

Generalized Fact-Finding

Jeff Pasternack

University of Illinois, Urbana-Champaign
201 North Goodwin
Champaign, Illinois 61801
jpaster2@illinois.edu

Dan Roth

University of Illinois, Urbana-Champaign
201 North Goodwin
Champaign, Illinois 61801
danr@illinois.edu

ABSTRACT

Once information retrieval has located a document, and information extraction has provided its contents, how do we know whether we should actually believe it? Fact-finders are a state-of-the-art class of algorithms that operate in a manner analogous to [2]’s Hubs and Authorities, iteratively computing the trustworthiness of an information *source* as a function of the believability of the *claims* it makes, and the believability of a claim as a function of the trustworthiness of those sources asserting it. However, as fact-finders consider only “who claims what”, they ignore a great deal of relevant background and contextual information. We present a framework for “lifting” (generalizing) the fact-finding process, allowing us to elegantly incorporate knowledge such as the confidence of the information extractor and the attributes of the information sources. Experiments demonstrate that leveraging this information significantly improves performance over existing, “unlifted” fact-finding algorithms.

Categories and Subject Descriptors

H.3.3 [Information Systems]: Information Search and Retrieval—*Information filtering*; I.2.m [Computing Methodologies]: Artificial Intelligence—*Miscellaneous*

General Terms

Algorithms, Experimentation, Reliability

Keywords

Fact-finders, Graph Algorithms, Data Integration, Trust

1. INTRODUCTION

Once upon a time, current events came to us via newspapers, television or radio, and the primary repositories of human knowledge were heavy, cumbersome artifacts known as books. While not everything heard or seen could be trusted, the major publishers and broadcasters were nonetheless assumed to perform due diligence, fact-checking their work to ensure accuracy prior to dissemination. However, with the rise of the Internet and especially collaborative media such as wikis, message boards, and blogs, the previous editorial framework has become largely obsolete, with the con-

sequence that we are exposed to far more information, but with much less certainty of its veracity.

If two authors make conflicting claims, such as “Shakespeare was born on April 26th, 1564” and “Shakespeare was born on April 23rd, 1564”, who do we believe? We could take a vote, selecting the claim supported by the largest number of authors, but this relies upon the premise that all authors are equally trustworthy. Fact-finders avoid this assumption by simultaneously finding both the believability of claims and the trustworthiness of sources and, consequently, surpass voting by an often large margin. Still, fact-finders remain relatively ignorant; the only information they consider is whether or not a source makes a claim (with 100% certainty). They cannot, for example, model the case where the information extractor is unsure whether a claim is really asserted in a document, or where the source himself expresses uncertainty in his claim (e.g. “I am 80% sure Shakespeare was born on...”). Some uncertainty is often unavoidable: does “Shakespeare” refer to William Shakespeare, the playwright, John Shakespeare, his father, or Joseph Shakespeare, mayor of New Orleans, etc.? And we may know something about the source himself that may affect his trustworthiness, e.g. he is a member of the Shakespeare Historical Society. Lifted fact-finders are able to encode such data into the fact-finding process, enabling us to make a more comprehensive, more accurate trust decision.

2. FACT-FINDERS

Let S be the set of sources, and C be the set of claims. At each iteration i , a fact-finder calculates the trustworthiness of each source s , $T^i(s)$, in terms of $B^{i-1}(C_s)$ (where $C_s \subseteq C$ are the claims asserted by s) and calculates the believability of each claim c , $B^i(c)$, in terms of $T^i(S_c)$ (where $S_c \subseteq S$ are sources asserting c); the initial belief in each claim is given by the “prior”, $B^0(c)$. Claims may be *mutually exclusive* with one another (e.g. Shakespeare can only have one birthday), and the goal of the fact-finder is to determine which claim (if any) is true for each subset of mutually exclusive claims, such that at the final iteration f , $\arg\max_{c \in M_{\bar{c}}} B^f(c) = \bar{c}$, where \bar{c} is a true claim, and $M_{\bar{c}}$ are all claims mutually exclusive to it. Experimental accuracy is the percentage of mutual exclusion sets in which the true claim is so chosen.

As an example, Sums is a very simple fact-finder derived from [2]’s Hubs and Authorities and defined by: $T^i(s) = \sum_{c \in C_s} B^{i-1}(c)$, $B^i(c) = \sum_{s \in S_c} T^i(s)$, and $B^0(c) = 1$. We also use five other (far more sophisticated) fact-finders in our experiments: TruthFinder [4], 3-Estimates [1], Average-Log, Investment, and PooledInvestment [3].

Table 1: Experimental Results for Lifted Fact-Finders. All values are percent accuracy.

Experiment	Voting	Sums	3-Estimates	TruthFinder	Average-Log	Investment	PooledInvestment
Unlifted	81.49	81.82	81.49	84.42	80.84	87.99	80.19
Tuned Certainty	81.90	82.90	82.20	87.20	83.90	90.00	80.60
Best Certainty	81.82	83.44	82.47	87.66	86.04	90.26	81.49
Group Layer	N/A	84.74	N/A	84.09	84.42	89.61	84.74

3. LIFTED FACT-FINDING

A key observation is that fact-finders operate on an (unweighted) bipartite graph of sources and claims, where an edge between source s and claim c indicates that s asserts c . In lifted fact-finding, we are able to elegantly encode our additional knowledge by augmenting this graph, creating a weighted k-partite graph; all that then remains is enhancing each fact-finder to accept this graph as its input.

3.1 Weight-Encoded Knowledge

In the lifted model, $s \in S$ asserts $c \in C$ with weight $\omega(s, c) = [0, 1]$. A variety of phenomenon can be encoded into a single weight; indeed, we calculate $\omega(s, c)$ as $\omega_u(s, c) \times \omega_p(s, c) + \omega_\sigma(s, c) + \omega_g(s, c)$. $\omega_u(s, c)$ is the probability that s asserted c according to the information extractor (due to inherent ambiguity, OCR error, etc.), while $\omega_p(s, c)$ is the certainty expressed by the source in the claim (e.g. “I’m 70% certain that...”); their product can be viewed as the expected probability of c according to s . Additionally, $\omega_\sigma(s, c)$ encodes similarity among claims (a source objects less to a claim similar to the one he asserts) and $\omega_g(s, c)$ provides an alternate method of encoding source groups and attributes, although we omit further description here for want of space.

3.2 Layer-Encoded Knowledge

We can add new “layers” of groups and attributes to the existing two layers of sources and claims, where the first additional layer connects directly to the sources and higher layers model meta-groups and meta-attributes, making the graph k-partite. In the bipartite case, we had B and T functions; now, given layers $L_{1\dots k}$, we instead have $D_j^i(L_j)$ over $j = 1\dots k - 1$ and $U_j^i(L_j)$ over $j = 2\dots k$, where $D_k^i(L_k) = U_k^i(L_k)$ and $U_1^i(L_1) = D_1^{i-1}(L_1)$. D_j and U_j may vary for each layer; e.g. by using an existing fact-finder for sources and claims and other, novel U and D functions to mediate between sources and groups or attributes. At each i , we find $U_j^i(L_j)$ for layers $j = 2$ to k , then compute $D_j^i(L_j)$ for layers $j = k - 1$ to 1 . For example, Sums readily extends to k layers as: $U_j^i(e) = \sum_{f \in L_{j-1}} \omega(e, f) U_{j-1}^i(f)$ and $D_j^i(e) = \sum_{f \in L_{j+1}} \omega(e, f) D_{j+1}^i(f)$, where $\omega(e, f) = \omega(f, e)$ is the weight between nodes e and f (for sources, groups and attributes, $\omega(e, f) = 1$ if e has group or attribute f or vice-versa, and 0 otherwise).

3.3 Lifting the Fact-Finder Algorithm

Almost any fact-finder can be lifted to take advantage of a weighted graph by applying a small set of rewriting rules to their T and B functions; we omit the details for brevity, but these include: $|S_c| \Rightarrow \sum_{s \in S_c} \omega(s, c)$, $|C_s| \Rightarrow \sum_{c \in C_s} \omega(s, c)$, $B^{i-1}(c) \Rightarrow \omega(s, c) B^{i-1}(c)$, and $T^i(s) \Rightarrow \omega(s, c) T^i(c)$. Applying these rules to Sums gives us $T^i(s) = \sum_{c \in C_s} \omega(s, c) B^{i-1}(c)$ and $B^i(c) = \sum_{s \in S_c} \omega(s, c) T^i(s)$.

4. EXPERIMENTS

We use [3]’s Population dataset consisting of 44,761 claims of city populations (extracted from Wikipedia infoboxes) from 171,171 editors, with 308 true claims identified from census data as an evaluation set.

4.1 Tuned Assertion Certainty

One problem in extracting claims from infoboxes is the question of user intent: if a user modifies a field other than the population field, or somewhere else on the page entirely, does this imply that he saw and approved the population that was already listed? A simple method to account for this is to assign a certainty to each edit location (“population field”, “other field”, and “elsewhere on page”). As our labeled data was limited, we tuned these certainties over 208 randomly-selected true claims and tested on the remaining 100, repeating this 10 times to obtain the “tuned” results across six fact-finders (and basic voting), with substantial improvement over the “unweighted” case, where the unlifted fact-finder is used and only direct “population field” edits are counted. The “best” results come from tuning (and testing) over the entire evaluation set.

4.2 Groups via Additional Layers

Wikipedia editors can be split into three groups: administrators, blocked users, and everyone else. As administrators are elected by the community, they can be expected to have rather high trustworthiness; conversely, blocked users can be expected to be rather untrustworthy. We incorporate these groups as an additional layer, creating a tripartite fact-finding graph. As it is not readily extended multiple layers, 3-Estimates is omitted; however, we find that (with the exception of TruthFinder) using lifted fact-finders with knowledge of the sources’ groups again produced significantly better results than the unlifted variants.

Acknowledgements

This research sponsored by the Army Research Laboratory (ARL) under agreement W911NF-09-2-0053. Any opinions, findings, conclusions or recommendations are those of the authors and do not necessarily reflect the view of the ARL.

5. REFERENCES

- [1] A. Galland, S. Abiteboul, A. Marian, and P. Senellart. Corroborating information from disagreeing views. In *WSDM*, 2010.
- [2] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5):604–632, 1999.
- [3] J. Pasternack and D. Roth. Knowing What to Believe (when you already know something). In *COLING*, 2010.
- [4] X. Yin, P. S. Yu, and J. Han. Truth Discovery with Multiple Conflicting Information Providers on the Web. *IEEE TKDE*, 20(6):796–808, 2008.