then by replacing $\bigcup_{1 \leq i \leq n} (A_i)$ with A in the previous statement we have that there exists $\gamma' \in \{\{A \setminus B\} \cap \mathbf{K}\}$ such that $(\nu(\gamma')) \leq (\nu(\beta'))$ for every $\beta' \in \{\{A \setminus C\} \cap \mathbf{K}\}$. By definition of \mathcal{V}^{\max} this is $\mathcal{V}^{\max}(\{A \setminus B\}) \leq \mathcal{V}^{\max}(\{A \setminus C\})$.

- **Pertinence:** for all $A \subseteq \bigcup_{1 \leq i \leq n} KB_{\mathcal{A}_i}$ and $\beta \in A$ it holds that if $\beta \notin \prod_{\mathbb{A}}$ then $\mathcal{V}^{\max}(\{A\}) = \mathcal{V}^{\max}(\{A \cup \{\beta\}\})$.

From Definition 18 we have that

$$\mathcal{V}^{\max}(\{A \cup \{\beta\}\}) = (\nu(\alpha))_{\alpha \in \{A \cup \{\beta\}\} \cap \mathbf{K}}$$

Since $\beta \notin \prod_{\mathbb{A}}$, then $\beta \notin \mathbf{K}$, and by extension $\beta \notin \{A \cup \{\beta\}\} \cap \mathbf{K}$. Thus, it holds that $\{A \cup \{\beta\}\} \cap \mathbf{K} = \{A\} \cap \mathbf{K}$, and

$$(\nu(\alpha))_{\alpha \in \{A \cup \{\beta\}\} \cap \mathbf{K}} = (\nu(\alpha))_{\alpha \in \{A\} \cap \mathbf{K}}.$$

From Definition 18 this is $\mathcal{V}^{\max}(\{A\}) = \mathcal{V}^{\max}(\{A \cup \{\beta\}\})$.

□

To sum up, we will introduce an example of particular incision functions based on the semantics introduced in this section, and the merging operators that arise from their use.

Example 16 (Particular Incision functions and Merging operators).

Consider the Datalog[±] ontology KB_2 presented in Example 2, which is the union of the particular knowledge bases of the agents in the community \mathbb{A}_5 presented in Example 5.

To see how incision functions based on accumulation and absorption accrual functions behave, we will once again focus on cluster C_2 . First, it is easy to see that every subset of C_2 satisfies the first two conditions in the definition of incision functions. Then, all that remains is to search for a subset such that its removal preserves the most credibility, according to the chosen semantics. To do this, incision functions should calculate the values of every possible solution for the cluster (that is, every subset that intersects all kernels in it) and find out which one retains the most value once removed. Looking into Example 14 we can see that such sets can be as follows:

- $\mathcal{V}^\uparrow(C_2 \setminus \{\nu_1\}) = \mathcal{V}(\{a_5\}) + \mathcal{V}(\{a_6\}) = 0.5 + 0.6 = 1.1$
- $\mathcal{V}^\uparrow(C_1 \setminus \{a_1\}) = \mathcal{V}(\{a_2\}) + \mathcal{V}(\{a_3\}) + \mathcal{V}(\{a_4\}) + \mathcal{V}(\{\tau_1\}) + \mathcal{V}(\{\sigma_1\}) + \mathcal{V}(\{\sigma_2\}) + \mathcal{V}(\{\sigma_3\}) = 0.6 + 0.6 + 0.5 + 0.5 + 0.2 + 0.3 + 0.6 = 3.3$
- $\mathcal{V}^{\max}(C_2 \setminus \{a_5\}) = \max(\{\mathcal{V}(\{\nu_1\}), \mathcal{V}(\{a_6\})\}) = 0.6$
- $\mathcal{V}^{\max}(C_1 \setminus \{\sigma_2\}) = \max(\{\mathcal{V}(\{a_1\}), \mathcal{V}(\{a_2\}), \mathcal{V}(\{a_3\}), \mathcal{V}(\{a_4\}), \mathcal{V}(\{\tau_1\}), \mathcal{V}(\{\sigma_1\}), \mathcal{V}(\{\sigma_3\})\}) = 0.6$

Finally we have that formulas chosen for removal are

- $I_{\text{accum}} = \{a_1, \nu_1\}$.
- $I_{\text{max}} = \{a_5, \sigma_2\}$.

Note that there exists other possibilities as well, particularly for Maximality, that arise different merging operators.

If we remove these formulas from KB_2 , then a consistent merged Datalog[±] ontology is obtained. By applying Υ , the result is the following merged knowledge base:

$$KB_2 \setminus I_{\text{accum}} = \left\{ \begin{array}{l} \text{D} : \{a_2 : \text{unemployment}(p1), a_3 : \text{incSBZ}(p1), \\ a_4 : \text{incOTB}(p1), a_5 : \text{fund}(c1, p1), a_6 : \text{fund}(c1, p2)\} \\ \Sigma_{NC} : \{\tau_1 : \text{investDA}(P) \wedge \text{investBZ}(P) \rightarrow \perp\} \\ \Sigma_E : \{\} \\ \Sigma_T : \{\sigma_1 : \text{expST}(P) \wedge \text{incOTB}(P) \rightarrow \text{investBZ}(P), \\ \sigma_2 : \text{expST}(P) \rightarrow \text{investDA}(P), \\ \sigma_3 : \text{unemployment}(P) \wedge \text{incSBZ}(P) \rightarrow \text{investBZ}(P)\} \end{array} \right\}$$

$$KB_2 \setminus I_{max} = \left\{ \begin{array}{l} D : \{a_1 : expST(p1), a_2 : unemployment(p1), a_3 : incSBZ(p1), \\ a_4 : incOTB(p1), a_6 : fund(c1, p2)\} \\ \Sigma_{NC} : \{\tau_1 : investDA(P) \wedge investBZ(P) \rightarrow \perp\} \\ \Sigma_E : \{\nu_1 : fund(E, P) \wedge fund(E, P') \rightarrow P = P'\} \\ \Sigma_T : \{\sigma_1 : expST(P) \wedge incOTB(P) \rightarrow investBZ(P), \\ \sigma_3 : unemployment(P) \wedge incSBZ(P) \rightarrow investBZ(P)\} \end{array} \right\}$$

Note that these knowledge bases are different among them. Nevertheless this is clearly intended since we aim to maximize credibility with respect to different particular credibility accrual functions.

In closing this section, it is interesting to analyze if, according to the chosen accumulation semantics, the resulting KB is the most credible one, and to verify this fact, it is necessary to check whether the merging operators based on incision functions arising from the use of this particular semantics satisfy **Optimality**. To do that we begin by considering the result of $\Upsilon(\mathbb{A}) = KB_2 \setminus I_{accum}$ and every KB' that is a superset of the merged KB, *i.e.*, $KB_2 \setminus I_{accum} \subsetneq KB' \subsetneq KB_2$. However, the only way to get a superset of the resulting merged knowledge base is to include some of the removed formulas. Nevertheless, if that is the case, then KB' is clearly inconsistent. So, we do not need to account for it, since it is of no interest as it falls out of the hypothesis in **Optimality**.

On the other hand, every $KB' \subsetneq KB_2$ such that KB' is consistent (and $KB' \neq \Upsilon(\mathbb{A})$) implies that $KB' \subsetneq \Upsilon(\mathbb{A})$ since the only way to obtain another consistent knowledge base will be to remove additional formulas from the original community knowledge. Let us consider any such set; since accumulation functions aggregate the credibility of the formulas in a set, it is easy to see that it cannot be the case that $\mathcal{V}(\Upsilon(\mathbb{A})) < \mathcal{V}(KB')$, which goes to show that a ck-merging operator using an accumulation accrual function as the basis of its incision function will indeed satisfy **Optimality**.

Now, we move on to the case of the merging operator using Maximality credibility accrual functions. It is easy to show that indeed the operator satisfies **Optimality** by design, since for every cluster at least the most valued formula is retained. Then, such formula is part of the merged KB. Plus, every formula that is in the original knowledge and is not part of a kernel (*i.e.*, is not involved in any conflict) is also retained in the merged knowledge. Given that scenario, to assume that there may be another subset of the original knowledge that is both consistent and has more credibility is an absurd, because that will imply that the selection in a cluster was not optimal, which is clearly not the case.

5. Formal comparison with Kernel Merge [19]

We have formally introduced up to this point a general approach to merge ontologies, where each of these ontologies represents the knowledge of an agent in the community. Moreover, each agent explicitly states how much it trusts another agent (when it has that information available about that other agent), and this trust degree is the cornerstone in our conflict resolution strategy.

We will now look into some formal properties of our approach. In particular, we will formally relate it with the work of Falappa *et al.* [19] that presented the Kernel Merge operator approach. We will adapt to our multi-agent, ontology-based approach, the notions introduced in [19].

Among the numerous works on knowledge merging in the literature, one that is particularly interesting to us is that of [19]. The rationale behind this focus is that in there, the authors also present merging operators based on Hansson's incision functions called *Kernel Merge* operators. For reasons of space, we will skip now the formal definition of Kernel Merge operators, and directly explore into their properties and how they relate to our approach, by exploiting the kernel merge operators' representation theorem introduced in [19]. Before that, however, we will give an intuitive presentation of the operator's behavior, referring the reader to the source for further clarification. Succinctly, the process used by kernel merge operators is as follows. To begin with, the operators consider as their epistemic input two belief bases K and A . The applied mechanism is then to perform a set union of K and A , and then proceed to eliminate from the result all possible inconsistencies in it, which is known in the Belief Revision literature as external revision [25]. To solve such conflicts kernel merge operators use incision functions over each minimally inconsistent subset of $K \cup A$.

We now move on to a formal analysis of the relation between our approach and Falappa *et al.*'s. We can show that operators defined in definition 11 satisfy some of the postulates in Theorem 3 in [19], adapted to our scenario; this result is formalized next.

Proposition 5. *Let $\mathcal{A}_1 = (KB_{\mathcal{A}_1}, \mathbb{V}_{\mathcal{A}_1})$ and $\mathcal{A}_2 = (KB_{\mathcal{A}_2}, \mathbb{V}_{\mathcal{A}_2})$ be two agents such that $KB_{\mathcal{A}_1}, KB_{\mathcal{A}_2}$ are the Datalog $^\pm$ ontologies for agents \mathcal{A}_1 and \mathcal{A}_2 , respectively; and $\mathbb{V}_{\mathcal{A}_1}, \mathbb{V}_{\mathcal{A}_2}$ are the credibility vectors for agents \mathcal{A}_1 and \mathcal{A}_2 , respectively. Also, let Υ be a merging operator defined as in Definition 11. Then, Υ satisfies the following postulates:*

- **Inclusion:** $\Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\}) \subseteq KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}$.
- **Strong Consistency:** $\Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$ is consistent.
- **Global Core Retainment:** If $\alpha \in (KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}) \setminus \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$ then there exists a set C such that $C \subseteq (KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2})$, C is consistent and $C \cup \{\alpha\}$ is inconsistent.

Proof Let $\mathcal{A}_1 = (KB_{\mathcal{A}_1}, \mathbb{V}_{\mathcal{A}_1})$ and $\mathcal{A}_2 = (KB_{\mathcal{A}_2}, \mathbb{V}_{\mathcal{A}_2})$ be two agents, then $KB_{\mathcal{A}_1}, KB_{\mathcal{A}_2}$ are the Datalog $^\pm$ ontologies for agents \mathcal{A}_1 and \mathcal{A}_2 and $\mathbb{V}_{\mathcal{A}_1}, \mathbb{V}_{\mathcal{A}_2}$

are the credibility vectors for agents \mathcal{A}_1 and \mathcal{A}_2 , respectively. Also, for some $S \subseteq (KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2})$. Let $S_D = S \cap (D_{\mathcal{A}_1} \cup D_{\mathcal{A}_2})$ and $S_\Sigma = S \cap (\Sigma_{\mathcal{A}_1} \cup \Sigma_{\mathcal{A}_2})$ be the set of atoms and the set of rules, respectively. Finally, let Υ be a ck-merging operator defined as in Definition 11.

- **Inclusion:** $\Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\}) \subseteq KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}$.

Follows from Definition 11, since

$$\Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\}) = (KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}) \setminus \varrho(\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}})$$

- **Strong Consistency:** $\Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$ is consistent.

To prove that Υ is strongly consistent it is enough to show that all minimal conflicts are attended to; *i.e.*, that there is no $X \in \coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}}$ such that $X \subseteq \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$.

Consider any arbitrary $X \in \coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}}$. From Definition 8 it follows that there exists $Y \in \coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}}$ such that $X \subseteq Y$ (see Proposition 2 in [12]). From the definition of ϱ it follows that $\varrho(\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}}) \cap X \neq \emptyset$. Thus, since $\Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\}) = (KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}) \setminus \varrho(\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}})$ we have that $X \not\subseteq \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$. Since this holds for any arbitrary $X \in \coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}}$ then there is no $A \subseteq \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$ such that $\text{mods}(A_D, A_\Sigma) = \emptyset$. Thus, it holds that $\text{mods}(\Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})_D, \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})_\Sigma) \neq \emptyset$, *i.e.*, $\Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$ is consistent.

- **Global Core Retainment:** If $\alpha \in (KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}) \setminus \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$ then there exists a set C such that $C \subseteq (KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2})$, C is consistent and $C \cup \{\alpha\}$ is inconsistent.

Let $\alpha \in (KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}) \setminus \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$. Then, $\alpha \notin \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$. From Definition 11 we have then that $\alpha \in \varrho(\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}})$, and from Definition 10 follows that $\alpha \in X$ for some $X \in \coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}}$. Let $C = X \setminus \{\alpha\}$. From Definition 6 we have $C \subseteq (KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2})$ (since $X \subseteq (KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2})$). Also, since $C \subset X$ then $\text{mods}(C_D, C_\Sigma) \neq \emptyset$ and C is consistent. Finally, $\text{mods}(X_D, X_\Sigma) = \emptyset$, and since $C \cup \{\alpha\} = X$ then $C \cup \{\alpha\}$ is inconsistent.

□

It is important to remark that, following Hansson's original paper's spirit, the definition of incision functions in [19] are left as general, not providing thus any particular way to define *what* the functions choose for deletion. As we have shown, we provide further conditionings over the incision functions to exploit the credibility associated with agents in the community, thus refining operators to perform deletions informed under that light. This characteristic becomes essential since it is the trigger behind the following result, which states that our operators do not satisfy the **Reversion** postulate in Theorem 3 in [19], thus separating the approaches.

Observation 1. Let $\mathcal{A}_1 = (KB_{\mathcal{A}_1}, \mathbb{V}_{\mathcal{A}_1})$, $\mathcal{A}_2 = (KB_{\mathcal{A}_2}, \mathbb{V}_{\mathcal{A}_2})$, $\mathcal{A}_3 = (KB_{\mathcal{A}_3}, \mathbb{V}_{\mathcal{A}_3})$ and $\mathcal{A}_4 = (KB_{\mathcal{A}_4}, \mathbb{V}_{\mathcal{A}_4})$ be agents such that $KB_{\mathcal{A}_i}$ is the Datalog[±] ontology for

agent \mathcal{A}_i ; and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i ($1 \leq i \leq 4$). Also, let Υ be a *ck-merging operator* defined as in Definition 11. Υ does not generally satisfy the following postulate:

- **Reversion:** If $KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}$ and $KB_{\mathcal{A}_3} \cup KB_{\mathcal{A}_4}$ have the same set of minimal inconsistent subsets then:

$$(KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}) \setminus \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\}) = (KB_{\mathcal{A}_3} \cup KB_{\mathcal{A}_4}) \setminus \Upsilon(\{\mathcal{A}_3, \mathcal{A}_4\}).$$

The proof for the statement in Observation 1 comes from the fact that, even if the conflicts in the knowledge bases for agents \mathcal{A}_1 and \mathcal{A}_2 are the same than those for agents \mathcal{A}_3 and \mathcal{A}_4 , it can be the case that the credibility vectors of the agents differ in such a way that different resolutions for the conflicts are triggered, thus deleting different formulas.

However, it is clear that if credibilities associated with agents are such that conflicting formulas are valued in a way such that the order is considered, then the approaches do coincide. In the light of that last statement, we can establish a stronger version of **Reversion** tailored for our scenario, asking that not only the minimal conflicts coincide but also the credibility associated with agents.

We first need to define when credibility vectors in some community of agents are equivalent. Intuitively, sets of credibility vectors (and thus so communities) are equivalent with respect to some particular accrual function when they induce the same order among formulas for the particular accrual function in consideration.

Definition 19 (Communities' equivalence). Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ and $\mathbb{B} = \{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_n\}$ be communities of agents on the form $\mathcal{A}_i = (KB_{\mathcal{A}_i}, \mathbb{V}_{\mathcal{A}_i})$ (respectively, $\mathcal{B}_i = (KB_{\mathcal{B}_i}, \mathbb{V}_{\mathcal{B}_i})$) such that $KB_{\mathcal{A}_i}$ (respectively, $KB_{\mathcal{B}_i}$) is the *Datalog[±] ontology* and $\mathbb{V}_{\mathcal{A}_i}$ (respectively, $\mathbb{V}_{\mathcal{B}_i}$) is the *credibility vector* for agent \mathcal{A}_i (respectively, \mathcal{B}_i), be such that $\coprod_{\mathbb{A}} = \coprod_{\mathbb{B}}$, and \mathcal{V} be a *credibility accrual function*.

We say that \mathbb{A} and \mathbb{B} are *equivalent with respect to \mathcal{V}* , denoted $\mathbb{A} \cong_{\mathcal{V}} \mathbb{B}$, iff for all $B, C \subseteq \coprod_{\mathbb{A}}$ and $B', C' \subseteq \coprod_{\mathbb{B}}$ such that $B = B'$ and $C = C'$ it holds that $\mathcal{V}(B) \leq \mathcal{V}(C)$ iff $\mathcal{V}(B') \leq \mathcal{V}(C')$.

Note that because of **Neutrality**, we only need to be concerned with conflicting formulas instead of asking for the complete ontologies to coincide. If two communities are equivalent with respect to some accrual function \mathcal{V} , then they both will induce the same \mathcal{V} -based incision functions, and thus the deletions will coincide. That is, even if agents in the community does not have the exact same credibility vectors if they are equivalent then some set A that is deemed as less valuable than another set B in a community will be also less valuable in the other one (even if the values for A and B are not precisely the same for both communities). This fact helps us to propose the following property of **Strong Reversion**. Then, we can show that operators defined in our approach belong to this *strong kernel merge* class of operators.

Proposition 6. Let $\mathbb{A} = \{\mathcal{A}_1, \mathcal{A}_2\}$ and $\mathbb{A}' = \{\mathcal{A}_3, \mathcal{A}_4\}$ be such that $\mathcal{A}_1 = (KB_{\mathcal{A}_1}, \mathbb{V}_{\mathcal{A}_1})$, $\mathcal{A}_2 = (KB_{\mathcal{A}_2}, \mathbb{V}_{\mathcal{A}_2})$, $\mathcal{A}_3 = (KB_{\mathcal{A}_3}, \mathbb{V}_{\mathcal{A}_3})$ and $\mathcal{A}_4 = (KB_{\mathcal{A}_4}, \mathbb{V}_{\mathcal{A}_4})$ are agents such that $KB_{\mathcal{A}_i}$ is the Datalog[±] ontology for agent \mathcal{A}_i ; and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i ($1 \leq i \leq 4$). Also, let Υ be a merging operator defined as in Definition 11. Then, Υ satisfies the following postulate:

- **Strong Reversion:** If $KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}$ and $KB_{\mathcal{A}_3} \cup KB_{\mathcal{A}_4}$ have the same set of minimal inconsistent subsets and $\mathbb{A} \cong_{\mathbb{V}} \mathbb{A}'$ then $(KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}) \setminus \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\}) = (KB_{\mathcal{A}_3} \cup KB_{\mathcal{A}_4}) \setminus \Upsilon(\{\mathcal{A}_3, \mathcal{A}_4\})$.

Proof Let $\mathcal{A}_1 = (KB_{\mathcal{A}_1}, \mathbb{V}_{\mathcal{A}_1})$, $\mathcal{A}_2 = (KB_{\mathcal{A}_2}, \mathbb{V}_{\mathcal{A}_2})$, $\mathcal{A}_3 = (KB_{\mathcal{A}_3}, \mathbb{V}_{\mathcal{A}_3})$ and $\mathcal{A}_4 = (KB_{\mathcal{A}_4}, \mathbb{V}_{\mathcal{A}_4})$ be agents such that $KB_{\mathcal{A}_i}$ is the Datalog[±] ontology for agent \mathcal{A}_i ; and $\mathbb{V}_{\mathcal{A}_i}$ is the credibility vector for agent \mathcal{A}_i ($1 \leq i \leq 4$). Also, let Υ be a merging operator defined as in Definition 11.

- **Strong Reversion:** If $KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}$ and $KB_{\mathcal{A}_3} \cup KB_{\mathcal{A}_4}$ have the same set of minimal inconsistent subsets and $\mathbb{A} \cong_{\mathbb{V}} \mathbb{A}'$ then

$$(KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}) \setminus \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\}) = (KB_{\mathcal{A}_3} \cup KB_{\mathcal{A}_4}) \setminus \Upsilon(\{\mathcal{A}_3, \mathcal{A}_4\}).$$

Let $KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}$ and $KB_{\mathcal{A}_3} \cup KB_{\mathcal{A}_4}$ be such that they have the same set of minimal inconsistent subsets and $\mathbb{A} \cong_{\mathbb{V}} \mathbb{A}'$.

Consider $I \subset ((KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}) \setminus \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\}))$. Then, $I \not\subseteq \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\})$. From Definition 11 we have then that $I \subset \varrho(\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}})$.

Consider any arbitrary $\varsigma \in \coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}}$, and let $I' \subset \varrho(\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}})$ be such that $I' = I \cap \varsigma$, i.e., the set of formulas removed from ς by ϱ . From Definition 10 we have that for all $F \subset \varsigma$ where F satisfies the two first conditions in Definition 10 it holds that $\mathcal{V}(\varsigma \setminus F) \leq \mathcal{V}(\varsigma \setminus I')$ (and we do not need to be concerned with those sets that do not satisfy the conditions because of **Strong Consistency**, since their deletion does not solve the conflict).

Since $KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}$ and $KB_{\mathcal{A}_3} \cup KB_{\mathcal{A}_4}$ are such that they have the same set of minimal inconsistent subsets, then from Definition 6 we have that $\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}} = \coprod_{\{\mathcal{A}_3, \mathcal{A}_4\}}$, and then from Definition 8 that $\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}} = \coprod_{\{\mathcal{A}_3, \mathcal{A}_4\}}$. Thus, $\varsigma \in \coprod_{\{\mathcal{A}_3, \mathcal{A}_4\}}$, and $I' \subset \varsigma$. Also, since $\mathbb{A} \cong_{\mathbb{V}} \mathbb{A}'$ then it holds that for all $F \subset \varsigma$ where F satisfies the two previous conditions in the definition it holds that $\mathcal{V}(\varsigma \setminus F) \leq \mathcal{V}(\varsigma \setminus I')$ for \mathbb{A}' as well. Then, it holds that $I' \subset \varrho(\coprod_{\{\mathcal{A}_3, \mathcal{A}_4\}})$.

So, we have that for $\varsigma \in \coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}}$ (and hence $\varsigma \in \coprod_{\{\mathcal{A}_3, \mathcal{A}_4\}}$) it holds that if $I' \subset \varrho(\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}})$ then $I' \subset \varrho(\coprod_{\{\mathcal{A}_3, \mathcal{A}_4\}})$, i.e., the conflict resolution for the cluster coincides. Since this holds for any arbitrary $\varsigma \in \coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}}$ and $\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}} = \coprod_{\{\mathcal{A}_3, \mathcal{A}_4\}}$ then we have that $\varrho(\coprod_{\{\mathcal{A}_1, \mathcal{A}_2\}}) = \varrho(\coprod_{\{\mathcal{A}_3, \mathcal{A}_4\}})$. From Definition 11 follows then that $(KB_{\mathcal{A}_1} \cup KB_{\mathcal{A}_2}) \setminus \Upsilon(\{\mathcal{A}_1, \mathcal{A}_2\}) = (KB_{\mathcal{A}_3} \cup KB_{\mathcal{A}_4}) \setminus \Upsilon(\{\mathcal{A}_3, \mathcal{A}_4\})$. □

6. Related work

In Section 5, we have already presented a formal comparison between our work and some related works that are particularly close. We will now take a more comprehensive look into the relation of our approach with the literature of both Belief Revision (especially knowledge merging) and Social Choice Theory.

6.1. Merging knowledge bases

Many significant works have contributed to the study the problem of merging potentially conflicting knowledge coming from multiple sources (see, for instance, [35, 30, 28, 31, 47, 19, 17, 44, 45]) which have provided inspiration and interesting foundations for further efforts on belief revision. We will look closer into similarities and differences between those works and our approach.

One of the cornerstones in the area of knowledge merging is the work presented by Konieczny and Pino-Pérez in [30]. Although we follow the same motivations by studying the problem of merging conflicting information, there are some differences between their work and ours. Unlike our approach, they state that the merged knowledge base will be consistent if the set of integrity constraints used to “guide” the merging process is consistent, and they do not study the alternative case. In our work, however, we do not impose assumptions over these “integrity constraints”, that is, NCs and EGDs whose semantics in some way restrict the knowledge. As a result, we can have an incoherent set of constraints, which is somehow equivalent to constraints inconsistency in ontological settings (see [12]), and the merged knowledge base will be consistent. Another difference is that they focus on proposing a logical characterization of merging operators and do not consider in their formalisms a model of trust defined over the knowledge bases. In contrast, we propose an approach for multi-agent scenarios where every agent attaches to the rest of the agents in the community a credibility value, defining in this manner a preference order over these agents.

Meyer, Lee, and Booth [39] use two well-known strategies for knowledge integration for propositional inconsistency management, adapting them to a Description Logic setting. The proposed approach takes knowledge bases and generates a disjunctive knowledge base (DKB) as the result of the integration. We propose some strategies to deal with inconsistency when integrating knowledge, but the resulting merged knowledge base is a regular Datalog[±] ontology instead of a disjunctive knowledge base. Also, while not explicitly stated here for reasons of space, in our approach is possible to deal with incoherent ontologies: it will simply render new incoherence-related kernels. Then, the conflict resolution machinery will take them and solve them, either by removing atoms, TGDs, or constraints. On the contrary, Meyer *et al.* set aside problems related to incoherence in the integration process as further research.

In [28], Hunter and Liu propose an approach for generating a stratified knowledge base from a set of input knowledge bases that is based on the support degree a formula receives. Following this, they propose some merging operators.

They also define the notion of most primed formulas and the most entailed formulas introducing methods to select them. In a development that keeps some similarity to these, we establish different methods to calculate the value attached to a particular formula in the merging process. However, in contrast, the main difference with their approach, besides the goal (degree of support vs. credibility accrual), is that they focused on how to obtain information from the original sources from which formula gathered more support, and we centered on agent credibilities to deal with the problem of pieces of conflicting information from different source of information in a multi-agent context. In fact, we use accrual functions to decide which formula needs to be removed to solve conflicts, which clearly provides a useful tool that goes further the work [28].

In the Logic Programming area, a number of relevant research works also have focused on dealing with the problem of knowledge merging. In [27], Hué *et al.* have introduced a merging approach based on stable model semantics using the logic of *Here-and-There* introduced in [49]. The authors consider a merging strategy based on the method of *Removed Sets Fusion* to decide the set of formulas to remove in order to restore inconsistency. To achieve this, they assume a pre-order among the candidates for deletion that are obtained by a given strategy. Although it falls outside of the goals of our work, it would be reasonably easy to adapt our approach to using a preference criterion over sets of formulas selected from conflicting information sets arising from the use of accrual functions, similarly as to how it is done on [13].

Delgrande *et al.* [14] introduced another important work in the Logic Programming area. In there, the authors propose two different operators for merging logic programs under answer set semantics. The first one follows the approach of arbitration, selecting those models of the programs that depart the least from the models of the other programs. They study unsatisfiable logic programs, and their strategy to solve the unsatisfiability leaves the unsatisfiable program out of consideration for the merging process without trying to solve the arising conflicts as we propose. The second operator can be seen as an instance of the approach introduced in [30], and it selects those models of a program P_0 that are closest to the models of the programs to be merged P_1, \dots, P_n , where P_0 is a particular program that can be considered analogous to a set of integrity constraints. Certainly, integrating our approach to [14] would add several attractive characteristics to this work. A significant benefit of introducing our trust model does not only favor to resolve conflicts using accrual functions but also the use of different types of information associated with the sources that will be merged, *e.g.*, the account of sources supporting pieces of information, or credibility assigned to such sources as we proposed.

A closely related work is the one reported in [17], where the authors present a proposal of new merging operators that can be defined from existing ones, and also study under which conditions IC merging operators can be composed in such a way that the resulting operation preserves rationality postulates. In particular, since this work does not consider in the formalization of the proposal any mechanism to deal with problems related to the set of integrity constraints (ICs), they start with the approach proposed in [30] where sources of informa-

tion are assumed consistent. Unlike them in our work, we do not make the assumption that the set of ICs should be conflict-free when considering multi-agent scenarios. Therefore, a significant advantage of our approach is that it allows us to cover cases where it is possible to join two or more bases without worrying about the possible conflicts that may arise.

In [44], Schwind *et al.* have studied the merging process taking into account rationalized belief bases with respect to integrity constraints. More precisely, they consider the merging operators presented in [30] extending them by the introduction of a class of rationalization-based merging operators where physically infeasible scenarios that are preferable to discard are considered, *i.e.*, “repairing” the input bases by discarding the models of them which do not satisfy the integrity constraints. Three types of rationalizations are considered for computing the merged result; merging operators are then applied to the rationalized sources of information to obtain the merged result. As in this research, we propose several methods to deal with the inconsistency problem. However, besides the difference in the research focus between that work and ours, it is important to remark that we also provide a complete inconsistency resolution process exploiting the credibility relation between agents in the community to define how the consistency should be restored. This process obtains the best-valued among the consistent bases by employing optimal incision functions, which represents one of the main benefits of applying our credibility approach.

In collaborative multi-agent scenarios where every agent plays a role of a source of information to others, the credibility assigned by each agent to its peers is an interesting aspect that has been studied in the literature [47, 46]. For instance, in [47], Tamargo *et al.* have introduced an approach to deal with the knowledge dynamics of a multi-agent system where the weigh of beliefs is established by using a credibility order among agents. This order is especially important when agents can obtain new information from other agents. Four change operators were introduced: expansion, contraction, prioritized revision, and nonprioritized revision. A belief is then revised when new contradictory incoming information arrives from a highly credible informant. Likewise, in our proposal, we introduce a credibility relation between agents; however, here, we intend to use this relation to associate a (computed) credibility degree to each piece of information.

Following the intuitions of [47], the authors in [46] present a formalism of multi-source multiple belief revision where the credibility attached to informant agents is considered in the revision process. Given that the trust the agent assigns to each source may differ, every agent establishes its own credibility partial order among its informants. In this work, a change operator considering information sources to decide which information should be accepted when performing a belief revision is presented. Indeed, using a partial order defined over the set of sources is a more general approach than ours since it gives the ability to represent the case where some sources are incomparable to others. However, the main difference with our work is that we use the credibility relation between agents to deal with the problem of conflicting information arising when merging knowledge bases. A significant advantage of using accrual functions is that they

provide us a useful computational device to define different conflict resolution strategies.

Finally, Schwind *et al.*'s work [43] is more closely related to our approach, where a formalization for modeling the belief dynamics of a group of agents called Belief Revision Games (BRGs) is introduced. In this approach, each agent iteratively updates its beliefs by applying a revision policy-based belief merging operators that take into account its current belief states and the belief states of its neighbors. Though the idea of interacting agents is similar to our approach, we do not focus on belief dynamics, *i.e.*, we do not study how the beliefs of a group of agents evolve depending on how agents share their beliefs. An attractive characteristic of our proposal is that the framework proposed in [43] could be extended incorporating our credibility model based on accrual functions.

6.2. Social Choice

Social Choice Theory is a scientific inquiry in an interdisciplinary area that studies the aggregation of individual preferences towards a collective choice [8]. In particular, the design and theoretical evaluation of voting rules and questions originally stemming from social choice have motivated interesting developments in different contexts [9].

Although coming from a different area, the notion of belief accrual shares a lot with methods for collective decision making studied in Social Choice Theory [8]. Computational social choice is concerned with the application of mechanisms developed in Computer Science for solving questions stemming from social choice. For instance, to determine how preferences of individual agents can be aggregated into collective preference relations is a question that is utterly important in many research fields; in particular, in multi-agent systems, it allows us to study how a community of agents can be treated as a single rational decision-maker. Preferences are not only a characteristic that members of an agent society may want to aggregate, but different methods of the social choice theory have also been applied to other types of information, such as private members' beliefs that, when considered together, could be inconsistent.

Although social choice theory studies and belief merging operators belong to distinct areas, there are several similarities between them, as discussed in [8]. Several merging operators proposed in the literature were inspired by voting methods studied in social choice theory. For instance, in [20], the authors study the voting problem under the perspective of belief revision and belief merging. They introduce a form of encoding a voting scenario in a belief revision-based logical framework. To resolve the conciliation problem of contradictory voting interests, they introduce a model-based belief merging operator. Other approaches have been proposed to use a merging operator in order to define judgment aggregation methods (*e.g.*, see [16]). The aggregation of individual judgments regarding the truth (or falsehood) of a given set of related statements is called judgment aggregation. In [16], the notion of distance is showed to relate judgment aggregation, belief merging, and social choice theory. The authors maintain that the use of a distance notion establishes a link between

belief merging and social choice theory, and show how a distance-based merging operator can be applied to deal with classic social choice problems, providing a link between distance-based approaches and social choice theory. Several important works in the literature have focused on dealing typical social choice problems employing a distance-based merging operator [31]. Links between the areas of belief merging and social choice have also been studied in [16].

As we shall see, several aspects of our approach provide benefits and useful advantages in comparison with other formalisms. For instance, accrual function-based merging operators can be related to those works mentioned; however, in there, they are used with a different purpose. It is interesting to note that these operators aggregate knowledge coming from equally reliable sources, in our approach the sources of a piece of information have different reliability, and this characteristic is exploited to define a measure of the strength of the piece of information itself. Here, this additional information plays an essential role in our formalization, allowing us to define different conflict resolution strategies when all information sources are considered at once during the merging process. An important point to note is the possibility of dealing with the problem addressed in this paper as a preference aggregation problem importing positive results from social choice theory. For this, two different types of approaches could be adopted. On the one hand, a qualitative approach could be used, where the different degrees of credibility assigned by each agent is mapped into a preference relation, and a collective preference relation is obtained from these individual relations. From this resulting relation, it is possible to identify the most desirable formulas. On the other hand, several social choice-based qualitative mechanisms could also be applied in our proposal, such as scoring rules [15] or weighted goals [32] among others. Despite being similar in spirit to aggregation functions traditionally studied in social choice theory, a significant advantage of introducing accrual functions for merging knowledge bases is to provide a natural way to identify agents, which are the sources of conflicting information. This characteristic does not only favor the use of features attached to that information but also the use of other types of information associated with the sources. Besides the credibility that each agent assigns to each other in our multi-agent scenario, the merging process could use, for instance, the availability, the satisfaction, or the number of informant agents in the community. Our research focus has been set in taking advantage of the accrual functions to better exploit the additionally available information for scenarios with multiple informants. Particularly, we propose to integrate different knowledge bases focusing on enforcing consistency, applying different credibility criteria.

7. Conclusions and Future Work

Over the years, a handful of important contributions have been advanced to merge information coming from different knowledge bases [35, 30, 31, 19, 44, 45, 27]. In these works, it is customary to consider just the formulas in the knowledge bases as all the information that is available for such a process; nevertheless, some of them also consider additional information that represents

the preference relation among formulas in the knowledge base. As pointed out in Section 1, having computational tools to define how much particular pieces of information are desirable based on trust models, and using such desirability is a useful advantage in collaborative multi-agent scenarios since it effectively increases the capacity of these agents to establish which formulas need to be removed to resolve conflicts in an agreed, consensual manner.

In this research, we have addressed the issue of merging in a single knowledge base the distributed knowledge of a community of agents. In this scenario, every agent has its knowledge (*i.e.*, its epistemic state) represented as a Datalog[±] ontology; then, a final ontology is assembled representing (as much as it is feasible) the knowledge of the community as a whole. The main contribution of this work is a formalization that characterizes a framework where every member of the community assigns a credibility value to its peers, and these values are used to compute the best-valued knowledge during the merging process. To measure the credibility of a particular set of formulas, we exploited the notion of accrual functions, tailoring them to our particular multi-agent credibility scenario defining in this manner a general credibility accrual functions that use the general contexts introduced in [13] specializing them to conflicts in the community. In particular, the formalizations introduced in Section 3, and theoretical results in Section 5, are two of the main strengths of this proposal.

Credibility represents a crucial part of our framework. An essential advantage of introducing credibility accrual functions to multi-agent settings is to provide more flexibility to define different conflict information strategies when all agents are considered at once, obtaining the most significant benefit. As we have explained before, different trust models that could be adopted have been proposed in the literature; for instance, one possibility is to consider an approach based on the number of removed formulas. Nevertheless, and as we have shown in Section 3, the optimality in our proposal is inherently connected to the credibility of agents. In this sense, our proposed accrual functions contribute to maximizing credibility instead of minimizing the number of formulas removed.

Regarding the merging operation *per se*, our approach adapts several definitions introduced in [12] to our setting in order to define a merging operator capable of returning a knowledge base that represents as much as possible the original ones, focusing on resolving inconsistency. This consistency solving is based on the formalization of credibility based incision functions, which handle inconsistencies through grouping related conflicts through clusters. The main idea behind incision functions presented here is to resolve conflicts using a credibility accrual function to determine which formulas can be deleted for every cluster considered. We have characterized our merging operator through two properties, showing that the merged knowledge base is always consistent and optimal (regarding the particular accrual functions being considered by the operator). As another formal result, we compared our approach with Falappa *et al.*'s Kernel Merge approach [19].

The introduction of a general mechanism for calculating the credibility value of a set of formulas allowed us to explore different particularizations of our general credibility accrual function. In Section 4, we presented particular strategies

to measure the credibility attached to a set of formulas. We also showed that these functions satisfy the properties of the general credibility accrual functions: Transitivity, Domination, and Pertinence. These properties are of particular interest in our formalism because they are directly involved in the optimality of the conflict resolution (see [13]). Although an analysis of the computational complexity of the approach is out of the scope of this work, we would like to point out that a naive approach to conflict resolution based on a comparison of every subset of a cluster is indeed computationally expensive. Nevertheless, it is worth noting that we can exploit particular semantics to reduce the search space, greatly lowering the complexity.

As for future work, there are several research lines that we are currently pursuing. On the other hand, we are also interested in studying new developments to tackle the limitations that we discuss in the following paragraphs.

Several works from the literature of Belief Revision and Merging (*e.g.*, [40]) have adopted for their proposals an approach based on additional information concerning the formulas in the knowledge base, such as a preference relation defined by domain experts that can be used to decide over conflicts. In this direction, we want to explore constructions based on exploiting preference relations among the formulas and credibility degrees assigned by every agent to other agents in order to define different strategies to choose which formulas to delete. Mainly, our future goal is to study two different situations: the relation between merging operators based on preference relations with respect to the one presented in this work, and how different semantics defined here affect their behavior.

As we have mentioned, our proposal corresponds to a multi-agent systems scenario; however, here, we have focused on exploring the theoretical aspects of the presented approach. Thus, computational aspects such as time efficiency optimization are of great importance for our long-term goal, but they fall out of this work scope since we have focused on exploring the presented approach's theoretical aspects. A line of work that we are already working on is the definition of different particular accrual strategies to measure the credibility attached to a set of formulas, such as the ones introduced in Section 4. In particular, as explained before, we aim to explore the computational aspects of the strategies presented here and to investigate those that will enable a computationally efficient conflict resolution since they can be calculated over a subset of the space of possible solutions. While we expect that the full, general credibility approach to be rather costly regarding execution time since it needs to look into every subset of a cluster, we also expect that some particularizations may prove to be more easily computed because there are significantly fewer possible solutions. An example of such a function would be absorption credibility accrual functions such as the one introduced in this work; that is, those that evaluate sets of formulas as equal to the most/least credible one in them. It is clear that in such a case, we only need to check on singleton subsets of the solution space, which could be done in linear time.

In [13], information about the application domain related to formulas is modeled through values that account for interesting features. In the formalism

we have introduced it is possible to define a feature to model those aspects that are to be considered in a merging process that can be seen as a combination of two or more features; however, our framework does not provide a systematic way to deal with information conflicts handling more than one feature. For this reason, another line of work that we are particularly interested in is to extend the approach presented here to include multiple features representing different dimensions, besides credibility, which can be considered to choose how to deal with conflicts. In this context, it will be of interest in studying the impact of different kinds of features in the merging process.

Finally, we are currently in the process of developing an implementation of our merging operator to produce an efficient application using Datalog[±] ontologies. We intend to perform analysis of both the computational complexity aspects over the defined accrual functions in order to determine which one are the most suitable to develop, and also we will analyze the execution time under different conditions.

Acknowledgements

This work has been partially supported by EU H2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 690974 for the project MIREL: MIning and REasoning with Legal texts, and by funds provided by CONICET, Universidad Nacional del Sur (UNS), Argentina, by PGI-UNS (grants 24/ZN34), PICT-2016-0215 Tipo III-A - Raices, and Universidad Nacional de Entre Ríos (UNER), by PID-UNER (grants 7053).

Appendix A. Notations

The notation used throughout the paper is summarized in the following table.

Name	Notation	Meaning
Relational schema	\mathcal{R}	A relational schema
Sublanguage	$\mathcal{L}_{\mathcal{R}}$	the sublanguage generated by \mathcal{R}
Chase	$chase(D, \Sigma_T)$	the chase for a database D and a set of TGDs Σ_T
Models	$mods(D, \Sigma)$	the set of models of D and Σ
Credibility degree	(\mathcal{A}, c)	a credibility c associated to agent \mathcal{A}
Agent's credibility Vector	$\mathbb{V}_{\mathcal{A}}$	the credibilities that agent \mathcal{A} assigns to other agents in the community
Kernels	$\mathbb{K}_{\mathbb{A}}$	the set of kernels in the union of knowledge bases of all agents in community \mathbb{A}

Name	Notation	Meaning
Clusters	$\mathbb{III}_{\mathbb{A}}$	the set of clusters in the union of knowledge bases of all agents in community \mathbb{A}
Credibility accrual function	$\mathcal{V}_{[cred]}(\{A\})$	An accrual function that assigns a value to the set of formulas A
Cluster incision function	$\varrho(\mathbb{III}_{\mathbb{A}})$	An incision function selecting formulas from clusters in $\mathbb{III}_{\mathbb{A}}$ for deletion
ck-Merging Operator	$\Upsilon(\mathbb{A})$	a merge operator for the community \mathbb{A}
Source	$\mathcal{S}_{\mathbb{A}}(\alpha)$	the set of agents that has formula α in their KB
Credibility Set	$\widetilde{\mathcal{C}}(\mathcal{A}_i)$	the set of credibilities assigned to an agent
Credibility of a particular agent	$\widetilde{\mathcal{C}}$	the minimal credibility associated to an agent
Credibility of a set of agents	$\mathcal{C}(\mathbb{A})$	the credibility associated to a set of agents
Credibility of a formula	$\nu(\alpha)$	the credibility associated to α
Accumulation Accrual function	$\mathcal{V}^{\uparrow\uparrow}(A)$	accumulation-based accrual of credibilities
Maximality Accrual function	$\mathcal{V}^{\max}(A)$	maximality-based accrual of credibilities
Communities' equivalence	$\mathbb{A} \cong_{\nu} \mathbb{B}$	Two equivalent communities <i>w.r.t.</i> formula ordering

- [1] Alchourrón, C., Gärdenfors, P., Makinson, D., 1985. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic* 50 (2), 510–530.
- [2] Alchourrón, C., Makinson, D., 1981. Hierarchies of Regulation and Their Logic. *New Studies in Deontic Logic*, 125–148.
- [3] Anam, S., Kim, Y. S., Kang, B. H., Liu, Q., 2016. Adapting a knowledge-based schema matching system for ontology mapping. In: *Proceedings of the Australasian Computer Science Week Multiconference*. ACM, p. 27.
- [4] Ardjani, F., Bouchiha, D., Malki, M., 2015. Ontology-alignment techniques: survey and analysis. *International Journal of Modern Education and Computer Science* 7 (11), 67.
- [5] Arenas, M., Bertossi, L. E., Chomicki, J., 1999. Consistent query answers in inconsistent databases. In: *Proceedings of the Eighteenth ACM Symposium on Principles of Database Systems (PODS)*. pp. 68–79.

- [6] Bell, D. A., Qi, G., Liu, W., 2007. Approaches to inconsistency handling in description-logic based ontologies. In: OTM Workshops (2). pp. 1303–1311.
- [7] Black, E., Hunter, A., Pan, J. Z., 2009. An argument-based approach to using multiple ontologies. In: SUM. pp. 68–79.
- [8] Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A. D., 2016. Handbook of computational social choice. Cambridge University Press.
- [9] Brandt, F., Conitzer, V., Hemaspaandra, L. A., Laslier, J., Zwicker, W. S. (Eds.), 2010. Computational Foundations of Social Choice, 07.03. - 12.03.2010. Vol. 10101 of Dagstuhl Seminar Proceedings. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany.
- [10] Cali, A., Gottlob, G., Lukasiewicz, T., 2012. A general Datalog-based framework for tractable query answering over ontologies. *Journal of Web Semantic* 14, 57–83.
- [11] Cholvy, L., 1998. A general framework for reasoning about contradictory information and some of its applications. In: Prade, H. (Ed.), *ECAI Workshop Conflicts among Agents*. England.
- [12] Deagustini, C. A. D., Martinez, M. V., Falappa, M. A., Simari, G. R., 2016. Datalog+- ontology consolidation. *Journal of Artificial Intelligence Research* 56, 613–656.
- [13] Deagustini, C. A. D., Martinez, M. V., Falappa, M. A., Simari, G. R., 2019. Belief base contraction by belief accrual. *Artif. Intell.* 275, 78–103.
- [14] Delgrande, J. P., Schaub, T., Tompits, H., Woltran, S., 2009. Merging logic programs under answer set semantics. In: *ICLP*. pp. 160–174.
- [15] Dietrich, F., 2014. Scoring rules for judgment aggregation. *Social Choice and Welfare* 42 (4), 873–911.
- [16] Eckert, D., Pigozzi, G., 2005. Belief merging, judgment aggregation and some links with social choice theory. In: Delgrande, J. P., Lang, J., Rott, H., Tallon, J. (Eds.), *Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics*. Vol. 05321 of Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, pp. 1–14.
- [17] Everaere, P., Konieczny, S., Marquis, P., 2012. Compositional belief merging. In: Brewka, G., Eiter, T., McIlraith, S. A. (Eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the 13th International Conference, KR*. AAAI Press, Italy, pp. 603–607.
- [18] Falappa, M. A., García, A. J., Kern-Isberner, G., Simari, G. R., 2013. Stratified belief bases revision with argumentative inference. *Journal of Philosophical Logic* 42 (1), 161–193.
<https://doi.org/10.1007/s10992-011-9217-z>

- [19] Falappa, M. A., Kern-Isberner, G., Reis, M. D. L., Simari, G. R., 2012. Prioritized and non-prioritized multiple change on belief bases. *Journal of Philosophical Logic* 41 (1), 77–113.
- [20] Gabbay, D., Pigozzi, G., Rodrigues, O., 2007. Common foundations for belief revision, belief merging and voting. In: Bonanno, G., Delgrande, J., Lang, J., Rott, H. (Eds.), *Formal Models of Belief Change in Rational Agents*. No. 07351 in Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, pp. 1–16.
- [21] Gärdenfors, P., 1982. Rule for rational changes of belief. *Philosophical Essay Dedicated To Lennart Åqvist on his Fiftieth Birthday*, 88–101.
- [22] Grandison, T., Sloman, M., 2000. A survey of trust in internet applications. *IEEE Communications Surveys & Tutorials* 3 (4), 2–16.
- [23] Hansson, S. O., 1991. *Belief Base Dynamics*. Ph.D. thesis, Uppsala University, Department of Philosophy, Uppsala, Sweden.
- [24] Hansson, S. O., 1994. Kernel contraction. *Journal of Symbolic Logic* 59 (3), 845–859.
- [25] Hansson, S. O., 2001. *A Textbook of Belief Dynamics: Solutions to Exercises*. Kluwer Academic Publishers, Norwell, MA, USA.
- [26] Harper, W., 1975. Rational Belief Change, Popper Functions and Counterfactuals. *Synthese* 30, 221–262.
- [27] Hué, J., Papini, O., Würbel, E., 2009. Merging belief bases represented by logic programs. In: *ECSQARU*. pp. 371–382.
- [28] Hunter, A., Liu, W., 2009. Knowledge base stratification and merging based on degree of support. In: Sossai, C., Chemello, G. (Eds.), *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, 10th European Conference, ECSQARU. Proceedings. Vol. 5590 of *Lecture Notes in Computer Science*. Springer, Italy, pp. 383–395.
- [29] Kfir-Dahav, N. E., Tennenholtz, M., 1996. Multi-agent belief revision. In: *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge*, De Zeeuwse Stromen, The Netherlands, March 17-20 1996. pp. 175–194.
- [30] Konieczny, S., Pérez, R. P., 2002. Merging information under constraints: A logical framework. *Journal of Logic and Computation* 12 (5), 773–808.
- [31] Konieczny, S., Pérez, R. P., 2011. Logic based merging. *Journal of Philosophical Logic* 40 (2), 239–270.

- [32] Lafage, C., Lang, J., 2000. Logical representation of preferences for group decision making. In: Proceedings of the Seventh International Conference on Principles of Knowledge Representation and Reasoning. Morgan Kaufmann Publishers Inc., pp. 457–468.
- [33] Lenzerini, M., 2002. Data integration: A theoretical perspective. In: PODS. pp. 233–246.
- [34] Levi, I., 1977. Subjunctives, Dispositions and Chances. *Synthese* 34 (4), 423–455.
- [35] Lin, J., Mendelzon, A. O., 1999. Knowledge base merging by majority. In: Dynamic worlds. Springer, pp. 195–218.
- [36] Liu, W., Williams, M.-A., 2001. A framework for multi-agent belief revision. *Studia Logica* 67 (2), 291–312.
- [37] Lukasiewicz, T., Martinez, M. V., Simari, G. I., 2012. Inconsistency handling in datalog+/- ontologies. In: 20th European Conference on Artificial Intelligence (ECAI). pp. 558–563.
- [38] Martinez, M., Pugliese, A., Simari, G., Subrahmanian, V., Prade, H., 2007. How Dirty is Your Relational Database? An Axiomatic Approach. In: 9th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU), Tunisie. Springer, Berlin, pp. 103–114.
- [39] Meyer, T., Lee, K., Booth, R., 2005. Knowledge integration for description logics. In: COMMONSENSE. AAAI Press, pp. 645–650.
- [40] Meyer, T. A., Labuschagne, W. A., Heidema, J., 2000. Refined epistemic entrenchment. *Journal of Logic, Language and Information* 9 (2), 237–259.
- [41] Nisar, T., Yeung, M., 02 2018. Twitter as a tool for forecasting stock market movements: A short-window event study. *The Journal of Finance and Data Science* 4.
- [42] Sabater, J., Sierra, C., 2005. Review on computational trust and reputation models. *Artif. Intell. Rev.* 24 (1), 33–60.
- [43] Schwind, N., Inoue, K., Bourgne, G., Konieczny, S., Marquis, P., 2015. Belief revision games. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, USA. pp. 1590–1596.
- [44] Schwind, N., Konieczny, S., Marquis, P., 2018. Belief base rationalization for propositional merging. *J. Log. Comput.* 28 (7), 1601–1634.
- [45] Schwind, N., Marquis, P., 2018. On consensus in belief merging. In: Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence

- (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018. pp. 1949–1956.
- [46] Tamargo, L. H., Deagustini, C. A. D., García, A. J., Falappa, M. A., Simari, G. R., 2019. Multi-source multiple change on belief bases. *Int. J. Approx. Reasoning* 110, 145–163.
- [47] Tamargo, L. H., García, A. J., Falappa, M. A., Simari, G. R., 2012. Modeling knowledge dynamics in multi-agent systems based on informants. *Knowledge Eng. Review* 27 (1), 87–114.
- [48] Tamargo, L. H., García, A. J., Falappa, M. A., Simari, G. R., 2014. On the revision of informant credibility orders. *Artificial Intelligence* 212, 36–58.
- [49] Turner, H., 2003. Strong equivalence made easy: nested expressions and weight constraints. *Theory and Practice of Logic Programming* 3 (4-5), 609–622.
- [50] Verhodubs, O., Grundspenkis, J., 2013. Ontology merging in the context of a semantic web expert system. In: *International Conference on Knowledge Engineering and the Semantic Web*. Springer, pp. 191–201.