Automated Fine-grained Trust Assessment in Federated Knowledge Bases

Andreas Nolle[1], Melisachew Wudage Chekol[2], Christian Meilicke[2],
German Nemirovski[1], and Heiner Stuckenschmidt[2]

[1] Albstadt-Sigmaringen University, Germany
{nolle,nemirovskij}@hs-albsig.de
[2] Research Group Data and Web Science, University of Mannheim, Germany
{mel,christian,heiner}@informatik.uni-mannheim.de

Abstract. The federation of different data sources gained increasing attention due to the continuously growing amount of data. But the more data are available from heterogeneous sources, the higher the risk is of inconsistency. To tackle this challenge in federated knowledge bases we propose a fully automated approach for computing trust values at different levels of granularity. Gathering both the conflict graph and statistical evidence generated by inconsistency detection and resolution, we create a Markov network to facilitate the application of Gibbs sampling to compute a probability for each conflicting assertion. Based on which, trust values for each integrated data source and its respective signature elements are computed. We evaluate our approach on a large distributed dataset from the domain of library science.

1 Introduction

The permanent growing amount of data published in the Linked Open Data (LOD) cloud opens new challenges in data integration. Additionally the use of different schema makes the task of federating several data sources a difficult problem. The federation of various data sources implies typically the amalgamation of ambiguous and possibly conflicting information and often leads to inconsistencies. The resolution of conflicts in federated large scale knowledge bases (KBs) is studied in [17]. Their approach is based on the generation and evaluation of federated clash queries, which are known to be complete for inconsistency detection in *DL-Lite*$_\mathcal{A}$ KBs. They apply a majority voting scheme to determine a partial repair. This approach does not aim at finding a global optimal repair, but applies an efficient heuristic where each step in the algorithm corresponds to a reasonable decision.

However, resolving conflicts by removing (or ignoring) a subset of the given assertions may result in loss of information. An alternative approach is to determine the trustworthiness of *individual assertions*, data source specific *signature*

We refer the interested reader to an extended version of this paper available at http://www.researchgate.net/publication/318722371.

elements (concept, role, or attribute names) and *data sources* integrated in the federated KB. Grandison and Sloman [8] define "trust as the belief in the competence of an entity to act dependably, securely, and reliably within a specified context". In our work we are concerned with data sources and their competences to provide reliably information with respect to a given assertion or with respect to the set of all assertions that have the same predicate (signature element of the TBox) in common. In that sense our definition of trust also builds on the notion of context dependency, while we understand context as reference to a given fact or reference to a predicate.[3]

We use the statistical evidence gathered by calculating a repair, as prior knowledge for the calculation of trust values at different levels of granularity. In particular, we consider the conflict graph, generated by clashing assertions, as a Markov network that can be used to determine the probability for each conflicting assertion via Gibbs sampling. With the aid of these probabilities, specific trust values for signature elements and data sources can be computed to estimate the probabilities of non-conflicting assertions. Consequently, our approach requires neither a full trusted data source (as in [5]) nor any manual assignments (or user interactions) and relies solely on the identified conflicts. Unlike other approaches [7,12–14,20,22] that in principle rely on the determination of source reliability, we additionally compute individual trust measures on the assertion and signature level of each integrated data source. Our main contribution is a fully automated approach of fine-grained trust assessment and consequently the transformation of a conventional (federated) KB into a probabilistic one.

In Section 2 we briefly introduce some fundamental terms and definitions. After introducing the generation of a conflict graph and its repair in Section 3 we propose our approach for assessing fine-grained trust values in Section 4. Subsequently, we present and discuss results of our experiments in Section 5. Before concluding in Section 7, we discuss related work in Section 6.

2 Preliminaries

We briefly introduce the definition of federated *DL-Lite*$_{\mathcal{A}}$ KBs, basic notions related to inconsistency in description logic (DL) KBs, and Markov networks.

2.1 Federated *DL-Lite*$_{\mathcal{A}}$ Knowledge Bases

DL-Lite is a family of languages in which checking KB satisfiability can be done in PTime in the size of the TBox and query answering in AC^0 in the size of the ABox. We consider the subfamily *DL-Lite*$_{\mathcal{A}}$, which has been designed for efficiently dealing with huge amounts of extensional information (ABox). We refer the reader to [3, 18] for a detailed discussion of the syntax and semantics. In general, subsumption axioms in *DL-Lite*$_{\mathcal{A}}$ can be normalized, i.e., each axiom

[3] The measure of trust essentially indicates the probability of an assertion to be true. While the term trust is used more on the data source level, probability is more often used with respect to a specific assertion. We use both terms interchangeably.

comprise only one element on the left of the subsumption relation ($\sqsubseteq$) and one element on the right hand side.

The signature Σ (also known as alphabet or vocabulary) of a KB is a finite set of concept, role and attribute names. Furthermore, for ABox assertions of the form $A(x)$, $R(x, y)$ and $U(x, v)$ we refer to the concept, role and attribute names of assertions, i.e., $A, R, U \in \Sigma$, as *signature elements*. In the context of federated KBs, where each integrated data source uses different terminologies (signatures) that are linked by an intermediary (central) schema, we can define a federated *DL-Lite*$_{\mathcal{A}}$ KB as well as federated ABox assertions as follows.

Definition 1. *A federated DL-Lite*$_{\mathcal{A}}$ *knowledge base is a DL-Lite*$_{\mathcal{A}}$ *knowledge base* $\mathcal{K}$ *with* $\mathcal{K} = \langle \mathcal{T}_c \cup \bigcup_{i \in \mathbb{F}} \mathcal{T}_i, \bigcup_{i \in \mathbb{F}} \mathcal{A}_i \rangle$ *where* $\mathcal{T}_c$ *is a central TBox, each* $\mathcal{T}_i$ *is a TBox and* $\mathcal{A}_i$ *is an ABox in data source* i *and* $\mathbb{F}$ *is a set of indices that refers to the federated data sources. A federated ABox assertion is a pair* $\langle \alpha, i \rangle$ *where* α *denotes an ABox assertion stated in* $\mathcal{A}_i$*.*

For compact presentation we write $\mathcal{T}$ instead of $\mathcal{T}_c \cup \bigcup_{i \in \mathbb{F}} \mathcal{T}_i$ and $\mathcal{A}$ instead of $\bigcup_{i \in \mathbb{F}} \mathcal{A}_i$. Besides, without loss of generality, in the remainder of this paper we assume that there is only one central schema $\mathcal{T}$ which might be the union of some data source specific schema and an intermediary one comprising mappings between the data source specific vocabularies. Furthermore, we do not address integration problems related to incoherency, i.e., we assume that $\mathcal{T}$ is coherent. Note that there are other works that deal with debugging issues on the terminological level, e.g., [9].

2.2 Inconsistency in Description Logics

In DL, an interpretation $\mathcal{I}$ that satisfies all assertions in $\mathcal{T} \cup \mathcal{A}$ of KB $\mathcal{K}$ is called a *model*. The set of all models for $\mathcal{K}$ is denoted by *Mod*($\mathcal{K}$). $\mathcal{K}$ is called *satisfiable* or *consistent*, if *Mod*($\mathcal{K}$) $\neq \emptyset$ [2, 6]. Otherwise $\mathcal{K}$ is called *inconsistent*. $\mathcal{K} \models \phi$ denotes that $\mathcal{K}$ logically entails or satisfies a closed first-order logic sentence (formula) ϕ, provided that $\phi^{\mathcal{I}}$ is true for every $\mathcal{I} \in$ *Mod*($\mathcal{K}$). If a set F of closed sentences is entailed by $\mathcal{K}$, we can also write $\mathcal{K} \models F$ [19].

An *explanation* (or justification) for $\mathcal{K} \models \phi$ is a subset $\mathcal{K}'$ of $\mathcal{K}$ such that $\mathcal{K}' \models \phi$ while $\mathcal{K}'' \not\models \phi$ for all $\mathcal{K}'' \subset \mathcal{K}'$ [10]. Consequently, an explanation can be interpreted as a minimal reason that explains why ϕ follows from $\mathcal{K}$. Given an inconsistent KB $\mathcal{K}$, an explanation for the inconsistency is called a minimal inconsistent subset (*MIS*) and is denoted by the subset $\mathcal{K}'$ of $\mathcal{K}$ such that $\mathcal{K}'$ is inconsistent while $\mathcal{K}''$ is consistent for all $\mathcal{K}'' \subset \mathcal{K}'$. A subset $\mathcal{R} \subseteq \mathcal{K}$ is called a *repair* (or repair plan) of an inconsistent KB $\mathcal{K}$, if $\mathcal{K} \setminus \mathcal{R}$ is consistent.

Assuming that all of the terminological axioms are (semantically) *correct*, we are only interested in the subset of a MIS that comprises only ABox assertions. We refer to such a subset of a MIS as a MISA (minimal inconsistency preserving sub-ABox). Please notice that in *DL-Lite*$_{\mathcal{A}}$ each MISA comprise at most two ABox assertions due to the normalized form of subsumption axioms (see Section 2.1). As we will show in Section 4.2, the conflict graph obtained from the MISAs can be represented as a Markov network.

2.3 Markov Networks

Graphical models are used to compactly describe a complex distribution over a multi-dimensional space as a graph and provide a central framework to reason on uncertain information. A *Markov network* or *Markov random field* is a probabilistic model that represents the joint probability distribution over a set of random variables $X = (x_1, x_2, ..., x_n)$ as an undirected graph [11]. Each variable is represented by a node and a direct probabilistic interaction between two nodes is represented by an edge. For each clique D comprising the set of nodes X_D there exists a real-valued weight w_D and a feature f_D mapping a possible state $\mathbf{x}_D$ of that clique to a real value. A clique of a graph is a set of nodes which are fully connected. The joint distribution of a Markov network can be defined as a log-linear model of the form

$$p(X = \mathbf{x}) = \frac{1}{Z} \exp\Bigg(\sum_D w_D f_D(\mathbf{x}_D)\Bigg), \tag{1}$$

where $\mathbf{x}$ is a vector, comprising the state of the variables X and Z is a normalization constant, called partition function. The *Markov blanket* B_x of a variable (node) x is defined as the minimal set of variables (nodes) that renders x independent from the rest of the graph, which is simply all neighboring nodes of x. We consider binary discrete variables, hence, the *state* of a variable is its truth value, i.e, either 1 or 0. The conditional probability of a variable x when its Markov blanket B_x is in a state $\mathbf{b}_x$ is given by:

$$\begin{aligned} & p(x = \mathbf{x} | B_x = \mathbf{b}_x) \qquad (2) \\ & = \frac{\exp\Bigg(\sum_{f_x \in F_x} w_x f_x(x = \mathbf{x}, B_x = \mathbf{b}_x)\Bigg)}{\exp\Bigg(\sum_{f_x \in F_x} w_x f_x(x = 0, B_x = \mathbf{b}_x)\Bigg) + \exp\Bigg(\sum_{f_x \in F_x} w_x f_x(x = 1, B_x = \mathbf{b}_x)\Bigg)}, \end{aligned}$$

where $\mathbf{b}_x$ is a vector that denotes the state of the Markov blanket B_x of node x, F_x is the set of features in which x appears and the feature f_x is a real value of the state, given $\mathbf{x}$ and $\mathbf{b}_x$. In this paper we focus on binary features $f(\mathbf{x}) \in \{0, 1\}$. We will use formula (2) to compute the probabilities of conflicting assertions as shown in Section 4.2.

3 Conflict Graph and Repair Generation

To illustrate how MISAs are used to generate a conflict graph, we introduce an example that is used throughout the remainder of this paper. Let $\mathcal{T}$ be a central schema that comprises the following axioms.

$$\begin{array}{cc} \mathit{Book} \sqcup \mathit{Paper} \sqsubseteq \mathit{Publication} & \mathit{Paper} \sqsubseteq \neg \mathit{Book} \\ \mathit{Proceedings} \sqsubseteq \mathit{Book} & \mathit{Publication} \sqsubseteq \neg \mathit{SlideSet} \\ \exists \mathit{isPartOf} \sqsubseteq \mathit{Paper} & \exists \mathit{isPartOf}^- \sqsubseteq \mathit{Proceedings} \\ \exists \mathit{hasSlideSet} \sqsubseteq \mathit{Paper} & \exists \mathit{hasSlideSet}^- \sqsubseteq \mathit{SlideSet} \end{array}$$

And let $\mathcal{A}_1$, $\mathcal{A}_2$, and $\mathcal{A}_3$ denote three distributed data sources that contain assertions shown in the following table.

$\mathcal{A}_1$		$\mathcal{A}_2$		$\mathcal{A}_3$	
$\mathit{Paper}(\mathbf{I1})$	(α_1)	$\mathit{Paper}(\mathbf{I1})$	(β_1)	$\mathit{SlideSet}(\mathbf{I1})$	(γ_1)
$\mathit{isPartOf}(\mathbf{I1}, \mathbf{C1})$	(α_2)	$\mathit{Proceedings}(\mathbf{I1})$	(β_2)	$\mathit{SlideSet}(\mathbf{I2})$	(γ_2)
$\mathit{Paper}(\mathbf{I2})$	(α_3)	$\mathit{isPartOf}(\mathbf{C1}, \mathbf{I1})$	(β_3)	$\mathit{hasSlideSet}(\mathbf{I3}, \mathbf{I2})$	(γ_3)
$\mathit{Paper}(\mathbf{I4})$	(α_4)	$\mathit{isPartOf}(\mathbf{I4}, \mathbf{C2})$	(β_4)	$\mathit{SlideSet}(\mathbf{I4})$	(γ_4)
$\mathit{isPartOf}(\mathbf{C2}, \mathbf{I5})$	(α_5)	$\mathit{Proceedings}(\mathbf{C2})$	(β_5)	$\mathit{hasSlideSet}(\mathbf{C2}, \mathbf{I4})$	(γ_5)
$\mathit{isPartOf}(\mathbf{I6}, \mathbf{C3})$	(α_6)	$\mathit{isPartOf}(\mathbf{I6}, \mathbf{C3})$	(β_6)	$\mathit{Proceedings}(\mathbf{C3})$	(γ_6)
$\mathit{Paper}(\mathbf{I6})$	(α_7)	$\mathit{Proceedings}(\mathbf{C3})$	(β_7)	$\mathit{Proceedings}(\mathbf{C4})$	(γ_7)
$\mathit{Paper}(\mathbf{I7})$	(α_8)	$\mathit{Paper}(\mathbf{C4})$	(β_8)	$\mathit{hasSlideSet}(\mathbf{I6}, \mathbf{C4})$	(γ_8)

For example, the assertion that **I1** is a *Paper* (α_1 in $\mathcal{A}_1$) and the assertion that **I1** is a *SlideSet* (γ_1 in $\mathcal{A}_3$) are obviously in contradiction due to the axiom $\mathit{Paper} \sqsubseteq \neg \mathit{SlideSet}$ originated from the axiom $\mathit{Publication} \sqsubseteq \neg \mathit{SlideSet}$ in $\mathcal{T}$. In addition, as the assertion $\mathit{Paper}(\mathbf{I1})$ is also found in $\mathcal{A}_2$ (β_1), it is also contradictory to $\mathcal{A}_3$. Furthermore, we can entail this assertion in $\mathcal{A}_1$ from $\mathit{isPartOf}(\mathbf{I1}, \mathbf{C1})$ (α_2) and the axiom $\exists \mathit{isPartOf} \sqsubseteq \mathit{Paper}$ in $\mathcal{T}$.

Note that our example can easily be extended to the case where the integrated data sources use different terminologies that are linked by equivalence or subsumption axioms by an intermediary schema. Relying on a previous work [17], we can efficiently detect and resolve inconsistency in federated *DL-Lite*$_{\mathcal{A}}$ KBs. The complete set of conflicts respectively the corresponding MISAs is generated by so-called federated clash queries. Hence, for the above KB, the complete set $\mathcal{C}$ of identified conflicts (MISAs) is given by $\{\, \{\alpha_1, \beta_2\}, \{\alpha_1, \beta_3\}, \{\alpha_1, \gamma_1\}, \{\alpha_2, \beta_2\}, \{\alpha_2, \beta_3\}, \{\alpha_2, \gamma_1\}, \{\beta_1, \beta_2\}, \{\beta_1, \beta_3\}, \{\beta_1, \gamma_1\}, \{\beta_2, \gamma_1\}, \{\beta_3, \gamma_1\}, \{\alpha_3, \gamma_2\}, \{\alpha_3, \gamma_3\}, \{\alpha_4, \gamma_4\}, \{\alpha_4, \gamma_5\}, \{\alpha_5, \beta_4\}, \{\alpha_5, \beta_5\}, \{\alpha_5, \gamma_5\}, \{\beta_4, \gamma_4\}, \{\beta_4, \gamma_5\}, \{\beta_5, \gamma_5\}, \{\beta_8, \gamma_7\}, \{\beta_8, \gamma_8\}, \{\gamma_7, \gamma_8\}\, \}$. The corresponding conflict graph comprising four independent subgraphs is shown in Fig. 1. Each federated assertion is represented by a node and a contradiction between two assertions is represented by an edge.

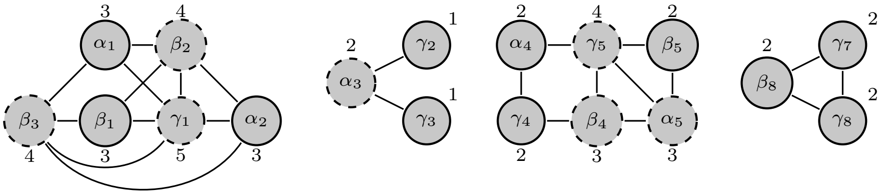

Fig. 1. Conflict Graph

Majority Voting Approach. Once all logical conflicts have been collected, the resolution of the identified contradictions is based on the assumption that the more data sources are integrated, the higher is the probability that correct assertions occur redundantly. Conversely, the probability that an assertion is incorrect correlates with the number of contradictions in which the assertion is involved. Based on this assumption, a majority voting scheme is applied on the assertion cardinalities, which are given by the number of involved MISAs for each assertion as illustrated in Fig. 1. MISAs comprising assertions with different cardinalities are iteratively resolved by adding the assertion with higher cardinality to the repair. Note that MISAs with minimum cardinality are resolved first, to reduce the impact (of wrong decisions) on subsequent decisions. Applying this heuristic to the conflict graph of our example will produce the repair $\{\beta_2, \beta_3, \gamma_1, \alpha_3, \alpha_5, \beta_4, \gamma_5\}$, depicted as dashed nodes.

Obviously, this heuristic may not resolve all logical conflicts, i.e., MISAs whose assertions having the same cardinalities (like β_8, γ_7 and γ_8). As a consequence, the heuristic generates a unique but not a full nor a global optimal repair. The application of this approach to a federated setting comprising four LOD data sources has shown that 39.5% of the detected conflicts could be solved with a precision up to 97% [17]. One possibility to get a full repair leading to a consistent KB could be for example to choose a random repair for all remaining contradictions. However, the resolution of conflicts implies the removal of all assertions that are part of the repair. To avoid loss of information we will now use the result of this approach to compute trust values respectively probabilities for individual assertions as well as for each data source and its individual signature elements.

4 Fine-grained Trust assessment

Since the evaluation of the approach for inconsistency resolution shows a high precision, we use the gathered statistical evidence as a basis for a fine-grained assessment of trust values at the level of assertions, signatures and data sources.

4.1 Signature Accuracy

We determine the *signature accuracy*[4] for each signature element with respect to a data source based on conflicting assertions and assertions that are 'correct'. Correct means in this case solely assertions whose individuals occur in a non-conflicting assertion in at least one other integrated data source. The set of conflicting assertions and correct assertions can be treated as an adequate sample of all assertions that use the same signature element. Furthermore, conflicting assertions can be defined into the following three subcategories:

- likely false assertions (assertions that are in the majority voting based repair),
- likely true assertions (conflicting assertions that would become conflict-free if the repair is removed),
- always conflicting assertions (assertions that are part of unresolvable MISAs).

[4] We intentionally avoid here the terms 'trust' or 'probability' to prevent any confusion with the calculated signature trusts later on.

Accordingly, we can define the accuracy for an signature element $\sigma \in \Sigma$ of a specific data source i formally as follows:

Definition 2. *Given a federated knowledge base $\mathcal{K} = \langle \mathcal{T}, \bigcup_{i \in \mathbb{F}} \mathcal{A}_i \rangle$, the set X of all conflicting assertions and the set $\mathcal{C} = \{\boldsymbol{c}_1, ..., \boldsymbol{c}_n\}$ of all conflicts (MISAs) in $\mathcal{K}$ with $\boldsymbol{c} := \{x_k, x_l\}$, $k \neq l$ and $x_k, x_l \in X$, a repair $\mathcal{R}$ computed by a majority voting approach, and the set $\mathcal{G}$ of non-conflicting assertions comprising individuals that occur in more than one data source (correct assertions). Let σ be either a concept, a property or an attribute in the signature Σ of $\mathcal{K}$, and let $\Psi \subseteq \bigcup_{i \in \mathbb{F}} \mathcal{A}_i$ be a set of federated assertions, then $sas(\sigma, \Psi, i)$ is defined as the subset of assertions in Ψ that use σ and originate from $\mathcal{A}_i$. The signature accuracy acc of σ with respect to $\mathcal{A}_i$ is defined as*

$$acc(\sigma, i) = 1 - \frac{|sas(\sigma, \mathcal{R}, i)| + \sum\limits_{x \in sas(\sigma, a \in \boldsymbol{c}: \boldsymbol{c} \in \mathcal{C}; \boldsymbol{c} \cap \mathcal{R} = \emptyset, i)} \frac{1}{|\boldsymbol{c} \in \mathcal{C} : x \in \boldsymbol{c}|}}{|sas(\sigma, X \cup \mathcal{G}, i)|}, \tag{3}$$

where x is an assertion in $\mathcal{A}_i$ that uses the signature element σ and is part of a MISA not resolved by $\mathcal{R}$. The accuracy of a signature is between 0 and 1, i.e., $0 < acc(\sigma, i) < 1$. Accuracy values that are outside of this range, i.e., for $acc(\sigma, i) = 0$ and $ar(\sigma, i) = 1$, the accuracy is set to a fixed value: 0.001 *and* 0.999 *respectively.*

Informally, the accuracy for a signature element of a specific data source is defined by '1− the ratio of incorrect assertions with respect to the total number of conflicting assertions and correct assertions'. The numerator in formula (3) is the number of incorrect assertions. It is given by the number of likely false assertions ($|sas(\sigma, \mathcal{R}, i)|$) and the probability of being true for each always conflicting assertion, which in turn is given by 1 divided by the number of contradicting assertions (number of involved conflicts).

Example 1. From the example of Section 3, the set of conflicting assertions comprises α_1, α_3 and α_4 with respect to the signature element *Paper* in data source $\mathcal{A}_1$, where only α_3 is part of the repair. On the other hand, α_7 is a correct assertion because it is verified by β_6 and not in conflict with any other assertion. Further, α_8 is neither a correct assertion nor part of any MISA. According to Definition 2 the accuracy for signature element *Paper* in data source $\mathcal{A}_1$ is given by $acc(Paper, 1) = 1 - \frac{1+0}{4} = 0.75$. The accuracy values for all signature elements are shown below.

$$\begin{array}{rrr}
acc(Paper, 1) = 0.75 & acc(Paper, 2) = 0.33 & acc(SlideSet, 3) = 0.67 \\
acc(isPartOf, 1) = 0.67 & acc(Proceedings, 2) = 0.67 & acc(hasSlideSet, 3) = 0.44 \\
 & acc(isPartOf, 2) = 0.67 & acc(Proceedings, 3) = 0.67
\end{array}$$

Based on the above definition, we can now use the calculated signature accuracy for a specific signature element with respect to a data source, to compute precise probabilities for conflicting assertions.

4.2 Assertion Trusts

We consider a conflict graph as a Markov network, where $X = \{x_1, \ldots, x_m\}$ represents the set of all federated assertions that are involved in some conflict and $\mathcal{C} = \{\boldsymbol{c}_1, ..., \boldsymbol{c}_n\}$ the set of all conflicts (edges) with $\boldsymbol{c} := \{x_k, x_l\}$, $k \neq l$. Since the edges in the conflict graph are undirected, we have chosen Markov Network as undirected graphical model. Each assertion x represents a binary random variable $x \in \{0, 1\}$, i.e., either true $(x = 1)$ or false $(x = 0)$. For each assertion x we have a feature $f_a \in F$ such that $f_a(x = 0) = 0$ and $f_a(x = 1) = 1$, i.e., $f_a(x) = x$. Moreover, in order to obtain a consistent *possible world* the condition $!(\bigwedge_{x \in \boldsymbol{c}} x)$ has to be satisfied for each conflict $\boldsymbol{c} \in \mathcal{C}$. A possible world is an assignment of truth values to all the variables. Consequently, each such condition is also treated as a feature $f_{\boldsymbol{c}} \in F$ such that the Markov network specifies a probability distribution over all possible worlds $\mathcal{X}$.

Since each condition $f_{\boldsymbol{c}}$ is a hard constraint that has to be satisfied in each possible world $\mathbf{x} \in \mathcal{X}$, the corresponding weight is $w_{\boldsymbol{c}} \to \infty$. If any constraint is violated, the joint distribution (given by equation (1)) is $\lim_{w \to \infty} p(X = \mathbf{x}) = 0$. Further, if at least one variable $b_x \in B_x$ in the Markov blanket of x is true, the conditional probability (given by equation (2)) of x is $\lim_{w \to \infty} p(x = 0 | B_x = \mathbf{b}_x) = 1$ and $\lim_{w \to \infty} p(x = 1 | B_x = \mathbf{b}_x) = 0$. This is because, the feature $f_x(x = 0, B_x = \mathbf{b}_x)$ (resp. $f_x(x = 1, B_x = \mathbf{b}_x)$) can only be true (resp. false) iff all its neighbors $b_x \in B_x$ are false, i.e., $b_x = 0$ (resp. true $b_x = 1$).

In order to compute the marginal probability of an assertion that uses signature element σ with respect to a data source i, we make use of the calculated signature accuracies. Hence we determine the weight w_a of a feature f_a for an assertion x in $\mathcal{A}_i$ as the log odds between a world in which an assertion x of $\mathcal{A}_i$ that uses σ is true and a world in which it is false, given by

$$w_a = \ln\left(\frac{acc(\sigma(x), i)}{1 - acc(\sigma(x), i)}\right), \tag{4}$$

where $\sigma(x)$ is the signature element of assertion x.

The complexity of computing the marginal probabilities is time exponential in the number of nodes. Thus, to perform approximate inference in Markov networks, Markov chain Monte Carlo (MCMC) particularly Gibbs sampling [11] is one of the most commonly used methods. In Gibbs sampling each node is sampled randomly in turn, given its Markov blanket using Equation (2). An approximation of the marginal probabilities, which is also called marginal inference, can be done by simply counting over the samples. Flipping the state of a node (e.g., changing its truth value from true to false) can be treated as a 'transition' between different worlds $\mathbf{x} \in \mathcal{X}$ (possible worlds of X). Because of the conditions $f_{\boldsymbol{c}} \in F$, a change of the state of an assertion x according to its conditional probability is only performed, iff all its neighbors $b_x \in B_x$ are false, i.e., $b_x = 0$ (denoted by $B_x = 0$ for short) and consequently the flip would not lead to an inconsistent world. Otherwise the state of an assertion remains unchanged. Given that $B_x = 0$, in Equation (2) all constraint features $f_{\boldsymbol{c}}$ in which x appears are zero $(f_{\boldsymbol{c}} = 0)$ and there remains one feature $f_a(x)$ whose value depends on the state of x. As we have already computed a possible repair based on a majority voting approach, we use it as a starting point for the Gibbs sampling. However, as there is no guarantee that all conflicts are resolved, a repair for all remaining contradictions is chosen randomly. Provided that we jump only between consistent possible worlds, there remain solely two cases for the conditional probability of node x, representing an assertion in $\mathcal{A}_i$ from which we have to sample:

1. if the current world contains $x = 0$ and $B_x = 0$, then the probability that x is true in the next possible world is given by:

$$p(x = 1|B_x = 0) = \frac{\exp\left(\ln\left(\frac{acc(\sigma(x), i)}{1 - acc(\sigma(x), i)}\right)\right)}{\exp(0) + \exp\left(\ln\left(\frac{acc(\sigma(x), i)}{1 - acc(\sigma(x), i)}\right)\right)} = acc(\sigma(x), i), \quad (5)$$

2. if the current world contains $x = 1$ and $B_x = 0$, the probability that x is false in the next possible world is given by:

$$p(x = 0|B_x = 0) = \frac{\exp(0)}{\exp(0) + \exp\left(\ln\left(\frac{acc(\sigma(x), i)}{1 - acc(\sigma(x), i)}\right)\right)} = 1 - acc(\sigma(x), i). \quad (6)$$

Consequently, the calculated accuracy of a signature element σ is exactly the conditional (prior) probability of an assertion $x \in \mathcal{A}_i$ comprising σ, given that all neighbors (contradicting assertions) are false. Since we start with a consistent world and ensure that an inconsistent world is never reached, the flipping of states causes that in some circumstances too many assertions are false (part of the repair), which is absolutely legitimate in terms of an acceptable repair. In terms of performance optimization, the sampling is applied to each independent subgraph of the conflict graph in parallel. After the sampling the approximate marginal probability (trust) of each assertion x can be calculated according to the following definition:

Definition 3. *Given a federated knowledge base $\mathcal{K} = \langle \mathcal{T}, \bigcup_{i \in \mathbb{F}} \mathcal{A}_i \rangle$, the set X of all conflicting assertions in $\mathcal{K}$, the set $\mathcal{G}$ of (non-conflicting) correct assertions and the set M containing, for each conflicting assertion $x \in X$, the number of Gibbs sampling states in which $x = 1$. Then, the assertion trust $p(x)$ for each federated assertion x in $\mathcal{A}$ of $\mathcal{K}$ is given by*

$$p(x) = \begin{cases} 1.0, & x \in \mathcal{G}, \\ \frac{M_{x=1}}{N}, & x \in X, \\ \varnothing, & \textit{otherwise}, \end{cases} \quad (7)$$

where $M_{x=1}$ is the number of states in which $x = 1$, N is the number of samples and $\varnothing$ denotes undefined.

Probabilities cannot be assessed for all assertions in $\mathcal{A}$ of $\mathcal{K}$, i.e., for those assertions that are not correct and are not involved in some MISAs. For such assertions, we determine (in Section 4.3) trust values for different signature elements with respect to a specific data source, called *signature trusts.*

Example 2. From the example in Section 3, $\alpha_6, \alpha_7, \beta_6, \beta_7$ and γ_6 are correct assertions and hence get a probability of 1.0. Using the accuracies of Example 1 and calculating assertion trusts using Gibbs sampling with $N = 10{,}000$ as described above, will result in the following assertion trusts:

$$\begin{array}{llllll}
p(\alpha_1) = 0.66 & p(\alpha_5) = 0.31 & p(\beta_1) = 0.58 & p(\beta_5) = 0.41 & p(\gamma_1) = 0.05 & p(\gamma_5) = 0.07 \\
p(\alpha_2) = 0.58 & p(\alpha_6) = 1.0 & p(\beta_2) = 0.07 & p(\beta_6) = 1.0 & p(\gamma_2) = 0.42 & p(\gamma_6) = 1.0 \\
p(\alpha_3) = 0.37 & p(\alpha_7) = 1.0 & p(\beta_3) = 0.03 & p(\beta_7) = 1.0 & p(\gamma_3) = 0.28 & p(\gamma_7) = 0.34 \\
p(\alpha_4) = 0.48 & p(\alpha_8) = \varnothing & p(\beta_4) = 0.15 & p(\beta_8) = 0.35 & p(\gamma_4) = 0.32 & p(\gamma_8) = 0.14
\end{array}$$

Only for assertion α_8 no probability is assessed ($\varnothing$), because it is not part of any conflict nor is correct.

4.3 Signature Trusts

Based on the previously computed probabilities of assertions, we can now define the trust for a signature element σ of a specific data source i as shown below.

Definition 4. *Given a federated knowledge base $\mathcal{K} = \langle \mathcal{T}, \bigcup_{i \in \mathbb{F}} \mathcal{A}_i \rangle$, the set X of all conflicting assertions in $\mathcal{K}$ and the set $\mathcal{G}$ of correct assertions in $\mathcal{K}$. Then, the signature trust $p(\sigma, i)$, for each signature element $\sigma \in \mathcal{A}_i$ of data source i in $\mathcal{K}$, is given by*

$$p(\sigma, i) = \begin{cases} \dfrac{\sum\limits_{a \in sas(\sigma, X \cup \mathcal{G}, i)} p(a)}{|sas(\sigma, X \cup \mathcal{G}, i)|}, & sas(\sigma, X, i) \neq \emptyset, \\ \varnothing, & otherwise. \end{cases} \tag{8}$$

Roughly, the signature trust with respect to a data source is defined by the average of all its assertion trusts. As a result we can now use the calculated signature trusts as the probability of assertions for which no trust value is assessed.

Example 3. In order to calculate the trust value of the signature element *Paper* in data source $\mathcal{A}_1$, we have to consider the probabilities of $\alpha_1, \alpha_3, \alpha_4$ and α_7. Following Definition 4, the signature trust of *Paper* is given by $p(\mathit{Paper}, 1) = \frac{0.66+0.37+0.48+1.0}{4} = 0.63$. Since for α_8 no assertion trust was computed, the signature trust of *Paper* in data source $\mathcal{A}_1$ is used as its probability. The calculated trusts for all signature elements with respect to the corresponding data sources are shown below:

$$\begin{array}{rrr}
p(\mathit{Paper}, 1) = 0.63 & p(\mathit{Paper}, 2) = 0.47 & p(\mathit{SlideSet}, 3) = 0.26 \\
p(\mathit{isPartOf}, 1) = 0.63 & p(\mathit{Proceedings}, 2) = 0.49 & p(\mathit{hasSlideSet}, 3) = 0.16 \\
 & p(\mathit{isPartOf}, 2) = 0.39 & p(\mathit{Proceedings}, 3) = 0.67
\end{array}$$

4.4 Data Source Trusts

Obviously, if there is no conflicting assertion that uses the signature element σ in a specific data source i, the signature trust value for σ with respect to data source i cannot be assessed. For this reason we in turn determine trust values for each data source in $\mathcal{K}$. Based on the definition of signature trusts, the trust value for a specific data source can be formally defined as:

Definition 5. *Given a federated knowledge base $\mathcal{K} = \langle \mathcal{T}, \bigcup_{i\in\mathbb{F}} \mathcal{A}_i \rangle$, the signature Σ of $\mathcal{K}$ and the complete set X of conflicting assertions in $\mathcal{K}$. Then, the trust value $p(i)$ for data source i in $\mathcal{K}$ is given by*

$$p(i) = \begin{cases} \dfrac{\sum\limits_{\sigma\in\Sigma:sas(\sigma,X,i)\neq\emptyset} p(\sigma,i) * |sas(\sigma,\mathcal{A}_i,i)|}{\sum\limits_{\sigma\in\Sigma:sas(\sigma,X,i)\neq\emptyset} |sas(\sigma,\mathcal{A}_i,i)|}, & \textit{if } \mathcal{A}_i \cap X \neq \emptyset, \\ \varnothing, & \textit{otherwise.} \end{cases} \tag{9}$$

Roughly, the trust in data source i is given by the average of the weighted sum of its signature trusts. Each signature trust is weighted by the number of assertions that uses the corresponding signature element in data source i. As there still might be some signature elements and consequently some assertions without an assessed probability, the trust value of the respective data source is used instead. Of course, if a data source contains no conflicting assertions the trust value for this data source cannot be computed. In this case a default or user-defined trust value could be used.

Example 4. With respect to the calculated signature trusts of Example 3 and using Definition 5, the data source trust for $\mathcal{A}_1$ is given by $p(\mathcal{A}_1) = \frac{0.63*5+0.63*3}{8} = 0.63$. The calculation of the data source trusts for $\mathcal{A}_2$ and $\mathcal{A}_3$ yields $p(\mathcal{A}_2) = 0.45$ and $p(\mathcal{A}_3) = 0.33$ respectively. If $\mathcal{A}_1$ would contain an additional assertion *SlideSet*(**I8**), the signature trust of *SlideSet* with respect to $\mathcal{A}_1$ and consequently the assertion trust for *SlideSet*(**I8**) would be the data source trust $p(\mathcal{A}_1) = 0.63$.

5 Experimental Evaluation

In order to evaluate our approach we have used a large distributed LOD dataset from the domain of library science, comprising four LOD data sources. Namely, FacetedDBLP ($\mathcal{A}_1$), BibSonomy ($\mathcal{A}_2$), RKB Explorer ePrints Open Archives ($\mathcal{A}_3$), and RKB Explorer DBLP ($\mathcal{A}_4$). Since the OWL 2 QL profile is based on *DL-Lite*, we have used it as specification language of our central TBox that includes the TBoxes of each data source. In order to ensure that the federated TBox is coherent and to gain a higher overlapping of the data sources, we have applied some small modifications of the data source specific TBoxes as well as its datasets (ABoxes). For more detail, we refer the interested reader to [17]. The collection of the central TBox as well as the referenced TBoxes is available online[5]. For legal reasons we are currently not able to publish the final dataset of each integrated data source. Please contact us if you are interested in these datasets. We run the implementation of our trust assessment approach on a CentOS 6.7 virtual machine consisting of 6x Intel Xeon CPUs (à 4 cores @ 2.50 GHz) and 128 GB of RAM.

[5] http://www.researchgate.net/publication/299852903

Accuracy and Trust Computation. The federated KB contains 284,355,894 assertions. The evaluation of 44,072 generated clash queries resulted in 18,146,950 MISAs[6]. The majority voting approach proposed in [17] could resolve 7,166,005 (39.5%) MISAs and generated a repair of 1,993,136 assertions. Note that the number of resolved MISAs is significantly higher than the size of the repair and indicates a high overlap of the MISAs. Based on this repair, the signature accuracy values are calculated using the formula in Definition 2. The distribution of the resulting values are depicted in Fig. 2(a). As shown in the figure, there exist one signature element with an accuracy < 0.1 with respect to $\mathcal{A}_1$ and $\mathcal{A}_3$. We had already observed that assertions involving the attribute *volume* are misused in $\mathcal{A}_1$ and $\mathcal{A}_3$, i.e., *volume* attributes are in both data sources not used at the level of collections like proceedings, journals or books, but on the level of articles published in a collection. Hence, it is not surprising that we get a low signature accuracy < 0.1 for *volume* with respect to $\mathcal{A}_1$ and $\mathcal{A}_3$.

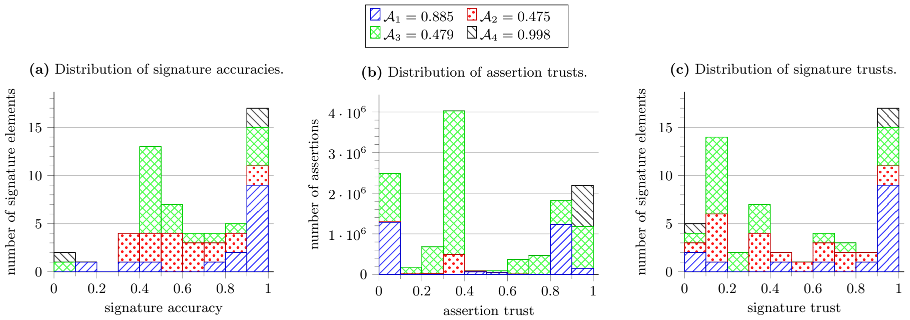

(a) Distribution of signature accuracies.

(b) Distribution of assertion trusts.

(c) Distribution of signature trusts.

Fig. 2. Signature accuracy and trust values of assertions and signature elements.

Since the generated repair resolves only 39.5% of the conflicts, we choose randomly a repair for all the remaining conflicts such that the starting point of the Gibbs sampling represents a possible world. The application of our approach for fine-grained trust assessment based on this repair and the calculated signature accuracy values result in the data source trusts, distributions of assertion trusts and signature trusts depicted in Fig. 2(b) and (c). In Fig. 2(c), if we consider data source $\mathcal{A}_4$, we see that it contains two signature trusts > 0.9 (for the signature elements *article−of−journal* and *Journal*); and one trust < 0.1 (for signature element *title*). Due to the negligible number of assertions in $\mathcal{A}_4$ with a low trust, the trust value of this data source is close to 1.0. Nevertheless, we cannot trust $\mathcal{A}_4$ with respect to the signature element *title*.

Runtime and Convergence Performance. The runtime, with increasing samples N (with a step size of 200) as well as the corresponding convergence of the trust values, is shown in Fig. 3. After a burn-in period of 1,000 samples in which the variables state may not exactly represent the desired distribution, the runtime increases linearly with the number of samples. After sampling each node 10,000

[6] Clashes of incorrect datatypes are not considered since its resolution is trivial.

times, the maximal deviation of a trust value compared to the previous sample is 0.019. Thus, the probabilities converge towards their true values as N increases.

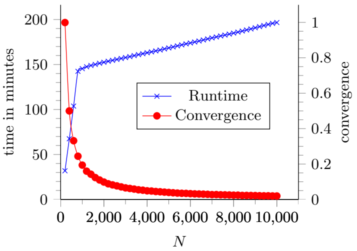

Fig. 3. Runtime and convergence.

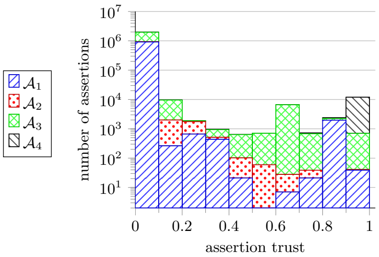

Fig. 4. Repair assertion trusts.

Comparison of Accuracy and Trust Measures. To give more insight into the generated trusts, we have done further analysis. We inspect in ascending order all conflicting assertions with low trust values. After removing all assertions with a trust value ≤ 0.5 (overall 7,465,415 assertions) all MISAs are resolved and the KB is consistent. Additionally, the distribution of assessed trusts for assertions that are part of the majority voting repair is depicted in Fig. 4. Notice that the y-axis is scaled logarithmically for presentation purposes. As shown in the figure, solely 22,497 (1.13%) assertions of the repair have trust values ≥ 0.5. This indicates that our approach performs very well (which is in line with the high precision (97%) of the repair shown in [17]).

Comparing the calculated signature accuracy values with the assessed signature trusts shows that the prior probabilities comprise 20 signature elements with a value ≤ 0.5 whereas the signature trusts have 30 elements. Table 1 shows the top 5 signature elements with a high deviation between the signature accuracy and the signature trust. For example, if we look at the signature element *Proceedings* of data source $\mathcal{A}_1$ and $\mathcal{A}_2$, it shows that most of the MISAs are not resolved by the majority voting. Moreover, the signature elements of conflicting assertions in resolved MISAs are in many cases different from the signature elements of conflicting assertions in unresolved MISAs. Since the accuracy for signature element *Proceedings* is less than the accuracy of conflicting signature elements in unresolved MISAs, the resulting trust for assertions that use *Proceedings* is low.

Table 1. Top 5 of Signature Elements with High Deviation

data source	$\sigma \in \Sigma$	**Signature Accuracy**	**Signature Trust**
$\mathcal{A}_1$	*http://swrc.ontoware.org/ontology♯Proceedings*	0.761	0.029
$\mathcal{A}_2$	*http://swrc.ontoware.org/ontology♯Proceedings*	0.578	0.178
$\mathcal{A}_3$	*http://purl.org/ontology/bibo/EditedBook*	0.591	0.224
$\mathcal{A}_3$	*http://purl.org/ontology/bibo/Book*	0.534	0.185
$\mathcal{A}_3$	*http://purl.org/ontology/bibo/Website*	0.500	0.163

Qualitative Analysis of Trust Values. To evaluate the quality of assessed assertion trusts, we randomly selected 100 assertions from the repair with a trust ≥ 0.8, representing a set of assertions that are probably mistaken for being part of the repair by the majority voting. Because of the already evaluated high precision of the repair, we omit the evaluation of assertions from the repair with a low trust value. The selected subset of assertions is manually evaluated by a domain expert. Since 81% of the assertions are annotated as correct, the evaluation indicates a high precision of the assessed probabilities and substantiate that the approach is reasonable. Besides, this precision score confirms that the calculation of signature accuracy values used as prior probability is a valid premise and enables a high precision of the assessed trust values.

6 Related Work

The notion of trust has been used in a heterogeneous way within the semantic web community (surveyed in [1]). The referred works are often based on the assumption that an external criteria is used to estimate trusts or that initial trust estimations are already given. Contrary to that, our work is based on the assumption that each data source and each assertion has the same level of trust prior to the majority voting. Moreover, our method is based on the idea that we have to readjust the initial assumption by analyzing and leveraging the logical contradictions of given assertions. Note also, that we could extend our approach, by starting with varying trust values based on an analysis of data provenance.

Beside addressing issues like the correction of mistakes that stem from the knowledge extraction process or with the aggregation of different values, dealing with contradictory assertions is one of the central tasks of knowledge fusion [4]. Given an inconsistent KB, one possible approach is to resolve all conflicts by eliminating at least one of the conflicting statements. However, conflict resolution often results in loss of information. Contrary to this, paraconsistent (inconsistency-tolerant) logics are used for reasoning in KBs that contain conflicts. To represent and reason on uncertain (i.e., imprecise) knowledge, there exist several approximation approaches. An overview of such approaches is for example given in [15].

In addition to paraconsistent logics, one straightforward approach is to eliminate conflicts by applying a majority voting scheme as shown in [17]. In order to consider the quality of different data sources, truth discovery techniques are proposed in [7,12,14,20,22]. A comprehensive survey on truth discovery is given by Li et al. [13]. The principle of truth discovery is to estimate the reliability of each source, i.e., the more frequently true information is provided, the higher is the trust in that source. Consequently, the information of a reliable source is considered as trustworthy. One shortcoming of (most of) these approaches is that they do not asses the quality of a data source with respect to some specific information or information type (signature element). As a consequence, all assertions of a data source have the same probability, yet assertions with respect to a specific signature element whose trust differ widely from the data source trust are neglected. So the trust calculation for an assertion on which the truth discovery is based upon is computed by means of the assumed data source trust (top-down), whereas in our approach the data source trust is determined by the signature trust and consequently by the individual assertion trusts (bottom-up). Another approach proposed by Ma et al. [16] considers the varying reliability of sources among different topics by automatically assigning topics to a question and estimating the topic-specific expertise of a source. Closer to our approach is the work proposed by Zhao et al. [23], since they calculate probabilistic values on facts (assertions) by using a Bayesian model and Gibbs sampling. Contrary to our approach, Zhao et al. base their notion of conflicting facts on direct contradictions that origin from a closed-world assumption instead of using a TBox that allows to find both explicit and implicit conflicts while still preserving the assumption that non-stated facts do not correspond to the claim of their negation.

In addition to the estimation of source reliability only by the accuracy of the provided information, there exist methodologies and frameworks for assessing data quality respectively its source by considering diverse quality dimensions and metrics, e.g., accessibility, performance, reputation, timeliness and others. Zaveri et al. [21] proposed a systematic review of such approaches that evaluate the quality of LOD sources and provide under a common classification scheme a comprehensive list of dimensions and metrics.

Our proposed approach is different from the approaches mentioned above in two aspects. First, we exploit the composition of the conflict graph, which is constructed based on a well-defined semantics, as well as the statistical evidence, gathered by inconsistency resolution, to compute individual probabilities for conflicting assertions. Second, the intention is not to use the computed probabilities for truth discovery but to enable the representation of uncertain knowledge and thereby the application of probabilistic reasoning and paraconsistent logics as well as the computation of the most probable consistent KB. To the best of our knowledge there is currently no other approach in this direction.

7 Conclusion

In this paper we proposed an automated approach for fine-grained trust assessment at different levels of granularity. In particular, by exploiting the statistical evidence generated by inconsistency resolution via majority voting and considering the conflict graph as a Markov network, we facilitate the application of Gibbs sampling to compute a probability for each conflicting assertion. Based on which, specific trust values for signature elements and data sources are computed to estimate the probabilities of non-conflicting assertions. We evaluated our approach on a large distributed dataset and could measure a high precision of the calculated probabilities.

Beside an evaluation against related truth discovery approaches, one further aspect will be to examine whether and to what extent it is possible to improve the calculated probabilities, by considering the entailment relation between several assertions (according to the given TBox) within the Gibbs sampling.

References

1. Artz, D., Gil, Y.: A survey of trust in computer science and the semantic web. Web Semantics: Science, Services and Agents on the WWW 5(2), 58–71 (2007)
2. Baader, F.: The description logic handbook: theory, implementation, and applications. Cambridge: Cambridge University Press (2003)
3. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Tractable reasoning and efficient query answering in description logics: The DL-Lite family. Journal of Automated Reasoning 39(3), 385–429 (2007)
4. Dong, X.L., Gabrilovich, E., Heitz, G., Horn, W., Murphy, K., Sun, S., Zhang, W.: From data fusion to knowledge fusion. PVLDB 7(10), 881–892 (2014)
5. Dong, X.L., Gabrilovich, E., Murphy, K., Dang, V., Horn, W., Lugaresi, C., Sun, S., Zhang, W.: Knowledge-based trust: Estimating the trustworthiness of web sources. PVLDB 8(9), 938–949 (2015)
6. Flouris, G., Huang, Z., Pan, J.Z., Plexousakis, D., Wache, H.: Inconsistencies, negations and changes in ontologies. AAAI 21(2), 1295–1300 (2006)
7. Galland, A., Abiteboul, S., Marian, A., Senellart, P.: Corroborating information from disagreeing views. In: WSDM. pp. 131–140. ACM (2010)
8. Grandison, T., Sloman, M.: A survey of trust in internet applications. IEEE Communications Surveys & Tutorials 3(4), 2–16 (2000)
9. Ji, Q., Haase, P., Qi, G., Hitzler, P., Stadtmüller, S.: Radonrepair and diagnosis in ontology networks. In: ESWC, pp. 863–867. Springer (2009)
10. Kalyanpur, A., Parsia, B., Horridge, M., Sirin, E.: Finding all justifications of OWL DL entailments. In: The Semantic Web. pp. 267–280. Springer (2007)
11. Koller, D., Friedman, N.: Probabilistic Graphical Models: Principles and Techniques. MIT Press (2009)
12. Li, X., Dong, X.L., Lyons, K., Meng, W., Srivastava, D.: Truth finding on the deep web: Is the problem solved? PVLDB 6(2), 97–108 (2012)
13. Li, Y., Gao, J., Meng, C., Li, Q., Su, L., Zhao, B., Fan, W., Han, J.: A survey on truth discovery. SIGKDD Explor. Newsl. 17(2), 1–16 (2016)
14. Liu, W., Liu, J., Duan, H., He, X., Wei, B.: Exploiting source-object network to resolve object conflicts in linked data. ESWC (2017)
15. Lukasiewicz, T., Straccia, U.: Managing uncertainty and vagueness in description logics for the semantic web. Journal of Web Semantics 6(4), 291–308 (2008)
16. Ma, F., Li, Y., Li, Q., Qiu, M., Gao, J., Zhi, S., Su, L., Zhao, B., Ji, H., Han, J.: Faitcrowd: Fine grained truth discovery for crowdsourced data aggregation. In: ACM SIGKDD. pp. 745–754. ACM (2015)
17. Nolle, A., Meilicke, C., Chekol, M.W., Nemirovski, G., Stuckenschmidt, H.: Schema-based debugging of federated data sources. ECAI pp. 381–389 (2016)
18. Poggi, A., Lembo, D., Calvanese, D., De Giacomo, G., Lenzerini, M., Rosati, R.: Linking data to ontologies. Journal on data semantics X pp. 133–173 (2008)
19. Rudolph, S.: Foundations of description logics. In: Reasoning Web. Semantic Technologies for the Web of Data, pp. 76–136. Springer (2011)
20. Yin, X., Han, J., Philip, S.Y.: Truth discovery with multiple conflicting information providers on the web. IEEE TKDE 20(6), 796–808 (2008)
21. Zaveri, A., Rula, A., Maurino, A., Pietrobon, R., Lehmann, J., Auer, S.: Quality assessment for linked data: A survey. Semantic Web 7(1), 63–93 (2016)
22. Zhao, B., Han, J.: A probabilistic model for estimating real-valued truth from conflicting sources. Proc. of QDB (2012)
23. Zhao, B., Rubinstein, B., Gemmell, J., Han, J.: A bayesian approach to discovering truth from conflicting sources for data integration. PVLDB 5(6), 550–561 (2012)