S&F: Sources and Facts Reliability Evaluation Method

Extended Abstract

Quentin Elsaesser
CRIL - CNRS
Université d'Artois, France
elsaesser@cril.fr

Patricia Everaere
CRIStAL
Université de Lille, France
patricia.everaere-caillier@univ-lille.fr

Sébastien Konieczny
CRIL - CNRS
Université d'Artois, France
konieczny@cril.fr

ABSTRACT

In this work we propose a family of methods that allow to conjointly compute the reliability of a set of information sources and the reliability of the facts on a set of objects in order to find the truth, by confronting the sources points of view. We use a (scoring-based) voting method for the evaluation of the trust of the sources, using Condorcet's Jury Theorem arguments in order to identify the truth and the reliable sources. We provide an experimental study that shows that we perform better than state of the art methods on the task of finding the truth among the possible facts, but we also show that we can, at the same time, adequately evaluate the reliability (trust) of the sources of information.

KEYWORDS

Reliability; Truth Tracking; Evidence-based Trust Evaluation; Voting

ACM Reference Format:

Quentin Elsaesser, Patricia Everaere, and Sébastien Konieczny. 2023. S&F: Sources and Facts Reliability Evaluation Method: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

There are many applications where (conflicting) information are received from multiple sources and an opinion needs to be formed from this information. In this situation, a standard way of resolving conflicts is to believe the most reliable sources.

Truth Discovery methods aim to resolve these conflicts and find the truth among this information [10, 14, 17, 18]. To achieve this task, these methods follow the idea that trustworthy sources claim believable facts.

In this paper, we propose methods that allow to identify the correct answers, but also to evaluate the reliability (truthfulness) of the sources. It is, as far as we know, the first approach that allows to perform these two tasks conjointly. We also provide an experimental evaluation. The idea is to test if we manage to achieve these tasks of evaluating the reliability of sources and facts in practice.

2 S&F METHODS

We consider three sets $\mathcal{S}$, $\mathcal{F}$ and $\mathcal{O}$ respectively called *Sources*, *Facts* and *Objects* (see Figure 1). *Sources* represent the human or artificial agents that provide the information. *Objects* are the questions on which we would like to have an answer, and the *Facts* are the possible answers. For each object, only one fact can be chosen (facts are distinct and exclusive).

We keep the same vocabulary used in previous works ([16–18]).

Definition 2.1. Let $G = (V, E)$ be a directed graph with $V = \mathcal{S} \cup \mathcal{F} \cup \mathcal{O}$ and $E \subseteq (\mathcal{S} \times \mathcal{F}) \cup (\mathcal{F} \times \mathcal{O})$, such that:

- There cannot be more than one path between a source $s \in \mathcal{S}$ and an object $o \in \mathcal{O}$.
- For each fact $f \in \mathcal{F}$ there is an unique object $o \in \mathcal{O}$ with $(f, o) \in E$.

$(s, f) \in E$ means that the source s claims that the fact f is the correct answer for its corresponding object. It is possible that a fact is not claimed by any source.

We suppose that we initially have no information about the reliability of the sources and we define an iterative procedure to determine their reliability.

At the beginning, we assign the same reliability to all the sources, then we compare the answers to the different questions. In order to find the true information and reward the sources, we rely on the idea of Condorcet's Jury Theorem [5], which states that it is more likely that the majority of the individuals will choose the correct solution. The hypothesis of this theorem (all the individuals have a reliability greater than 0.5) can be more-or-less relaxed ([1–4, 7–9, 11, 13, 15]). More precisely, at each iteration, the sources give strength to the facts they claim on the different objects. With the sum of the obtained strengths, we got the confidence of each fact.

We want to reward the sources that provide pieces of information that are confirmed by others, and then that are more likely true. To reward the sources, the objects take part to a vote where they rank their related facts from most reliable to least reliable ones. We use scoring-based voting rules in order to associate a number to each rank of facts. Then the new reliability of each source is computed by combining all these scores. But, we wish to give the reliability of the source, i.e. an estimation of the probability of this source to find the true facts. So, we have to normalize the reliability of the sources to ensure that this reliability is between 0 and 1. There are at least two ways to normalize the reliability. The first one (normalization A) favors sources that provide the most of correct answers. The reliability of the source is divided by the total number of objects in the graph. The second (normalization C) favors sources that are more careful and do not fail often. The reliability of the source is divided by the number of objects on which it claims a fact.

Then a new iteration begins with the updated reliability of each source. The algorithm stops when the process converges, i.e. when the cosine similarity between the reliability of the sources of the last and the current iteration is smaller than $1 - \epsilon$ where $\epsilon = 0.001$.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

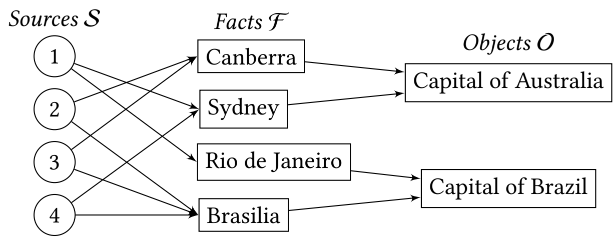

Figure 1: Sources, Facts & Objects

In Figure 1, the majority claims that *Brasilia* is the *Capital of Brazil* but there is a tie for the *Capital of Australia*. The sources that claim *Brazilia* will get a reward for proposing the most popular answer (that is then considered as the most plausible one). Then, at the next iteration, the reliability of *Camberra* will be better than the reliability of *Sydney*.

3 EXPERIMENTAL STUDY

We generated synthetic data sets to be able to perform an experimental evaluation. All the generated graphs are composed of 10 objects and 4 facts by object. For each object, we randomly choose one of the facts to be the true value of this object. This will be our ground truth to evaluate our methods with the metrics. After the generation, we know the *a posteriori* probability of choosing a true fact for all the sources. This value represent the true reliability of these sources. In the tests, we rank the experiments with respect to the average reliability of the sources. This allows us to see what happens when the sources are globally more or less reliable. In the graph, an average reliability of x% means that there is x% of links between sources and true facts. Each point on the graph corresponds to the mean value obtained by generating 1000 graphs.

In figures, PlA and PlC stand respectively for the S&F method with plurality vote and the normalization A or C. Bor$A$ and Bor$C$ correspond to the methods with the Borda rule and the normalization A and C. We compare the results of our methods against related methods of the literature (tf for Truth Finder[18], hna for Hubs and Authorities[10], Sums[14] and usums for Unbounded-Sums[17]).

Facts Credibility - Truth Discovery. We can see in Figure 2 (and the zoom in Figure 3) that the S&F method with the plurality rule and normalization A is better for *precision* [6, 12] that the other methods until they all find 100% of the true facts (when the average reliability is greater than 52%). Note that the results for methods with normalization C are significantly worse despite their natural meaning and justifications. This normalization should thus be avoided if efficiency of *Truth Discovery* is the main concern.

Reliability of the sources. Now we want to evaluate our methods on the task of estimating the reliability of the sources. To our knowledge, there is no other work that computes the reliability of the sources conjointly with the credibility of facts (Truth Discovery). As a consequence, we can not compare our methods with any other method for this task. We will focus on the averaged difference, that is the difference between the computed reliability and the (*a posteriori*) probability of choosing the true fact for every objects. So this distance measures how far the estimated reliability of the sources

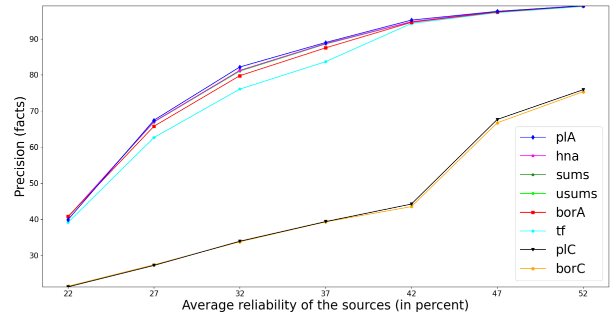

Figure 2: Precision - 10 sources

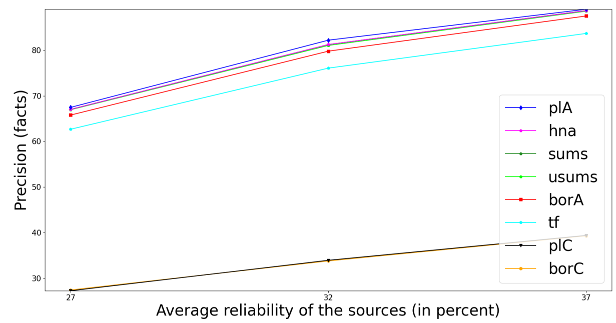

Figure 3: Precision - 10 sources (zoom between 27-37%)

is from the true one (the *a posteriori* probability). We can see in Figure 4 that the computed reliability is close to the *a posteriori* probability. Especially, the plurality rule and the normalization A gets exactly and quickly the true reliability. Because the Borda rule gives points to all the sources, it is impossible to obtain the exact reliability but the results are still good.

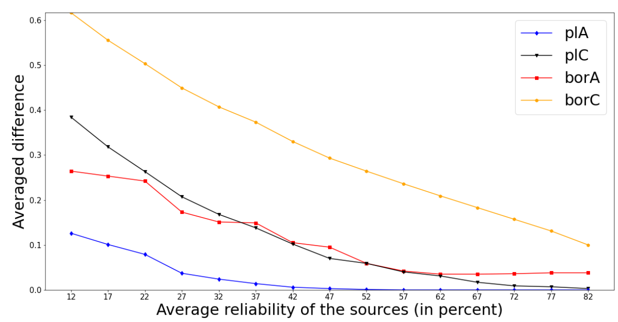

Figure 4: Sources reliability - Averaged difference - 10 sources

4 CONCLUSION

In this paper we have introduced the S&F methods for evaluating the reliability of the sources conjointly to the credibility of the facts in an information-based multi-agent system. We have performed some experimental evaluations and we show that our methods (in particular the methods with normalization A) outperform methods from the literature in identifying the true facts. But we also show that our methods allow to correctly estimate the reliability of the sources. It is, as far as we know, the first approach that allows to perform these two tasks together. We only have space to give here the results on generated benchmarks, but we also test our methods with similar results on two real benchmarks.

REFERENCES

[1] David Austen-Smith and Jeffrey Banks. 1996. Information Aggregation, Rationality, and the Condorcet Jury Theorem. *American Political Science Review* 90 (1996), 34–45.

[2] Ruth Ben-Yashar and Jacob Paroush. 2000. A nonasymptotic Condorcet jury theorem. *Social Choice and Welfare* 17, 2 (2000), 189–199.

[3] Ruth Ben-Yashar and Mor Zahavi. 2011. The Condorcet jury theorem and extension of the franchise with rationally ignorant voters. *Public Choice* 148, 3/4 (2011), 435–443.

[4] Daniel Berend and Jacob Paroush. 1998. When is Condorcet's Jury Theorem valid? *Social Choice and Welfare* 15, 4 (1998), 481–488.

[5] Marquis de Condorcet. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix.* Imprimerie royale Paris.

[6] Ralph Donaldson, Rosemary Dyer, and Michael Kraus. 1975. An objective evaluator of techniques for predicting severe weather events. In *Preprints, Ninth Conference on Severe Local Storms, Norman, Oklahoma, American Meteorological Society*, Vol. 321326.

[7] David Estlund. 1994. Opinion leaders, independence, and Condorcet's Jury Theorem. *Theory and Decision* 36, 2 (1994), 131–162.

[8] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. 2010. The epistemic view of belief merging: can we track the truth?. In *Nineteenth European Conference on Artificial Intelligence (ECAI'10)*. 621–626.

[9] Patrick Hummel. 2010. Jury theorems with multiple alternatives. *Social Choice and Welfare* 34, 1 (2010), 65–103.

[10] Jon M Kleinberg. 1999. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)* 46, 5 (1999), 604–632.

[11] Christian List and Robert E. Goodin. 2001. Epistemic Democracy: Generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy* 9, 3 (2001), 277–306.

[12] David L Olson and Dursun Delen. 2008. *Advanced data mining techniques.* Springer Science & Business Media.

[13] Guillermo Owen, Bernard Grofman, and Scott L. Feld. 1989. Proving a distribution-free generalization of the Condorcet Jury Theorem. *Mathematical Social Sciences* 17, 1 (1989), 1 – 16.

[14] Jeff Pasternack and Dan Roth. 2010. Knowing what to believe (when you already know something). In *Proceedings of the 23rd International Conference on Computational Linguistics (COLING 2010)*. 877–885.

[15] Bezalel Peleg and Shmuel Zamir. 2012. Extending the Condorcet Jury Theorem to a general dependent jury. *Social Choice and Welfare* 39, 1 (2012), 91–125.

[16] Joseph Singleton and Richard Booth. 2020. An axiomatic approach to truth discovery. In *Autonomous Agents and Multi-Agent Systems (AAMAS'20)*. 2011–2013.

[17] Joseph Singleton and Richard Booth. 2022. Towards an axiomatic approach to truth discovery. *Journal of Autonomous Agents and Multi-Agent Systems* 36, 2 (2022), 1–49.

[18] Xiaoxin Yin, Jiawei Han, and Philip S. Yu. 2008. Truth Discovery with Multiple Conflicting Information Providers on the Web. *IEEE Transactions on Knowledge and Data Engineering* 20, 6 (2008), 796–808.