An Axiomatic Approach to Truth Discovery

Paper #1015

ABSTRACT

The problem of *truth discovery*, i.e., of trying to find the true facts concerning a number of objects based on reports from various information sources of unknown trustworthiness, has received increased attention recently. The problem is made interesting by the fact that the relative believability of facts depends on the trustworthiness of their sources, which in turn depends on the believability of the facts the sources report. Several algorithms for truth discovery have been proposed, but their evaluation has mainly been performed experimentally by computing accuracy against large datasets. Furthermore, it is often unclear how these algorithms behave on an intuitive level. We develop a general framework for truth discovery which allows comparison and evaluation of algorithms based instead on their theoretical properties. To do so we pose truth discovery as a social choice problem, and formulate various *axioms* that any reasonable algorithm should satisfy. Along the way we provide an axiomatic characterisation of the baseline 'Voting' algorithm – which leads to an impossibility result showing that a certain combination of the axioms cannot hold simultaneously – and check which axioms some well-known existing algorithms satisfy. We find that, surprisingly, our more fundamental axioms do not hold, and propose modifications to the algorithms to partially fix these problems.

KEYWORDS

Truth discovery; Axioms; Trust and reputation; Social choice theory

1 INTRODUCTION

There is an increasing amount of data available in today's world, particularly from the web, social media platforms and crowdsourcing systems. The openness of such platforms makes it simple for a wide range of users to share information quickly and easily, potentially reaching a wide international audience. It is inevitable that amongst this abundance of data there are *conflicts*, where data sources disagree on the truth regarding a particular object or entity. This can be caused by low-quality sources mistakenly providing erroneous data, or by malicious sources aiming to misinform.

Resolving such conflicts and determining the true facts is therefore an important task. Truth discovery has emerged as a set of techniques to achieve this by considering the *trustworthiness* of sources [11, 14]. The general principle is that true facts are those claimed by trustworthy sources, and trustworthy sources are those that claim believable facts. Application areas include real-time traffic navigation [8], drug side-effect discovery [17], crowdsourcing and social sensing [16, 20, 27].

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 2020, Auckland, New Zealand

https://doi.org/doi

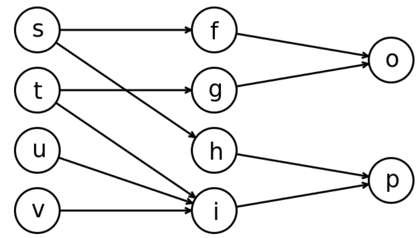

Figure 1: Illustrative example of a dataset to which truth discovery can be applied with sources $\{s, t, u, v\}$, facts $\{f, g, h, i\}$ and objects $\{o, p\}$.

For a simple example of a situation where trust can play an important role, consider Fig. 1 which shows data sources, 'facts' and objects from left to right. Suppose object p represents the question 'how much money does the UK send to the EU per week?', for this particular question has certainly had at least two answers put forward recently, and suppose object o represents another matter such as 'what is the current stock price for Google?'. Without considering trust information, the stock price question seems to be a tie, with both options f and g receiving one vote from sources s and t respectively.

Taking a trust-aware approach, however, we can look beyond object o to consider the *trustworthiness* of s and t. Indeed, when it comes to object p, t agrees with two extra sources u and v, whereas s disagrees with everyone. In principle there could be *hundreds* of extra sources here instead of just two, in which case the effect would be even more striking. We may conclude that s is *less trustworthy* than t. Returning to o, we see that g is supported by a more trustworthy source, and conclude that it should be accepted over f.

Many truth discovery algorithms have been proposed in the literature with a wide range of techniques used, e.g. iterative heuristic-based methods [10, 18], probabilistic models [24], maximum likelihood estimation and optimisation-based methods [15]. It is common for such algorithms to be evaluated empirically by running them against some large dataset for which the true facts are already known; this allows *accuracy* and other metrics to be calculated, and permits comparison between algorithms. This may be accompanied by some theoretical analysis, such as calculating run-time complexity [11], proving convergence of an iterative algorithm [25], or proving convergence to the 'true' facts under certain assumptions on the distribution of source trustworthiness [21, 22].

A limitation of this kind of analysis is that the results apply only to a single algorithm (or class of algorithm) due to the assumptions made. In this work we aim to take a more general view and study truth discovery without reference to any specific methodology or probabilistic framework. To do so we note the similarities between truth discovery and related problems such as judgment aggregation [9], voting theory [30] ranking and recommendation systems [1–3, 19] in which the *axiomatic approach* of social choice has been successfully applied.

In taking the axiomatic approach one aims to formulate *axioms* that encode intuitively desirable properties that an algorithm may possess. The interaction between these axioms can then be studied; typical results include *impossibility results*, where it is shown that a set of axioms cannot hold simultaneously, and *characterisation results*, where it is shown that a set of axioms are uniquely satisfied by a particular algorithm. Such analysis facilitates principled comparison between algorithms based on intuitive behaviour, in contrast to the somewhat opaque accuracy scores that result from empirical evaluation.

With this in mind we develop a general framework for truth discovery in which axioms can be formulated, and go on to give both an impossibility result and an axiomatic characterisation of a baseline voting algorithm. We also analyse an existing truth discovery algorithm, *Sums* [18], with respect to the axioms.

The paper is organised as follows. The next section informally outlines our framework and axioms. The formal definitions for the framework are then given in Sect. 2. Sect. 3 provides examples of truth discovery algorithms from the literature expressed in the framework. In Sect. 4 we formally introduce the axioms and present an impossibility result showing a subset of these cannot all be satisfied simultaneously. The examples of Sect. 3 are then revisited in Sect. 5, where we analyse them with respect to the axioms and propose modifications to satisfy some axioms that fail. Finally we provide closing remarks and directions for future work in Sect. 6.

We note that proofs are omitted from the paper due to space limitations, and are instead made available online at https://bit.ly/2qg7syf

1.1 Overview of the framework and results

In this section we provide a high-level summary of the framework and our key axioms and results. We only aim to sketch the ideas behind the concepts at this stage, and refer the reader to sections 2 and 4 for the formal definitions. This style of presentation is inspired by [3].

The input to a truth discovery problem (termed a *truth discovery network*) is modelled as a directed tripartite graph whose nodes consist of sources, facts and objects. Fig. 1 provides an example, with sources, facts and objects shown from left to right. We model the output of the truth discovery process as a pair $(\sqsubseteq, \preceq)$, where $\sqsubseteq$ is an ordering of the sources and $\preceq$ an ordering of the facts. We take $s_1 \sqsubseteq s_2$ to mean s_2 is ranked *at least as trustworthy* as s_1; similarly $f_1 \preceq f_2$ means f_2 is ranked *at least as believable* as f_1.

The axioms are summarised below. Example networks in which the axioms apply are shown in Fig. 2.

Coherence: The two rankings $\sqsubseteq$ and $\preceq$ should *cohere* with one another in the following sense: if the sources for two facts f_1, f_2 can be paired up in such a way that the sources for f_1 are less trustworthy than those for f_2, then f_2 should be seen as more believable than f_1. Conversely, if the facts for a source s_1 are pairwise less believable than those for s_2, then s_2 should be seen as the more trustworthy source.

Symmetry: The rankings for isomorphic networks should themselves be isomorphic. That is, the rankings depend only on the structure of the network and not on the 'names' of the sources, facts and objects present.

Unanimity and Groundedness: Facts claimed by all sources should rank maximally (Unanimity), and facts claimed by no sources should rank minimally (Groundedness).

Monotonicity: Increasing support for a fact f by adding a claim from a new source should increase f's ranking relative to other facts.

Independence axioms: The rankings of two sources s_1, s_2 should not be affected by 'irrelevant' details of the network that do not concern s_1 or s_2 (and similarly for the rankings of facts).

Independence actually covers a range of axioms, with different notions of 'irrelevance' leading to different axioms. In this paper we consider three independence axioms, first using irrelevance criteria inspired by equivalent axioms in social choice, and then adopting a truth discovery-specific approach. We will show that the first two of the resulting axioms lead to undesirable behaviour with regards to the fact rankings, and are not compatible with Coherence.

2 A FRAMEWORK FOR TRUTH DISCOVERY

With an overview given in the previous section, we now define our formal framework, which provides the key definitions required for theoretical discussion and analysis of truth discovery methods.

We consider fixed finite and mutually disjoint sets $\mathcal{S}$, $\mathcal{F}$ and $\mathcal{O}$ throughout, called the *sources*, *facts* and *objects* respectively. All definitions and axioms will be stated with respect to these sets.

2.1 Truth discovery networks

A core definition of the framework is that of a *truth discovery network*, which represents the input to a truth discovery problem. We model this as a tripartite graph with certain constraints on its structure, in keeping with approaches taken throughout the truth discovery literature [11, 24].

Definition 2.1. A *truth discovery network* (hereafter a *TD network*) is a directed graph $N = (V, E)$ where $V = \mathcal{S} \cup \mathcal{F} \cup \mathcal{O}$, and $E \subseteq (\mathcal{S} \times \mathcal{F}) \cup (\mathcal{F} \times \mathcal{O})$ has the following properties:

1. For each $f \in \mathcal{F}$ there is a unique $o \in \mathcal{O}$ with $(f, o) \in E$, denoted $\text{obj}_N(f)$. That is, each fact is associated with exactly one object.
2. For $s \in \mathcal{S}$ and $o \in \mathcal{O}$, there is at most one directed path from s to o. That is, sources cannot claim multiple facts for a single object.
3. $(\mathcal{S} \times \mathcal{F}) \cap E$ is non-empty. That is, at least one claim is made.

We will say that s *claims* f when $(s, f) \in E$. Let $\mathcal{N}$ denote the set of all TD networks.

Note that there is no requirement that a source makes a claim for *every* object, or even that a source makes any claims at all. This reflects the fact that truth discovery datasets are in practise extremely sparse, i.e. each individual source makes few claims. Conversely, we allow for facts that receive no claims from any sources.

Also note that the object associated with a fact f is not fixed across all networks. This is because we view facts as *labels* for information that sources may claim, not the pieces of information themselves. Similarly, we consider objects simply as labels for real-world entities. Whilst a particular piece of information has a fixed entity to which it pertains, the labels do not. For example, when

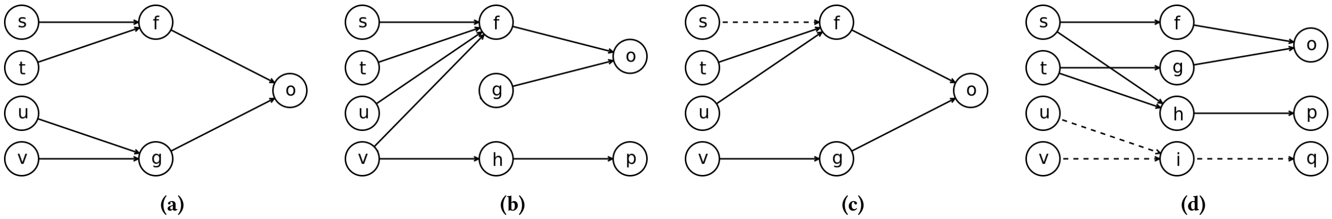

Figure 2: Truth discovery networks where our axioms may be applied. a) Coherence: if $s \sqsubseteq u$ and $t \sqsubset v$, then Coherence dictates $f \prec g$. b) Unanimity and Groundedness: f must rank maximally and g must rank minimally. c) Monotonicity: If $g \preceq f$ before the dashed line is added, we must have $g \prec f$ afterwards. d) Independence: the addition or removal of the dashed edges should not affect the relative rankings of s and t, nor the rankings of f, g and h.

implementing truth discovery algorithms in practise it is common to assign integer IDs to the 'facts' and 'objects'; the algorithm then operates using only the integer IDs. In this case there is no reason to require that fact 17 is always associated with object 4, for example, and the same principle applies in our framework.

A special case of our framework is the binary case in which every object has exactly two associated facts. This brings us close to the setting studied in *judgment aggregation* [9] and, specifically (since sources do not necessarily claim a fact associated to every object) to the setting of *binary aggregation with abstentions* [6, 7].

In a network N, the implicit assumption is that, for each object o, one of the facts in $\text{obj}_N^{-1}(o)$ is the *true* fact associated with that object. However we do not assume any *constraints* on the possible configurations of true facts across *different* objects. That is, any combination of facts is feasible. In judgment aggregation such an assumption has the effect of neutralising the impossibility results that arise in that domain (see, e.g., [6]). We shall see that that is not the case in our setting.

To simplify the notation in what follows, for a network $N = (V, E)$ we write $\text{facts}_N(s) = \{f \in \mathcal{F} : (s, f) \in E\}$ for the set of facts claimed by a source s, and $\text{src}_N(f) = \{s \in \mathcal{S} : (s, f) \in E\}$ for the set of sources claiming a fact f.

2.2 Truth discovery operators

Having defined the input to a truth discovery problem, the output must be defined. Contrary to many approaches across the existing truth discovery literature, we consider the output to be *rankings* of the sources and facts. This is because we are primarily interested in *ordinal properties* rather than quantitative values. Indeed, for the theoretical analysis we wish to perform it is only important that a source is *more trustworthy* than another; the particular numeric scores produced by an algorithm are irrelevant.

Moreover, the scores produced by existing algorithm often have no semantic meaning [18], and so referring to numeric values is not meaningful when comparing across algorithms. In this case it is only the rankings of sources and facts that can be compared, which is further motivation for our choice. This point of view is also common across the social choice literature.

For a set X, let $\mathcal{L}(X)$ denote the set of all total preorders on X, i.e. the set of transitive, reflexive and complete binary relations on X. We now define a *truth discovery operator* as a mapping from the space of inputs to outputs.

Definition 2.2. An *ordinal truth discovery operator* T (hereafter *TD operator*) is a mapping $T : \mathcal{N} \to \mathcal{L}(\mathcal{S}) \times \mathcal{L}(\mathcal{F})$. We shall write $T(N) = (\sqsubseteq_N^T, \preceq_N^T)$, i.e. $\sqsubseteq_N^T$ is a total preorder on $\mathcal{S}$ and $\preceq_N^T$ is a total preorder on $\mathcal{F}$.

Intuitively, the relation $\sqsubseteq_N^T$ is a measure of *source trustworthiness* in the network N according to T, and $\preceq_N^T$ is a measure of *fact believability*. The notation $\sqsubset_N^T$ and $\simeq_N^T$ will be used to denote the strict order and symmetric closure induced by $\sqsubseteq_N^T$ respectively. For fact rankings, $\prec_N^T$ and $\approx_N^T$ are defined similarly.

Whilst rankings of sources and facts are a core component of our framework, it is true that many existing truth discovery algorithms do not produce rankings directly, but instead assign each source a numeric *trust score* and each fact a *belief score* [10, 18, 24, 26, 28, 29]. This is captured in the following definition.

Definition 2.3. A *numerical TD operator* is a mapping $T : \mathcal{N} \to \mathbb{R}^{\mathcal{S} \cup \mathcal{F}}$, i.e. T assigns to each TD network N a function $T(N) = T_N : \mathcal{S} \cup \mathcal{F} \to \mathbb{R}$. For $s \in \mathcal{S}$, $T_N(s)$ is the *trust score* for s in the network N according to T; for $f \in \mathcal{F}$, $T_N(f)$ is the *belief score* for f. The set of all numerical TD operators will be denoted by $\mathcal{T}_{Num}$.

Note that any numerical operator T naturally induces an ordinal operator $\hat{T}$, where $s_1 \sqsubseteq_N^{\hat{T}} s_2$ iff $T_N(s_1) \leq T_N(s_2)$, and $f_1 \preceq_N^{\hat{T}} f_2$ iff $T_N(f_1) \leq T_N(f_2)$. Henceforth we shall write $\sqsubseteq_N^T$, $\preceq_N^T$ without explicitly defining the induced ordinal operator $\hat{T}$.

It is worth noting that yet other truth discovery algorithms output neither rankings nor numeric scores for facts, but only a single 'true' fact for each object [15]. This is also the approach taken in judgment aggregation, where an *aggregation rule* selects which formulas are to be taken as true. Such algorithms can be modelled in our framework as numerical operators where $T_N(f) = 1$ for each identified 'true' fact f, and $T_N(g) = 0$ for other facts g. To go in the reverse direction and obtain the 'true' facts according to an operator, one may simply select the set of facts for each object that rank maximally.

3 EXAMPLES OF TRUTH DISCOVERY OPERATORS

Our framework is general enough to capture a wide range of operators that have been proposed in the truth discovery literature. In this section we provide two concrete examples: *Voting*, which is a simple approach commonly used as a baseline method, and *Sums*

[18]. We go on to outline the class of *recursive operators* – of which *Sums* is an instance – which contains many more examples from the literature.

3.1 Voting

In *Voting*, we consider each source to cast 'votes' for the facts they claim, and facts are ranked according to the number of votes received. Clearly this method disregards the source trustworthiness aspect of truth discovery, as a vote from one source carries as much weight as a vote from any other. As such, *Voting* cannot be considered a serious contender for truth discovery. It is nonetheless useful as a simple baseline method against which to compare more sophisticated operators.

Definition 3.1. *Voting* is the numerical operator defined as follows: for any network $N \in \mathcal{N}$, $s \in \mathcal{S}$ and $f \in \mathcal{F}$, $T_N(s) = 1$ and $T_N(f) = |\mathrm{src}_N(f)|$.

Consider the network N shown in Fig. 1. Facts f, g and h each receive one vote, whereas i receives 3. The fact ranking induced by *Voting* is therefore $f \approx g \approx h \prec i$. On the other hand, all sources receive a trust score of 1 and therefore rank equally.

3.2 Sums

Sums [18] is a simple and well-known operator adapted from the *Hubs and Authorities* [12] algorithm for ranking web pages. The algorithm operates iteratively and recursively, assigning each source and fact a score for each $n \in \mathbb{N}$; the final score is taken as the *limit* of this sequence of scores.

Initially, scores are fixed at a constant value of 1/2. The trust score for each source is then updated by summing the belief score of its associated facts. Similarly, belief scores are updated by summing the trust scores of the associated sources. To prevent these scores from growing without bound as the algorithm iterates, the authors normalise at each iteration by dividing each trust score by the maximum across all sources (belief scores are normalised similarly).

Expressed in our framework, we have that if T is the (numerical) operator giving the scores at iteration n, then the pre-normalisation scores at iteration $n + 1$ are given by T', where

$$T'_N(s) = \sum_{f \in \mathrm{facts}_N(s)} T_N(f); \quad T'_N(f) = \sum_{s \in \mathrm{src}_N(f)} T'_N(s) \tag{1}$$

Consider again the network N shown in Fig. 1. It can be shown that, with T denoting the limiting scores from *Sums* with normalisation, we have $T_N(s) = 0$, $T_N(t) = 1$, and $T_N(u) = T_N(v) = \sqrt{2}/2$. The induced ranking of sources is therefore $s \sqsubset u \simeq v \sqsubset t$.

For fact scores, we have $T_N(f) = 0$, $T_N(g) = \sqrt{2} - 1$, $T_N(h) = 0$ and $T_N(i) = 1$, so the ranking is $f \approx h \prec g \prec i$. Note that fact g fares better under *Sums* than *Voting*, due to its association with the highly-trusted source t.

3.3 Recursive truth discovery operators

The iterative and recursive aspect of *Sums* is hoped to result in the desired mutual dependence between trust and belief scores: namely that sources claiming high-belief facts are seen as trustworthy, and vice versa. In fact, this recursive approach is near universal across the truth discovery literature (see for instance [8, 10, 15, 23, 28, 29]). As such it is appropriate to identify the class of *recursive operators* as an important subset of $\mathcal{T}_{Num}$. To make a formal definition we first define an *iterative operator.*

Definition 3.2. An *iterative operator* is a sequence $(T^n)_{n \in \mathbb{N}}$ of numerical operators. An iterative operator is said to *converge* to a numerical operator T^* if $\lim_{n \to \infty} T^n_N(z) = T^*_N(z)$ for all networks N and $z \in \mathcal{S} \cup \mathcal{F}$. In such case the iterative operator can be identified with the ordinal operator induced by its limit T^*.

Note that it is possible that an iterative operator $(T^n)_{n \in \mathbb{N}}$ converges for only a subset of networks. In such case we can consider $(T^n)_{n \in \mathbb{N}}$ to converge to a 'partial operator' and identify it with the induced partial ordinal operator; that is, a partial function $\mathcal{N} \to \mathcal{L}(\mathcal{S}) \times \mathcal{L}(\mathcal{F})$.

Recursive operators can now be defined as those iterative operators where T^{n+1} can be obtained from T^n alone.

Definition 3.3. An iterative operator $(T^n)_{n \in \mathbb{N}}$ is said to be *recursive* if there is a function $U : \mathcal{T}_{Num} \to \mathcal{T}_{Num}$ such that $T^{n+1} = U(T^n)$ for all $n \in \mathbb{N}$.

In this context the mapping $U : \mathcal{T}_{Num} \to \mathcal{T}_{Num}$ is called the *update function*, and the initial operator T^1 is called the *prior operator.* For a prior operator T and update function U, we write $\mathrm{rec}(T, U)$ for the associated recursive operator; that is, $T^1 = T$ and $T^{n+1} = U(T^n)$.

Returning to *Sums*, we see that (1) defines a mapping $\mathcal{T}_{Num} \to \mathcal{T}_{Num}$ and consequently an update function U^{Sums}. The normalisation step can be considered a separate update function norm which maps any numerical operator T to T', where[1]

$$T'_N(s) = \frac{T_N(s)}{\max_{x \in \mathcal{S}} |T_N(x)|}, \quad T'_N(f) = \frac{T_N(f)}{\max_{y \in \mathcal{F}} |T_N(y)|}$$

It can then be seen that *Sums* is the recursive operator $\mathrm{rec}(T^{Fixed}, \mathrm{norm} \circ U^{Sums})$, where $T^{Fixed}_N \equiv 1/2$.

Many other existing algorithms proposed in the literature can also be realised as recursive operators in the framework, such as *Investment*, *PooledInvestment* [18], *TruthFinder* [24], LDT [28] and others. For brevity, we shall not define any more here.

4 AXIOMS FOR TRUTH DISCOVERY

Having laid out the formal framework, we now introduce axioms for truth discovery. Each axiom encodes a theoretical property that we believe any 'reasonable' operator T should satisfy. Many axioms are adapted from the social choice literature, although modifications are necessary in places to match the semantics of truth discovery. Afterwards, we shall revisit the specific operators of the previous section to check which axioms are satisfied.

4.1 Coherence

As mentioned previously, a guiding principle of truth discovery is that sources claiming highly believed facts should be seen as trustworthy, and that facts backed by highly trusted sources should be seen as believable.

[1] If $\max_{x \in \mathcal{S}} |T_N(x)| = 0$ then the above is ill-defined; we set $T'_N(s) = 0$ for all s in this case. Fact belief scores are defined similarly if the maximum is 0.

Whilst this intuition is hard to write down formally, it is possible to do so in particular cases where there are obvious means by which to compare the set of facts for two sources (and vice versa). This situation is considered in the axiomatic analysis of ranking and reputation systems under the name *transitivity* [2, 19], and we shall adapt it to truth discovery presently. First, some preliminary definitions are required.

Definition 4.1. Let T be a TD operator, N be a TD network and $Y, Y' \subseteq \mathcal{F}$. We shall say Y is *less believable* than Y' with respect to N and T if there is a bijection $\phi : Y \to Y'$ such that $f \preceq_N^T \phi(f)$ for each $f \in Y$, and $\hat{f} \prec_N^T \phi(\hat{f})$ for some $\hat{f} \in Y$.

For $X, X' \subseteq \mathcal{S}$ we define X *less trustworthy* than X' with respect to N and T in a similar way.

In plain English, Y less believable than Y' means that the facts in each set can be paired up in such a way that each fact in Y' is at least as believable as its counterpart in Y, and at least one fact in Y' is strictly more believable. Now, consider a situation where $\text{facts}_N(s_1)$ is less believable than $\text{facts}_N(s_2)$. In this case the intuition outlined above tells us that s_2 provides 'better' facts, and should thus be seen as more trustworthy than s_1. A similar idea holds if $\text{src}_N(f_1)$ is less trustworthy than $\text{src}_N(f_2)$ for some facts f_1, f_2. We state this formally as our first axiom.

Axiom 1 (Coherence). *For any network N, $\text{facts}_N(s_1)$ less believable than $\text{facts}_N(s_2)$ implies $s_1 \sqsubset_N^T s_2$, and $\text{src}_N(f_1)$ less trustworthy than $\text{src}_N(f_2)$ implies $f_1 \prec_N^T f_2$.*

Coherence can be broken down into two sub-axioms: *Source-Coherence*, where the first implication regarding source rankings is satisfied; and *Fact-Coherence*, where the second implication is satisfied. We take Coherence to be a fundamental axiom for TD operators.

4.2 Symmetry

Our next axiom requires that rankings of sources and facts should not depend on their 'names', but only on the structure of the network. To state it formally, we need a notion of when two networks are essentially the same but use different names.

Definition 4.2. Two TD networks N and N' are *equivalent* if there is a graph isomorphism π between them that preserves sources, facts and objects, i.e., $\pi(s) \in \mathcal{S}$, $\pi(f) \in \mathcal{F}$ and $\pi(o) \in \mathcal{O}$ for all $s \in \mathcal{S}$, $f \in \mathcal{F}$ and $o \in \mathcal{O}$. In such case we write $\pi(N)$ for N'.

Axiom 2 (Symmetry). *Let N and $N' = \pi(N)$ be equivalent networks. Then for all $s_1, s_2 \in \mathcal{S}$, $f_1, f_2 \in \mathcal{F}$, we have $s_1 \sqsubseteq_N^T s_2$ iff $\pi(s_1) \sqsubseteq_{N'}^T \pi(s_2)$ and $f_1 \preceq_N^T f_2$ iff $\pi(f_1) \preceq_{N'}^T \pi(f_2)$.*

In the theory of voting in social choice, Symmetry as above is expressed as two axioms: *Anonymity*, where output is insensitive to the names of voters, and *Neutrality*, where output is insensitive to the names of alternatives [30]. Analogous axioms are also used in judgment aggregation.

Correspondingly, Symmetry can by broken down into sub-axioms where the above need only hold for a subset of permutations π satisfying some condition: *Source-Symmetry* (where π must leave facts and objects fixed) and *Fact-Symmetry* (where π leaves sources and objects fixed). For truth discovery we have the additional notion of objects and thus *Object-Symmetry* can defined similarly.

4.3 Fact ranking axioms

Next, we introduce axioms that dictate the ranking of particular facts in cases where there is an 'obvious' ordering. *Unanimity* and *Groundedness* express the idea that if all sources are in agreement about the status of a fact, then an operator should respect this in its verdict. Two obvious ways in which sources can be in agreement are when *all* sources believe a fact is true, and when *none* believe a fact is true.

Axiom 3 (Unanimity). *Suppose $N \in \mathcal{N}$, $f \in \mathcal{F}$, and $\text{src}_N(f) = \mathcal{S}$. Then for any other $g \in \mathcal{F}$, $g \preceq_N^T f$.*

Axiom 4 (Groundedness). *Suppose $N \in \mathcal{N}$, $f \in \mathcal{F}$, and $\text{src}_N(f) = \emptyset$. Then for any other $g \in \mathcal{F}$, $f \preceq_N^T g$.*

That is, f cannot do better than to be claimed by all sources when T satisfies Unanimity, and cannot do worse than to be claimed by none when T satisfies Groundedness.

Unanimity here is a truth discovery rendition of the same axiom in judgment aggregation, and can also be compared to the *weak Paretian* property in voting [5]. Groundedness is a version of the same axiom studied in the analysis of collective annotation [13].

The next axiom is a basic monotonicity property, which states that if f receives extra support from a new source s, then its ranking should receive a strictly positive boost.

Axiom 5 (Monotonicity). *Suppose $N \in \mathcal{N}$, $s \in \mathcal{S}$, $f \in \mathcal{F} \setminus \text{facts}_N(s)$. Write E for the set of edges in N, and let N' be the network in which s claims f; i.e. the network with edge set*

$$E' = \{(s, f)\} \cup E \setminus \{(s, g) : g \neq f, \text{obj}_N(g) = \text{obj}_N(f)\}$$

Then for all $g \neq f$, $g \preceq_N^T f$ implies $g \prec_{N'}^T f$.

Note that Monotonicity does not imply anything about the ranking of the source s.

4.4 Independence axioms

An important idea in social choice is that of *independence*. In voting, this takes the form of Independence of Irrelevant Alternatives (IIA) [4], which requires that the ranking of two alternatives A and B depends only on the individual assessments of A and B, not on some 'irrelevant' alternative C. That is, if the voter preferences are changed such that the individual rankings of A versus B remain unchanged, the social ranking of A and B should remain unchanged.

To translate this principle into an axiom for truth discovery, we need to decide which properties of a network N should be considered relevant to the ranking of two facts (or two sources). There is no canonical choice here, since the role of objects is unique to truth discovery and can be handled in various ways.

We start by considering the case where facts f_1, f_2 relate to the same object o. If one aims to construct an object-aware version of IIA, it is reasonable to suggest that only the other facts for o, and the sources claiming them, are relevant to the ranking of f_1 and f_2. This leads to the following axiom.

Axiom 6 (Per-object Independence (POI)). *Let $o \in \mathcal{O}$. Suppose N_1, N_2 are networks such that $F_o = \text{obj}_{N_1}^{-1}(o) = \text{obj}_{N_2}^{-1}(o)$ and $\text{src}_{N_1}(f) = \text{src}_{N_2}(f)$ for each $f \in F_o$. Then the restrictions of $\preceq_{N_1}^T$ and $\preceq_{N_2}^T$ to F_o are equal; that is, $f_1 \preceq_{N_1}^T f_2$ iff $f_1 \preceq_{N_2}^T f_2$ for all $f_1, f_2 \in F_o$.*

Whilst this axiom would make sense in a voting-like context where there are no dependencies between objects, one may question whether it is truly a desirable property for truth discovery. Indeed, the intuition behind our motivating example in Sect. 1 was that the ranking of facts for object o was decided by the trustworthiness of s and t, which in turn were decided based on their claims for object p. However, POI excludes the possibility of using trust information from claims for other objects.

It appears then that POI may force an operator to discard the source-trustworthiness aspect of truth discovery, as we have seen *Voting* does. It is clear that *Voting* satisfies POI, and in fact there is also a relationship the other way around: POI *forces Voting*-like behaviour within $\mathrm{obj}_N^{-1}(o)$ for each $o \in \mathcal{O}$, when combined with our less controversial requirements of Symmetry and Monotonicity. We note that, for the special case of binary networks, similar results have been shown in the literature on binary aggregation with abstentions [6].

Theorem 4.3. *Let T be any operator satisfying Symmetry, Monotonicity and POI. Then for any $N \in \mathcal{N}$, $o \in \mathcal{O}$ and $f_1, f_2 \in \mathrm{obj}_N^{-1}(o)$ we have $f_1 \preceq_N^T f_2$ iff $|\mathrm{src}_N(f_1)| \le |\mathrm{src}_N(f_2)|$.*

Proof (sketch). We will sketch the main ideas of the proof here with some technical details omitted. Let N be a network, o be an object and $f_1, f_2 \in \mathrm{obj}_N^{-1}(o)$. Consider modifying N by removing all claims for objects other than o. By POI, we have $f_1 \preceq_N^T f_2$ iff $f_1 \preceq_{N'}^T f_2$. Since $|\mathrm{src}_N(f_j)| = |\mathrm{src}_{N'}(f_j)|$ also ($j \in \{1,2\}$), it is sufficient for the proof to show that $f_1 \preceq_{N'}^T f_2$ iff $|\mathrm{src}_{N'}(f_1)| \le |\mathrm{src}_{N'}(f_2)|$.

For the 'if' direction, first suppose $|\mathrm{src}_{N'}(f_1)| = |\mathrm{src}_{N'}(f_2)|$. Let π be the permutation which swaps f_1 with f_2 and swaps each source in $\mathrm{src}_{N'}(f_1)$ with one in $\mathrm{src}_{N'}(f_2)$; then we have $\pi(N') = N'$, and Symmetry of T gives $f_1 \approx_{N'}^T f_2$. In particular $f_1 \preceq_{N'}^T f_2$ as required.

Otherwise, $|\mathrm{src}_{N'}(f_2)| - |\mathrm{src}_{N'}(f_1)| = k > 0$. Consider N'' where k sources from $\mathrm{src}_{N'}(f_2)$ are removed, and all other claims remain. By Symmetry as above, $f_1 \approx_{N''}^T f_2$. Applying Monotonicity k times we can produce N' from N'' and get $f_1 \prec_{N'}^T f_2$ as desired.

For the 'only if' statement, suppose $f_1 \preceq_{N'}^T f_2$ but, for contradiction, $|\mathrm{src}_{N'}(f_1)| > |\mathrm{src}_{N'}(f_2)|$. Applying Monotonicity again as above gives $f_1 \succ_{N'}^T f_2$ and the required contradiction. □

Recall that Coherence formalises the idea that source-trustworthiness should inform the fact ranking, and vice versa. Clearly *Voting* does not conform to this idea, and in fact even the object-wise voting patterns in Thm. 4.3 are incompatible with Coherence. This can easily be seen in the network in Fig. 1 where, regarding object p, we have $|\mathrm{src}_N(h)| < |\mathrm{src}_N(i)|$ (hence $h \prec_N^T i$) and, regarding object o, we have $|\mathrm{src}_N(f)| = |\mathrm{src}_N(g)|$ (hence $f \approx_N^T g$). Hence $\mathrm{facts}_N(s)$ is less believable than $\mathrm{facts}_N(t)$. If Coherence held this would give $s \sqsubset_N^T t$, but then $\mathrm{src}_N(f)$ is less trustworthy than $\mathrm{src}_N(g)$, giving $f \prec_N^T g$ – a contradiction. From this discussion and Thm. 4.3 we obtain as a corollary the following first impossibility result for truth discovery.

Theorem 4.4. *There is no TD operator satisfying Coherence, Symmetry, Monotonicity and POI.*

Given that Thm. 4.3 characterises the fact ranking of *Voting* for facts relating to a single object, it is natural to ask if there is a stronger form of independence that guarantees this behaviour across *all* facts. As our next result shows, the answer is *yes*, and the necessary axiom is obtained by ignoring the role of objects altogether for fact ranking.

Axiom 7 (Strong Independence). *For any networks N_1, N_2 and facts f_1, f_2, if $\mathrm{src}_{N_1}(f_j) = \mathrm{src}_{N_2}(f_j)$ for each $j \in \{1,2\}$ then $f_1 \preceq_{N_1}^T f_2$ iff $f_1 \preceq_{N_2}^T f_2$.*

That is, the ranking of two facts f_1 and f_2 is determined solely by the sources claiming f_1 and f_2. In particular, the fact-object affiliations and claims for facts other than f_1, f_2 are irrelevant when deciding on f_1 versus f_2. We have the following result.

Theorem 4.5. *Suppose $|\mathcal{O}| \ge 3$. Then an operator T satisfies Strong Independence, Monotonicity and Symmetry if and only if for any network N and $f_1, f_2 \in \mathcal{F}$ we have*

$$f_1 \preceq_N^T f_2 \iff |\mathrm{src}_N(f_1)| \le |\mathrm{src}_N(f_2)|$$

The proof of Thm. 4.5 is similar in spirit to that of Thm. 4.3, but uses a different transformation to obtain a modified network N' in the first step.

Clearly neither POI nor Strong Independence are satisfactory axioms for truth discovery, and the notion of 'irrelevance' within a network needs to be reconsidered. Fig. 1 can help us once again in this regard. Whereas POI and Strong Independence would say that facts h and i are irrelevant to f, the argument with Coherence for Thm. 4.4 suggests otherwise due the indirect links via the sources. We therefore propose that only when there is no (undirected) path between two nodes can we consider them to be truly irrelevant to each other. That is, nodes are relevant to each other iff they lie in the same *connected component* of the network.[2]

Our final rendering of independence states that the ordering of two facts in the same connected component does not depend on any claims outside of the component, and similarly for sources.

Axiom 8 (Independence). *For any TD networks N_1, N_2 with a common connected component G, the restrictions of $\sqsubseteq_{N_1}^T$ and $\sqsubseteq_{N_2}^T$ to $G \cap \mathcal{S}$ are equal, and the restrictions of $\preceq_{N_1}^T$ and $\preceq_{N_2}^T$ to $G \cap \mathcal{F}$ are equal; that is, $s_1 \sqsubseteq_{N_1}^T s_2$ iff $s_1 \sqsubseteq_{N_2}^T s_2$ and $f_1 \preceq_{N_1}^T f_2$ iff $f_1 \preceq_{N_2}^T f_2$ for $s_1, s_2 \in G \cap \mathcal{S}$ and $f_1, f_2 \in G \cap \mathcal{F}$.*

In analogy with Source/Fact Coherence and Source/Fact Symmetry, it is possible to split the two requirements of Independence into sub-axioms Source-Independence and Fact-Independence.

5 SATISFACTION OF THE AXIOMS

With the axioms formally defined, we can now consider whether they are satisfied by the example operators of Sect. 3. It will be seen that neither *Voting* nor *Sums* satisfy all our desirable axioms, but it is possible to modify each operator to gain some improvement with respect to the axioms.

[2] We have found that existing datasets with multiple connected components do indeed occur in practise; for example the *Book* and *Restaurant* datasets found at the following web page each contain two connected components: http://lunadong.com/fusionDataSets.htm

5.1 Voting

Being the simplest operator, we will consider *Voting* first. The following theorem shows that all axioms except Coherence are satisfied. Since Coherence is a fundamental principle of truth discovery, and we actually consider POI and Strong Independence to be *undesirable*, this rules out *Voting* as a viable operator.

Theorem 5.1. Voting *satisfies Symmetry, Unanimity, Groundedness, POI, Strong Independence, Independence and Monotonicity.* Voting *does not satisfy Coherence.*

Note that once Symmetry, Monotonicity and POI are shown, the fact that *Voting* fails Coherence follows from our impossibility result (Thm. 4.4), and Fig. 1 serves as an explicit counterexample.

5.2 Sums

Next we look towards *Sums*. Coherence and the simpler axioms are satisfied here, and the undesirable independence axioms are not. However, Monotonicity and Independence do *not* hold. Since Independence is one of our most important axioms that we expect any reasonable operator to satisfy, this potentially limits the usefulness of *Sums* in practise.

Theorem 5.2. Sums *satisfies Coherence, Symmetry, Unanimity and Groundedness.* Sums *does not satisfy POI, Strong Independence, Independence or Monotonicity.*

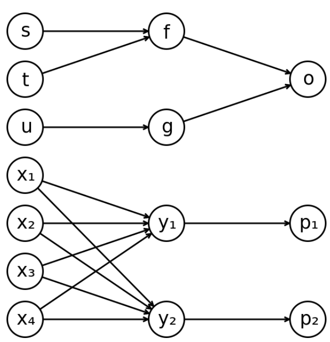

Figure 3: Network which yields counterexamples for POI, Strong Independence, Independence and Monotonicity for *Sums*.

Fig. 3 shows a network from which counterexamples for POI, Strong Independence, Independence and Monotonicity can be obtained. In this network the lower connected component is 'dense', in the sense that each source x_i claims all facts in the component, and each fact y_j is claimed by all sources in the component. Moreover, sources elsewhere in the network claim fewer facts than the x_i, and facts elsewhere are claimed by fewer sources than the y_j.

Since scores are obtained by a simple sum, this results in the scores for the x_i and y_j dominating those of the other sources and facts. The normalisation step then divides these scores by the (comparatively large) maximum. As the next result shows, under certain conditions this causes scores to decrease *exponentially* and become 0 in the limit. In particular, we can generate examples such as Fig. 3 where a whole connected component receives scores of 0 in the limit, which leads to failure of Monotonicity and the independence axioms.

Lemma 5.3. *Let N be a network. Suppose there is $X \subseteq \mathcal{S}$, $Y \subseteq \mathcal{F}$ such that* $\text{facts}_N(x) = Y$ *for each $x \in X$ and* $\text{src}_N(y) = X$ *for each $y \in Y$. Additionally suppose that for each $s \in \mathcal{S} \setminus X$ we have* $\text{facts}_N(s) \cap Y = \emptyset$ *and* $|\text{facts}_N(s)| \le |Y|/2$*, and that for each $f \in \mathcal{F} \setminus Y$ we have* $\text{src}_N(f) \cap X = \emptyset$ *and* $|\text{src}_N(f)| \le |X|/2$*. Then, with $(T^n)_{n\in\mathbb{N}}$ denoting* Sums, $T^n_N(s) \le 2^{1-n}$ *and* $T^n_N(f) \le 2^{1-n}$ *for all $s \in \mathcal{S} \setminus X$, $f \in \mathcal{F} \setminus Y$ and $n > 1$.*

In the case of the network N in Fig. 3, we get that $f \approx^{T^*}_N g$. Letting N' denote the network containing claims from the top connected component only, each of POI, Strong Independence and Independence would imply $f \preceq^{T^*}_N g$ iff $f \preceq^{T^*}_{N'} g$. However, it is easily seen that $T^*_{N'}(f) = 1 > 0 = T^*_{N'}(g)$, so $g \prec^{T^*}_{N'} f$ – this contradicts each of the independence axioms.

For Monotonicity, consider removing the edge (u, g) from N to obtain N''. Applying the above lemma, we get $f \approx^{T^*}_{N''} g$; in particular, $f \preceq^{T^*}_{N''} g$. Monotonicity would therefore imply that $f \prec^{T^*}_N g$, but we have seen that this is not true.

5.3 Modifying *Voting* and *Sums*

So far we have seen that neither of the basic operators *Voting* or *Sums* are completely satisfactory with respect to the axioms of section 4. Armed with the knowledge of how and why certain axioms fail, one may wonder whether it is possible to modify the operators accordingly so that the axioms *are* satisfied. Presently we shall show that this is partially possible both in the case of *Voting* and *Sums*.

5.3.1 Voting. A core problem with *Voting* is that it fails Coherence. Indeed, all sources are ranked equally regardless of the 'votes' for facts, so in some sense it is obvious that the two rankings do not cohere with each other. An easy improvement is to ensure the rankings cohere in at least one direction: we can aim for *Source*-Coherence by constructing the source rankings based on the fact ranking of *Voting*.

Definition 5.4. For a network N, define a binary relation $\lhd_N$ on $\mathcal{S}$ by $s_1 \lhd_N s_2$ iff $\text{facts}_N(s_1)$ is less-believable than $\text{facts}_N(s_2)$ with respect to *Voting*. The numerical operator *SC-Voting* (Source-Coherence Voting) is defined by

$$T^{SCV}_N(s) = |\{t \in \mathcal{S} : t \lhd_N s\}|, \quad T^{SCV}_N(f) = |\text{src}_N(f)|$$

It is easily seen that *SC-Voting* has Source-Coherence (although not Fact-Coherence; e.g. consider Fig. 1), which is a significant improvement over regular *Voting*. Since $\lhd_N$ relies on 'global' properties on N, however, this unfortunately comes at the expense of Source-Independence. Satisfaction of the other axioms is inherited from *Voting*.

Theorem 5.5. SC-Voting *satisfies Source-Coherence, Symmetry, Unanimity, Groundedness, Monotonicity, Fact-independence, POI and Strong Independence. It does not satisfy Fact-Coherence or Source-Independence.*

We note at this stage that the idea behind *SC-Voting* can be generalised beyond *Voting*: it is possible to define $\lhd_N$ in terms of *any* operator T, and to construct a new operator T' whose source ranking is given according to $\lhd_N$ as above, and whose fact ranking coincides with that of T. This ensures T' satisfies Source-Coherence whilst keeping the existing fact ranking from T.

Moreover we can go in the other direction and ensure *Fact*-Coherence whilst retaining the source ranking of T by defining a relation $\blacktriangleleft_N$ on $\mathcal{F}$ in a analogous manner to $\lhd_N$, and proceeding similarly.

5.3.2 Sums. Our main concern with *Sums* is the failure of Independence and Monotonicity. We have seen that this is some sense caused by the normalisation step: in Fig. 3 the scores of s, t, u etc go to 0 in the limit after dividing the 'global' maximum score across the network, which happens to come from a different connected component.

A natural fix for Independence is to therefore divide by the maximum score *within each component.* In this case the score for a source s depends only on the structure of the connected component in which it lies, which is exactly what is required for Independence.

However, this approach does not negate the undesirable effects of lemma 5.3. Indeed, suppose the network in Fig. 3 was modified so that fact y_1 is associated with object o instead of p_1. Clearly lemma 5.3 still applies after this change, and all sources and facts shown now belong to the same connected component. Therefore the 'per-component *Sums*' operator gives the same result as *Sums* itself, and in particular our Monotonicity counterexample still applies. Perhaps even worse, one can show that Coherence fails for this operator. We consider the loss of Coherence too high a price to pay for Independence.

Instead, let us take a step back and consider why and if normalisation is necessary. On the one hand, without normalisation the trust and belief scores are unbounded and therefore do not converge. On the other, we are not interested in the numeric scores for their own sake, but rather for the *rankings* that they induce. It may be possible that whilst the scores diverge without normalisation, the induced rankings *do* converge to a fixed one, which we may take as the 'ordinal limit'. This is in fact the case.

Definition 5.6. UnboundedSums is the recursive operator $\text{rec}(T^{Prior}, U^{Sums})$ where $T_N^{Prior}(s) = 1$, $T_N^{Prior}(f) = |\text{src}_N(f)|$ and U^{Sums} is defined as in Sect. 3.2.[3]

Theorem 5.7. UnboundedSums *is ordinally convergent, in the sense that there is an ordinal operator T^* such that for each network N there exists $J_N \in \mathbb{N}$ such that $T_N^n(s_1) \le T_N^n(s_2)$ iff $s_1 \sqsubseteq_N^{T^*} s_2$ for all $n \ge J_N$ and $s_1, s_2 \in \mathcal{S}$ (and similarly for facts).*

That is, the rankings on $\mathcal{S}$ and $\mathcal{F}$ in N are given by T^* after J_N iterations.

With the normalisation problems aside, *UnboundedSums* provides an example of a principled operator satisfying our two key axioms: Coherence and Independence. We conjecture that Monotonicity is also satisfied, but this remains to be proven.

Theorem 5.8. UnboundedSums *satisfies Coherence, Symmetry, Unanimity, Groundedness and Independence.* UnboundedSums *does not satisfy POI and Strong Independence.*

[3] Note that to simplify proof of ordinal convergence we use a different prior operator to *Sums*, but this does not effect the operator in any significant way.

	Voting	SC-Voting	Sums	U-Sums
Coherence	X	X	✓	✓
Symmetry	✓	✓	✓	✓
Unanimity	✓	✓	✓	✓
Ground.	✓	✓	✓	✓
Mon.	✓	✓	X	?
POI	✓	✓	X	X
Str. Indep	✓	✓	X	X
Indep.	✓	X	X	✓

Table 1: Satisfaction of the axioms for the various operators.

To summarise this section, Table 1 shows which axioms are satisfied by each of the operators.

6 CONCLUSION

In this paper we formalised a mathematical framework for truth discovery which is applicable to many algorithms in the literature. This provided the setting for the axiomatic methods of social choice to be applied. To our knowledge, this is the first such axiomatic treatment in this context.

It was possible to adapt many axioms from social choice theory and related areas. In particular, the *transitivity* axiom studied in the context of ranking systems [2, 19] took on new life in the form of Coherence, which we consider a core axiom for TD operators.

We proceeded to provide an impossibility result and an axiomatic characterisation of the baseline *Voting* method, before turning to more practical matters and analysing the existing TD algorithm *Sums.* We found that, surprisingly, it fails Independence. This is a serious issue for *Sums* which has not been discussed in the literature to date, and its discovery here highlights the benefits of the axiomatic method. To resolve this, we suggested a modification to *Sums* for which Independence *is* satisfied.

A restriction of our analysis is that only one 'real-world' algorithm was considered. Further axiomatic analysis of algorithms provides a deeper understanding of how algorithms operate on an intuitive level, but is made difficult by the complexity of the state-of-the-art truth discovery methods. New techniques for establishing the satisfaction (or otherwise) of axioms would be helpful in this regard.

There is also scope for extensions to our model of truth discovery in the framework itself. For example, we make the somewhat simplistic assumption that there are a fixed finite number of available sources, objects and facts. Regarding sources, this is at odds with the overwhelming number of individuals and devices that may participate in truth discovery, which is for all practical purposes infinite. For facts, it means we can only consider categorical values; modelling an object whose true fact is a real number, for example, is not straightforward in our framework.

Finally, our model does not account for any associations or constraints between objects, whereas in reality the belief in a fact for one object may strengthen or weaken our belief in other facts for related objects. These types of constraints or correlations have been studied both on the theoretical side (e.g. in judgment aggregation) and practical side in truth discovery [23].

REFERENCES

[1] Alon Altman and Moshe Tennenholtz. 2005. Ranking systems: the PageRank axioms. In *Proceedings of the 6th ACM conference on Electronic commerce*. ACM, 1–8.

[2] Alon Altman and Moshe Tennenholtz. 2008. Axiomatic Foundations for Ranking Systems. *J. Artif. Int. Res.* 31, 1 (March 2008), 473–495. http://dl.acm.org/citation.cfm?id=1622655.1622669

[3] Reid Andersen, Christian Borgs, Jennifer Chayes, Uriel Feige, Abraham Flaxman, Adam Kalai, Vahab Mirrokni, and Moshe Tennenholtz. 2008. Trust-based Recommendation Systems: An Axiomatic Approach. In *Proceedings of the 17th International Conference on World Wide Web (WWW '08)*. ACM, 199–208. https://doi.org/10.1145/1367497.1367525

[4] Kenneth J. Arrow. 1952. Social Choice and Individual Values. *Ethics* 62, 3 (1952), 220–222.

[5] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. 2016. Introduction to Computational Social Choice. In *Handbook of Computational Social Choice* (1st ed.), Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, New York, NY, USA, Chapter 1.

[6] Zoé Christoff and Davide Grossi. 2017. Binary Voting with Delegable Proxy: An Analysis of Liquid Democracy. In *Proc. TARK 2017*.

[7] Elad Dokow and Ron Holzman. 2010. Aggregation of binary evaluations with abstentions. *Journal of Economic Theory* 145 (2010), 544–561.

[8] Y. Du, Y. Sun, H. Huang, L. Huang, H. Xu, Y. Bao, and H. Guo. 2019. Bayesian Co-Clustering Truth Discovery for Mobile Crowd Sensing Systems. *IEEE Transactions on Industrial Informatics* (2019), 1–1. https://doi.org/10.1109/TII.2019.2896287

[9] Ulle Endriss. 2016. Judgment Aggregation. In *Handbook of Computational Social Choice* (1st ed.), Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, New York, NY, USA, Chapter 17.

[10] Alban Galland, Serge Abiteboul, Amélie Marian, and Pierre Senellart. 2010. Corroborating Information from Disagreeing Views. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining (WSDM '10)*. ACM, New York, NY, USA, 131–140. https://doi.org/10.1145/1718487.1718504

[11] Manish Gupta and Jiawei Han. 2011. Heterogeneous Network-based Trust Analysis: A Survey. *SIGKDD Explor. Newsl.* 13, 1 (Aug. 2011), 54–71. https://doi.org/10.1145/2031331.2031341

[12] Jon M. Kleinberg. 1999. Authoritative Sources in a Hyperlinked Environment. *J. ACM* 46, 5 (Sept. 1999), 604–632. https://doi.org/10.1145/324133.324140

[13] Justin Kruger, Ulle Endriss, Raquel Fernandez, and Ciyang Qing. 2014. Axiomatic Analysis of Aggregation Methods for Collective Annotation. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems (AAMAS '14)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1185–1192. http://dl.acm.org/citation.cfm?id=2617388.2617437

[14] Yaliang Li, Jing Gao, Chuishi Meng, Qi Li, Lu Su, Bo Zhao, Wei Fan, and Jiawei Han. 2016. A Survey on Truth Discovery. *SIGKDD Explor. Newsl.* 17, 2 (2016), 1–16. https://doi.org/10.1145/2897350.2897352

[15] Yaliang Li, Qi Li, Jing Gao, Lu Su, Bo Zhao, Wei Fan, and Jiawei Han. 2016. Conflicts to Harmony: A Framework for Resolving Conflicts in Heterogeneous Data by Truth Discovery. *IEEE Transactions on Knowledge and Data Engineering* 28, 8 (Aug 2016), 1986–1999. https://doi.org/10.1109/TKDE.2016.2559481

[16] Fenglong Ma, Yaliang Li, Qi Li, Minghui Qiu, Jing Gao, Shi Zhi, Lu Su, Bo Zhao, Heng Ji, and Jiawei Han. 2015. FaitCrowd: Fine Grained Truth Discovery for Crowdsourced Data Aggregation. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '15)*. ACM, New York, NY, USA, 745–754. https://doi.org/10.1145/2783258.2783314 event-place: Sydney, NSW, Australia.

[17] Fenglong Ma, Chuishi Meng, Houping Xiao, Qi Li, Jing Gao, Lu Su, and Aidong Zhang. 2017. Unsupervised Discovery of Drug Side-Effects from Heterogeneous Data Sources. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*. ACM, New York, NY, USA, 967–976. https://doi.org/10.1145/3097983.3098129

[18] 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 '10)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 877–885. http://dl.acm.org/citation.cfm?id=1873781.1873880

[19] Moshe Tennenholtz. 2004. Reputation Systems: An Axiomatic Approach. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence (UAI '04)*. AUAI Press, Arlington, Virginia, United States, 544–551. http://dl.acm.org/citation.cfm?id=1036843.1036909

[20] Dong Wang, Lance Kaplan, Hieu Le, and Tarek Abdelzaher. 2012. On Truth Discovery in Social Sensing: A Maximum Likelihood Estimation Approach. In *Proceedings of the 11th International Conference on Information Processing in Sensor Networks (IPSN '12)*. ACM, 233–244. https://doi.org/10.1145/2185677.2185737 event-place: Beijing, China.

[21] Houping Xiao. 2018. *Multi-sourced Information Trustworthiness Analysis: Applications and Theory.* Ph.D. Dissertation. University at Buffalo, State University of New York.

[22] Houping Xiao, Jing Gao, Zhaoran Wang, Shiyu Wang, Lu Su, and Han Liu. 2016. A Truth Discovery Approach with Theoretical Guarantee. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*. ACM, New York, NY, USA, 1925–1934. https://doi.org/10.1145/2939672.2939816

[23] Yi Yang, Quan Bai, and Qing Liu. 2019. A probabilistic model for truth discovery with object correlations. *Knowledge-Based Systems* 165 (2019), 360 – 373. https://doi.org/10.1016/j.knosys.2018.12.004

[24] 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 (June 2008), 796–808. https://doi.org/10.1109/TKDE.2007.190745

[25] Xiaoxin Yin and Wenzhao Tan. 2011. Semi-supervised Truth Discovery. In *Proceedings of the 20th International Conference on World Wide Web (WWW '11)*. ACM, 217–226. https://doi.org/10.1145/1963405.1963439 event-place: Hyderabad, India.

[26] Daniel Yue Zhang, Rungang Han, Dong Wang, and Chao Huang. 2016. On robust truth discovery in sparse social media sensing. In *2016 IEEE International Conference on Big Data (Big Data)*. 1076–1081. https://doi.org/10.1109/BigData.2016.7840710

[27] Daniel Yue Zhang, Rungang Han, Dong Wang, and Chao Huang. 2016-12. On robust truth discovery in sparse social media sensing. In *2016 IEEE International Conference on Big Data (Big Data)*. 1076–1081. https://doi.org/10.1109/BigData.2016.7840710

[28] Liyan Zhang, Guo-Jun Qi, Dong Zhang, and Jinhui Tang. 2018. Latent Dirichlet Truth Discovery: Separating Trustworthy and Untrustworthy Components in Data Sources. *IEEE Access* 6 (2018), 1741–1752. https://doi.org/10.1109/ACCESS.2017.2780182

[29] Shi Zhi, Bo Zhao, Wenzhu Tong, Jing Gao, Dian Yu, Heng Ji, and Jiawei Han. 2015. Modeling Truth Existence in Truth Discovery. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '15)*. ACM, New York, NY, USA, 1543–1552. https://doi.org/10.1145/2783258.2783339

[30] William S. Zwicker. 2016. Introduction to the Theory of Voting. In *Handbook of Computational Social Choice* (1st ed.), Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, New York, NY, USA, Chapter 2.