

An Evidential Path Logic for Multi-Relational Networks

Marko A. Rodriguez

Center for Nonlinear Studies
Los Alamos National Laboratory
Los Alamos, New Mexico 87545

Joe Geldart

Department of Computer Science
University of Durham
South Road, Durham, DH1 3LE

Abstract

Multi-relational networks are used extensively to structure knowledge. Perhaps the most popular instance, due to the widespread adoption of the Semantic Web, is the Resource Description Framework (RDF). One of the primary purposes of a knowledge network is to reason; that is, to alter the topology of the network according to an algorithm that uses the existing topological structure as its input. There exist many such reasoning algorithms. With respect to the Semantic Web, the bivalent, monotonic reasoners of the RDF Schema (RDFS) and the Web Ontology Language (OWL) are the most prevalent. However, nothing prevents other forms of reasoning from existing in the Semantic Web. This article presents a non-bivalent, non-monotonic, evidential logic and reasoner that is an algebraic ring over a multi-relational network equipped with two binary operations that can be composed to execute various forms of inference. Given its multi-relational grounding, it is possible to use the presented evidential framework as another method for structuring knowledge and reasoning in the Semantic Web. The benefits of this framework are that it works with arbitrary, partial, and contradictory knowledge while, at the same time, it supports a tractable approximate reasoning process.

Knowledge structures are used to represent facts about the world. The most common formal data structure to represent knowledge is the network. With respect to symbolic knowledge representation, the multi-relational network (also known as an edge labeled graph or semantic network) is widely used. A multi-relational network is composed of a set of vertices and a family of edge sets, where each edge set has a different nominal, or categorical, label. Formally, a multi-relational network can be represented as $M = (V, \mathbb{E})$, where V is the set of vertices and $\mathbb{E} = \{E_0, E_1, \dots, E_m \subseteq (V \times V)\}$ is the family of directed edge sets. In recent years, perhaps the most popular instance of a multi-relational data structure for knowledge representation is the Resource Description Framework (Dau 2006) of the Semantic Web initiative (Berners-Lee, Hendler, and Lassila 2001)¹. An edge

in an RDF network is called a statement, or triple, as it is composed of a subject, predicate, and object. For example, suppose the statement (i, k, j) . This statement denotes that i is related to j by means of an k -type relationship. Given the previous definition of M , this is equivalent to the directed edge $(i, j) \in E_k$. A particular instance of a statement is $(\text{marko}, \text{coauthor}, \text{joe})$. This statement denotes that Marko has a coauthorship relationship to Joe. Languages such as the RDF Schema (RDFS) and the Web Ontology Language (OWL) impose a set of constructs that serve to structure knowledge in a particular manner. The particularities of such a structure are used by an RDFS or OWL reasoner to infer new statements that can be added to the RDF network. The statements inferred by such reasoners are bivalent in that they are either true or false and moreover, their truth value is monotonic as it does not change once it has been asserted.

While RDFS and OWL are common languages in the Semantic Web, the flexibility of RDF can easily support other knowledge structures and reasoning algorithms. The purpose of this article is to present a non-bivalent, non-monotonic, evidential logic and reasoner for multi-relational networks that leverages many of the ideas from Non-Axiomatic Logic (NAL) (Wang 2006) and the Non-Axiomatic Reasoning System (NARS) (Wang 1993). The philosophical foundation of an evidential logic is that no statement is inherently true or false and that a statement only maintains levels of evidence to support or negate its claim. The notion of

experience-grounded semantics [is where] the truth value of a judgment indicates the degree to which the judgment is supported by the system's experience. Defined in this way, truth value is system-dependent and time-dependent. Different systems may have conflicting opinions, due to their different experiences. (Wang 1994)

The typical metaphor in an evidential logic system is that of an agent that perceives the world, represents its perceptions in an internal knowledge structure, and reasons on that structure to infer new knowledge (Wang 2004a). Moreover, it is assumed that this agent has limited computational resources in terms of both space and time and thus, does not maintain an objective knowledge structure nor does it necessarily have the ability to reason across its entire subjective knowledge structure. In other words, the agent has only so much information that it can store and process at any one

Copyright © 2009, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Other formal models of RDF include a bipartite graph (Hayes and Gutierrez 2004) and hypergraph (Morale and Serodio 2006) representation.

time. This notion is known as the Assumption of Insufficient Knowledge and Insufficient Resources (AIKIR). Non-axiomatic logic is contrasted to axiomatic logic, where truth is bivalent, is defined independent of the time and space requirements necessary to derive it, can be reasoned from a finite set of premises, and where all reasoning produces true, immutable conclusions.

The evidential logic presented in this article forms an algebraic ring over a multi-relational network (i.e. the knowledge structure) equipped with two binary operations (i.e. the atoms of the inferencing algorithms). Given the logic's multi-relational formulation, it is possible to comfortably represent this structure in RDF and thus, on the Semantic Web. The primary contribution of this article is the application of evidential logics to multi-relational networks and the formulation of an algebraic evidential reasoner.

Evidence in an Inheritance Network

With evidential logics, there does not exist an objective boolean truth value for every question that can be asked of the world as the world is not reasoned on directly (Wang 2004b). What is reasoned on is the agent's internal knowledge structure. For the agent, knowledge is gained as new evidence from the world is discovered (either through direct perception or through communication) or as knowledge is inferred given the agent's internal reasoning system. The Non-Axiomatic Reasoning System (NARS) is an example of an evidential reasoning system (Wang 1993). The data structure proposed for NARS version 2.2 is a directed evidence network denoted $G = (V, E, \lambda)$, where V is a set of vertices, $E \subseteq (V \times V)$ is a set of directed "inheritance" edges, and $\lambda : E \rightarrow \langle [0, 1], [0, 1] \rangle$ maps each edge to an evidence tuple. For example, an edge is denoted

$$(i, \langle w^+, w^- \rangle, j),$$

where i is the tail of the inheritance edge, j is the head of the inheritance edge, and $\langle w^+, w^- \rangle$ is the evidence tuple for that edge. Moreover, $\langle w^+, w^- \rangle \equiv \lambda(i, j)$. The meaning of $(i, \langle w^+, w^- \rangle, j)$ is that there is w^+ positive evidence supporting the claim that $i \text{ isA } j$ and w^- negative evidence that $i \text{ isA } j$. Another interpretation of this edge is that, upon examination, it appears, to the agent, that i has w^+ properties in common with j and w^- properties not in common with j . This idea is also expressed as i being the intension of j and j being the extension of i (Wang 1994).

In NARS, the evidence tuple of an edge is revised by means of external perception or internal reasoning. The act of perceiving i and j in the external world will return evidence of their relationship and thus, increment w^+ or w^- accordingly. With respect to internal reasoning, there are four syllogisms that can be applied to the agent's internal knowledge network: deduction, induction, abduction (Patzig 1968), and exemplification. For deduction,

$$(i, \langle w_1^+, w_1^- \rangle, j), (j, \langle w_2^+, w_2^- \rangle, k) \rightarrow (i, \langle w_3^+, w_3^- \rangle, k).$$

For induction,

$$(i, \langle w_1^+, w_1^- \rangle, j), (i, \langle w_2^+, w_2^- \rangle, k) \rightarrow (j, \langle w_3^+, w_3^- \rangle, k).$$

For abduction,

$$(i, \langle w_1^+, w_1^- \rangle, j), (k, \langle w_2^+, w_2^- \rangle, j) \rightarrow (i, \langle w_3^+, w_3^- \rangle, k).$$

Finally, there is a less widely used fourth syllogism known as exemplification (Bochenski 1970; Wang 1994). Exemplification is defined as

$$(i, \langle w_1^+, w_1^- \rangle, j), (j, \langle w_2^+, w_2^- \rangle, k) \rightarrow (k, \langle w_3^+, w_3^- \rangle, i).$$

The values for $\langle w_3^+, w_3^- \rangle$ depend upon the specific inference rules of the the evidential reasoner. In (Wang 2006), it is explicitly stated that the rules presented are not set in stone, but rather subject to revision themselves as more is understood about the design of evidential systems.

The evidence tuple $\langle w^+, w^- \rangle$ of an edge can be transformed into a normalized "truth value" consisting of a new tuple $\langle f, c \rangle \in \langle \emptyset \cup [0, 1], \emptyset \cup [0, 1] \rangle$. The first component f is the frequency of positive evidence and is defined as

$$f = \frac{w^+}{w^+ + w^-}.$$

The second component c is the confidence in the stability of f as k -more observations are made and is defined as

$$c = \frac{w^+ + w^-}{(w^+ + w^- + k)}.$$

The parameter $k \in \mathbb{R}_0^+$ is a user-defined, system constant. In words, as more positive evidence accumulates relative to negative evidence, f increases towards 1. As more total evidence accumulates relative to some constant k , c increases towards 1. If there is no evidence, then the edge has an f -component of \emptyset which means that the relationship is unknown. Finally, hard "truth" can be modeled with an evidence tuple of $\langle 1, 0 \rangle$ with $k = 0$ and thus, an f -tuple of $\langle 1, 1 \rangle$.

The contribution of this article is to extend the aforementioned evidential logic framework to multi-relational networks composed of both inheritance and non-inheritance edges. Moreover, this article contributes an algebraic ring formulation of evidential reasoning which situates the reasoner within well understood mathematics. From this multi-relational foundation, evidential reasoning can be comfortably executed in the RDF-rich world of the Semantic Web.

Evidential Reasoning using Path Expressions

A path algebra to map a multi-relational network to a single-relational form was originally presented in (Rodriguez and Shinavier 2008). The motivation behind the algebra was to provide a formal means by which the large class of single-relational network analysis algorithms could be applied to multi-relational networks in a meaningful way. With respect to this article, the binary operations of $+$ and \cdot are updated so as to work with evidence tuples.

A multi-relational, evidence network is defined as $M = (V, \mathbb{E} = \{E_0, E_1, \dots, E_m \subseteq (V \times V)\}, \lambda)$, where V is a set of vertices, \mathbb{E} is a family of edge sets, and $\lambda : E_k \rightarrow \langle [0, 1], [0, 1] \rangle$ maps each edge in $E_k : k \leq m$ to an evidence tuple. The algebraic formulation of the presented evidential

path algebra operates on an $n \times n \times m$ tensor representation of this network (Kolda, Bader, and Kenny 2005)². The evidence tensor \mathcal{A} is defined as

$$\mathcal{A}_{i,j}^k = \begin{cases} \lambda(i, j) & \text{if } (i, j) \in E_k \\ \langle 0, 0 \rangle & \text{otherwise.} \end{cases}$$

The two n dimensions represent the vertices and the single m dimension represents the various edge labels. Thus, there is a one-to-one mapping between a multi-relational, evidence network and an evidence tensor. The entries of the tensor denote the amount of positive (w^+) and negative (w^-) evidence for the edge $(i, j) \in E_k$, where $\langle 0, 0 \rangle$ denotes no evidence when no such edge exists. Inference on this tensor can be accomplished through the two binary operations

$$+ : \langle [0, 1], [0, 1] \rangle \times \langle [0, 1], [0, 1] \rangle \rightarrow \langle [0, 1], [0, 1] \rangle$$

and

$$\cdot : \langle [0, 1], [0, 1] \rangle \times \langle [0, 1], [0, 1] \rangle \rightarrow \langle [0, 1], [0, 1] \rangle.$$

The function rules of these operations are

$$\langle w_1^+, w_1^- \rangle + \langle w_2^+, w_2^- \rangle = \langle (w_1^+ + w_2^+), (w_1^- + w_2^-) \rangle$$

and

$$\langle w_1^+, w_1^- \rangle \cdot \langle w_2^+, w_2^- \rangle = \langle (w_1^+ \cdot w_2^+), (w_1^+ \cdot w_2^- + w_1^- \cdot w_2^+ + w_1^- \cdot w_2^-) \rangle.$$

These two operations form an algebraic ring over the evidence tensor (i.e. multi-relational, evidence network). The binary operation $+$ is associative, has an identity of $\langle 0, 0 \rangle$, and is commutative. The binary operation \cdot is associative with an identity of $\langle 1, 0 \rangle$. The operation $+$ supports the notion that evidence (from two independent experiences/inferences) can be summed together (Wang 1994). The operation \cdot supports the notion that positive evidence can be multiplied to form new positive evidence and that conflicting and negative evidence accounts for negative evidence.

The next subsections will formalize the various syllogisms on inheritance edges in a multi-relational network and then the following subsection will discuss the application of these operations to any arbitrary path through an evidence tensor. Note that what is presented is a set of operations that will operate on $n \times n$ matrix “slices” of the evidence tensor. In this respect, the presented operations are “global” computations and thus, are computationally inefficient and contrary to AIKIR. However, these operations can be implemented as “local” computations using various methods such as grammar walks (Rodriguez 2008) in specific areas of the network because evidential reasoning does not require a network-wide computation.

Inferring Inheritance Evidence

Inheritance relations are handled as a special case situation when reasoning in a multi-relational network. Example inheritance predicates in an RDF network are `rdf:type` and `rdfs:subClassOf`. Figure 1 presents a simple example network that will be used to demonstrate the inference rules of deduction, induction, abduction, and exemplification.

²The tensor representation can also be thought of as a set of m adjacency matrix “slices”, where each matrix has its own label.




Figure 1: An an inheritance network.

Deductive Inheritance Deduction is defined as a two step “walk” on the inheritance component of an evidential network. A two step walk can be computed by squaring a matrix. Suppose a standard square $\{0, 1\}$ -matrix denoted \mathbf{A} , where $\mathbf{A}_{i,j} = 1$ if there is an edge between vertex i and j , and 0 otherwise. The 2nd power of this matrix, as defined by ordinary matrix multiplication, \mathbf{AA} , will yield a new matrix where entry $(\mathbf{AA})_{i,j}$ denotes the total number of paths of length 2 starting from vertex i and ending on vertex j (Chartrand 1977). With respect to an evidence tensor, determining the product of two adjacency matrix “slices”, will yield a new adjacency matrix where the entry (i, j) denotes the total amount of deductive evidence supporting (i, j) .

Evidential matrix multiplication is defined as ordinary matrix multiplication, but respective of the rules of \cdot and $+$. Thus,

$$(\mathcal{A}^k \mathcal{A}^{k'})_{i,j} = \sum_{l \in V} \mathcal{A}_{i,l}^k \cdot \mathcal{A}_{l,j}^{k'} : k, k' \leq m.$$

In the degenerate case where all positive evidence is 1 and all negative evidence is 0, such that

$$\mathcal{A}_{i,j}^k = \begin{cases} \langle 1, 0 \rangle & \text{if } (i, j) \in E_k \\ \langle 0, 0 \rangle & \text{otherwise,} \end{cases}$$

evidential matrix multiplication will set w^+ to the total number of paths from vertex i to vertex j and 0 to w^- . In this form, the evidential path algebra yields results that are isomorphic to the original formulation of the path algebra in (Rodriguez and Shinavier 2008).

In Figure 2, deduction, as defined by $\mathcal{A}^{isA} \mathcal{A}^{isA}$, infers four new edges. Note that the evidence tuples are not presented in order to preserve diagram clarity.




Figure 2: Deduction in an inheritance network.

Inductive Inheritance Induction is the process of generalization given instances. In order to compute induction in an inheritance region of an evidence tensor, it is important to take the *converse transpose* of an adjacency matrix

“slice”. The operation of taking a statement like “a scholar is a writer” and reversing it to derive the statement “a writer is a scholar” is known as taking the converse of the statement. With respect to determining the evidence for the converse of a statement, all positive evidence for a “scholar is a writer” is considered positive evidence that a “writer is a scholar”. However, all negative evidence for “a scholar is a writer” should not be considered negative evidence for a “writer is a scholar” as there is no evidence in the original statement for writers not being scholars. Such statement converses can be expressed using the evidential algebra. For standard matrices, the transpose of a matrix is defined as $(\mathbf{A}^\top)_{i,j} = \mathbf{A}_{j,i}$ and denotes reversing the direction of an edge (i.e. taking the converse of a statement). However, for evidential edges, the converse transpose of an evidential matrix is defined as

$$\hat{\mathcal{A}}_{i,j}^k = \langle \gamma^+(\mathcal{A}_{j,i}^k), 0 \rangle,$$

where $\gamma^+ : ([0, 1], [0, 1]) \rightarrow [0, 1]$ maps an evidence tuple to its first component (i.e. w^+ positive evidence). This operation ensures that the converse of an evidence tuple maintains no negative evidence.

In Figure 3, induction, as defined by $\mathcal{A}^{isA} \hat{\mathcal{A}}^{isA}$, infers two new isA edges between `journalist` and `scholar`. As stated previous, and to stress the point, the negative evidence for these two new evidence tuples is 0.




Figure 3: Induction in an inheritance network.

Abductive Inheritance Abduction is the reverse of induction. In Figure 4, abduction, as defined by $\hat{\mathcal{A}}^{isA} \mathcal{A}^{isA}$, infers two new isA edges between `person` and `author`. Similar to induction, a converse transpose will generate no negative evidence for these two new evidence tuples.




Figure 4: Abduction in an inheritance network.

Exemplative Inheritance Exemplary inheritance paths can be determined by the multiplication of two converse transpose matrices. In Figure 5, exemplification, as defined by $\hat{\mathcal{A}}^{isA} \hat{\mathcal{A}}^{isA}$, infers four new evidence tuples.

This subsection presented the syllogisms of deduction, induction, abduction, and exemplification and their use




Figure 5: Exemplification in an inheritance network.

in an inheritance region of an evidence tensor. Note that this region may account for more than a single m -dimension as many labels can have an similar meaning to `isA` (e.g. `similarTo`, `equivalentTo`, `implies`, etc.). The next section will discuss reasoning using arbitrary paths through a multi-relational evidence network and thus, for those paths that may not necessarily contain `isA` edges.

Inferring Non-Inheritance Evidence

A multi-relational network may be composed of various types of relationships. Figure 6 diagrams an example multi-relational network that will be referred to in the examples of this subsection³.




Figure 6: A multi-relational knowledge network.

The network in Figure 6 is composed of a reference to this article, the denoted authors of this article, and two citations from this article to other articles. In a bivalent logic, these statements are true because they exist. However, in scholarly publishing it is rare, nearly impossible, for two people to “equally” write an article together. While ideas are shared and drafts are written, read, and edited, the article’s final form is always a biased reflection of the approach of some authors over others. With respect to the statements diagrammed in Figure 6, what is the evidence for these statements? The following descriptions are provided to expose, for each edge presented above, how much supporting or detracting evidence there is for i ’s m -type relationship to j :

(marko, wrote, this_article):

- w^+ : notation, writing style, diagram style, american spelling
- w^- : logic, reasoning, citations, philosophy

(marko, wrote, path_article):

- w^+ : notation, writing style
- w^- : algebra, no diagrams

(joe, wrote, this_article):

- w^+ : logic, reasoning, citations, philosophy, rings

³There is nothing that prevents the network in Figure 1 to be merged with the network in Figure 6 (e.g. marko and joe are both scholars). However, for diagram clarity, this is not represented.

- w^- : notation, writing style, diagram style, american spelling
(`this_article`, `cites`, `path_article`):
- w^+ : (Rodriguez and Shinavier 2008) citation, rings
- w^- : graph notation, philosophy, only a single algebra citation
(`this_article`, `cites`, `nars_article`):
- w^+ : (Wang 1993) citation, evidential notation, syllogisms
- w^- : path expressions, semantic web, rdf, owl, rdfs

The presented positive and negative evidence “metadata” (e.g. writing style, american spelling, citation patterns, etc.) can be represented in a multi-relational network. From this multi-relational encoding, it is possible, through automated means, to infer new evidence or revise existing evidence in the network with prescribed path expressions. In other words, a region of the network can provide further supporting and/or detracting evidence for another region of the network. In order to demonstrate two examples of this, the multi-relational network in Figure 6 will be used in this section to

1. infer new independent evidence supporting the claims that `marko` and `joe` wrote `this_article` according to the notion that self-citations are positive evidence supporting authorship, and
2. infer new evidence that `marko` and `joe` have a `coauthorship` edge between them.

Self-Citation Paths Suppose that the article citation evidence in Figure 6 was experienced by the agent (e.g. a repository feed) and added to its existing internal network. At the point of insertion, it is possible to revise the evidence tuple for `wrote` based upon the idea that self-citations in an article are considered evidence for authorship of that article (Rodriguez, Bollen, and Van de Sompel 2009). Furthermore, in order to ensure independent evidence, assume that the current evidence for `wrote` in Figure 6 was not determined using self-citation information⁴. Given this scenario, the following inference rule will update all `wrote` evidence according to inferred self-citation evidence:

$$\mathcal{A}_{(t+1)}^{\text{wrote}} = \left(\left(c(\mathcal{A}_{(t)}^{\text{wrote}}) \mathcal{A}_{(t)}^{\text{cites}} \mathcal{A}_{(t)}^{\text{wrote}\top} \right) \circ \mathbf{I} \right) + \mathcal{A}_{(t)}^{\text{wrote}},$$

where $t \in \mathbb{N}_0^+$ is the current time step, $c(\mathcal{A}^{\text{wrote}})$ “clips” the evidence in $\mathcal{A}^{\text{wrote}}$, \circ is the entry-wise Hadamard multiplication operation⁵, and \mathbf{I} is the evidential identity matrix

$$\mathbf{I}_{i,j} = \begin{cases} \langle 1, 0 \rangle & \text{if } i = j \\ \langle 0, 0 \rangle & \text{otherwise.} \end{cases}$$

In words, the self-citation inference rule states that evidence for $\mathcal{A}^{\text{wrote}}$ can be modulated by the total evidence for the path that goes from an author, to their written articles, to the

⁴Refer to (Wang 1993) for the definition and importance of independent evidence in evidential logics.

⁵For review, Hadamard entry-wise multiplication is defined as

$$\mathbf{A} \circ \mathbf{B} = \begin{bmatrix} \mathbf{A}_{1,1} \cdot \mathbf{B}_{1,1} & \cdots & \mathbf{A}_{1,j} \cdot \mathbf{B}_{1,j} \\ \vdots & \ddots & \vdots \\ \mathbf{A}_{i,1} \cdot \mathbf{B}_{i,1} & \cdots & \mathbf{A}_{i,j} \cdot \mathbf{B}_{i,j} \end{bmatrix}.$$

articles that those articles cite, and then finally, to the authors of those cited articles. However, in order to ensure that those cited authors are the original author from the start of the path (i.e. self-citation), it is important to filter on the identity matrix \mathbf{I} . Hadamard entry-wise multiplication is used to apply a matrix filter to a path. Note that the transpose of an evidence matrix, not the converse transpose of the evidence matrix, is used when taking the converse of a non-inheritance statement. The reason for this is that the positive and negative evidence for the statement “marko wrote this article” is the same for “this article was written by marko”. Next, inferred evidence for `wrote` must be independent of the evidence used to calculate it. Thus, $\mathcal{A}^{\text{wrote}}$ is mapped to a $(\langle 1, 0 \rangle, \langle 0, 0 \rangle)$ -matrix using the clip c operation, where

$$c(\mathcal{A}^k)_{i,j} = \begin{cases} \langle 1, 0 \rangle & \text{if } \mathcal{A}_{i,j}^k \neq \langle 0, 0 \rangle \\ \langle 0, 0 \rangle & \text{otherwise.} \end{cases}$$

Finally, the total evidence for the self-citation path is summed with the current evidence for the `wrote` edge. Thus, the agent has used self-citations as further, revising evidence for `wrote`.

Coauthorship Paths Two people are considered coauthors if they have both written an article together. The evidence for coauthorship is determined by the total evidence across all their jointly written articles. In its algebraic form, the evidence for `coauthor` can be determined by

$$\mathcal{A}_{(t+1)}^{\text{coauthor}} = \left(\left(\mathcal{A}_{(t)}^{\text{wrote}} \mathcal{A}_{(t)}^{\text{wrote}\top} \right) \circ n(\mathbf{I}) \right) + \mathcal{A}_{(t)}^{\text{coauthor}},$$

where $n(\mathbf{I})$ “nots” the evidential identity matrix such that every $\langle 1, 0 \rangle$ is a $\langle 0, 0 \rangle$ and every $\langle 0, 0 \rangle$ is a $\langle 1, 0 \rangle$. The reason for the $n(\mathbf{I})$ filter is that to represent a `coauthor` path from a person to their authored papers and then to *other* authors of those papers, the path must exclude the original author as an author is not a coauthor of themselves. In other words, it must filter out the identity evidence matrix. The two inferred `coauthor` edges between `marko` and `joe` are diagrammed in Figure 7.




Figure 7: Coauthoring in an inheritance network.

The article (Rodriguez and Shinavier 2008) provides an in-depth review of different inferences that can be made with arbitrary paths, various filters, and how the theorems of the general path algebra can be applied to derive equivalent, yet more computationally efficient paths. The examples presented in (Rodriguez and Shinavier 2008) can be applied to an evidence tensor as long as the definitions of $+$ and \cdot , as defined in this article, are respected.

Conclusion

Reasoning with axiomatic logics is computationally expensive (Donini 2002; Fensel and van Harmelen 2007). With respect to the Semantic Web, and with the integration capabilities brought forth by the Linked Data initiative, such reasoning is intractable. The original assumption driving the development of NARS is the Assumption of Insufficient Knowledge and Insufficient Resources (AIKIR) (Wang 1993). Given space and time constraints, an agent cannot reason over the entire Semantic Web, and potentially, not even over its entire internal knowledge network. The benefit of the inheritance-based syllogisms are that they do not require a global analysis of the knowledge network, can be executed independently of each other, and at their core, are very simple and computationally efficient. With respect to inferencing with arbitrary path expressions, the efficiency is dependent on the length of the path and the number of applied filters. It is important to note that the matrix formalism presented in this article is very much intractable as the best known algorithm to compute ordinary matrix multiplication is approximately $O(|V|^{2.807})$ (Strassen 1969). As stated previously, this matrix model can be approximated using various techniques such as grammar walk algorithms which do not compute the inferences over the entire network, but instead, on local subgraphs (i.e. paths starting from particular vertices) (Rodriguez 2008). With evidential logic, such walks can be executed when resources are available and only in those areas of the knowledge network where it is deemed necessary (e.g. $f \sim 0.5$ and/or low c areas).

Finally, to actually represent a multi-relational, evidence network in RDF and on the Semantic Web, some form of reification can be used. A popular technique is the quad-form of a “triple” where a statement maintains a fourth component known as a named graph (Carroll et al. 2005). With reification, statements can be attached to statements and thus, a $\langle w^+, w^- \rangle$ evidence tuple can be assigned to an RDF statement.

This article presented a non-axiomatic evidential logic that can be implemented within the constructs of RDF and thus, can be used as an evidential reasoning system in the Semantic Web. The benefit of this system is that it works with arbitrary, partial, and contradictory knowledge while, at the same time, in a non-matrix implementation, supports a tractable approximate reasoning process.

Acknowledgements

Vadas Gintautas provided useful insight during the development of these ideas. This work was funded by a Director’s Fellowship from the Los Alamos National Laboratory, EP-SRC grant EP/D504376/1, and British Telecom plc.

References

Berners-Lee, T.; Hendler, J. A.; and Lassila, O. 2001. The Semantic Web. *Scientific American* 34–43.

Bochenski, I. M. 1970. *A History of Formal Logic*. New York, USA: Chelsea Publishing Company.

Carroll, J. J.; Bizer, C.; Hayes, P.; and Stickler, P. 2005. Named graphs, provenance and trust. In *The Fourteenth International World Wide Web Conference (WWW05)*, 613–622. Chiba, Japan: ACM Press.

Chartrand, G. 1977. *Introductory Graph Theory*. Dover.

Dau, F. 2006. RDF as graph-based, diagrammatic logic. *Lecture Notes in Computer Science* 4203.

Donini, F. 2002. *The Description Logic Handbook*. Cambridge University Press. chapter Complexity of Reasoning, 101–141.

Fensel, D., and van Harmelen, F. 2007. Unifying reasoning and search to web scale. *IEEE Internet Computing* 11(2):94–96.

Hayes, J., and Gutierrez, C. 2004. Bipartite graphs as intermediate model for RDF. In *Proceedings of the International Semantic Web Conference ISWC04*, 47–61.

Kolda, T. G.; Bader, B. W.; and Kenny, J. P. 2005. Higher-order web link analysis using multilinear algebra. In *Proceedings of the Fifth IEEE International Conference on Data Mining ICDM’05*. IEEE.

Morale, A. A. M., and Serodio, M. E. V. 2006. A directed hypergraph model for RDF. In Simperl, E.; Diederich, J.; and Schreiber, G., eds., *Proceedings of the Knowledge Web PhD Symposium (KWEPSY07)*.

Patzig, G. 1968. *Aristotle’s Theory of The Syllogism*. Boston, Massachusetts: D. Reidel Publishing Company.

Rodriguez, M. A., and Shinavier, J. 2008. Exposing multi-relational networks to single-relational network analysis algorithms. Technical Report LA-UR-08-03931, Los Alamos National Laboratory.

Rodriguez, M. A.; Bollen, J.; and Van de Sompel, H. 2009. Automatic metadata generation using associative networks. *ACM Transactions on Information Systems* 27(2).

Rodriguez, M. A. 2008. Grammar-based random walkers in semantic networks. *Knowledge-Based Systems* 21(7):727–739.

Strassen, V. 1969. Gaussian elimination is not optimal. *Numerische Mathematik* 13:354–356.

Wang, P. 1993. Non-axiomatic reasoning system (version 2.2). Technical Report 75, Center for Research on Concepts and Cognition at Indiana University.

Wang, P. 1994. From inheritance relation to non-axiomatic logic. *International Journal of Approximate Reasoning* 11:281–319.

Wang, P. 2004a. Cognitive logic versus mathematical logic. In *Proceedings of the Third International Seminar on Logic and Cognition*.

Wang, P. 2004b. Experience-grounded semantics: a theory for intelligent systems. *Cognitive Systems Research* 6(4):282–302.

Wang, P. 2006. *Rigid Flexibility*. Springer.