



Modeling News Trust: A Defeasible Logic Programming Approach

Fernando M. Sagui^{†§} Ana G. Maguitman^{†§}
Carlos I. Chesñevar^{†§} Guillermo R. Simari[†]

[†]Laboratorio de Investigación y Desarrollo en Inteligencia Artificial
Departamento de Ciencias e Ingeniería de la Computación, Universidad Nacional del Sur
Av. Alem 1253 , (8000) Bahía Blanca, Argentina

[§]Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET)
{fms,agm,cic,grs}@cs.uns.edu.ar

Abstract Deciding whether to trust an information sources on the Web has been recognized as one of the main problems in today's Information Society. In particular, assessing the credibility of news is a major research challenge. Typically, criteria such as freshness, relevance and viewer profile have been used by news services to rank news. However, these services do not deal with credibility from a qualitative perspective, and do not provide mechanisms to cope with controversial news reports. To fill this gap, this paper proposes a novel framework that brings the notions of trust and pluralism into play. In our proposal, we integrate dialectical reasoning into a news recommender system. The system is based on a set of basic principles characterizing the nature of trust. We use Defeasible Logic Programming (DeLP), a general-purpose defeasible argumentation formalism based on logic programming, to model the notion of trust. Our approach helps identify antagonism among sources of news and facilitates the analysis of opposing positions.

Keywords: personalized news, Web, argumentation, trust, credibility, trustworthiness, pluralism.

1 Introduction

According to several studies by Nielsen/NetRatings [4], reading news has become one of the most important activities on the Web. The number of visitors to news websites has steadily increased over the last years, and the abundant supply of online news is a clear indication of users' urge to be informed. Moreover, the availability of multiple sources of news provides an opportunity to access pluralistic opinions, which can be regularly found on the media.

Nowadays, there are several commercial multi-source news providers on the Web, such as Google News [1], Yahoo! News [2], etc. Although none of them has disclosed the technical details underlying the way news are selected, aggregated and ranked, it is evident that factors such as freshness, sources and popularity are taken into account. The information provided in news reports may not always be fully verifiable and therefore another important factor that can help select news is *trust* or *credibility*.

Research on multi-source news has generally overlooked the dynamics of news credibility, or if considered, it has been studied through quantitative approaches (e.g. [20]). There is also documented evidence [9] of Google News's plans to build a database of news source credibility based on information such as average story length, number of staff a news source employs, the volume of internet traffic to its website and the number of countries accessing the site. Google's approach to dealing with news credibility is to

take all these and other parameters to create a single value used to rank the results of any news search. However, a foolproof approach to deal with news trust has not been developed yet and multi-source news services remain vulnerable to credibility breaches.¹

A purely quantitative perspective to news credibility has several limitations. On the one hand, the absence of a logic-based model underlying quantitative approaches makes it hard to provide viewers with a justification of why certain news should be trusted. Because quantitative approaches are not equipped with inference capabilities, much of the implicit information remains undiscovered. On the other hand, they are incapable of dealing formally with the defeasible nature of trust. In addition, because trust is to a great degree subjective, quantifying trust by combining measures coming from a pool of credibility assessments may not be entirely realistic.

The goal of this research paper is to define a qualitative and personalized trust-based news service. The service will allow news viewers to access and compare the trustworthiness of news sources and their reports. Viewers' trust statements on sources and reports can be based on the viewers' subjective beliefs or, when absent, trust assumptions can be obtained indirectly from other viewers' beliefs. However, in order to derive trust from other viewers, a trust relationship between viewers must exist.

In this proposal, trust is modeled using DeLP, a defeasible argumentation framework based on logic programming [11]. This allows us to integrate dialectical reasoning into a news service, which will provide a reasoned basis for the news presented to the user.

The paper is organized as follows. The next section presents the background and motivations for our proposal, reviewing DeLP and discussing the problem of ranking and trusting Web news. Section 3 proposes a set of principles for news trust and shows how to represent them by means of a set of DeLP rules. In section 4 we illustrate the proposal with an example. Finally, section 5 overviews related work and section 6 outlines our conclusions.

2 Background and Motivations

2.1 Defeasible logic programming

Defeasible logic programming (DeLP) [11] is a general-purpose defeasible argumentation formalism based on logic programming, intended to model inconsistent and potentially contradictory knowledge. A defeasible logic program has the form $\mathcal{P} = (\Pi, \Delta)$, where Π and Δ stand for *strict* and *defeasible* knowledge, respectively. The set Π involves *strict rules* of the form $P \leftarrow Q_1, \dots, Q_k$ and *facts* (strict rules with empty body), and it is assumed to be *non-contradictory* (i.e., no complementary literals P and $\sim P$ can be inferred, where $\sim P$ denotes the contrary of P). The set Δ involves *defeasible rules* of the form $P \rhd Q_1, \dots, Q_k$, which stand for " Q_1, \dots, Q_k provide a tentative reason to believe P ." Rules in DeLP are defined in terms of *literals*. A literal is an atom A or the strict negation ($\sim A$) of an atom. Default negation (denoted *not* A) is also allowed in the body of defeasible rules (see [11] for details).

Deriving literals in DeLP results in the construction of *arguments*. An argument \mathcal{A} for a literal Q (denoted $\langle \mathcal{A}, Q \rangle$) is a (possibly empty) set of ground defeasible rules that together with the set Π provide a proof for a given literal Q , satisfying the additional constraints of *non-contradiction* (i.e., an argument should not allow the derivation of contradictory literals) and *minimality* (i.e., the set of defeasible information used to derive Q should be minimal). Note that arguments are obtained by a mechanism similar to the usual query-driven SLD derivation from logic programming, performed by backward chaining on *both* strict and defeasible rules; in this context a negated literal $\sim P$ is treated just as a new predicate name *no_P*. In DeLP, arguments provide tentative support for claims (literals). Clearly, as a program \mathcal{P} represents incomplete and tentative information, an argument $\langle \mathcal{A}, Q \rangle$ may be *attacked* by other arguments also derivable from \mathcal{P} . An argument $\langle \mathcal{B}, R \rangle$ is a *counter-argument* for $\langle \mathcal{A}, Q \rangle$ whenever a subargument $\langle \mathcal{A}', Q' \rangle$ (with $\mathcal{A}' \subseteq \mathcal{A}$) in $\langle \mathcal{A}, Q \rangle$ can be identified, such that $\langle \mathcal{B}, R \rangle$ and $\langle \mathcal{A}', Q' \rangle$ cannot be simultaneously accepted since their joint acceptance would allow contradictory conclusions to be inferred from $\Pi \cup \mathcal{A}' \cup \mathcal{B}$. If the attacking argument $\langle \mathcal{B}, R \rangle$ is preferred over $\langle \mathcal{A}', Q' \rangle$, then $\langle \mathcal{B}, R \rangle$ is called a *defeater* for $\langle \mathcal{A}, Q \rangle$. The preference criterion commonly used is *specificity* [11], preferring those arguments which are more direct or more informed, although other criteria could be adopted.

¹A famed example of this vulnerability problem is that of a fifteen-year old teenager who put out a fake release saying he had been hired by Google and within hours, his release was picked up in Google News.

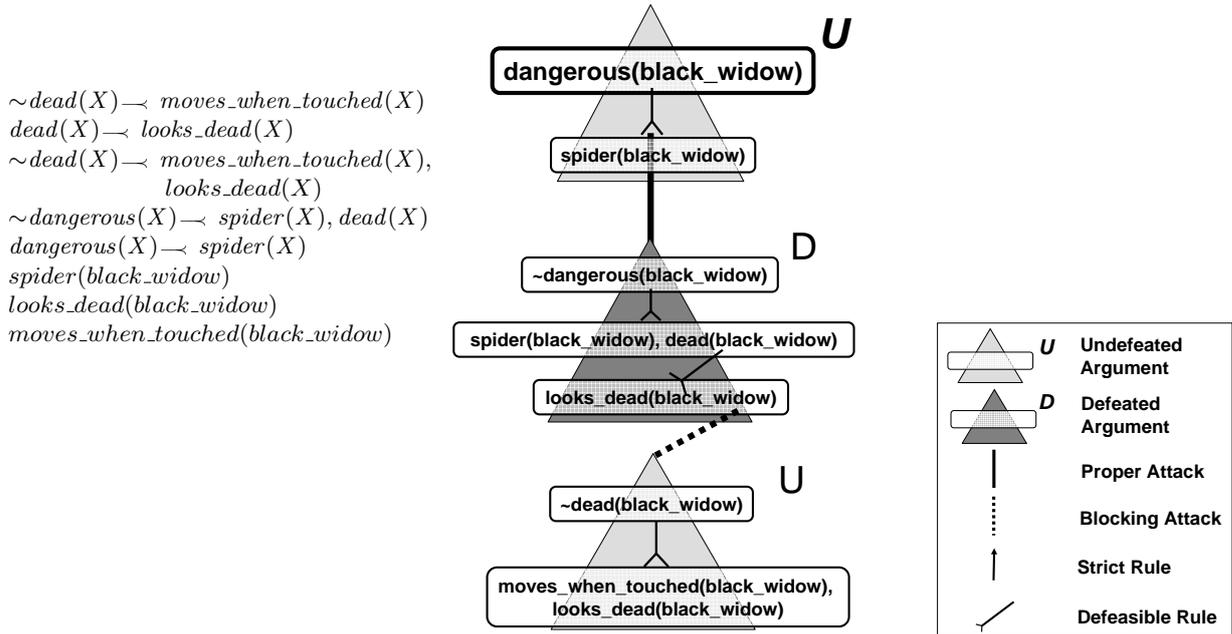


Figure 1: A sample DeLP program about spiders, and the dialectical analysis associated with solving the query dangerous(black_widow)

In DeLP the search for defeaters for a given argument $\langle \mathcal{A}, Q \rangle$ prompts a recursive process, resulting in the generation of a *dialectical tree*: the root node of this tree is the original argument at issue, and every children node in the tree is a defeater for its parent. Additional restrictions help to avoid circular situations when computing branches in a dialectical tree, guaranteeing that every dialectical tree is finite (see [11] for details). Nodes in the tree can be marked either as *defeated* (*D*-nodes) or as *undefeated* (*U*-nodes). The marking of the dialectical tree is performed as in an AND-OR trees: leaves are always marked as undefeated nodes (as they have no defeaters); inner nodes can be marked either as undefeated (if and only if every of its children nodes is marked as defeated) or as defeated (whenever at least one of its children has been marked as undefeated). The original argument $\langle \mathcal{A}, Q \rangle$ (the root of tree) is deemed as ultimately acceptable or *warranted* whenever it turns out to be marked as undefeated after applying the above process.

Figure 1 shows an example of how DeLP can be used to represent commonsense knowledge about spiders. In this sample program there are four defeasible rules (e.g. spiders are usually dangerous, spiders which are dead are usually not dangerous, etc.). The program also includes some facts about a particular spider in a given situation (e.g. a black widow spider which looks dead, but moves when touched). By performing the query dangerous(black_widow), DeLP allows us to conclude that we have a warranted argument supporting the claim that this spider is dangerous. Note that this involves computing a dialectical tree, which in this particular example involves just a single branch. As the query dangerous(black_widow) is supported by a warranted argument, the answer provided by DeLP is *yes*.

Note also that the computation of the dialectical tree is performed automatically by the DeLP interpreter on the basis of the program available. This process is based on an abstract machine which extends Warren’s abstract machine for PROLOG [11]. Given a DeLP program \mathcal{P} , solving a query Q with respect to \mathcal{P} may result in four possible answers: YES (there is at least one warranted argument \mathcal{A} for Q); NO (there is at least one warranted argument \mathcal{A} for $\sim Q$); UNDECIDED (none of the previous cases hold); and UNKNOWN (Q is not present in the program signature). The emerging semantics is skeptical, computed by DeLP on the basis of the goal-directed construction and marking of dialectical trees, which is performed in a depth-first fashion. Additional facilities (such as visualization of dialectical trees, zoom-in/zoom-out view of arguments, etc.) are integrated in the DeLP environment to facilitate user interaction when solving queries.

2.2 Ranking and Trusting News

The problem of ranking Web news has attracted much attention in recent years. There are several reasons why measures of page authority such as PageRank [6] cannot be directly applied at the moment of ranking Web news. Differently from what happens with webpages, the Internet newspapers rarely use linking. Moreover, breaking news usually have priority over previous news because viewers prefer to see information about news events as soon as they take place. However, fresh news usually have very few incoming links, which precludes the application of link analysis algorithms to favor fresh news over stale ones. A ranking model that gives high priority to fresh news, however, will have some deficiencies. Reports on fresh news tend to be incomplete and many stories presented as breaking news are revised when additional information becomes available.

Another approach to rank news could be based on news popularity, estimated by monitoring the number of viewers accessing a report or by a system of voting on favorite stories. However, news popularity may not reflect the real value a news report has for individual viewers.

Usually a viewer has to decide whether a news report is worth reading and whether the facts described in the report are credible. Unfortunately, developing an algorithm for Web news selection and ranking is very difficult because it needs to combine many, sometimes conflicting, aspects. The level of trust a viewer has on a piece of news is not necessarily associated with measures of news authority or popularity, and it may even be negatively correlated with news freshness. A news service that uses a trust-management system can support the viewer in making the decision by selecting reports from trusted sources or based on another trustworthy viewer's opinion.

There are important aspects of trust that need to be considered in order to develop a realistic modeling of trust, as well as reliable and usable services based on this notion:

- **Trust should be Justified.** Most existing news ranking services act as “black boxes”, because they refuse to disclose how they select certain news or rank them in certain ways. This results in trustworthiness issues because they do not provide viewers with a justification of why certain news should be trusted.
- **Trust is Defeasible.** News reports trusted by some viewer can be superseded by other reports carrying more authority, say from CNN or some other trusted source. In the meantime, news agencies are subject to time constraints, which results in the publication of reports with incomplete or inaccurate information. Trust on such reports could be revoked by the release of other more recent ones.
- **Trust is Subjective.** Like many aspects of the Web, news browsing is becoming a collaborative activity. However, judgment of news credibility is idiosyncratic. Therefore, models that deal with trust as an objective notion are unrealistic. In this case, the opinion from the “wisdom of the crowds” may not be as useful as the viewer's personal opinion, or the opinion of another trusted viewer.

The rest of the paper is aimed at providing a framework for dealing with news credibility on the Web, in which the above characteristics of trust are taken into consideration.

3 A Formal Framework for Reasoning about News Trust

A system that reasons about news trust should take a number of ingredients into consideration:

- **Reports.** A report or news article is a written communication of a news event prepared by a specific news agency (source). When a report is made available on the Web, we can identify fields such as *title*, *source*, *timestamp*, *description*, *category* and *link to news content*. Other information related to the report such as *author* can also be derived in certain situations.
- **Sources.** The source of a news article is the agency in charge of supplying the report to be used by the media.
- **Viewers.** A viewer is a user of the news service. The system maintains a pool of viewers. Viewers can provide trust statements about reports, sources and other viewers.

- **Trust/Distrust Statements.** A trust (distrust) statement is an explicit assertion of the fact that a viewer trusts (distrusts) a report, a source or another viewer. These statements allow to infer implicit trust relations, which are useful to provide recommendations to the viewer based on trust.

3.1 Principles for Trust Statements

This subsection presents a set of principles for trust (and distrust) statements. The set of principles will embody the general intuitions about the way trust and distrust statements could be derived from existing ones.

Principle 1. A report coming from a trusted source will typically be trusted.

Principle 2. A report coming from a distrusted source will typically be distrusted.

Principle 3. A report trusted by a trusted viewer will typically be trusted.

Principle 4. A report distrusted by a trusted viewer will typically be distrusted.

Principle 5. A source trusted by a trusted viewer will typically be trusted.

Principle 6. A source distrusted by a trusted viewer will typically be distrusted.

An interesting situation will arise when two conflicting conclusions can be reached. For example, a trusted viewer distrusts a report, but the report was released by a trusted source. Assuming that in general we prefer to base our opinion on information provided by trusted sources, we can add the following two principles:

Principle 7. A report coming from a trusted source will typically be trusted, even if it is distrusted by a trusted viewer.

Principle 8. A report coming from a distrusted source will typically be distrusted, even if it is trusted by a trusted viewer.

Additional principles could be added. For example, if some viewer has a very good reputation for fact-checking, we will prefer to trust this viewer's opinion even if it conflicts with other viewers' opinions. Other principles that could be added to the list may include references to the timestamp of the report. For instance, a more recent report will be trusted over an outdated one, unless the report is "just out" (due to the eagerness to publish the story some fresh news reports may not be as reliable as old ones). The list of principles could be extended indefinitely, including references to news author, country of origin of the source, news category, etc. Some principles could be personalized, because different viewers may disagree on the conclusion that should be adopted given certain facts. For the sake of simplicity, we will take principles 1 to 8 as the core principles for our trust-management system.

3.2 Using DeLP to Reason about News Trust and Pluralism

Principles 1 to 8 can be naturally modeled using the DeLP rules as shown in Table 1.

$trust_report(V, R)$	\leftarrow	$report_source(R, S), trust_source(V, S)$	(R_1)
$\sim trust_report(V, R)$	\leftarrow	$report_source(R, S), \sim trust_source(V, S)$	(R_2)
$trust_report(V, R)$	\leftarrow	$trust_viewer(V, V1), trust_report(V1, R)$	(R_3)
$\sim trust_report(V, R)$	\leftarrow	$trust_viewer(V, V1), \sim trust_report(V1, R)$	(R_4)
$trust_source(V, S)$	\leftarrow	$trust_viewer(V, V1), trust_source(V1, S)$	(R_5)
$\sim trust_source(V, S)$	\leftarrow	$trust_viewer(V, V1), \sim trust_source(V1, S)$	(R_6)
$trust_report(V, R)$	\leftarrow	$report_source(R, S), trust_source(V, S),$ $trust_viewer(V, V1), \sim trust_report(V1, R)$	(R_7)
$\sim trust_report(V, R)$	\leftarrow	$report_source(R, S), \sim trust_source(V, S),$ $trust_viewer(V, V1), trust_report(V1, R)$	(R_8)

Table 1: DeLP Rules

We envision a trust-management system with built-in rules R_1 to R_8 . We should remark that viewers will not need to deal directly with DeLP rules. However, through a user-friendly question-answering

interface it will be possible to extend or adjust the built-in core rules based on the viewer's preferences. Trust and distrust statements about reports, sources and other viewers will be added to the system whenever the viewer rates these entities.

For a particular viewer v , and based on the corresponding DeLP rules and facts, news reports will be classified into three sets:

Trusted Reports: those reports r_i for which there exists at least one warranted argument supporting $trust_report(v, r_i)$.

Distrusted Reports: those reports r_i such that there is a warranted argument supporting $\sim trust_report(v, r_i)$.

Undecided: those reports r_i for which there is no warranted argument for $trust_report(v, r_i)$ or $\sim trust_report(v, r_i)$.

This classification will allow the viewer to focus on those reports considered trustworthy, and to be warned about the non trustworthy ones.

4 A Worked Example

Assume Joe is a viewer, whose personalized trust-management system contains rules R_1 to R_8 together with the following facts:

$report_source(true_news, the_truthteller)$	(F_1)
$report_source(false_news, the_corker)$	(F_2)
$report_source(some_news, the_incog)$	(F_3)
$trust_source(joe, the_truthteller)$	(F_4)
$\sim trust_source(tom, the_corker)$	(F_5)
$\sim trust_report(ann, true_news)$	(F_6)
$trust_report(ann, false_news)$	(F_7)
$\sim trust_report(ann, some_news)$	(F_8)
$trust_report(tom, some_news)$	(F_9)
$trust_viewer(joe, ann)$	(F_{10})
$trust_viewer(joe, tom)$	(F_{11})

Suppose that the reports “True News”, “False News” and “Some News” need to be classified based on their trustworthiness status.

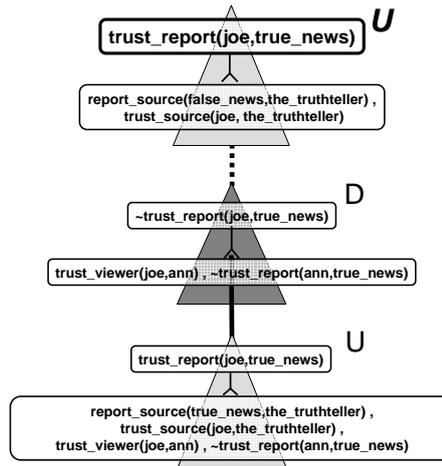


Figure 2: A DeLP dialectical tree supporting the conclusion that the “True News” report should be trusted by Joe.

Figure 2 shows that “True News” can be trusted by Joe, as there exists a warranted argument supporting $trust_report(joe, true_news)$. On the other hand, the existence of a warranted argument for

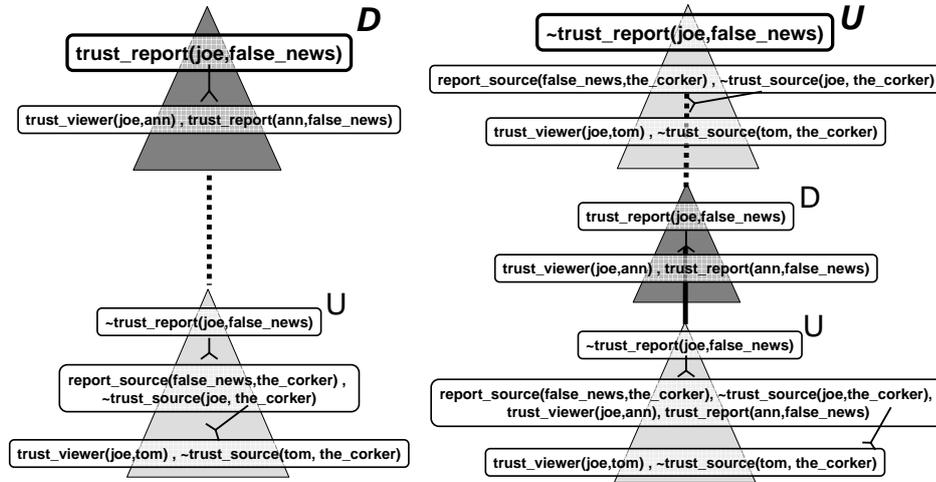


Figure 3: DeLP dialectical trees supporting the conclusion that the “False News” report should be distrusted by Joe.

$\sim trust_report(joe, false_news)$, as shown in figure 3, allows the system to conclude that “False News” must be distrusted. Finally, figure 4 shows that “Some News” can neither be trusted nor distrusted as there is no warranted argument for $trust_report(joe, some_news)$ or $\sim trust_report(joe, some_news)$.

An important by-product of using an argument-based approach to classify news reports according to their trustworthiness status is that the viewer will be able to inspect the reasons that lead the system to provide a conclusion. In order to benefit from this facility a special user interface would be required. This interface will be in charge of providing a user-friendly mechanism to interact with the dialectic tree.

5 Related Work

A variety of methods have been proposed and a number of systems have been developed to facilitate access to news on the Web. NewsInEssence [23] is a system that searches and clusters related news. QCS [8] is a tool that facilitates the task of grouping and categorizing news. In [15] a method is proposed to search web articles while TV news are on the air. Other tools (e.g. [24]) automatically extract domain-oriented news from websites. Velthune [14] is a search engine that extracts information both from the Web and from newsfeeds. Another system, NewsJunkie [10], has the purpose of identifying novel news in the context of stories previously reviewed by the user. Compare&Contrast [18] is a Case-Based Reasoning system that uses the Web as a knowledge base to discover comparable cases for news stories. Many techniques have been proposed to organize news in topics. Most of them try to recognize, track and summarize stories [21, 19, 17, 5]. ArgueNet, a system previously proposed by some of the authors [7], is an argument-based framework for ranking web results. The ArgueNet system differs from this proposal in being a general approach for identifying relevant results and not attempting to provide an underlying formal model of the notion of trust.

A few systems deal with the notions of credibility and media bias. In [20] a method is proposed to rate the credibility of news documents. The method uses algorithms that compare the content of different news sources. PolyNews [22] is a news service framework that tries to mitigate the effect of media bias by the creation of multiple classified viewpoints. NewsTrust [3] is a service created to evaluate news where users can rank news reports, news writers and news sources. A multi-layer recommendation system based on trust is proposed in [16] and a general framework for the analysis of the propagation of trust and distrust is presented in [13].

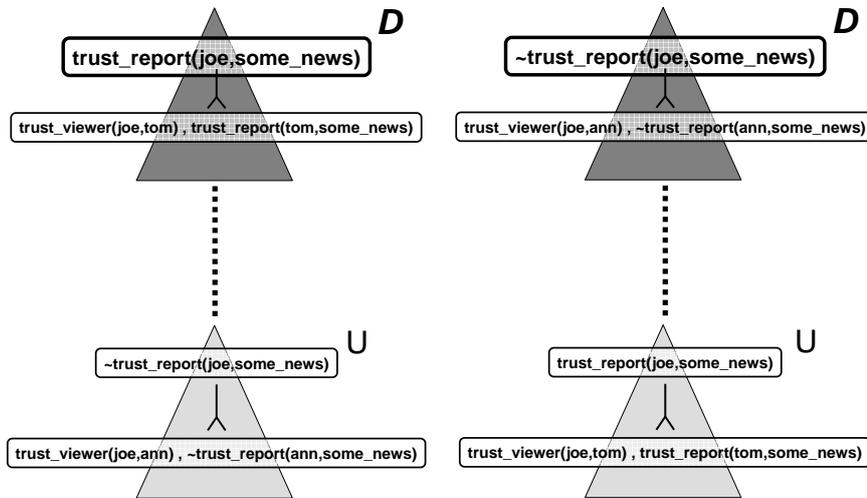


Figure 4: DeLP dialectical trees showing that it is not possible to conclude that the “Some News” report should be trusted or distrusted by Joe and therefore it will be classified as undecided.

6 Conclusions and Future Work

In this paper we have proposed a formal framework to deal with news trust and distrust. A set of principles has been proposed and modeled using DeLP rules.

An implementation of the proposed trust-based system is in progress. As a first step we have implemented a multi-source news service that monitors several newsfeeds and populates a database of news. We are currently developing algorithms for clustering news reports based on the news events the reports describe. In the near future, we expect to integrate this environment with information about viewers and their trust statements, as well as to equip the system with reasoning capabilities based on DeLP as described in this article. The DeLP environment is available online to test at http://lidia.cs.uns.edu.ar/delp_client.

The proposed tool relies on the user declaring his or her preference criteria, which the system codifies as facts and rules. This process could be complemented by the application of techniques for defeasible rule discovery as described in [12]. As part of our future work we expect to study more powerful models of trust. A user is much more likely to believe statements from a trusted acquaintance than from a stranger. Therefore, trust could be quantified and trust values could be transitively obtained from other viewer and scaled down depending on the viewer providing the information. In this sense we expect to study ways of extending our model of trust by adding possibilistic and probabilistic reasoning.

We believe our proposal presents a number of advantages over existing news recommendation services. An argument-based approach allows the exploration of arguments that challenge the viewer perspective. Using an argument-based framework can help the viewer review previously maintained trust statement, as it is possible to analyze the justification in favor and against trusting some report. This could also help discover events in which biased sources present facts in a convincing manner, but neglect to include other important facts that are against their position. Because conclusions are justified by the system, the viewer is able to go back and figure out where errors were revealed. Clearly, in practice most of such facts will remain unverifiable, and there will be room for ideology and interpretation. The proposed framework intends to be faithful to this aspect of reality.

7 Acknowledgements

Authors acknowledge the “Agencia Nacional de Promoción Científica y Tecnológica” (PICT 2005, N 32373), “Consejo Nacional de Investigaciones Científicas y Técnicas” (PIP 5050) and “Universidad Nacional del Sur” (24/N016 and 2/ZN13).

References

- [1] <http://news.google.com/>.
- [2] <http://news.yahoo.com/>.
- [3] <http://www.newstrust.net/>.
- [4] <http://www.nielsen-netratings.com/>.
- [5] James Allan, Rahul Gupta, and Vikas Khandelwal. Temporal summaries of new topics. In *SIGIR '01: Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 10–18, New York, NY, USA, 2001. ACM Press.
- [6] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1–7):107–117, 1998.
- [7] Carlos Chesñevar and Ana Maguitman. ArgueNet: An argument-based recommender system for solving Web search queries. In *Proceedings of the International IEEE Conference on Intelligent Systems (IS 2004)*, pages 282–287. IEEE, June 2004.
- [8] Daniel M. Dunlavy, John Conroy, and Dianne P. O’Leary. Qcs: a tool for querying, clustering, and summarizing documents. In *NAACL '03: Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology*, pages 11–12, Morristown, NJ, USA, 2003. Association for Computational Linguistics.
- [9] Barry Fox. Google searches for quality not quantity. *New Scientist magazine*, 2497:24. 30 April 2005.
- [10] Evgeniy Gabrilovich, Susan Dumais, and Eric Horvitz. Newsjunkie: providing personalized newsfeeds via analysis of information novelty. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 482–490, New York, NY, USA, 2004. ACM Press.
- [11] A. García and G. Simari. Defeasible Logic Programming: An Argumentative Approach. *Theory and Practice of Logic Programming*, 4(1):95–138, 2004.
- [12] G. Governatori and A. Stranieri. Towards the application of association rules for defeasible rules discovery. In *Legal Know. & Inf. Sys.*, pages 63–75. JURIX, IOS Press, 2001.
- [13] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *International World Wide Web Conference (WWW2004)*, 2004.
- [14] A. Gulli. The anatomy of a news search engine. In *WWW '05: Special interest tracks and posters of the 14th international conference on World Wide Web*, pages 880–881, New York, NY, USA, 2005. ACM Press.
- [15] Monika Henzinger, Bay-Wei Chang, Brian Milch, and Sergey Brin. Query-free news search. In *WWW '03: Proceedings of the 12th international conference on World Wide Web*, pages 1–10, New York, NY, USA, 2003. ACM Press.
- [16] Claudia Hess. Trust-based recommendations for publications: A multi-layer network approach. *TCDL Bulletin*, 2(2), 2006.

-
- [17] Andrew J. Kurtz and Javed Mostafa. Topic detection and interest tracking in a dynamic online news source. In *JCDL '03: Proceedings of the 3rd ACM/IEEE-CS joint conference on Digital libraries*, pages 122–124, Washington, DC, USA, 2003. IEEE Computer Society.
- [18] Jiahui Liu, Earl Wagner, and Larry Birnbaum. Compare&contrast: using the web to discover comparable cases for news stories. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 541–550, New York, NY, USA, 2007. ACM Press.
- [19] Masaki Mori, Takao Miura, and Isamu Shioya. Topic detection and tracking for news web pages. In *WI '06: Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence*, pages 338–342, Washington, DC, USA, 2006. IEEE Computer Society.
- [20] Ryosuke Nagura, Yohei Seki, Noriko Kando, and Masaki Aono. A method of rating the credibility of news documents on the web. In *SIGIR '06: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 683–684, New York, NY, USA, 2006. ACM Press.
- [21] Ramesh Nallapati, Ao Feng, Fuchun Peng, and James Allan. Event threading within news topics. In *CIKM '04: Proceedings of the thirteenth ACM international conference on Information and knowledge management*, pages 446–453, New York, NY, USA, 2004. ACM Press.
- [22] Souneil Park, Seungwoo Kang, and Junehwa Song. Polynews: Delivering multiple aspects of news to mitigate media bias. Technical report, KAIST, 2006.
- [23] Dragomir Radev, Jahna Otterbacher, Adam Winkel, and Sasha Blair-Goldensohn. Newsinessence: summarizing online news topics. *Commun. ACM*, 48(10):95–98, 2005.
- [24] D. C. Reis, P. B. Golgher, A. S. Silva, and A. F. Laender. Automatic web news extraction using tree edit distance. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 502–511, New York, NY, USA, 2004. ACM Press.