

Trusting Information Sources One Citizen at a Time

Yolanda Gil and Varun Ratnakar

USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292
{gil, varunr}@isi.edu

Abstract. This paper describes an approach to derive assessments about information sources based on individual feedback about the sources. We describe TRELIS, a system that helps users annotate their analysis of alternative information sources that can be contradictory and incomplete. As the user makes a decision on which sources to dismiss and which to believe in making a final decision, TRELIS captures the derivation of the decision in a semantic markup. TRELIS then uses these annotations to derive an assessment of the source based on the annotations of many individuals. Our work builds on the Semantic Web and presents a tool that helps users create annotations that are in a mix of formal and human language, and exploits the formal representations to derive measures of trust in the content of Web resources and their original source.

1 Introduction

The Semantic Web can be described as a substrate to support advanced functions for collaboration (human-human, computer-human, computer-computer), sharing of Web resources, and reasoning about their content [3]. The markup languages that are being proposed for the Semantic Web will be the basis to develop reasoners, proof checking and derivation tools, and many other functions such as Web services. The Semantic Web will also be the basis for the Web of Trust, which will provide mechanisms to handle authentication, permission, and validation of attribution in a Web where, by design, anyone can contribute content, links, and services.

A lot of current emphasis on the Web of Trust is in accessing resources, specifically authentication and permission issues. Digital signatures and public keys support authentication. Proofs are another important technology in the Web of Trust, since permission schemes are often described with rules and statements (e.g., anyone working for company C should be allowed to access D) and will need to rely on proofs that can reason about the rules and conclude whether access should be granted. An important issue with respect to both authentication and permission is checking that a document can be attributed to the source specified. For example, if Joe Doe writes an article and publishes it claiming Henry Kissinger as the author, it should be possible to check the truth about the document's authorship.

Our work addresses a different issue on the Web of Trust regarding whether to trust the *content* of a Web resource depending on its source. It seems that people reach some times informal consensus on how and when to trust what a source says. Many qualifiers about sources seem to be common knowledge only to those familiar with the topic. Some sources are generally considered more trustworthy or reliable than others (e.g., [19]). Some sources are considered authoritative in specific topics (e.g., [14,22]). Some sources are preferred to others depending on the specific context of use of the information (e.g., student travelers may prefer [17], families may find [12] more preferable, and business people may go with [6]). Some sources are considered pretty accurate but it is understood they are not necessarily up to date (e.g., [5]). Finally, specific statements by traditionally authoritative sources can be proven wrong in light of other information, while the source's reputation will still hold. In this sense, there is a finer grain of detail in attributing trust to a source with respect to specific statements made by it.

These kinds of observations about sources are the result of informal consensus and should be captured by the Web of Trust. Individual users should be provided with tools to annotate their views and opinions on resources available on the Web, as well as to qualify and justify these views if they choose to. The Semantic Web will provide the markup languages and reasoners to derive consensus on how to assess a source based on these individual annotations about sources.

This paper describes our initial work on TRELLIS to enable users to express their trust on a source and the statements made by it, and to combine individual views into an overall assessment of each source of information. TRELLIS enables users to annotate how they analyze and use information for some decision making purpose. As the user considers information from different sources relevant to their purpose, they annotate their views as they find contradictory or complementary statements, make an opinion on what to believe in the absence of complete information. TRELLIS includes a vocabulary and markup language for semantic annotations of decisions and tradeoffs, and allows users to extend this vocabulary with domain specific terms or constructs that are useful to their particular task. We have used TRELLIS with a variety of scenarios to annotate tradeoffs and decisions (e.g., military planning), organize materials (e.g., search results), analyze disagreements and controversies on a topic (e.g., intelligence analysis), and handle incomplete and conflicting information (e.g., genealogy research).

The paper begins with an overview of TRELLIS as an information analysis tool. Then we describe how users can specify source descriptions and qualifications in TRELLIS. We show how TRELLIS derives ratings for each source, averaged over many users and many analyses. We discuss how these ratings can be presented in useful ways to users to help them assess sources in subsequent analysis with TRELLIS. We conclude with related work and a discussion of future directions.

2 Information Analysis in TRELIS

TRELIS is an interactive tool that helps users annotate the rationale for their decisions, hypotheses, and opinions as they analyze information from various sources. This section gives a brief overview of TRELIS, more details can be found in [13]. TRELIS is available on-line at trellis.semanticweb.org.

Our work is inspired by military intelligence analysts, but we believe that this kind of information analysis is a common task for many users in the information age. Intelligence analysis carries over to political, strategic, and competitive (business) areas. Someone doing genealogy research looks at various sources of information that may be incomplete and contradictory, make plausible hypotheses in the light of what is known, and create a consistent story about what happened in the family. Someone planning a vacation makes a lot of decisions after consulting many airlines and hotel possibilities, their friends traveling to the same destination may want to consult the same sources though perhaps making different choices.

Our goal is to enable users to create annotations of their analysis of alternative sources of information as they make a decision or reach a conclusion based on their analysis. Once this rationale is recorded, it can be used to help users justify, update, and share the results of their analysis. Users need support after they have made a decision, reached a conclusion, or made a recommendation, since they are often required to: 1) explain and justify their views to others, 2) update the decision in light of additional information or new data, 3) expose the intermediate products of the final recommendation to others that may be analyzing related information to make similar decisions.

TRELIS includes a language for annotating information analysis, which can be extended by users to suit their needs. The language uses the following basic components. A *statement* is a piece of information or data relevant to an analysis, such as "Cuban pilots were trained in Czechoslovakia", "Prince Larry did not inhale". A statement may have been extracted, summarized, or concluded from a document. Documents are often an existing Web resource (text, imagery, or any other format) indicated by a URI, but can also be a user-provided document such as an email message or a note relating a conversation in which case TRELIS helps users turn them into Web resources as well. The user can also create a statement to introduce a hypothesis, a conclusion, or an observation that will be used in the analysis, which may or may not be backed up by (associated with) a document. Every document has a source description, describing its creator, publisher, date, format, etc. Each statement and its source can have a source qualification specified as a degree of reliability and credibility. The next section explains in detail how sources are described in TRELIS.

A *compound statement* is composed of several statements related by a construct. *Constructs* reflect how individual statements are related in the analysis. For example, a causal construct is used to form the compound statement: "Cuban pilots were trained in Czechoslovakia" and "Arrival of latest Soviet warplanes in Cuba" **results in** "A major upgrade of the Cuban Air Forces". Statements can have a *likelihood qualification*, which is a subjective informal indication of the

analyst's reaction to a statement (or compound statement). This can indicate surprise, dismissal, saliency, accuracy, etc. A *reason* can be used to justify a compound statement, a source qualification, and a likelihood qualification.

These basic components are used to create *units*. The basic structure of a unit is:

statement {and *statement*}* *construct* {and *statement*}*
 is {not} *likelihood-qualifier* because
 according to *source-description* which is
reliability-qualifier because *statement* and
credibility-qualifier because *statement*

An example of a unit, taken from a Special Operations planning analysis, is:

water temperature unsustainable for SDV divers
is elaborated in
average March water temperature is 55-60 degrees
and
platoon requires minimum water temperature of 65 degrees
according to source
Cmdr Smith *which is*
completely reliable (A)
because Cmdr Smith has 15 years experience with JSOC
and
probably true
because Cmdr Smith has been platoon cmdr for 3 years

The user may or may not provide all the components of a unit, only a statement is required to form a unit.

An analysis can be done with an overarching *purpose* (or topic), which is often a question or request that the information analyst starts with. An analysis is composed of many such units. They can be linked as subunits of one another. Units or statements can be left with no links to the overall analysis, and in that case they can be specified as attachments to the analysis. This is useful to indicate that they have been considered by the user but do not appear in the derivation of the final analysis (for lack of time, or because the analyst found better options to justify their conclusions). An analysis is represented as a tree, but can be a set of trees (especially if the user has not reached a final conclusion yet).

Our default set of constructs used in units is drawn from argumentation and discourse relations, logic connectives, action representations, causality, and object representations (see [13] for details). Some examples are {P}* contradicts {Q}*, {P}* is supported by {Q}*, {P}* before {Q}*, {P}* is summarized by {Q}*. Our emphasis is not on the completeness or semantics of these constructs, but rather on offering users a sensible set of constructs that they can draw from in forming an argument. Our experience has been that users often cannot pin down the logic behind their argument but would still like to reflect their conclusion and some sort of (informal) justification of it. Users can add additional constructs to suit their specific needs using the TRELLIS interface.

A very important feature of TRELIS is that users can include in the analysis references to sources and indicate that they were dismissed and why, typically because they contradict other information used in the analysis or simply because other sources provided more precise or detailed data.

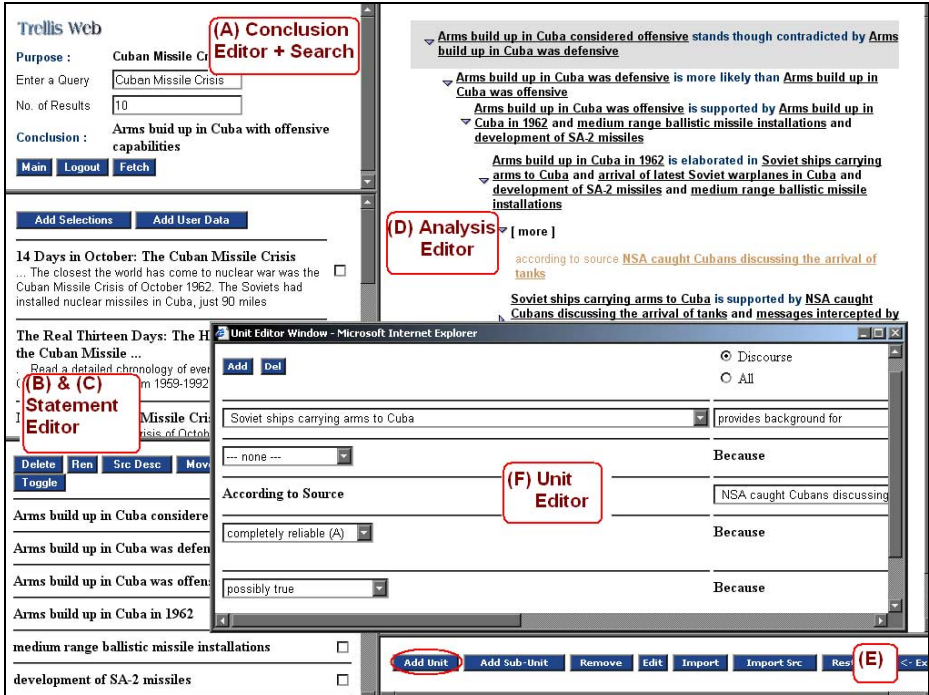


Fig. 1. A Snapshot of the TRELIS user interface. From top-left counter clock wise the system shows the purpose and conclusions of the analysis (A), the original documents and associated statements (B&C), the units of the analysis (E,F), and the overall analysis (D).

Figure 1 shows the current TRELIS user interface. The example is inspired on the Cuban missile crisis, one of the most thoroughly studied cases of intelligence analysis. The purpose of the analysis and the final conclusion are shown in Frame (A). Analysis and opinions revolve around facts, statements, and hypotheses. With Frame (B), users can search the Web and mark documents to be indexed by TRELIS, or can add their own documents which are then converted to a Web resource and given a URI. Each resource is then associated with a short statement entered by the user in Frame (C). Users can specify several statements per resource, each summarizing a salient piece of information described within the resource in terms that are suitable to the user. The 'Src Desc' button in Frame (C) allows users to enter meta-data about the resource. Frame (E) invokes the Unit Editor (F). The overall analysis is composed using the Analysis Editor, shown in Frame (D).

TRELIS generates annotations of the user's analysis in several markup languages (XML, RDF, and DAML+OIL). The constructs used in units as well as the source descriptions are represented in the corresponding schema languages and ontologies.

Each analysis is turned into an annotated Web document and contains links to all the resources referenced within. When users add new constructs to the language, the corresponding schemas and ontologies are updated.

```

<element rdf:ID="elem29">
  <title>Unstability of Cuban and Soviet governments</title>
  <link>/trellis_web/cached/demo_dir7_1011051896.html</link>
</element>
<statement rdf:ID="st1">
  <left_elem rdf:resource="#elem1" />
  <relation rdf:resource="http://excalibur.isi.edu:8888/trellis_web/library#stands_though_contradicted_by" />
  <right_elem rdf:resource="#elem2" />
  <qualifier />
  <sub_st rdf:resource="#st2" />
  <sub_st rdf:resource="#st15" />
</statement>
<statement rdf:ID="st2">
  <left_elem rdf:resource="#elem2" />
  <relation rdf:resource="http://excalibur.isi.edu:8888/trellis_web/guest/library#is_more_likely_than" />
  <right_elem rdf:resource="#elem3" />
  <qualifier />
  <sub_st rdf:resource="#st3" />
  <sub_st rdf:resource="#st11" />
  <sub_st rdf:resource="#st14" />
</statement>
<statement rdf:ID="st9">
  <left_elem rdf:resource="#elem3" />
  <relation rdf:resource="http://excalibur.isi.edu:8888/trellis_web/library#is_supported_by" />
  <right_elem rdf:resource="#elem4" />

```

Fig. 2. A part of semantic RDF markup of the analysis in Fig. 1. TRELLIS provides an RDF schema as well as a DAML ontology for the terms used in the markup.

TRELLIS allows users to view the markup of the annotations, as well as the schemas and ontologies for the underlying constructs. Figure 2 shows part of the semantic RDF markup annotations of the analysis shown in Figure 1. These annotations are posted as a Web document as we mentioned above, and are available to other tools for search, reference, and reasoning.

TRELLIS also supports sharing and collaboration. The user can search and view the analyses entered by other users, import relevant portions into their own analysis, and change the imported portions as they see fit (this is done through the "Import" button on the bottom right of the screen).

3 Source Attribution and Description

TRELLIS allows users to annotate the source attribution for each statement used in the analysis, to describe the source, and to make qualifications about it. Figure 3 shows an example, discussed in detail throughout this section.

For each document indexed in TRELLIS, the user can annotate meta-data regarding its attribution using the Dublin Core [8]. The Dublin Core (*dc:*) was developed as a standard to describe resources (e.g., documents). A document is described with 15 main attributes: *dc:title*, *dc:creator*, *dc:subject*, *dc:description*, *dc:publisher*, *dc:contributor*, *dc:date*, *dc:type*, *dc:format*, *dc:identifier*, *dc:source*, *dc:language*, *dc:relation*, *dc:coverage*, and *dc:rights*. Five of them are concerned with attribution

of information. The *dc:creator* is an entity primarily responsible for making the content of the resource. The *dc:publisher* is an entity responsible for making the resource available. The *dc:contributor* is an entity responsible for making contributions to the content of the resource. The *dc:source* is a reference to a resource from which the present resource is derived. The *dc:relation* is a reference to a related resource.

Consider a pseudo-fictitious example of a document in an FA Times article from the CREUTERS agency that reports on drug problems in the Monarchy and that contains the statement "At a press conference last Monday, Duckingham Palace was adamant that Prince Larry did not inhale". In this case, the CREUTERS journalist would be the creator, FA Times the publisher, the Duckingham Palace spokesperson would be a contributor, the original CREUTERS article would be the source, and the tapes of the press conference could be specified as a relation.

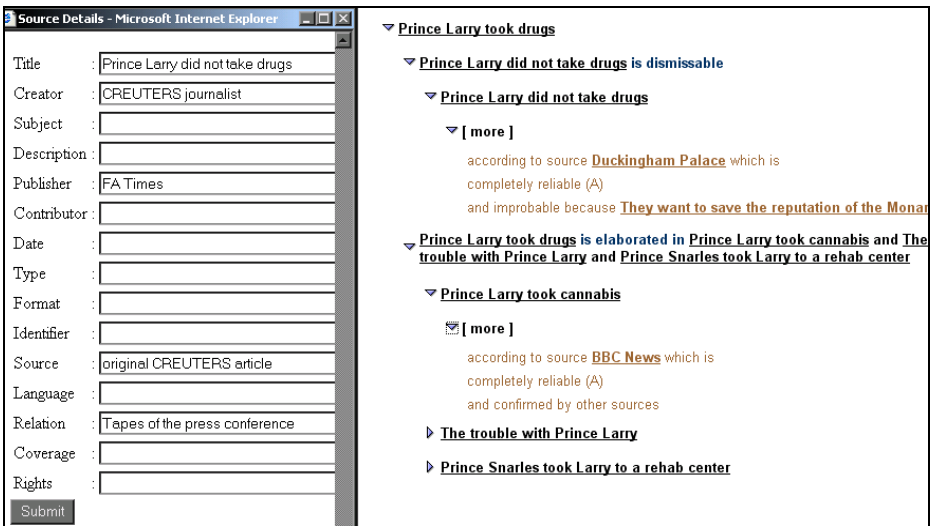


Fig. 3. An Example of attribution of a statement in TRELLIS, where the source stated in the analysis (shown on the right) may or may not be one of the entities mentioned in the DC metadata description (shown on the left).

In TRELLIS, document used in an analysis is first indexed with a short statement, as a way to summarize the particular aspect of the document used in the analysis. The statement points to the document, and must become part of an analysis unit. In the unit, the user can specify a *source* for that statement. This TRELLIS source can be any of the five fields in the Dublin Core metadata that are related to attribution and that we mentioned above, but can also be any other entity that is not indicated in it. TRELLIS gives the user this flexibility because the user may trust (or distrust) any of these sources enough to take some stand about the statement. In the rest of the discussion, we will refer to sources as those chosen by a TRELLIS user to be associated with a statement within a unit.

In our example, the user could use the statement "Prince Larry did not take drugs" to index the news article, but specify as its source one of: the FA Times, the publisher of the newspaper, CREUTERS, Duckingham Palace, or perhaps Prince Larry himself depending on the nature and the context of the analysis. In Figure 3, the source chosen by the user for that statement chosen by the user in their analysis (shown on the right) is Duckingham Palace, which is the dc:contributor (the DC metadata are shown on the left).

In TRELLIS, the user can also qualify the source of a statement by its *reliability* and *credibility*. Reliability is typically based on credentials and past performance of the source. Credibility specifies the analyst's view of probable truth of a statement, given all the other information available to the analyst. Reliability and credibility are not the same, for example a completely reliable source may provide some information that may be judged to be not very credible given other known information. In our example, Duckingham Palace may be judged to be a pretty reliable source based on its reputation, yet the user may decide the source to be not very credible in making this particular statement. In TRELLIS, we use the default ratings of reliability and credibility that are typically found in military intelligence manuals. Reliability is specified by a six-valued scale ranked A to F (completely reliable, usually reliable, fairly reliable, not usually reliable, unreliable, and not possible to judge). Credibility can have one of six values on a scale (confirmed by other sources, probably true, possibly true, doubtfully true, improbable, and not possible to judge).

When users introduce hypotheses in an analysis, they are designated as the source and are not allowed to give themselves any ratings.

Because an important goal is to give users maximum flexibility, TRELLIS does not require users to specify or qualify the sources for any of the statements.

In summary, the user can specify the source of a statement used in the analysis, which may or may not be one of the fields used in the Dublin Core meta-data for the document that supports the statement. Users can also qualify the source attribution according to its reliability and credibility.

4 Deriving an Assessment about a Source

As many users create multiple analyses that refer to common sources, TRELLIS creates an overall consensus assessment about each source as we explain in this section. We describe here our initial approach to get a rough approximation of the relative ratings of each source. In future work we plan to formalize, evaluate, and extend this approach to consider additional factors.

For each purpose or topic analyzed by the user, we derive a rating for each of the sources referenced in that analysis. First, we derive a rating for each statement associated to each source.

The first criterion for rating the sources is provided by the user in terms of R (reliability) and C (credibility) on a 1 to 6 scale. If the statement is unqualified in the unit, then an average value is taken (the fact that the user did not specify R and C is taken into account by one of the other rating variables described below).

For those sources where C and R are not specified in the unit (i.e., the unit just says "A according to source S", we track the use of the statement along the analysis tree as an indicator of the user's view on the source in the current analysis context. We describe the source to be one of the following depending on the status of 'A'.

✂ Used: This means that the user has found the statement relevant to the analysis and is using it to make the final conclusion. So if the user makes another upper level statement "'B' is supported by 'A'", then we track both 'B' and 'A' up the tree, since in some sense 'S' now transitively vouches for both 'A' and 'B'. In this way, when the final top-level conclusion is reached, we assume that the source 'S' was considered by the user to be a contributor to reaching a conclusion and thus the user had some level of trust on the source.

✂ Tainted: The system marks a statement by a source as *tainted* or *dismissed* if somewhere along the analysis tree the user refers to it with a pejorative construct (such as *dismiss*, or *contradicts*). Therefore, if any upper level statement uses such construct then the source 'S' is marked as tainted.

✂ Not used: The statement is not used in reaching the conclusion. This happens, as we mentioned, when users leave statements as part of the analysis in case they are useful in future updates of the analysis. Therefore, if the upper level statements do not transitively use 'A' such that the top level conclusion is never reached, then the source is marked as 'Unused'.

Therefore, each source-statement pair (s,a) has the following rating attributes:

1. C(s,a) - Credibility - based on credibility qualifications over all the units that specify this qualifier for the source. On a scale of 1 to 6.
2. R(s,a) - Reliability - based on reliability qualifications over all the units that specify this qualifier for the source. On a scale of 1 to 6.
3. U(s,a) - Used - a count of the number of times it was used to reach a conclusion.
4. T(s,a) - Tainted - a count of the number of times it was found to be *tainted*.
5. N(s,a) - Not Used - a count of the number of times it was unused in reaching the final conclusion.

The system's overall rating O(s,a) of the source-statement pair is derived from the normalized averages of the above ratings:

$$O(s,a) = k_1 * (C + R) + k_2 * U - k_3 * T - k_4 N$$

The coefficients k₁, k₂, k₃, and k₄ reflect the relative weight given to each individual factor. Here U, T and N are normalized with respect to the total number of times that the (s,a) pair is used. There might be some cases, as previously mentioned, where the user does not provide a credibility or a reliability and where TRELIS assigns 3 (the average value) for C and R, then the formula for O(s,a) allows other factors (if it is used in reaching a conclusion or not, if it is dismissed or not, if it is not considered in reaching a conclusion) to influence the final value for the rating of the source by adding to its rating with U and taking from its rating T and to a lesser extent N.

In the future, we would like advanced users to be able to change the coefficients k1, k2, k3, k4 so that they can control how much they would like to weigh different uses of the source by other users. The current values of the coefficients are set to 2/3, 2, 2, and 1 respectively.

The screenshot shows the Trelis Web interface. On the left, there is a search bar with the query "Rubber Duck operation" and a "Go for it" button. Below the search bar is a list of items with checkboxes, including "Military Collectibles", "R.Smg TORELLI", "WWF stamps with online payment - Local FDC", "Avg Wind speed in March around 14 Knots", "METOC Rota (9.8/10)", "Temperature ok", "Wave height ok", "Gibraltar avg temperature in March is 60 F", "Avg. wave height in March is 4 ft", "Unofficial Special Operations Site (8.6/10)", and "Wind factor not ok".

On the right, there is a detailed view of the selected item, showing a list of reasons for its selection. The reasons are:

- Fixed wing aircrafts are much more silent
- Helicopter rotor noise may compromise secrecy
- Wind speed ok for fixed wing is supported by Fixed wing aircrafts cannot be used in gusty winds (+30 Knots) and Winds from the east at 10 Knots and Avg Wind speed in March around 14 Knots
- Fixed wing aircrafts cannot be used in gusty winds (+30 Knots)
- Winds from the east at 10 Knots
- Avg Wind speed in March around 14 Knots
- Drop point 5 miles west of Gibraltar is elaborated in Drop point west of the peninsula and a visible shoreline is preferred and Visibility is around 6 miles and Critical Meteorological data ok
- Drop point west of the peninsula is supported by Winds from the east at 10 Knots and Approach into the wind preferred
- a visible shoreline is preferred
- Visibility is around 6 miles
- Critical Meteorological data ok is elaborated in Temperature ok and Wave height ok
- Temperature ok is supported by Temperature must be between 90 and 50 F and Gibraltar avg temperature in March is 60 F and Gibraltar temperature 60 F
- Temperature must be between 90 and 50 F
- [more]
- according to source Unofficial Special Operations Site which is usually reliable (E) and confirmed by other sources
- Gibraltar temperature 60 F
- Gibraltar avg temperature in March is 60 F
- Wave height ok is supported by Wave height should be less than 8 ft

At the bottom of the right pane, there are buttons for "Add Unit", "Add Sub-Unit", "Remove", "Edit", "Import", "Import Src", "Restore", and "< Ex".

Fig. 4. An intelligence analyst is using TRELIS to annotate the reasons why he chose a certain drop site and a certain mode of transport in a hypothetical operation.

The overall rating O(s) for a single source is the average of all the ratings of its associated statements:

$$O(s) = 1 / N \sum_{(i=1 \text{ to } N)} O(s, a_i)$$

TRELIS derives and updates these ratings automatically as users enter different analyses that rely on those sources. The next section shows how these ratings are used and shown to the users to help them make decisions about what sources to trust.

5 Helping Users Select Sources

As a user is considering a topic for an analysis, he or she may wonder what sources were considered by other users on topics relevant to their analysis, as well as how those sources were rated by users in light of what they were considering and in light of their expertise on the topic. TRELIS allows users to search for sources on specific topics, see how they rank based on their overall ratings, and view the details of a source's ratings based on the individual factors considered in deriving the ratings.

Figure 4 shows an intelligence analyst reasoning about the choice of a drop site and the mode of transportation for a hypothetical “rubber duck” operation in the Gibraltar area. A rubber duck operation is used by the military to insert SEALs into a target area. A rubber craft is dropped into the sea, which is then used by the SEALs to move towards the target. The intelligence analyst knows sources that provide critical meteorological data for the operation, as well as some manuals that describe the techniques preferred in carrying out the operation. He points to weather sources to indicate whether conditions match the critical meteorological data, and if the preferred techniques in the manual can be carried out.

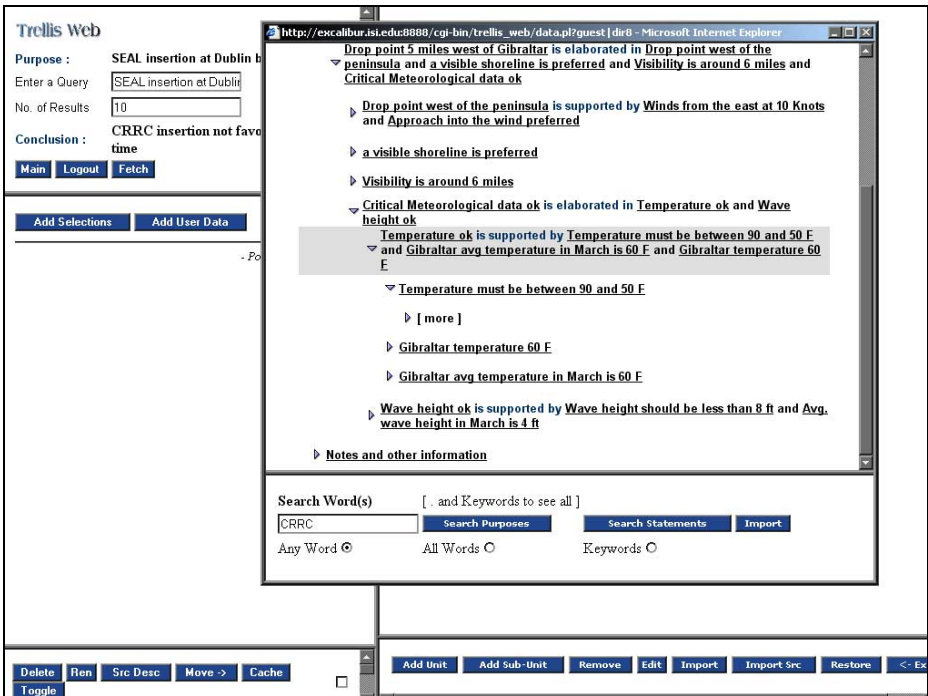


Fig. 5. An intelligence analyst considering whether a hypothetical mission is feasible is browsing the analysis done by another intelligence analyst for a mission that was shown in Figure 4.

Figure 5 shows a relatively inexperienced intelligence analyst trying to analyze the feasibility of a hypothetical “rubber duck” mission to Dublin. The analyst uses the

“Import” button to bring up the Analysis Browser window. He searches for “CRRC” (which is the craft used in a “rubber duck” operation) and finds the analysis of the previous intelligence analyst for a “rubber duck” operation to Gibraltar. He can now select a portion of this analysis and “import” it into his own analysis, or he can simply browse it and figure out what kind of sources will be needed. For example, the analyst finds out that according to a reasonably accurate source, the water temperature should be between 50 F and 90 F for the CRRC to be safe.

Now the analyst needs to find out the average water temperature in the Dublin area. He now invokes the “source query tool” by pressing the ‘Import Src’ button in the bottom right frame of the main window. Figure 6 shows the user query for “temperature” and the results that are returned. TRELIS shows the rating of all sources that are related to the topic (here “temperature”). The analyst then selects the sources that he considers appropriate and imports them to his selection of statements and sources in the ‘Statement Editor’ (bottom-left frame of the main TRELIS window).

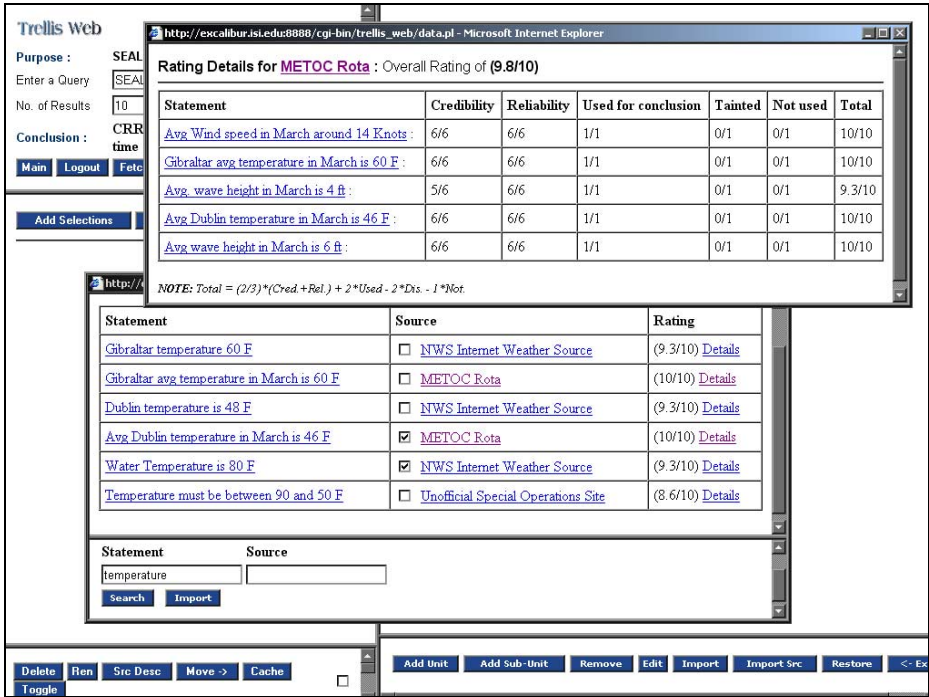


Fig. 6. TRELIS shows users its assessment of a source based on previous analysis by other users, showing both an overall rating and the details about how the rating was derived.

The user can also see further details about the ratings of a source, shown at the top of Figure 6. This shows the detailed factors and ratings of the source for all statements that it has been used with.

6 Related Work

Annotea [15] is a general metadata annotation infrastructure that enables users to comment on resources (also using the Dublin Core as well as domain-specific markup in RDF), share and summarize these comments, and manage reply threads. Annotea's goals and emphasis are complementary to the work on TRELLIS, which has a more specific focus on information analysis for decision making.

Web users are often allowed to offer ratings that can be used by other users in deciding what or who to trust or perhaps prefer [20]. This kind of feedback schemas can be used to rate other users (e.g., [9]), products (e.g., [1]), or any other item of entertainment or informative value (e.g., [11,18]). These ratings capture a user's opinion about an item, but not about its source or producer. They also only have very limited ways for users to articulate their justification. TRELLIS also captures the context of an opinion based on the use of the source.

SEAS [16] uses an alternative approach to support similar kinds of analysis in the military and competitive intelligence arena. Users define argument templates that contain questions to support the argument and an inference structure to derive the conclusion from the answers to the questions. The approach emphasizes the use of shared patterns as well as support for automated inference on demand, and does not capture or generate ratings for the sources used. TRELLIS has more support for assessing sources but does not provide as much support for automation nor domain-specific standard patterns to facilitate sharing.

Other research is aimed to capture argumentation in collaborative environments, where the focus is decision making and consensus within a group by supporting collaboration and sharing of decision rationale [7]. Some tools are also geared to capture decisions and their rationale in specific contexts, such as software or engineering design [21]. They are customized to their domain and are not as general as TRELLIS, but they provide expressive representations to express design tradeoffs adequately.

7 Conclusions

We have shown an approach to capturing assessments of users about their trust on individual information sources as they are deciding whether and how to use information from each source in a specific analysis or decision making process. Users may express a range of opinions, such as dismissing the information in light of other sources, or express the credibility and reliability of the source explicitly or implicitly in their analysis. Our approach captures the trust of individual users based on an actual context of use of the source as well as their expertise on the topic as they go through the analysis. Other users benefit of these indications of trust as they decide which sources to use in new analyses and decisions.

In future work, we would like to use the approach presented as a basis to extend service descriptions with meta-data about whether other customers have found ser-

vices trustworthy or reliable. Current languages for Web services as well as e-commerce applications [4,2,10] support the advertisement of services, but should be extended to allow qualifications of what is advertised with the customer's view on the actual performance of the provider. Our approach could be used to ensure truth in advertising by collecting feedback from individuals together with justifications for their opinions.

Many of the mechanisms necessary in the Web of Trust will be imposed by strict protocols of authentication, accessibility, and attribution. Our work opens the way for individuals to provide their assessment of truth in more subjective way that contributes to create a collective consensus of trust on information sources.

Acknowledgements

This work was funded by the US Air Force Office of Scientific Research (AFOSR) under grant F49620-00-1-0337. We would like to acknowledge Jim Blythe, Jihie Kim, Larry Kite, and Fred Bobbitt for their comments and suggestions about this work.

References

1. Amazon Web Site, <http://www.amazon.com>.
2. Ankolekar, A., Burstein, M., Hobbs, J.R., Lassila, O., Martin, D.L., McIlraith, S.A., Narayanan, S., Paolucci, M., Payne, T.R., Sycara, K., Zeng, H.: DAML-S: Semantic Markup for Web Services. In: International Semantic Web Workshop (2001).
3. Berners-Lee, T., Hendler, J., Fensel, D.: The Semantic Web. In: Scientific American 78(3) (2001): 20-88.
4. Christensen, E., Curbera, F., Meredith, G., Weerawarana, S.: Web Services Description Language (WSDL) 1.1. <http://www.w3.org/TR/2001/NOTE-wsdl-20010315> (2001).
5. CIA Factbook Web Site, <http://www.cia.gov/cia/publications/factbook> (2001).
6. Concierge Web Site, <http://www.concierge.com> (2002).
7. Conklin, J. and Begeman, M., gIBIS: A Hypertext Tool for Exploratory Policy Discussion, ACM Transactions on Office Information Systems, Vol. 6, pp. 303-331 (1988).
8. DCES : The Dublin Core Element Set. <http://dublincore.org/documents/dces/> (1999).
9. eBay Web Site, <http://www.ebay.com>.
10. ebXML Web Site. <http://www.ebxml.org/>.
11. Epinions Web Site. <http://www.epinions.com>.
12. Fodors Web Site. <http://www.fodors.com>.
13. Gil, Y. and Ratnakar, V. : TRELIS: An Interactive Tool for Capturing Information Analysis and Decision Making. Internal Project Report (2001).
14. Janes Web Site. <http://www.janes.com>.
15. Koivunen, M.R. and Swick, R.: Metadata Based Annotation Infrastructure offers Flexibility and Extensibility for Collaborative Applications and Beyond. In: Proceedings of the K-CAP 2001 Workshop on Knowledge Markup and Semantic Annotation, Victoria, British Columbia (2001).

16. Lawrence, J. D., Harrison, I.W., Rodriguez, A. C.: Capturing Analytic Thought. In: Proceedings of the First International Conference on Knowledge Capture (K-CAP)
17. Letsgo Web Site. <http://www.letsgo.com>.
18. Movielens Web Site. <http://www.movielens.umn.edu>.
19. The New York Times Web Site. <http://www.nyt.com>.
20. Sarwar, B M., Karypis, G., Konstan, J.A., Riedl J.: Item-based collaborative filtering recommendation algorithms. In: To appear in Proceedings of the 10th International World Wide Web Conference (WWW10), Hong Kong (2001).
21. Shum, S.B.: Design Argumentation as Design Rationale. In: Encyclopedia of Computer Science and Technology (M.Dekker Inc: NY) (1996).
22. World Wide Web Consortium Web Site. <http://www.w3c.org>.