

Fusing Data with Correlations

Ravali Pochampally
University of Massachusetts
ravali@cs.umass.edu

Anish Das Sarma
ClearList Inc.
anish.dassarma@gmail.com

Xin Luna Dong
Google Inc.
lunadong@google.com

Alexandra Meliou
University of Massachusetts
ameli@cs.umass.edu

Divesh Srivastava
AT&T Labs-Research
divesh@research.att.com

ABSTRACT

Many applications rely on Web data and extraction systems to accomplish knowledge-driven tasks. Web information is not curated, so many sources provide inaccurate, or conflicting information. Moreover, extraction systems introduce additional noise to the data. We wish to automatically distinguish correct data and erroneous data for creating a cleaner set of integrated data. Previous work has shown that a naïve voting strategy that trusts data provided by the majority or at least a certain number of sources may not work well in the presence of copying between the sources. However, correlation between sources can be much broader than copying: sources may provide data from complementary domains (*negative correlation*), extractors may focus on different types of information (*negative correlation*), and extractors may apply common rules in extraction (*positive correlation, without copying*). In this paper we present novel techniques modeling correlations between sources and applying it in truth finding. We provide a comprehensive evaluation of our approach on three real-world datasets with different characteristics, as well as on synthetic data, showing that our algorithms outperform the existing state-of-the-art techniques.

1. INTRODUCTION

The Web is an incredibly rich source of information, which is growing at an unprecedented pace and is amassed by a plethora of contributors. An increasing number of users and applications rely on online data as the main resource to satisfy their information needs. Web data is not curated and sources may often provide erroneous or conflicting information. Additionally, a lot of Web data is largely unstructured, lacking a predefined schema or consistent format. As a result, knowledge-driven applications in various domains (*e.g.*, finance, technology, advertisement, *etc.*) rely on information extraction systems to retrieve structured relations from online sources. However, extraction systems have less than perfect accuracy, invariably introducing more noise to the data.

Our goal is to automatically distinguish correct data and erroneous data for creating a cleaner set of data. A naïve approach to achieve this goal is majority voting: we trust the data provided by the majority, or at least a certain number of sources. However, such

a strategy may perform badly for two reasons. First, sources may provide data from complementary domains (*e.g.*, information on scientific books vs. on biographies) and extractors may focus on different types of information (*e.g.*, extracting from the Infobox or the texts of Wikipedia pages); blindly requiring agreement among sources may miss correct data and cause *false negatives*. Second, sources may easily copy and share data [2] and extractors may apply common rules; blindly trusting agreement among sources may enforce erroneous data and cause *false positives*. Such correlation or anti-correlation between sources makes it especially hard to tell the truth from wrong statements or extractions, as illustrated next.

EXAMPLE 1.1. *Figure 1 depicts example data extracted from the Wikipedia page for Barack Obama, using five different extraction systems. Extracted data consist of knowledge triples in the form of {subject, predicate, object}; for example, {Obama, spouse, Michelle} states that the spouse of Obama is Michelle. Some extracted triples are incorrect. For example, triple t_2 is false: extraction systems S_1 and S_2 derived the triple from a sentence referring to Barack Obama Sr, rather than the current US president.*

Various types of correlations exist among the five sources. First, S_1, S_4 and S_5 implement similar extraction patterns and extract similar sets of triples; there is a positive correlation between these sources. Second, S_3 extracts triples from the Infobox of the Wikipedia page while S_1 (similarly, S_4 and S_5) extracts triples from the text; their extracted triples are largely complementary and there is a negative correlation between them.

*Figure 1c shows the precision, recall, and F-measure of voting techniques: Union- k accepts a triple as true if at least $k\%$ of the extractors extract it; *e.g.*, Union-25 accepts triples provided by at least 2 extractors: it has high recall (missing only one triple), but makes a lot of mistakes (extracting 4 false triples) because of the common mistakes by the correlated sources S_1, S_4 , and S_5 . Union-75 accepts triples provided by at least 4 extractors; it misses a lot of true triples since S_3 is anti-correlated with three other sources.*

Resolving conflicts with consideration of source copying has been studied in *data fusion* [6, 5]. There are two limitations of previous approaches. First, they focus on copying of data between sources and are based on the intuition that common mistakes are strong evidence of copying; correlation is much broader: it can be positive or negative and can be caused by different reasons. Previous approaches are effective in detecting positive correlation on false data, but are not effective with positive correlation on true data or negative correlation. Second, their model relies on the *single-truth* assumption such as everyone has a unique birthplace; however, in practice there can be multiple truths for certain “facts”, such as someone may have multiple professions (*e.g.*, triples t_1, t_3 , and t_{10} in Figure 1 are all correct).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

ID	Web document	KnowledgeTriple	Correct?	S ₁	S ₂	S ₃	S ₄	S ₅
t ₁	wiki/Barack_Obama	{Obama,profession,president}	Yes	✓	✓		✓	✓
t ₂	wiki/Barack_Obama	{Obama,died,1982}	No	✓	✓			
t ₃	wiki/Barack_Obama	{Obama,profession,lawyer}	Yes			✓		
t ₄	wiki/Barack_Obama	{Obama,religion,Christian}	Yes		✓	✓	✓	✓
t ₅	wiki/Barack_Obama	{Obama,age,50}	No		✓	✓		
t ₆	wiki/Barack_Obama	{Obama,support,White Sox}	Yes	✓			✓	✓
t ₇	wiki/Barack_Obama	{Obama,spouse,Michelle}	Yes	✓	✓	✓		
t ₈	wiki/Barack_Obama	{Obama,administered by,John G. Roberts}	No	✓	✓		✓	✓
t ₉	wiki/Barack_Obama	{Obama,surgical operation,05/01/2011}	No	✓	✓		✓	✓
t ₁₀	wiki/Barack_Obama	{Obama,profession,community organizer}	Yes	✓		✓	✓	✓

(a) Data extracted by five different extractors from the Wikipedia page for *Barack Obama*. The ✓ symbols indicate which extraction systems produce each knowledge triple; for example, t_3 is extracted by S_3 , but not by any other extractor.

	Precision	Recall
S ₁	0.57	0.67
S ₂	0.43	0.5
S ₃	0.8	0.67
S ₄	0.67	0.67
S ₅	0.67	0.67

	Joint prec	Joint rec
S ₂ S ₃	0.67	0.33
S ₁ S ₃	1	0.33
S ₁ S ₂ S ₄	0.33	0.167
S ₁ S ₄ S ₅	0.6	0.5

	Precision	Recall	F-measure
Union-25	0.56	0.83	0.67
Union-50	0.71	0.83	0.77
Union-75	0.6	0.5	0.55

(b) Precision and recall for each extractor, and joint precision and joint recall for some combinations of extractors.

(c) Naïve fusion approaches based on voting do not achieve very good results, as they do not account for correlations among the extractors.

Figure 1: Example 1.1: (a) knowledge triples derived by 5 extractors, (b) extractor quality and correlations, (c) voting results.

In this paper, we address the problem of finding truths among data provided by multiple sources, which may contain complex correlations. We make the following contributions.

- We propose measuring the quality of a source as its *precision* and *recall* and measuring the correlation between a subset of sources as their *joint precision* and *joint recall*. We express them in terms of conditional probability (Section 2).
- We present a novel technique that derives the probability of a triple being true from the precision and recall of the sources using Bayesian analysis under the independence assumption (Section 3). Our experiments show that even before incorporating correlations, our basic approach often outperforms existing state-of-the-art techniques.
- We extend our approach to handle correlations between the sources. We first present an *exact solution* that is exponential in the number of data sources. We then present two approximation schemes: the *aggressive approximation* reduces the computational complexity from exponential to linear, but sacrifices the accuracy of the predictions; our *elastic approximation* provides a mechanism to trade efficiency for accuracy (Section 4).
- We conduct a comprehensive evaluation of our techniques against three real-world data sets, as well as synthetic data. Our experiments show that our methods can significantly improve the results by considering correlation without adding too much overhead for efficiency (Section 5).

2. THE FUSION PROBLEM

In this section, we introduce our data model and its semantics, we provide a formal definition of the problem of fusing data from sources with unknown correlations, and we present a high-level overview of our approach. We summarize notations in Figure 2.

2.1 Data model

We consider a set of data sources $\mathcal{S} = \{S_1, \dots, S_n\}$. Each source provides some data and we call each unit of data a *triple*; a triple can be considered as a cell in a database table in the form of *[row-entity, column-attribute, value]* (e.g., in a table about politicians, a row can represent *Obama*, a column can represent attribute *profession*, and the corresponding cell can have value *president*),

or an RDF triple in the form of *{subject, predicate, object}*, such as *{Obama, profession, president}*. We denote with O_i the triples provided by source $S_i \in \mathcal{S}$; interchangeably, we denote with $S_i \models t$ or $t \in O_i$ that S_i provides triple t . Our data model consists of $\mathcal{S} = \{S_1, \dots, S_n\}$ and the collections of their output triples $\mathcal{O} = \{O_1, \dots, O_n\}$. In a slight abuse of notation, we write $t \in \mathcal{O}$ to denote that $\exists O_i \in \mathcal{O}$ such that $t \in O_i$. We use \mathcal{O}_t to represent the subset of outputs in \mathcal{O} that involve triple t ; note that \mathcal{O}_t contains the observation that a source S_i does not provide t only if S_i provides other data in the domain of t , so we do not unnecessarily penalize data missing from irrelevant sources.

We consider deterministic sources: a source either outputs a triple, or it does not. In practice, a source $S_i \in \mathcal{S}$ may provide a confidence score associated with each triple $t \in O_i$; we can consider that S_i outputs t if the assigned confidence score exceeds a certain threshold. As in previous work [6, 25], we assume that schema mapping and reference reconciliation have been applied so we can compare the triples across sources.

Our goal is to purge the output of all incorrect triples to obtain a high-quality data set $R = \{t : t \in \mathcal{O} \wedge t \text{ is true}\}$. We say that a triple t is *true* if it is consistent with the real world, and *false* otherwise; for example, *{Obama, profession, president}* is true whereas *{Obama, died, 1982}* is false. We next show an instantiation of our data model for the data extraction scenario.

EXAMPLE 2.1. *Figure 1 shows triples extracted by five extractors from the Wikipedia page for Barack Obama and we need to determine which triples are correctly extracted. We consider that each extractor corresponds to a source; for example, S_1 corresponds to the first extractor and it provides (among others) triple $t_1 : \{\text{Obama, profession, president}\}$. We denote this as $S_1 \models t_1$, meaning that the extractor believes that t_1 is a fact that appears on the Wikipedia page. Accordingly, $O_1 = \{t_1, t_2, t_6, t_7, t_8, t_9, t_{10}\}$.*

Based on \mathcal{S} and \mathcal{O} , we decide whether each triple t_i is true ($i \in [1, 10]$). In this scenario, the extractor input (the processed web page) represents the “real world”, against which we evaluate the correctness of the extractor outputs. We consider a triple to be true (i.e., correctly extracted) if the web page indeed provides the triple.

Semantics: In this paper, we make two assumptions about semantics of the data: First, we assume *triple independence*: the truth-

fulness of each triple is independent of that of other triples. For example, whether the page indeed provides triple t_1 is independent of whether the page provides triple t_2 . Second, we assume *open-world* semantics: a source considers any triple in its output as *true*, and any triple not in its output as *unknown* (rather than *false*). For example, in Figure 1, S_1 provides t_1 and t_2 but not t_3 , meaning that it considers t_1 and t_2 as being provided by the page, but does not know whether t_3 is also provided. Note that this is in contrast with the *conflicting-triple*, *closed-world* semantics in [6]; under this semantics, $\{Obama, religion, Christian\}$ and $\{Obama, religion, Muslim\}$ would be considered conflicting with each other, as we typically assume one can have at most one religion and a source claiming the former implicitly claims that the latter is false.

We make these assumptions for two reasons. The first reason is that they are suitable for many application scenarios. One application scenario is data extraction, as shown in our motivating example: when an extractor derives two different triples from a Web page (often from different sentences or phrases), the correctness of the two extractions are independent; if an extractor does not derive a triple from a Web page, it usually indicates that the extractor *does not know* whether the page provides the triple, rather than that it *believes* that the page does not provide the triple. Another scenario is attributes that can accept multiple truths. For example, a person can have multiple professions: the correctness of each profession is largely independent of other professions¹, and a source that claims that *Obama* is a president does not necessarily claim that *Obama cannot* be a lawyer. The second reason is that, to the best of our knowledge, all previous work that studies correlation of sources focuses on the conflicting-triple, closed-world semantics; the independent-triple, open-world semantics allows us to fill the gap in the existing literature. Note that we can apply strategies for conflicting-triple and closed-world semantics in the case of independent-triple and closed-world semantics, or in the case of conflicting-triple and open-world semantics. We leave combination of all semantics for future work.

2.2 Measuring truthfulness

The objective of our framework is to distinguish true and false triples in a collection of source outputs. A key feature of our approach is that it does not assume any knowledge of the inner workings of the sources and how they derive the data that they provide. First, this is indeed the case in practice for many real-world data sources — they provide the data without telling us how they obtain it. Second, even when some information on the data derivation process is available, it may be too complex to reason about; for example, an extractor often learns thousands (or even more) of patterns (e.g., distance supervision [18]) and uses internal coding to present them; it is hard to understand all of them, let alone to reason about them and compare them across sources.

Next, we show which key evidence we consider in our approach and then formally define our problem.

Source quality

The quality of the sources affects our belief of the truthfulness of a triple. Intuitively, if a source S has high *precision* (i.e., most of its provided triples are true), then a triple provided by S is more likely to be true. On the other hand, if S has a high *recall* (i.e., most of the true triples are provided by S), then a triple not provided by S is more likely to be false.

¹Arguably, it is unlikely for a person to be a doctor, a lawyer, and a plumber at the same time as they require very different skills; we leave such joint reasoning with *a priori* knowledge for future work.

Notation	Description
\mathcal{S}	Set of sources $\mathcal{S} = \{S_1, \dots, S_n\}$
O_i	Set of output triples of source S_i
\mathcal{O}	$\mathcal{O} = \{O_1, \dots, O_n\}$
\mathcal{O}_t	Subset of observations in \mathcal{O} that refer to triple t
p_i (resp. $p_{\mathcal{S}^*}$)	Precision of source S_i (resp. sources \mathcal{S}^*)
r_i (resp. $r_{\mathcal{S}^*}$)	Recall of source S_i (resp. sources \mathcal{S}^*)
q_i (resp. $q_{\mathcal{S}^*}$)	False positive rate of S_i (resp. \mathcal{S}^*)
$S_i \models t$	S_i outputs t ($t \in O_i$)
$\mathcal{S}^* \models t$	$\forall S_i \in \mathcal{S}^*, S_i \models t$
$\Pr(t \mathcal{O})$	Correctness probability of triple t
$\Pr(t), \Pr(-t)$	$\Pr(t = true)$ and $\Pr(t = false)$ respectively

Figure 2: Summary of notations used in the paper.

We define precision and recall in the standard way: the precision p_i of source $S_i \in \mathcal{S}$ represents the portion of triples in the output O_i that are true; the recall r_i of S_i represents the portion of all true triples that appear in O_i . These metrics can be described in terms of probabilities as follows.

$$p_i = \Pr(t | S_i \models t) \quad (1)$$

$$r_i = \Pr(S_i \models t | t) \quad (2)$$

EXAMPLE 2.2. Figure 1b shows the precision and recall of the five sources. For example, the precision of S_1 is $\frac{4}{7} = 0.57$, as only 4 out of the 7 triples in O_1 are correct. The recall is $\frac{4}{6} = 0.67$, as 4 out of the 6 correct triples are included in O_1 .

The recall of a source should be calculated with respect to the “scope” of its input. For example, if a source S provides only information about *Obama* but not about *Bush*, we may penalize the recall of S for providing only 1 out of the 3 professions of *Obama*, but should not penalize the recall of S for not providing any profession for *Bush*. For simplicity of presentation, in the rest of the paper we ignore the “scope” of each source in our discussion, but all of our techniques work with either version of recall calculation.

Correlation

Another key factor that can affect our belief of triple truthfulness is the presence of correlations between data sources. Intuitively, if we know that two sources S_i and S_j are nearly duplicates of each other, thus they are *positively correlated*, the fact that both provide a triple t should not significantly increase our belief that t is true. On the other hand, if we know two sources S_i and S_j are complementary and have little overlap, so are *negatively correlated*, the fact that a triple t is provided by one but not the other should not significantly reduce our belief that t is true. Note the difference between *correlation* and *copying* [6]: copying can be one reason for positive correlation, but positive correlation can be due to other factors, such as using similar extraction patterns or implementing similar algorithms to derive data, rather than copying.

We use *joint precision* and *joint recall* to capture correlation between sources. The joint precision of sources in \mathcal{S}^* , denoted by $p_{\mathcal{S}^*}$, represents the portion of triples in the output of all sources in \mathcal{S}^* (i.e., intersection) that are correct; the joint recall of \mathcal{S}^* , denoted by $r_{\mathcal{S}^*}$, represents the portion of all correct triples that are output by all sources in \mathcal{S}^* . If we denote by $\mathcal{S}^* \models t$ that a triple t is output by all sources in \mathcal{S}^* , we can describe these metrics in terms of probabilities as follows.

$$p_{\mathcal{S}^*} = \Pr(t | \mathcal{S}^* \models t) \quad (3)$$

$$r_{\mathcal{S}^*} = \Pr(\mathcal{S}^* \models t | t) \quad (4)$$

EXAMPLE 2.3. Figure 1b shows the joint precision and recall for selected subsets of sources. Take the sources $\{S_1, S_4, S_5\}$ as

an example. They provide similar sets of triples: they all provide t_1, t_6, t_8, t_9 , and t_{10} . Their joint precision is $\frac{3}{5} = 0.6$ and their joint recall is $\frac{3}{6} = 0.5$. Note that if the sources were independent, their joint recall would have been $r_1 \cdot r_4 \cdot r_5 = 0.3$, much lower than the real one (0.5); this indicates positive correlation.

On the other hand, S_1 and S_3 have little overlap in their data: they both provide triples t_7 and t_{10} . Their joint precision is $\frac{2}{2} = 1$ and their joint recall is $\frac{2}{6} = 0.33$. Note that if the sources were independent, their joint recall would have been $r_1 \cdot r_3 = 0.45$, higher than the real one (0.33); this indicates negative correlation.

We define positive and negative correlation formally in Section 4.

Problem definition

Our goal is to determine the truthfulness of each triple in \mathcal{O} . We model the truthfulness of t as the probability that t is true, given the outputs of all sources; we denote this as $\Pr(t | \mathcal{O})$. We can accept a triple t as true if this probability is above 0.5, meaning that t is more likely to be true than to be false. As we assume the truthfulness of each triple is independent, we can compute the probability for each triple separately conditioned on the provided data regarding t ; that is, \mathcal{O}_t . We frame our problem statement based on source quality and correlation. For now we assume the source quality metrics and correlation factors are given as input; we discuss techniques to derive them shortly. We formally define the problem as follows:

DEFINITION 2.4 (TRIPLE TRUTHFULNESS). *Given (1) a set of sources $\mathcal{S} = \{S_1, \dots, S_n\}$, (2) their outputs $\mathcal{O} = \{O_1, \dots, O_n\}$, and (3) the joint precision $p_{\mathcal{S}^*}$ and recall $r_{\mathcal{S}^*}$ of each subset of sources $\mathcal{S}^* \subseteq \mathcal{S}$, compute the probability for each output triple $t \in \mathcal{O}$, denoted by $\Pr(t | \mathcal{O}_t)$.*

Note that given a set \mathcal{S} of n sources, there is a total of $2(2^n - 1)$ joint precision and recall parameters. Since the input size is exponential in the number of sources, even a polynomial algorithm will be infeasible in practice. We show in Section 4 how we can reduce the number of parameters we consider in our model and solve the problem efficiently.

2.3 Overview

We start by studying the problem of triple probability computation under the assumption that sources are, indeed, independent. We will show that even in this case, there are challenges to overcome in order to derive the probability. We then extend our methods to account for correlations among sources. Here, we present an overview of some high-level intuitions that we apply in each of these two settings.

Independent sources (Section 3)

Fusion of data from multiple sources is challenging because the inner-workings of each source are not completely known. We present a method that uses source quality metrics (precision and recall) to derive the probability that a source provides a particular triple, and applies Bayesian analysis to compute the truthfulness of each triple. We describe how to derive the quality metrics if those are unknown. With this model, we are able to improve the F-measure to .86 (precision=.75, recall=1) for the motivating example.

Correlated sources (Section 4)

Sources are often correlated: they may copy data from each other, employ similar techniques in deriving the data, or analyze complementary portions of the raw data sets. Correlations can be positive or negative, and are generally unknown. We address two main challenges in the case of correlated sources.

- *Using correlations:* We start by assuming that we know concrete correlations between sources. We will see that the main insight into revising the probability of triples is to determine how likely it is for a particular triple to have appeared in the output of a given subset of sources but not in the output of any other source. Further, we use the inclusion-exclusion principle to express the correctness probability of a triple using the joint precision and joint recall of subsets of sources.
- *Exponential complexity:* The number of correlation parameters is exponential in the number of sources, which can make our computation infeasible. To counter this problem, we develop two approximation methods: our *aggressive approximation* reduces the computation from exponential to linear, but sacrifices accuracy; our *elastic approximation* provides a mechanism to trade efficiency for accuracy and improve the quality of our approximation incrementally.

Considering correlations, we can further improve the F-measure to 0.91 (precision=1, recall=0.83) for our motivating example, which is 18% higher than Union-50 (*i.e.*, majority voting).

3. FUSING INDEPENDENT SOURCES

In this section, we start with the assumption that the sources are independent. Our goal is to estimate the probability that an output triple t is true given the observed data: $\Pr(t | \mathcal{O}_t)$. We describe a novel technique to derive this probability based on the quality of each source (Sec. 3.1). Since these quality metrics are not always known in advance, we also describe how to derive them if we are given the ground truth on a subset of the extracted data (Sec. 3.2).

3.1 Estimating triple probability

Given a collection of output triples for each source O_i , our objective is to compute, for each $t \in \mathcal{O}$, the probability that t is true, $\Pr(t | \mathcal{O})$, based on the quality of each source. Due to the independent-triple assumption, $\Pr(t | \mathcal{O}) = \Pr(t | \mathcal{O}_t)$.

We use Bayes' rule to express $\Pr(t | \mathcal{O}_t)$ based on the inverse probabilities $\Pr(\mathcal{O}_t | t)$ and $\Pr(\mathcal{O}_t | \neg t)$, which represent the probability of deriving the observed output data conditioned on t being true or false respectively. In addition, we denote the *a priori* probability that t is true with $\Pr(t) = \alpha$.

$$\Pr(t | \mathcal{O}_t) = \frac{\alpha \Pr(\mathcal{O}_t | t)}{\alpha \Pr(\mathcal{O}_t | t) + (1 - \alpha) \Pr(\mathcal{O}_t | \neg t)} \quad (5)$$

The denominator in the above expression is equal to $\Pr(\mathcal{O}_t)$. The a-priori probability α can be derived from a training set (*i.e.*, a subset of the triples with known ground truth values, see Section 3.2).

We denote by \mathcal{S}_t the set of sources that provide t , and by $\mathcal{S}_{\bar{t}}$ the set of sources that do not provide t . Assuming that the sources are independent, the probabilities $\Pr(\mathcal{O}_t | t)$ and $\Pr(\mathcal{O}_t | \neg t)$ can then be expressed using the *true positive rate*, also known as *sensitivity* or recall, and the *false positive rate*, also known as the complement of *specificity*, of each source as follows:

$$\Pr(\mathcal{O}_t | t) = \prod_{S_i \in \mathcal{S}_t} \Pr(S_i \models t | t) \prod_{S_i \in \mathcal{S}_{\bar{t}}} (1 - \Pr(S_i \models t | t)) \quad (6)$$

$$\Pr(\mathcal{O}_t | \neg t) = \prod_{S_i \in \mathcal{S}_{\bar{t}}} \Pr(S_i \models t | \neg t) \prod_{S_i \in \mathcal{S}_t} (1 - \Pr(S_i \models t | \neg t)) \quad (7)$$

From Eq. (2), we know $r_i = \Pr(S_i \models t | t)$. We denote the false positive rate by $q_i = \Pr(S_i \models t | \neg t)$ and describe how we derive it in Section 3.2. Applying these to Eq. (5), we obtain the following theorem.

THEOREM 3.1 (INDEPENDENT SOURCES). *Given a set of independent sources $\mathcal{S} = \{S_1, \dots, S_n\}$, the recall r_i and the false positive rate q_i of each source S_i , the correctness probability of an output triple t is $\Pr(t | \mathcal{O}_t) = \frac{1}{1 + \frac{1-\alpha}{\alpha} \cdot \frac{1}{\mu}}$, where*

$$\mu = \prod_{S_i \in \mathcal{S}_t} \frac{r_i}{q_i} \prod_{S_i \in \mathcal{S}_{\bar{t}}} \left(\frac{1-r_i}{1-q_i} \right) \quad (8)$$

Intuitively, we compute the correctness probability based on the (weighted) contributions of each source for each triple. Each source S_i has contribution $\frac{r_i}{q_i}$ for a triple that it provides, and contribution $\frac{1-r_i}{1-q_i}$ for a triple that it does not provide. Given a triple t , we multiply the corresponding contributions of all sources to derive μ , and then compute the probability of the triple accordingly.

We say a source S_i is *good* if it is more likely to provide a true triple than a false triple; that is, $\Pr(S_i \models t | t) > \Pr(S_i \models t | \neg t)$ (i.e., $r_i > q_i$). Thus, a good source has a positive contribution for a provided triple — once it provides a triple, the triple is more likely to be true; otherwise, the triple is more likely to be false.

PROPOSITION 3.2. *Let $S' = S \cup \{S'\}$ and $\mathcal{O}' = \mathcal{O} \cup \{\mathcal{O}'\}$.*

- *If S' is a good source:*
 - *If $S' \models t$, then $\Pr(t | \mathcal{O}'_t) > \Pr(t | \mathcal{O}_t)$.*
 - *If $S' \not\models t$, then $\Pr(t | \mathcal{O}'_t) < \Pr(t | \mathcal{O}_t)$.*
- *If S' is a bad source:*
 - *If $S' \models t$, then $\Pr(t | \mathcal{O}'_t) < \Pr(t | \mathcal{O}_t)$.*
 - *If $S' \not\models t$, then $\Pr(t | \mathcal{O}'_t) > \Pr(t | \mathcal{O}_t)$.*

EXAMPLE 3.3. *We apply Theorem 3.1 to derive the probability of t_2 , which is provided by S_1 and S_2 but not by S_3, S_4 , or S_5 :*

$$\mu = \frac{r_1}{q_1} \cdot \frac{r_2}{q_2} \cdot \frac{1-r_3}{1-q_3} \cdot \frac{1-r_4}{1-q_4} \cdot \frac{1-r_5}{1-q_5}$$

Suppose we know that $q_1 = 0.5$, $q_2 = 0.67$, $q_3 = 0.167$, and $q_4 = q_5 = 0.33$, and we know the recall as shown in Figure 1b. Then we compute $\mu = 0.1$. With $\alpha = 0.5$, Theorem 3.1 gives $\Pr(t_2 | \mathcal{O}_{t_2}) = 0.09$, so we correctly determine that t_2 is false.

However, assuming independence can lead to wrong results: t_8 is provided by $\{S_1, S_2, S_4, S_5\}$, but not by S_3 . Using Eq. (8) produces $\mu = 1.6$ and $\Pr(t_8 | \mathcal{O}_{t_8}) = 0.62$, but t_8 is in reality false.

3.2 Estimating source quality

Theorem 3.1 uses the recall and false positive rate of each source to derive the correctness probability. We next describe how we compute them from a set of training data, where we know the truthfulness of each triple. Existing work [9] also relies on training data to compute source quality, while crowdsourcing platforms, such as Amazon Mechanical Turk, greatly facilitate the labeling process [17].

Computing the recall (r_i) relies on knowledge of the set of true triples, which is typically unknown *a priori*. Since we only need to decide truthfulness for each provided triple, we use the set of true triples that are provided by at least one source in the training data. Then, for each source $S_i, i \in [1, n]$, we count the number of true triples it provides and compute its recall according to the definition. In our motivating example (Figure 1), there are 6 true triples extracted by at least one extractor; accordingly, the recall of S_1 is $\frac{4}{6} = 0.67$ since it provides 4 true triples.

However, we cannot compute the false positive rate (q_i) in a similar way by considering only false triples in the training data. We next illustrate the problem using an example.

EXAMPLE 3.4. *Consider deriving the quality of S_1 from the training set $\{t_1, \dots, t_{10}\}$. Since the sources are all reasonably good, only 4 out of 10 triples are false. If we compute q_1 directly from the data, we have $q_1 = \frac{3}{4} = 0.75$. Since $q_1 > r_1 = 0.67$, we would (wrongly) consider S_1 as a bad source.*

Now suppose there is an additional source S_0 that provides 10 false triples $t_{11} - t_{20}$ and we include it in the training data. We would then compute $q_1 = \frac{3}{14} = 0.21$; suddenly, S_1 becomes a good source and much more trustable than it really is.

To address this issue, we next describe a way that derives the false positive rate from the precision and recall of a source. The advantage of this approach is that the precision of a source can be easily computed according to the training data and would not be affected by the quality of other sources. Using Bayes' Rule on $\Pr(t | S_i \models t)$ we obtain a formula similar to Eq. (5), and then we apply the conditional probability expressions for p_i, r_i , and q_i :

$$\Pr(t | S_i \models t) = \frac{\alpha \Pr(S_i \models t | t)}{\alpha \Pr(S_i \models t | t) + (1-\alpha) \Pr(S_i \models t | \neg t)}$$

$$\xrightarrow{(1)(2)} p_i = \frac{\alpha r_i}{\alpha r_i + (1-\alpha) q_i} \implies q_i = \frac{\alpha}{1-\alpha} \cdot \frac{1-p_i}{p_i} \cdot r_i$$

For our example, we would compute the precision of S as $\frac{4}{7} = 0.57$. Assuming $\alpha = 0.5$, we can derive $q_1 = \frac{0.5}{1-0.5} \cdot \frac{1-0.57}{0.57} \cdot 0.67 = 0.5$, implying that S_1 is a borderline source, with fairly low quality (recall that $r_1 = 0.67 > 0.5$). Note that for q_i to be valid, it needs to fall in the range of $[0, 1]$. The next theorem formally states the derivation and gives the condition for it to be valid.

THEOREM 3.5. *Let $S_i, i \in [1, n]$, be a source with precision p_i and recall r_i .*

- *If $\alpha \leq \frac{p_i}{p_i + r_i - p_i r_i}$, we have $q_i = \frac{\alpha}{1-\alpha} \cdot \frac{1-p_i}{p_i} \cdot r_i$*
- *If $p_i > \alpha$, S_i is a good source (i.e., $q_i < r_i$).*

Finally, we show in the next proposition that a triple provided by a high-precision source is more likely to be true, whereas a triple not provided by a good, high-recall source is more likely to be false, which is consistent with our intuitions.

PROPOSITION 3.6. *Let $S' = S \cup \{S'\}$ and $\mathcal{O}' = \mathcal{O} \cup \{\mathcal{O}'\}$. Let $S'' = S \cup \{S''\}$ and $\mathcal{O}'' = \mathcal{O} \cup \{\mathcal{O}''\}$. The following hold.*

- *If $r_{S'} = r_{S''}$, $p_{S'} > p_{S''}$, and $S' \models t$ and $S'' \models t$, then $\Pr(t | \mathcal{O}'_t) > \Pr(t | \mathcal{O}''_t)$.*
- *If $p_{S'} = p_{S''} > \alpha$, $r_{S'} > r_{S''}$, and $S' \not\models t$ and $S'' \not\models t$, then $\Pr(t | \mathcal{O}'_t) < \Pr(t | \mathcal{O}''_t)$.*

Comparison with LTM [25]

The closest work to our independent model is the *Latent Truth Model* (LTM) [25]; it treats source quality and triple correctness as latent variables and constructs a graphical model, and performs inference using Gibbs sampling. LTM is similar to our approach in that (1) it also assumes triple independence and open-world semantics, and (2) its probability computation also relies on recall and false positive rate of each source. However, there are three major differences. First, it derives the correctness probability of a triple from the Beta distribution of the recall and false positive rate of its providers; our model applies Bayesian analysis to maximize the *a posteriori* probability. Using the Beta distribution enforces assumptions about the generative process of the data, and when this model does not fit the actual dataset, LTM has a disadvantage against our non-parametric approach. Second, it computes the recall and false positive rate of a source as the Beta distribution of

the percentage of provided true triples and false triples; our model derives false positive rate from precision and recall to avoid being biased by very good sources or very bad sources. Third, it iteratively decides truthfulness of the triples and quality of the sources; we derive source quality from training data.

We compare LTM with our basic approach experimentally, showing that we have comparable results in general and sometimes better results; we also show that the correlation model we will present in the next section obtains considerably better results than LTM, which assumes independence of sources.

4. FUSING CORRELATED SOURCES

Theorems 3.1 and 3.5 summarize our approach for calculating the probability of a triple based on the precision and recall of each source. In this section, we extend the results to account for correlations among sources. Before we proceed, we first show several scenarios where considering correlation between sources can significantly improve the results.

EXAMPLE 4.1. Consider a set of n good sources $\mathcal{S} = \{S_1, \dots, S_n\}$. All sources in \mathcal{S} have the same recall r and false positive rate q , $r > q$. Given a triple t provided by all sources, Theorem 3.1 computes $\mu_{indep} = (\frac{r}{q})^n$.

Scenario 1 (Source copying): Assume that all sources in \mathcal{S} are replicas. Ideally, we want to consider them as one source; indeed, their joint recall is r and joint false positive rate is q . Thus, we compute $\mu_{corr} = \frac{r}{q} < \mu_{indep}$, which results in a lower probability for t ; in other words, a false triple would not get a high probability just because it is copied multiple times.

Scenario 2 (Sources overlapping on true triples): Assume that all sources in \mathcal{S} derive highly overlapping sets of true triples but each source makes independent mistakes (e.g., extractors that use different patterns to extract the same type of information). Accordingly, their joint recall is close to r and their joint false positive rate is q^n . Thus, we compute $\mu_{corr} \approx \frac{r}{q^n} > \mu_{indep}$, which results in a higher probability for t ; in other words, we will have much higher confidence for a triple provided by all sources.

Scenario 3 (Sources overlapping on false triples): Consider the opposite case: all sources have a high overlap on false triples but each source provides true triples independently (e.g., extractors that make the same kind of mistakes). In this case, the joint recall is r^n and the joint false positive rate is close to q . Thus, we compute $\mu_{corr} \approx \frac{r^n}{q} < \mu_{indep}$, which results in a lower probability for t ; in other words, considering correlations results in a much lower confidence for a common mistake.

Scenario 4 (Complementary sources): Assume that all sources are nearly complementary: their overlapping triples are rare but highly trustable (e.g., three extractors respectively focus on info-boxes, texts, and tables that appear on a Wikipedia page). Accordingly, the sources have low joint recall but very high joint precision; assume their joint recall is $r' \ll r$, and their joint false positive rate is q' , which is close to 0. Then, we compute $\mu_{corr} = \frac{r'}{q'} \approx \infty$; in other words, we highly trust the triples provided by all sources.

Under the same scenario, consider a triple t' provided by only one source $S \in \mathcal{S}$. Considering the negative correlation, the probability that a triple is provided only by S is r for true triples and q for false triples; thus, $\mu'_{corr} = \frac{r}{q} > \frac{r}{q} \cdot (\frac{1-r}{1-q})^{n-1} = \mu'_{indep}$, which results in higher probability for t' . In other words, considering the negative correlation, the correctness probability of a triple won't be penalized if only a single source provides the triple.

These scenarios exemplify the differences of our work and copy

detection in [5]. Copy detection can handle scenario 1 appropriately; in scenarios 2 and 3 it may incorrectly conclude with copying and compute lower probability for true triples; it cannot handle anti-correlation in Scenario 4.

We first present an *exact solution*, described by Theorem 4.2. However, exact computation is not feasible for problems involving a large number of sources, as the number of terms in the computation formula grows exponentially. In Section 4.2, we present an *aggressive approximation*, which can be computed in linear time by enforcing several assumptions, but may have low accuracy. Our *elastic approximation* (Section 4.3) relaxes the assumptions gradually, and can achieve both good efficiency and good results. Note that we can compute joint precision and joint recall, and derive joint false positive rate exactly the same way as we compute them for a single source (Section 3.2).

4.1 Exact solution

Recall that Eq. (6) and (7) compute $\Pr(\mathcal{O}_t | t)$ and $\Pr(\mathcal{O}_t | \neg t)$ by assuming independence between the sources. Now, we show how to compute them in the presence of correlations. Using \mathcal{S}_t to represent the set of sources that provide t , and $\mathcal{S}_{\bar{t}}$ to represent the set of sources that do not provide t , we can express $\Pr(\mathcal{O}_t | t)$ as:

$$\Pr(\mathcal{O}_t | t) = \Pr\left(\left(\bigwedge_{S \in \mathcal{S}_t} S \models t\right) \wedge \left(\bigwedge_{S' \in \mathcal{S}_{\bar{t}}} S' \not\models t\right) \mid t\right) \quad (9)$$

We apply the *inclusion-exclusion principle* to rewrite the formula using the joint recall of the sources:

$$\begin{aligned} \Pr(\mathcal{O}_t | t) &= \sum_{S^* \subseteq \mathcal{S}_{\bar{t}}} (-1)^{|S^*|} \Pr(\{\mathcal{S}_t \cup S^*\} \models t | t) \\ &= \sum_{S^* \subseteq \mathcal{S}_{\bar{t}}} (-1)^{|S^*|} r_{\mathcal{S}_t \cup S^*} \end{aligned} \quad (10)$$

Note that when the sources are independent, Eq. (10) computes exactly $\prod_{S_i \in \mathcal{S}_t} r_i \prod_{S_i \in \mathcal{S}_{\bar{t}}} (1 - r_i)$, which is equivalent to Eq. (6). We compute $\Pr(\mathcal{O} | \neg t)$ in a similar way, using the joint false positive rate of the sources, which can be derived from joint precision and joint recall as we described in Theorem 3.5:

$$\Pr(\mathcal{O}_t | \neg t) = \sum_{S^* \subseteq \mathcal{S}_{\bar{t}}} (-1)^{|S^*|} q_{\mathcal{S}_t \cup S^*} \quad (11)$$

Theorem 4.2 extends Theorem 3.1 for the case of correlated sources.

THEOREM 4.2. Given a set of sources $\mathcal{S} = \{S_1, \dots, S_n\}$, the joint recall and joint false positive rate for each subset of the sources, the probability of a triple t is $\Pr(t | \mathcal{O}) = \frac{1}{1 + \frac{1-\alpha}{\alpha} \cdot \frac{1}{\mu}}$, where

$$\mu = \frac{\Pr(\mathcal{O}_t | t)}{\Pr(\mathcal{O}_t | \neg t)} \quad (12)$$

and $\Pr(\mathcal{O}_t | t)$, $\Pr(\mathcal{O}_t | \neg t)$ are computed by Eq. (10) and (11).

COROLLARY 4.3. Given a set $\mathcal{S} = \{S_1, \dots, S_n\}$, where all sources are independent, the correctness probabilities computed using Theorems 3.1 and 4.2 are equal.

EXAMPLE 4.4. Triple t_8 of Figure 1a is provided by $\mathcal{S}_{t_8} = \{S_1, S_2, S_4, S_5\}$. We use notations $r_{\{S_1, S_2, S_4, S_5\}}$ and r_{1245} interchangeably. We can compute joint recall for a set of sources as we do for a single source (Section 3.2), but here we assume that all the joint recall and joint false positive rate parameters are given.

	S_1	S_2	S_3	S_4	S_5
C^+	$\frac{0.11}{0.67 \cdot 0.167} = 1$	1	0.75	1.5	1.5
C^-	$\frac{0.037}{0.5 \cdot 0.037} = 2$	1	1	3	3

Figure 3: Correlation parameters of the aggressive approximation computed for each source of Figure 1a.

We compute $\Pr(\mathcal{O}_t | t_8)$ and $\Pr(\mathcal{O}_t | \neg t_8)$, according to Eq. (10):

$$\Pr(\mathcal{O}_{t_8} | t_8) = r_{1245} - r_{12345} = 0.22 - 0.11 = 0.11$$

$$\Pr(\mathcal{O}_{t_8} | \neg t_8) = q_{1245} - q_{12345} = 0.22 - 0.037 = 0.185$$

Assuming a-priori probability $\alpha = 0.5$, we derive $\Pr(t_8 | \mathcal{O}) = \frac{1}{1 + \frac{0.185}{0.11}} = 0.37$. Note that although t_8 is provided by four out of the five sources, S_1, S_4 , and S_5 are correlated, which reduces their contribution to the correctness probability of t_8 . Using correlations allows us to correctly classify t_8 as false, whereas the independence assumption leads to the wrong result, as shown in Example 3.3.

Even though accounting for correlations can significantly improve accuracy, it increases the computational cost. The computation of $\Pr(\mathcal{O}_t | t)$ and $\Pr(\mathcal{O}_t | \neg t)$ is exponential in the number of sources that do not provide t , thus impractical when we have a large number of sources. We next describe two ways to approximate $\Pr(\mathcal{O}_t | t)$ and $\Pr(\mathcal{O}_t | \neg t)$.

4.2 Aggressive approximation

In this section, we present a linear approximation that reduces the total number of terms in the computation by enforcing a set of assumptions. We first present the main result for the approximation in Definition 4.5, and we show how we derive it later.

DEFINITION 4.5 (AGGRESSIVE APPROXIMATION). *Given a set of sources $\mathcal{S} = \{S_1, \dots, S_n\}$, the recall r_i and false positive rate q_i of each source S_i , and the joint recall and joint false positive rate for sets \mathcal{S} and $\mathcal{S} - S_i$, the aggressive approximation of the probability $\Pr(t | \mathcal{O}_t)$ is defined as: $\frac{1}{1 + \frac{1-\alpha}{\alpha} \cdot \frac{1}{\mu_{agg}}}$, where*

$$\mu_{agg} = \prod_{S_i \in \mathcal{S}_t} \frac{C_i^+ r_i}{C_i^- q_i} \prod_{S_i \in \mathcal{S}_{\bar{t}}} \left(\frac{1 - C_i^+ r_i}{1 - C_i^- q_i} \right) \quad (13)$$

$$C_i^+ = \frac{r_{1\dots n}}{r_i \cdot r_{12\dots(i-1)(i+1)\dots n}} \quad (14)$$

$$C_i^- = \frac{q_{1\dots n}}{q_i \cdot q_{12\dots(i-1)(i+1)\dots n}} \quad (15)$$

Eq. (13) differs from (8) in that it replaces r_i (resp. q_i) with $C_i^+ r_i$ (resp. $C_i^- q_i$). Intuitively, C_i^+ and C_i^- represent the correlation between S_i and the rest of \mathcal{S} , in the case of true and false triples respectively. Eq. (13) weighs r_i and q_i by these ‘‘correlation’’ parameters. When the sources are independent, $C_i^+ = C_i^- = 1$, and the approximation obtains the same result as Theorem 3.1. In contrast with Definition 2.4, aggressive approximation only uses $2n+1$ instead of $2(2^n - n - 1)$ correlation parameters.

COROLLARY 4.6. *Given a set $\mathcal{S} = \{S_1, \dots, S_n\}$, where all sources are independent, the correctness probabilities computed using Theorem 3.1 and Definition 4.5 are the same.*

EXAMPLE 4.7. *Consider triple t_8 in Figure 1a. Figure 3 shows the correlation parameters for each source and illustrates how they are computed for S_1 . These parameters indicate that S_1, S_4 and*

S_5 are positively correlated for false triples; for true triples, S_3 is anti-correlated with the rest of the sources, whereas S_4 and S_5 are correlated. Accordingly, we compute μ_{agg} as follows:

$$\mu_{agg} = \frac{0.67 \cdot 0.5 \cdot (1 - 0.75 \cdot 0.67) \cdot 1.5 \cdot 0.67 \cdot 1.5 \cdot 0.67}{2 \cdot 0.5 \cdot 0.67 \cdot (1 - 0.167) \cdot 3 \cdot 0.33 \cdot 3 \cdot 0.33} = 0.3$$

Thus, we compute $\Pr(t_8 | \mathcal{O}) = \frac{1}{1 + \frac{1}{\mu_{agg}}} = 0.23$, which is lower than the exact computation in Example 4.4. Both approaches correctly determine that t_8 is false.

Obviously, the computation is linear in the number of sources. Also, instead of having an exponential number of joint recall and false positive rate values, we only need C_i^+ and C_i^- for each $S_i, i \in [1, n]$, which can be derived from a linear number of joint recall and false positive rate values. However, as the following proposition shows, this aggressive approach can produce bad results for special cases with strong correlation (i.e., sources are replicas), or strong anti-correlation (i.e., sources are complementary to each other).

PROPOSITION 4.8. *If all sources in \mathcal{S} provide the same data, Definition 4.5 computes probability α for each provided triple.*

If every pair of sources in \mathcal{S} are complementary to each other, Definition 4.5 does not compute a valid probability for any triple.

Next, we proceed to describe the three major steps that lead to Definition 4.5.

I. Correlation factors

Accounting for correlations, the probability that a set \mathcal{S}^* of sources all provide a true triple is $r_{\mathcal{S}^*}$ instead of $\prod_{S_i \in \mathcal{S}^*} r_i$ (similarly for a false triple). We define two *correlation factors*: $C_{\mathcal{S}^*}$ and $C_{\mathcal{S}^*}^-$:

$$C_{\mathcal{S}^*} = \frac{\Pr(\mathcal{S}^* \models t | t)}{\Pr_{indep}(\mathcal{S}^* \models t | t)} = \frac{r_{\mathcal{S}^*}}{\prod_{S_i \in \mathcal{S}^*} r_i} \quad (16)$$

$$C_{\mathcal{S}^*}^- = \frac{\Pr(\mathcal{S}^* \models t | \neg t)}{\Pr_{indep}(\mathcal{S}^* \models t | \neg t)} = \frac{q_{\mathcal{S}^*}}{\prod_{S_i \in \mathcal{S}^*} q_i} \quad (17)$$

If the sources in \mathcal{S}^* are independent, then $C_{\mathcal{S}^*} = C_{\mathcal{S}^*}^- = 1$. Deviation from independence may produce values greater than 1, which imply *positive correlations* (e.g., for S_4 and S_5 in Figure 1a, $C_{45} = \frac{0.67}{0.67 \cdot 0.67} = 1.5 > 1$), or lower than 1, which imply *negative correlations*, also known as *anti-correlations* (e.g., for S_1, S_3 in Figure 1a, $C_{13} = \frac{0.33}{0.67 \cdot 0.67} = 0.75 < 1$).

Using separate parameters for true triples and false triples allows for a richer representation of correlations. In fact, two sources may be correlated differently for true and false triples. For example, sources S_2 and S_3 in Figure 1a are independent with respect to true triples ($C_{23} = 1$), and negatively correlated with respect to false triples ($C_{23}^- = 0.5 < 1$).

Using correlation factors, $\Pr(\mathcal{O} | t_8)$ in our running example can be rewritten as follows:

$$\Pr(\mathcal{O}_{t_8} | t_8) = C_{1245} r_1 r_2 r_4 r_5 - C_{12345} r_1 r_2 r_3 r_4 r_5$$

II. Assumptions on correlation factors

To transform the equations with correlation factors into a simpler form, we make *partial independence assumptions*. Before we formally state the assumptions, we first illustrate it using an example.

EXAMPLE 4.9. *Consider sources $\mathcal{S} = \{S_1 \dots S_5\}$ and assume S_4 is independent of the set of sources $\{S_1, S_2, S_3\}$ and of sources $\{S_1, S_2, S_3, S_5\}$. Then, we have*

$$r_{123} \cdot r_4 = r_{1234}$$

$$r_{12345} = r_{1235} \cdot r_4$$

Thus, $r_{123}r_{12345} = r_{1234}r_{1235}$. Using the definition of the correlation factors, it follows that $C_{123}C_{12345} = C_{1234}C_{1235}$.

Accordingly, we can rewrite the correlation factors: for example $C_{123} = \frac{C_{1234}C_{1235}}{C_{12345}}$, and similarly, $C_{23} = \frac{C_{1234}C_{1235}C_{2345}}{(C_{12345})^2}$. Combining Eqs (14) and (16), the following equations hold: $C_1^+ = \frac{C_{12345}}{C_{2345}}$, $C_4^+ = \frac{C_{12345}}{C_{1235}}$, $C_5^+ = \frac{C_{12345}}{C_{1234}}$. Using these, we can rewrite the following correlation factors: $C_{123} = \frac{C_{12345}}{C_4^+ C_5^+}$ and $C_{23} = \frac{C_{12345}}{C_1^+ C_4^+ C_5^+}$.

According to Eqs.(14–17), we can compute $C_i^+ = \frac{C_S}{C_{S \setminus \{s_i\}}}$ and $C_i^- = \frac{C_{\bar{S}}}{C_{\bar{S} \setminus \{s_i\}}}$. As illustrated in Example 4.9, partial independence assumptions lead to the following equations:

$$C_{S^*} = \frac{C_S}{\prod_{S_i \in S \setminus S^*} C_i^+} \quad \text{and} \quad C_{\bar{S}^*} = \frac{C_{\bar{S}}}{\prod_{S_i \in S \setminus S^*} C_i^-} \quad (18)$$

As a special form, when $S^* = \emptyset$, we have $C_{S^*} = C_{\bar{S}^*} = 1$, so,

$$C_S = \prod_{S_i \in S} C_i^+ \quad \text{and} \quad C_{\bar{S}} = \prod_{S_i \in S} C_i^- \quad (19)$$

Under these assumptions, $\Pr(\mathcal{O} | t_8)$ in our running example can be rewritten as follows:

$$\Pr(\mathcal{O}_{t_8} | t_8) = \frac{C_{12345}}{C_3^+} r_1 r_2 r_4 r_5 - C_{12345} r_1 r_2 r_3 r_4 r_5$$

III. Transformation

We are ready to transform the equation into a simpler form, which is the same as the one in Definition 4.5. We continue illustrating the main intuition with our running example:

$$\begin{aligned} \Pr(\mathcal{O}_{t_8} | t_8) &= \frac{C_{12345}}{C_3^+} r_1 r_2 r_4 r_5 (1 - C_3^+ r_3) \\ &\stackrel{(19)}{=} \frac{C_1^+ C_2^+ C_3^+ C_4^+ C_5^+}{C_3^+} r_1 r_2 r_4 r_5 (1 - C_3^+ r_3) \\ &= (C_1^+ r_1)(C_2^+ r_2)(1 - C_3^+ r_3)(C_4^+ r_4)(C_5^+ r_5) \quad (20) \end{aligned}$$

4.3 Elastic approximation

So far we have presented two solutions: the exact solution gives precise probabilities but is computationally expensive; the aggressive approximation enforces partial independence assumptions resulting in linear complexity, but in the worst case can compute probabilities independent of the quality of the sources. In this section, we present an elastic approximation algorithm that makes a tradeoff between efficiency and quality.

The key idea of the elastic approximation is to use the linear approximation as a starting point and gradually adjust the results by relaxing the assumptions in every step. We call the algorithm “elastic” because it can be configured to iterate over different levels of adjustments, depending on the desired level of approximation. We illustrate this idea with our running example.

EXAMPLE 4.10. Triple t_8 is provided by four sources $\mathcal{S}_{t_8} = \{S_1, S_2, S_4, S_5\}$ (Figure 1a). We will adjust the linear approximation of $\Pr(\mathcal{O}_{t_8} | t_8)$ from Eq. (20), by adding specific terms at every level. We refer to the degree of a term in the aggressive approximation, as the number of recall (or false positive rate) parameters associated with that term. The aggressive approximation for $\Pr(\mathcal{O}_{t_8} | t_8)$ contains two terms of degrees 4 and 5: $C_1^+ C_2^+ C_4^+ C_5^+ r_1 r_2 r_4 r_5$ and $C_1^+ C_2^+ C_3^+ C_4^+ C_5^+ r_1 r_2 r_3 r_4 r_5$ respectively (directly derived from Eq. (20)).

Algorithm 1 ELASTIC (Elastic approximation)

```

1:  $R \leftarrow r_{S_t} \prod_{S_i \in \mathcal{S}_{\bar{t}}} (1 - C_i^+ r_i)$ ;
2:  $Q \leftarrow q_{S_t} \prod_{S_i \in \mathcal{S}_{\bar{t}}} (1 - C_i^- q_i)$ ;
3: for  $l = 1 \rightarrow \lambda$  do  $\triangleright \lambda \geq 1$  is the desired adjustment level
4:   for all subsets  $S^* \subseteq \mathcal{S}_{\bar{t}}$  of size  $l$  do
5:      $S_l \leftarrow \{S_t \cup S^*\}$ ;
6:      $R \leftarrow R + (-1)^l (C_{S_l} - C_{S_l} \prod_{S_i \in S^*} C_i^+) \prod_{S_i \in S_l} r_i$ ;
7:      $Q \leftarrow Q + (-1)^l (C_{\bar{S}_l} - C_{\bar{S}_l} \prod_{S_i \in S^*} C_i^-) \prod_{S_i \in S_l} q_i$ ;
8: return  $\frac{R}{Q}$ ;

```

Elastic approximation makes corrections to the aggressive approximation based on terms of a given degree at every level. At level-0 we consider the terms with degree of $|\mathcal{S}_{t_8}| + 0 = 4$, i.e., the term $C_1^+ C_2^+ C_4^+ C_5^+ r_1 r_2 r_4 r_5$; the exact coefficient of the term is C_{1245} but we approximated it to $C_1^+ C_2^+ C_4^+ C_5^+$ based on the assumption that $C_{1245} = \frac{C_{12345}}{C_3^+} = C_1^+ C_2^+ C_4^+ C_5^+$. To remove the assumption, we need to replace $(C_1^+ r_1)(C_2^+ r_2)(C_4^+ r_4)(C_5^+ r_5)$ with $C_{1245} r_1 r_2 r_4 r_5 = r_{1245}$. Since $r_{1245} = q_{1245} = 0.22$, we have

$$\mu = \frac{0.22}{0.22} \cdot \frac{1 - 0.75 \cdot 0.67}{1 - 0.167} = 0.6$$

Note that the level-0 adjustment affects not only terms with degree 4, but actually all terms as we show next.

At level-1, we consider the terms with degree of $|\mathcal{S}_{t_8}| + 1 = 5$. After level-0 adjustment, the 5-degree term is $C_{1245} C_3^+ r_1 r_2 r_3 r_4 r_5$. We will replace $C_{1245} C_3^+$ with the exact coefficient C_{12345} , which will now give us the exact solution.

In summary, the μ_{agg} parameter calculated by the aggressive approximation, the level-0 adjustment, and the level-1 adjustment are 0.3, 0.6, and 0.59 respectively. Note that, as is the case in this example, we don’t need to compute all the levels; stopping after a constant number of levels can get close to the exact solution.

Our ELASTIC algorithm (Algorithm 1) contains the pseudo code of our elastic approximation. Lines 1–2 compute the initial values of the numerator R and denominator Q for μ . Note that they have already applied the level-0 adjustment. Then for each level l from 1 up to the required level λ (line 3), we consider each term with degree $|\mathcal{S}_t| + l$ (lines 4–5), and make up the difference between the exact coefficient and the approximate coefficient (lines 6–7). Finally, line 8 returns $\frac{R}{Q}$ as the value of μ .

PROPOSITION 4.11. Given a set of n sources, a set of m triples for probability computation, and an approximation level λ , ELASTIC takes times $O(m \cdot n^\lambda)$ and the number of required correlation parameters is in $O(m \cdot n^\lambda)$.

5. EVALUATION

This section describes a thorough evaluation of our models on three real-world datasets as well as synthetic data. Our experimental results show that (1) considering correlations between sources can significantly improve fusion results; (2) our elastic approximation can effectively estimate triple probability with much shorter execution time; and (3) even in presence of only independent sources, our model can outperform state-of-the-art data fusion approaches.

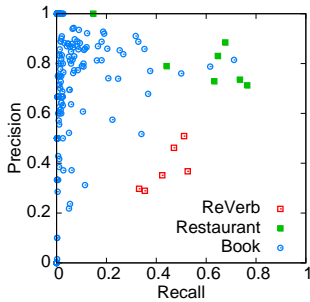
Datasets

We first describe the real-world datasets we used in our experiments; we describe our synthetic data generation in Section 5.2.

REVERB: We used the ReVerb ClueWeb Extraction dataset from [11], which samples 500 sentences from the Web using Yahoo’s random link service and uses 6 extractors to extract triples from these sentences. The gold standard contains 2407 extracted triples (616 true and 1791 false).

RESTAURANT: We used the restaurant dataset from [17], which consists of triples on the location of a collection of 1000 restaurants provided by 7 sources (*Yelp, Foursquare, OpenTable, Mechanical-Turk, YellowPages, CitySearch, MenuPages*). The gold standard contains 93 triples (68 true and 25 false), selected by majority vote over 10 Mechanical Turk responses.

BOOK: We used the book dataset from [6], collected by crawling *abebooks.com*. The dataset consists of 5900 unique book-author triples from 879 seller sources. The gold standard consists of 225 randomly sampled books for which the authors are manually identified from book covers; 482 authors are correctly provided for these books and 935 authors are wrongly provided. Note that our version of this dataset has more noise than the one used in [25], resulting in a more challenging setting.



We observe that these datasets display varied characteristics: the sources in RESTAURANT all have high precision, and most have high recall; the sources in REVERB have fairly low precision and recall; the sources in BOOK have large variations in precision, and most of them have low recall. Such differences allow us to evaluate our models in a variety of scenarios.

Comparisons

We compared our models with several state-of-the-art techniques that apply to the independent-triple and open-world semantics.

UNION- K : Considers a triple to be true if at least $K\%$ of the sources provide it. Union-50 is equivalent to majority voting.

3-ESTIMATE [13]: Iteratively computes trustworthiness of sources, trustworthiness of triples, and truthfulness of triples. This is the best model among the three proposed in [13], and we observed similar results from the other two models on our datasets.

LTM [25]: Constructs a graphical model and uses Gibbs sampling to determine source quality and truthfulness of each triple. We used the default parameters suggested by [25].

PRECREC (Section 3): Computes truthfulness of each triple from the precision and recall of each source. We set $\alpha = 0.5$ and computed source precision and recall according to the gold standard.

PRECCORR (Section 4): Extends PRECREC by considering correlation between sources. By default we report the results for the exact solution; however, as we show in Figure 5, we obtain similar results using level-3 elastic approximation. We computed joint precision and recall according to the gold standard. Note that BOOK is considerably larger than the other two datasets, which poses challenges for deriving the correlation parameters: (a) the number of correlation parameters is very large, and (b) there may not be enough support data to understand the correlation among the sources. We overcome this issue using a simple clustering approach: we divide sources into clusters based on their pairwise correlations, and assume that sources across clusters are independent.

We used a C# implementation of LTM and we implemented the other models in Java. For REVERB, RESTAURANT, and synthetic

data, we ran experiments on a Macbook Air with 4GB RAM, 1.7 GHz Intel Core i5 processor, and OSX Lion 10.7.5. The BOOK experiments were run on a m1.large Amazon EC2 server instance [1].

Metrics

We present results according to three metrics.

Precision/Recall/F1: We measure the correctness of binary decisions with three metrics. *Precision* measures among the returned true triples, how many are indeed true; *recall* measures among the provided true triples, how many are returned; *F-measure* computes their harmonic mean (*i.e.*, $F1 = \frac{2 \cdot \text{prec} \cdot \text{rec}}{\text{prec} + \text{rec}}$).

PR-curve/ROC-curve: We rank the provided triples in decreasing order of the computed truthfulness score (for UNION- K , we rank in decreasing order of the number of providers). As we add the triples gradually, *PR-curve* plots the precision versus the recall after adding each triple and *ROC-curve* plots the true positive rate versus the false positive rate. In addition, we compute the area under the curve, called *AUC-PR* and *AUC-ROC* respectively. These curves and measures allow us to examine whether the correctness probabilities we compute are consistent with the reality.

Execution time: We report execution time for each method.

5.1 Real-World Data

We first compare the different models on the three real-world data sets. Figure 4 reports the precision, recall, and F-measure of each method on each dataset. We also plot the PR-curve and ROC-curve of the methods on each data set. Note that the curves for UNION- K of different K are the same so we plot only one; also note that the results of 3-ESTIMATE are significantly worse than other methods, so we did not plot its curves to avoid cluttering.

Overall, we observe that among different datasets, most of the methods obtain higher quality results on RESTAURANT and BOOK, but lower quality on REVERB. This is not surprising given that the data sources in REVERB have fairly low precision and recall and they extract a lot of wrong triples. PRECCORR obtains the best results on all datasets: comparing with PRECREC, its F-measure is 5.2% higher on average, its AUC-PR is 10.3% higher on average, and its AUC-ROC is 3.3% higher on average. We note that although the improvement on F-measure is not that large, the improvement for AUC-PR and AUC-ROC is significant; this is because with consideration of correlations between the sources, we often compute a much higher probability for a true triple and a lower probability for a false triple, but this difference may be hindered when we apply the threshold and make binary decisions.

Among the methods that assume independence between sources, PRECREC obtains the best results: on average its F-measure is 14% higher than LTM and 41% higher than 3-ESTIMATE. For LTM, its F-measure is comparable to PRECREC on RESTAURANT and BOOK, but much lower on REVERB because of a very low precision. Its PR-curves and ROC-curves are not in a very good shape; indeed, its AUC-PR is 24% lower than PRECREC and its AUC-ROC is 20.8% lower on average. We observed that the probabilities it outputs typically fall in extreme ranges; for example, for most of the triples that it considers as true on RESTAURANT, it computes a probability very close to 1. 3-ESTIMATE obtains very low recall in all of the three datasets; as a result, its F-measure is the lowest among all methods.

For UNION- K , increasing K increases the precision but drops the recall. UNION-25 turns out to have the best F-measure, comparable to PRECREC on each data set, but lower than PRECCORR. However, its PR-curves and ROC-curves are in slightly worse shapes comparing with PRECREC; indeed, its AUC-PR and AUC-ROC is lower than that of PRECREC by up to 4.5%. As we show later




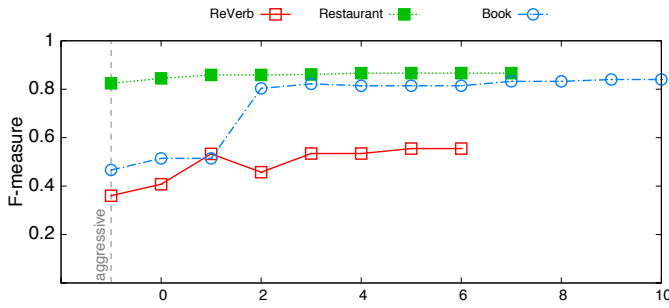
Figure 4: Our experiments show that PRECREC and PRECRECCORR result in better fusion results compared to other approaches. In the REVERB dataset both PRECREC and PRECRECCORR showed significant improvement in the F-measure compared to the state-of-the-art (3-ESTIMATE, and LTM). In the RESTAURANT and BOOK datasets, LTM and UNION-25 are comparable to the results of PRECREC, but the PR and ROC curves demonstrate that PRECRECCORR provides significantly better truthfulness estimates for triples.

on synthetic data, UNION- K is sensitive on source quality; for example, even UNION-25 can obtain very low F-measure when the sources have low precision or low recall.

Figure 5b shows the execution time of the different models. UNION- K is very efficient, while 3-ESTIMATE and PRECREC are the next most efficient, with runtimes up to one order of magnitude longer than UNION. We terminated LTM after 10 iterations; each iteration on average took 5.6 times longer than PRECREC. PRECRECCORR is one order of magnitude slower than PRECREC on average; however, the level-3 elastic approximation obtained similar results but finished in only half of the time. For our largest dataset (BOOK), level-3 approximated the exact solution in 40 minutes; we consider these runtimes reasonable, since this is an offline cleaning process. Parallelization can significantly improve the efficiency of PRECRECCORR, as the terms at different levels and across different clusters can be computed independently. With maximum parallelization PRECRECCORR terminates in 80 seconds, however a systematic study of these improvements is outside the scope of this paper.

Elastic approximation: Figure 5 demonstrates the behavior of our aggressive approximation and elastic approximation (Algorithm 1) over the three datasets. We observe that the aggressive estimate is much worse than the exact solution on REVERB and RESTAURANT, while comparable on BOOK; it is even worse than PRECREC, which does not consider correlation. Each line in the graph shows the progression of the approximation from the aggressive estimate to the exact computation. At every level, the elastic approximation refines the probability estimates of the earlier levels to gradually approach PRECRECCORR. Since the elastic approximation is heuristic in nature, there is no guarantee that the method improves the estimate with every level (*e.g.*, on REVERB the elastic approximation performs worse at level 2 than level 1). However, for all datasets, the elastic approximation comes close to the exact result within a small number of levels. We observe that on all three data sets, the result of level-3 approximation is already quite close to that of the exact solution, whereas the execution time is much shorter.

Discovered correlations: To better understand the improvement

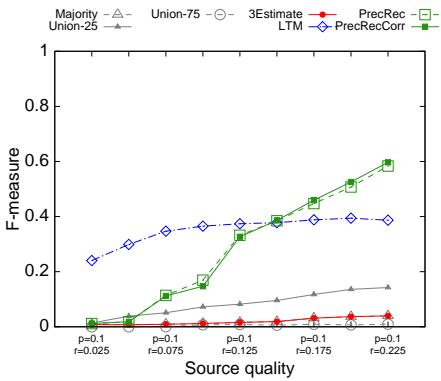


(a) Elastic approximation levels

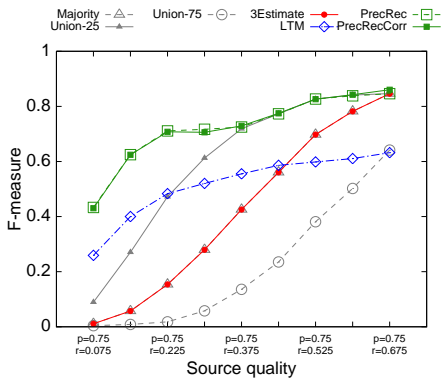
time(sec)	REVERB	RESTAURANT	BOOK
UNION-25	0.39	0.56	3.86
UNION-50	0.14	0.32	3.71
UNION-75	0.11	0.35	3.00
3-ESTIMATE	0.7	0.06	39
LTM (10 iter)	49	5.3	3791
PRECREC	2.6	0.3	35
PRECRECCORR	124	5.4	6786
PRECRECCORR-LVL3	79	2.25	2452

(b) Runtimes of algorithms (in seconds) for all datasets.


Figure 5: As expected, our elastic approximation gradually approaches the result of PRECRECCORR.



(a) Low precision sources, with low to fair recall, in a dataset of 25% true triples.



(b) High precision sources, with increasing recall, in a dataset of 50% true triples.



(c) Low recall sources, with increasing precision, in a dataset of 25% true triples.

Figure 6: Experimental results on synthetic data with independent sources. Our techniques are particularly effective with sources of low quality, and demonstrate significant gains in many configurations.

of PRECRECCORR over PRECREC, we examine in more detail the discovered correlations between the sources.

REVERB has 6 sources. With respect to true triples, we detect strong correlation on a group of 2 sources and on a group of 3 sources. With respect to false triples, 2 pairs of sources are strongly correlated, and one source is strongly anti-correlated with every other source. Of the 7 RESTAURANT sources, we detect strong correlation on a group of 4 sources and fairly strong anti-correlation on a pair of sources, with respect to true triples. For false triples, there is strong correlation on a group of 6 sources. Finally, for BOOK, there are 333 sources that provide triples in the gold standard. Recall that we cluster the sources according to their correlation. In terms of true triples, we obtain three clusters of size 22, 3, and 2. In terms of false triples, we obtain four clusters of size 22, 3, 2, and 2. Interestingly, except two sources between which we find strong correlation both on true triples and on false triples, the clusters for true triples and for false triples contain very different sources.

These observations indicate that our model of correlation is much richer than what can be captured by pure copying relationships, as in [6]. For our datasets, [6] applies only to BOOK dataset by considering the author list as a whole, but not the other datasets. In BOOK, this approach achieves high precision of 0.97 as it successfully detects copying and reduces the vote counts of false values. However, it has a low recall of 0.82, since it also discounts vote counts on true values and ignores other types of correlations. We leave an effective combination of that approach and ours for future work.

5.2 Synthetic Data

We generated synthetic data to evaluate our algorithms under a

large range of scenarios; in this section we present interesting cases that arise both in the case of independent sources, as well as in the case of correlations.

Our first set of experiments compares the different models on independent sources. We generated 5 sources providing data on 1000 triples according to a pre-configured precision and recall; we averaged 10 repetitions and show the results in Figure 6. Our results show that even without correlations, PRECREC provides significant improvements over existing approaches, while PRECRECCORR has similar performance. Figure 6a shows the performance of all the algorithms against a dataset of low quality sources. LTM is quite robust to variations in source quality, and performs well in this challenging setting; however, it does not benefit much from increases in source quality, and PRECREC quickly becomes better as recall increases over 0.15. In Figures 6b and 6c, we vary recall and precision respectively, while keeping the other constant. In both cases, our techniques perform remarkably well in comparison to the other algorithms. Note that UNION-25 is very sensitive to source quality and performs badly with low-quality sources.

Our second set of experiments considers correlated sources. Figure 7 demonstrates two cases: (1) a set of four sources are positively correlated on true triples, and (2) the sources are negatively correlated on false triples. In both cases, PRECRECCORR demonstrates significantly better performance than all the other approaches.

6. RELATED WORK

There has been extensive work in the area of data fusion (*i.e.*, resolving conflicts and finding the truth); [4, 8] surveyed early ap-




Figure 7: Experimental results on synthetic data with correlated sources. PRECRECCORR obtains better results compared to all other approaches.

proaches and [15] compared recent approaches on Deep Web data. Among these approaches, [6, 14, 19, 20, 21, 23, 24] jointly infer truth and source quality, but they assume the conflicting-triple, closed-world semantics. COSINE and 3ESTIMATE [13] can be applied under the independent-triple, open-world semantics. Instead of using precision and recall of sources, it considers a single quality metric—*accuracy* of a source; we compared with them in our experiments (Section 5). The model closest to ours is LTM [25]; we have made detailed comparisons in Section 3 and in experiments. All of these approaches assume independence between sources.

Correlation between sources are studied in two bodies of works. First, copy detection has been surveyed in [10] for various types of data and studied in [3, 5, 6, 7, 16] for structured data. Our approach is different in three aspects. First, in addition to copying, we consider broader scopes of correlations, including positive correlations not caused by copying (*e.g.*, extractors employing common extraction patterns), and also negative correlations. Second, instead of just discounting votes from copiers, we may boost contributions from providers correlated on true triples and reduce penalty from non-providers anti-correlated on true triples. Third, we assume independent-triple and open-world semantics, opposite to their conflicting-triple, closed-world semantics. We have compared with this approach in our experiments.

Second, there are other ways of measuring correlations. Qi et al. [22] constructed a graphical model that clusters dependent sources into groups and measures the quality of each group as a whole (instead of each individual source). Kappa measure [12] measures correlation by taking into account the agreement by chance. We measure correlations by the joint precision and recall for subsets of sources. Our measures have much higher expressiveness in that (1) they consider both positive and negative correlations; (2) they distinguish correlation on true data and on false data; and (3) they essentially consider correlation for every subset of sources.

7. CONTRIBUTIONS AND FUTURE WORK

In this paper we presented a novel technique for fusing data that contains correlations, which uses Bayesian analysis to derive the truthfulness of a fact based on the quality of sources that provide it. We evaluated our approach against other state-of-the-art techniques, and showed that our algorithms achieve significant improvements in the fusion results. The power of our approach lies in its generality: our algorithms do not need to have any knowledge of possible correlations, and all required parameters can be computed from a training set. As a result, PRECREC and PRECRECCORR perform well even in low quality datasets that prove challenging for other techniques.

There are still several interesting challenges in this problem. Our model uses independent-triple, open-world semantics, which al-

lows our techniques to consider multiple truth values for an entity (*e.g.*, a person may have multiple professions). However, this assumption may not always apply (*e.g.*, a person only has a single birth date). We consider modifications in our model to account for such scenarios in future work. Another challenge is that source quality may vary, based on the domain. For example, a source may have low overall precision, but may be particularly accurate with respect to Pizzerias, or restaurants in the Bay Area. In our model, we can consider domains separately, but deriving the proper domain subdivisions automatically is not straightforward.

Acknowledgements: We thank the authors of [25] for providing us with the implementation of LTM. This work was partially supported by NSF CCF-1349784 and a Google faculty research award.

8. REFERENCES

- [1] Amazon EC2 Instances. <http://aws.amazon.com/ec2/instance-types/>.
- [2] L. Berti-Equille, A. D. Sarma, X. Dong, A. Marian, and D. Srivastava. Sailing the information ocean with awareness of currents: Discovery and application of source dependence. In *CIDR*, 2009.
- [3] L. Blanco, V. Crescenzi, P. Merialdo, and P. Papotti. Probabilistic models to reconcile complex data from inaccurate data sources. In *CAISE*, 2010.
- [4] J. Bleiholder, K. Draba, and F. Naumann. FuSem-exploring different semantics of data fusion. In *VLDB*, pages 1350–1353, 2007.
- [5] X. L. Dong, L. Berti-Equille, Y. Hu, and D. Srivastava. Global detection of complex copying relationships between sources. *PVLDB*, 3(1-2):1358–1369, Sept. 2010.
- [6] X. L. Dong, L. Berti-Equille, and D. Srivastava. Integrating conflicting data: The role of source dependence. *PVLDB*, 2(1):550–561, 2009.
- [7] X. L. Dong, L. Berti-Equille, and D. Srivastava. Truth discovery and copying detection in a dynamic world. *PVLDB*, 2(1), 2009.
- [8] X. L. Dong and F. Naumann. Data fusion—resolving data conflicts for integration. *PVLDB*, 2009.
- [9] X. L. Dong, B. Saha, and D. Srivastava. Less is more: Selecting sources wisely for integration. *PVLDB*, 6(2):37–48, Dec. 2012.
- [10] X. L. Dong and D. Srivastava. Large-scale copying detection. In *SIGMOD (Tutorial)*, 2011.
- [11] A. Fader, S. Soderland, and O. Etzioni. Identifying relations for open information extraction. In *EMNLP*, 2011.
- [12] J. Fleiss. *Statistical methods for rates and proportions*. John Wiley and Sons, 1981.
- [13] A. Galland, S. Abiteboul, A. Marian, and P. Senellart. Corroborating information from disagreeing views. In *WSDM*, 2010.
- [14] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. In *SODA*, 1998.
- [15] X. Li, X. L. Dong, K. B. Lyons, W. Meng, and D. Srivastava. Truth finding on the deep web: Is the problem solved? *PVLDB*, 6(2), 2013.
- [16] X. Liu, X. L. Dong, B.-C. Ooi, and D. Srivastava. Online data fusion. *PVLDB*, 4(12), 2011.
- [17] A. Marian and M. Wu. Corroborating information from web sources. *IEEE Data Eng. Bull.*, 34(3):11–17, 2011.
- [18] M. Mintz, S. Bills, R. Snow, and D. Jurafsky. Distant supervision for relation extraction without labeled data. In *ACL*, 2009.
- [19] J. Pasternack and D. Roth. Knowing what to believe (when you already know something). In *COLING*, pages 877–885, 2010.
- [20] J. Pasternack and D. Roth. Making better informed trust decisions with generalized fact-finding. In *IJCAI*, pages 2324–2329, 2011.
- [21] J. Pasternack and D. Roth. Latent credibility analysis. In *WWW*, 2013.
- [22] G.-J. Qi, C. Aggarwal, J. Han, and T. Huang. Mining collective intelligence in groups. In *WWW*, 2013.
- [23] X. Yin, J. Han, and P. S. Yu. Truth discovery with multiple conflicting information providers on the web. In *SIGKDD*, 2007.
- [24] X. Yin and W. Tan. Semi-supervised truth discovery. In *WWW*, 2011.
- [25] B. Zhao, B. I. P. Rubinstein, J. Gemmell, and J. Han. A Bayesian approach to discovering truth from conflicting sources for data integration. *PVLDB*, 5(6):550–561, 2012.