Evaluating the Integration of Datasets

Matteo Paganelli
University of Modena and Reggio Emilia
Modena, Italy
matteo.paganelli@unimore.it

Francesco Del Buono
University of Modena and Reggio Emilia
Modena, Italy
francesco.delbuono@unimore.it

Francesco Guerra
University of Modena and Reggio Emilia
Modena, Italy
francesco.guerra@unimore.it

Nicola Ferro
University of Padua
Padua, Italy
ferro@dei.unipd.it

ABSTRACT

Evaluation is a bottleneck in data integration processes: it is performed by domain experts through manual onerous data inspections. This task is particularly heavy in real business scenarios, where the large amount of data makes checking all integrated tuples infeasible. Our idea is to address this issue by providing the experts with an unsupervised measure, based on word frequencies, which quantifies how much a dataset is representative of another dataset, giving an indication of how good is the integration process. The paper motivates and introduces the measure and provides extensive experimental evaluations, that show the effectiveness and the efficiency of the approach.

CCS CONCEPTS

• **Information systems** → **Mediators and data integration**; **Entity resolution**; **Deduplication**;

KEYWORDS

Entity Resolution, Evaluation

ACM Reference Format:
Matteo Paganelli, Francesco Del Buono, Francesco Guerra, and Nicola Ferro. 2022. Evaluating the Integration of Datasets. In *Proceedings of ACM SAC Conference (SAC'22).* ACM, New York, NY, USA, Article 4, 10 pages. https://doi.org/xx.xxx/xxx_x

1 INTRODUCTION

Data integration has always been considered as a key need for both research and industry. Traditionally the focus has been the integration of structured (typically relational) data sources where the information is divided into multiple tables. More recently, the attention paid to artificial intelligence and machine learning has led to the development of specific techniques for integrating datasets. From a technical perspective, these approaches typically implement a pipelined architecture, which consists of three major steps: schema alignment, entity resolution, and data fusion [6].

Despite the effort put by the research community (and partially reviewed in the Related Work Section), data and dataset integration is still far from being a solved problem and it is even less mature when applied in a real production context. Apart from the intrinsic complexity of the task, one of the barriers to fully empowering data integration is the human effort needed for evaluating and tuning the approaches. Indeed, you need to resort to controlled datasets, built on top of a manually created ground-truth, in order to compare your approach against this gold standard and score it accordingly. This is a long and economically demanding process, it presents serious challenges for scaling it up at the huge amounts of data needed in a real business scenario. Moreover, it is not able to keep pace with the quickly evolving data sources that you find in a real context and that call for a repeated over time and/or incremental integration process.

We address the problem from a completely different point of view, looking for an unsupervised way to measure how "good" is an integration process. By "good" we mean how much a data source is *representative* of another one, i.e. how much it preserves the informative content of another data source. Intuitively, the more a dataset can be represented by an integrated source, the less there is a loss of information when the integrated source is considered in place of the original one; we call this *input representativeness*. Viceversa, the more an integrated source can be represented by its datasets, the more it is consistent with them; we call this *output representativeness*.

Besides being an unsupervised measure, which reduces the required human effort and is suitable also for highly iterative and/or incremental real business scenarios, our approach considers the integration process as a whole and evaluates its quality after the data fusion step, which is what practitioners and domain experts are confronted with in a real context. On the contrary, most of the current literature [9, 12, 20] focuses on evaluating just the entity resolution step by using reference benchmarks and measures like precision or recall.

This paper extends and consolidates the preliminary results introduced in [14] by means of: (1) the introduction of an unsupervised metric for supporting the validation and verification of a dataset integration process; (2) the definition of scenarios demonstrating the ability of the measure to work in a real business context; (3) an extensive experimental evaluation[1] based on shared datasets

[1]See the project github at https://github.com/softlab-unimore/UEDI

SAC'22, April 25 –April 29, 2022, Brno, Czech Republic

ACM ISBN 978-1-4503-8713-2/22/04…$15.00
https://doi.org/xx.xxx/xxx_x

(a) D_1: the first data source.

entity id	authors	title	venue
freund1995a	yoav freund.	boosting a weak ...	in proceedings ...
haussler1994	haussler, d ...	rigorous learning ...	in proc. 7th ...
kearns1987	m. kearns, m. li ...	on the learnability ...	proceedings of ...
kearns1990	michael j. kearns.	the computational ...	
kearns1993b	m.j. kearns.	efficient noise-tolerant ...	in proc. 25th ...
schapire1996	r. e. schapire ...	learning sparse ...	j. of computer ...
kearns1994a	michael kearns, ...	on the learnability ...	proc. of the 26th ...
blum1994	avrim blum	weakly learning ...	in proceedings ...
freund1997a	yoav freund...	a decision-theoretic ...	journal of ...

(b) D_2: the second data source.

entity id	authors	title	venue
freund1995a	freund, y.	boosting a weak ...	in 'proceedings ...
haussler1994	haussler ...	rigorous learning ...	in proceedings ...
kearns1987	m. kearns ...	on the learn-ability ...	in proc. 19th stoc,
kearns1990	michael ...	the computational ...	
kearns1993b	m. kearns.	efficient noise-tolerant ...	in proceedings ...
haussler1994a	d. haussler, ...	bounds on the sample ...	machine learning,
kearns1988b	michael kearns.	thoughts on ...	(unpublished),
schapire1997	schapire, r.e ...	w.s.: boosting ...	proceedings of ...
rivest1989	r. l. rivest ...	inference of ...	in acm symposium ...

Table 1: Source datasets used in the motivating example.

(a) I_P: the *P*erfect integrated dataset.

id	entity id	...
1	*freund1995a*	...
2	*haussler1994*	...
3	*kearns1987*	...
4	*kearns1990*	...
5	*kearns1993b*	...
6	*schapire1996*	...
7	*schapire1997*	...
8	*blum1994*	...
9	*freund1997a*	...
10	*haussler1994a*	...
11	*kearns1988b*	...
12	*kearns1994a*	...
13	*rivest1989*	...

(b) I_C: low quality integrated dataset (concatenation).

id	entity id	...
1	*freund1995a*	...
2	*haussler1994*	...
3	*kearns1987*	...
4	*kearns1990*	...
5	*kearns1993b*	...
6	*schapire1996*	...
7	*schapire1997*	...
8	*blum1994*	...
9	*freund1997a*	...
10	*haussler1994a*	...
11	*kearns1988b*	...
12	*kearns1994a*	...
13	*rivest1989*	...
14	*kearns1987*	...
15	*kearns1990*	...
16	*kearns1993b*	...
17	*freund1995a*	...
18	*haussler1994*	...

(c) I_M: low quality integrated dataset (merging).

id	entity id	...
1	*freund1995a*	...
2	*haussler1994*	...
3	*kearns1987*	...
4	*kearns1990*	...
5	*kearns1993b*	...
6	*schapire1996, schapire1997*	...
7	*blum1994, rivest1989*	...
8	*haussler1994a, freund1997a*	...
9	*kearns1988b, kearns1994a*	...

Table 2: Three possible integrated datasets.

demonstrating the effectiveness and efficiency of the approach for high-dimensional and iterative / incremental integration processes.

The paper is organized as follows: Section 2 presents our approach; Section 3 introduces some relevant scenarios and reports experiments about them; Section 4 discusses related works; finally, Section 5 draws conclusions and outlooks for future work.

2 THE APPROACH

2.1 Motivating Example

Data integration in real scenarios is usually performed via try and error approaches, requiring several iterations, where domain experts evaluate the correctness of the integrated datasets produced at each step. The integration strategy is improved and tuned at each step until the experts are satisfied with the result obtained.

Clearly, this is a fully manual and very demanding task in terms of time, effort, and resources required. We provide here an example of how this process works in practice to motivate the need for automatic and unsupervised tools for supporting it.

We use the popular "Cora Citation Matching" data[2] to create two datasets of pubblications – D_1 and D_2 shown in Table 1 – where each publication is described by a unique identifier, authors, title, and venue. Table 2 shows some possible results from their integration. In particular, Table 2a shows I_P, the perfect integration according to the Cora ground truth. On the other hand, Table 2b shows I_C, a low-quality integration, obtained by just concatenating entities for the two sources. As a result, some merges are missing from it, i.e. some items from D_1 and D_2 are not recognized as referring to the same entity; for example, publication *haussler1994* is mapped to two separate entities – respectively, the second and the last entity– instead of the same one. Finally, Table 2c shows I_M, another low-quality integration, obtained merging each entity in D_1 with an entity in D_2. Five entities in I_M are the result of a correct integration process, since they are also in I_P. The remaining 4 entities (which were not merged in I_P) are here randomly integrated. For example, the last entity, that refers to the publication *kearns1988b*, contains also information from the pubblication *kearns1994a*, which is therefore not recognized as a distinct entity.

A domain expert would manually assess the quality of I_C, and I_M, by: 1) randomly sampling (or based on "sentinel" elements defined a priori) a number of entities to check; 2) verifying their correctness; and, 3) categorizing erroneous outputs to support the development of improvements in the integration approach. An expert, analyzing the integrated dataset I_M, may discover that the second entity has been correctly created while the sixth one contains an error since it merges two items referring to different real world entities, i.e. *schapire1996* in D_1 and *schapire1997* in D_2. On the other hand, I_C contains two separate entries for the entity *kearns1990* which actually refer to the same entity and therefore are a duplication.

The effort required for performing the error analysis is very huge due to the large size of the datasets typically involved. An accurate evaluation requires scanning the entire integrated dataset searching for duplicated and/or wrongly merged entities and a comparison with the input datasets to verify that every real-world entity has been included in the final result. Moreover, since the integrated dataset is obtained after several try and error iterations, the error analysis is repeated multiple times. Therefore, an automatic tool for analyzing the quality of an integration process would largely reduce the effort required for performing an integration task.

2.2 The model

We consider a dataset D as a collection of entities $D = \{e_1, \dots, e_N\}$. The integration of datasets is performed by means of an entity integration function, defined below.

Definition 2.1 (Entity Integration process). The Entity Integration process exploits an Entity Integration function (EI) to create an integrated dataset of entities $I = EI(\mathcal{D})$ from a collection of datasets $\mathcal{D} = \{D_1, \dots, D_k\}$. The *EI* function defines the logic for matching and merging the entities in the input dataset collection $\mathcal{D}$.

The integration approaches are usually evaluated with controlled datasets, pre-existing ground truths. Accuracy, and, more frequently, due to the unbalanced datasets, recall, precision, and F-measure are used to evaluate the quality of the integration result.

[2] https://people.cs.umass.edu/~mccallum/data.html

In business environments, the absence of a ground truth imposes to define a different procedure for the evaluation. The quality of the integration can be assessed through a *verification and validation process*. The verification process aims to check the formal correctness of the integrated dataset.

Definition 2.2 (Verified Entity Integrated Dataset). The Entity Integrated Dataset $I = EI(\mathcal{D})$, where EI is an entity integration function applied to a collection of datasets $\mathcal{D} = \{D_1, \ldots, D_k\}$, should be:

- **total**: each entity of every input dataset should be represented in I, i.e., $\forall e_i \in D_k, \exists\, e_j \in I, s.t.\ e_j$ and e_i refer to the same real-world entity;
- **minimal**: I should not contain duplicated entities, i.e., $\forall e_i, e_j \in I$, e_i and e_j refer to different real-world entities.

The *validation process* assesses the correspondence of the informative content of the integrated dataset with the input sources.

The unsupervided technique for evaluating EI processes proposed in this paper is based on a *representativeness function* that scores how much a dataset D_1 can be represented by a second dataset D_2 through the loss of information in using D_2 instead of D_1. We decided to implement the *representativeness function* by analyzing the word frequency distribution in the datasets.

Definition 2.3 (Word frequency distribution in datasets). Given a dataset D, let V be its vocabulary of terms. The word frequency distribution $freq_D(w) : V \to \mathbb{N}_0$ of the dataset D is a function which associates each term $w \in V$ with its frequency in D.

The simplest approach for the definition of a vocabulary of terms V for a dataset is to apply a tokenization algorithm to the concatenation of all tuples in D. Token splitting can be considered as a solved problem [21] and a large number of techniques are available in NLP code libraries.

Definition 2.4 (Dataset representativeness score). Given two datasets D_1 and D_2, the dataset representativeness $r_{D_1 \to D_2}$ quantifies the extent to which dataset D_1 represents D_2 by measuring how much the word frequency distribution $freq_{D_1}$ approximates $freq_{D_2}$.

In the next section, we propose a way to measure the approximation between two word frequency distributions in the context of a data integration process. The representativeness score should provide users with an assessment of how much datasets are represented by integrated sources by showing if there is any loss of information; vice-versa, it should quantify how much integrated sources are represented by the original datasets by showing if there is any redundancy or irrelevant content.

2.3 Scoring representativeness

When assessing the quality of the integration process, we need to consider the two sides of the coin, i.e. how well a source D is represented by the integration I and, vice-versa, how well the integration I is represented by a source D.

If the integration process is perfect, we expect that the content of D is completely "covered" by the content of I. This means that the vocabulary used in D should be included in the vocabulary used in I, and the word frequency distribution of words in D should be less than or equal to the one in I. The measure of the coverage of these word frequency distributions can provide a measure of the representativity of an integration source for a dataset. We call this measure *input representativeness* $r_{D \to I}$ and we define it in Equation (1).

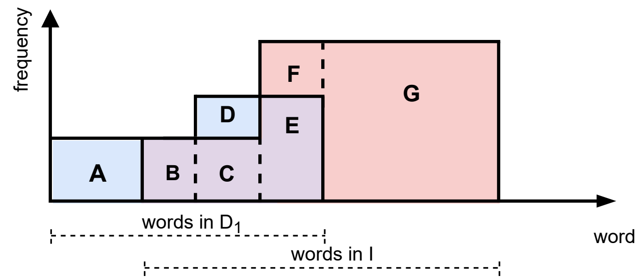

Figure 1: Example of word distributions.

Definition 2.5 (Input representativeness). Given two datasets D and I, where I is the integration of D according to some EI function, let V_D be the vocabulary of D and $freq_X(w)$ be the word frequency distribution of either D or I. We define the following representativeness score:

$$r_{D \to I} = 1 - \frac{1}{|V_D|} \sum_{w \in V_D} \frac{freq_D(w) - min(freq_D(w), freq_I(w))}{max(freq_D(w), freq_I(w))} \tag{1}$$

We expect that an integrated dataset contains more entities than an input dataset, due to the contribution of other datasets. Nevertheless, excluding stop words and other very generic words, we can suppose that the distribution of frequencies of words belonging to the intersection of the vocabularies of I and D is close. By measuring this closeness, we can evaluate how much the dataset can represent its integration for the shared words. We call this measure *output representativeness*, $r_{I \to D}$, and it is defined in equation (2).

Definition 2.6 (Output representativeness). Given two datasets D and I, where I is the integration of D according to some EI function, let V_D be the vocabulary of D and $freq_X(w)$ be the word frequency distribution of either D or I. We define the following representativeness score:

$$r_{I \to D} = 1 - \frac{1}{|V_D|} \sum_{w \in V_D} \frac{freq_I(w) - min(freq_D(w), freq_I(w))}{max(freq_D(w), freq_I(w))} \tag{2}$$

We observe as the output representativeness $r_{I \to D}$ is defined over the vocabulary V_D of the dataset D and not on the vocabulary of the integration I. Indeed, there is an intrinsic asymmetry in the integration process and we need to keep the focus on the dataset D, either considering how much it is represented by the integration I, i.e. $r_{D \to I}$, or how much it represents the integration I, i.e. $r_{I \to D}$, but without skewing the scores by including all the terms of V_I. Considering the whole vocabulary V_I, and not just its overlap with V_D, would just bring in all the other sources than D, whose vocabulary may differ a lot from V_D, and, as a result, these additional (and possibly unrelated) terms would mask how much D and I represent each other.

Example 1. Figure 1 shows a simplified word frequency distribution for a dataset D_1 and its integration I. The x-axis represents the words found in the data sources and the y-axis their respective distribution. Note that, for sake of simplicity, the heights of the frequency histograms are approximated to three possible values and the actual words are not reported on the x-axis. The areas A, B, C, D, E represent the word frequency distribution for D_1 and the areas B, C, E, F, G the one of I. A and G represent words belonging only to the input dataset and integrated dataset respectively. The words in B, C, D, E, F are common to both the sources and: (1) those of B have the same frequency distribution; (2) those of C and D have frequency distribution equal to C in the integration and frequency distribution equal to $C + D$ in the input dataset; (3) those of E and F have frequency distribution equal to E in the input dataset and frequency distribution equal to $E + F$ in the integration. To have a high value of representativeness, (1) the frequency of the common terms in the datasets should be similar (i.e. the regions D and F have to be as small as possible), and (2) a small number of terms should be used in a dataset only (i.e. the area of region A is limited). This is the behavior modeled by equations 1 and 2, which correspond to $r_{D\to I} \propto 1 - (A + \frac{D}{C+D})$, and $r_{I\to D} \propto 1 - (\frac{F}{E+F})$ when applied to the scenario represented in Figure 1.

2.4 Representativeness supporting the verification

The representativeness score can be used to *verify* an integration process, where the *input representativeness* score measures the *totality* of the integrated dataset; the *output representativeness* score the *minimality* of the integrated dataset.

Let I be obtained by the integration of D_1 and D_2. The input representativeness of I with respect to the input datasets D_1 and D_2 is obtained by averaging their input representativeness scores (i.e. $r_{D_1\to I}$ and $r_{D_2\to I}$). This aggregated score provides a measure of the totality of the integration process, since the more I represents the sources D_1 and D_2, the more the entities of D_1 and D_2 are also in I. On the other side, the output representativeness of D_1 and D_2 with respect to I, obtained by averaging $r_{I\to D_1}$ and $r_{I\to D_2}$, is a measure of the minimality of the integration process. Indeed, if D_1 and D_2 have high output representativeness, it follows that I does not contain duplicated entities.

2.5 Representativeness supporting the validation

An integration process can be *validated* by plotting the representativeness scores in a two-dimensional Cartesian plane. The x-axis reports the *input representativeness* $r_{D\to I}$, i.e. the *totality*, and shows the values obtained by the datasets with respect to the integration; the y-axis reports the *output representativeness* $r_{I\to D}$, i.e. the *minimality*, and shows the behavior of the integration with respect to the input sources. Values closest to the point $(1, 1)$ represent the best performance. We call the distance from (1, 1) *representativeness distance* and we claim that this is a measure of the validation of an integration approach. Indeed, the more we depart from (1, 1), the more the correspondences between entities in the input and integrated datasets decreases. Note that only in ideal scenarios, where the entities are represented in the input datasets with the same property values, the combined representativeness score of a verified and validated integrated dataset is (1, 1). Often, data representing the same entities are not the same, due to updates, mismatches and mistakes. This affects the word frequency distributions of the corresponding datasets which will have small differences and make representativeness values departing from (1,1).

Scenario	input repr.	output repr.
$D_1 \to I_C$	1	0.626
$D_2 \to I_C$	1	0.538
$D_1 \to I_M$	0.712	0.953
$D_2 \to I_M$	0.633	0.918
$D_1 \to I_P$	1	0.889
$D_2 \to I_P$	0.917	0.849

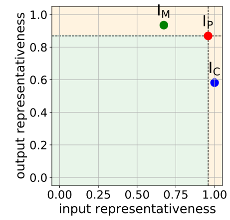

Figure 2: Input and output representativeness for the sources of the motivating example.

Example 2. Figure 2 shows the values of the representativeness scores obtained for the I_P, I_C, and I_M integrated datasets, described in Section 2.1. As expected, I_P is the best integrated dataset, being the closest to the point (1,1).

We observe that I_C is the integration that better represents the input datasets since it has the highest values for the *input representativeness.* It is the concatenation of the input datasets, so the resulting input representativeness value is 1, since the input word frequency distribution is completely included in the integrated dataset. Nevertheless, I_C obtains the worst value of *output representativeness*, thus meaning that it contains duplicated entries. I_M shows the highest results for the *output representativeness.* I_M has been built minimizing duplicated items (all entries in the input datasets have been merged). The worst values obtained for the *input representativeness* score means that the integrated dataset does not completely represent the input datasets. This is due to the wrong entity-merges that we have introduced. Note that input and output representativeness have to be jointly evaluated and the values assumed by the ground truth (I_P in the example) do not constitute an upper bound for the values that input and output representativeness can assume. In Figure 2, I_M and I_C are both located in the yellow area, which includes the elements with representativeness value greater than the one of the ground truth for at least one dimension. Nevertheless, even if I_M has a higher value of output representativeness, the quality of I_M (as the distance from (1,1) shows) is worst than the one of I_P due to the lower input representativeness. The same happens for the quality of I_C, which is worst than the one of I_P due to the lower output representativeness.

3 EXPERIMENTAL EVALUATION

We conduct a quantitative (in Sections 3.2 to 3.4) and qualitative (in Section 3.5) evaluation of the effectiveness of our proposed measures. Finally, in Section 3.6, we assess their efficiency.

3.1 Experimental setup

We use 12 publicly available use cases (see Table 3) from the benchmark of the Magellan tool[3], that is the main reference to evaluate entity matching approaches. The use cases consist each one of two datasets of entities and the ground truth contains pairs of entities, one for each dataset, labelled as matching and non matching items. According to the literature [7], we consider entities as referring to the same real world entity when the matching elements form

[3]https://github.com/anhaidgroup/deepmatcher/blob/master/Datasets.md

Use Case	Name	Input Datasets	Integrated Dataset	Shared Entities (%)	Unique Entities (%)
U1	Abt-Buy	$\|D_1\| = 949 - \|D_2\| = 920$	$\|I\| = 1174$	58.5	41.5
U2	Amazon-Google	$\|D_1\| = 1171 - \|D_2\| = 1843$	$\|I\| = 2232$	32.7	67.3
U3	Beer	$\|D_1\| = 237 - \|D_2\| = 233$	$\|I\| = 412$	14.1	85.9
U4	Fodors-Zagats	$\|D_1\| = 89 - \|D_2\| = 238$	$\|I\| = 422$	24.9	75.1
U5	iTunes-Amazon	$\|D_1\| = 272 - \|D_2\| = 278$	$\|I\| = 450$	20.9	79.1
U6	iTunes-Amazon	$\|D_1\| = 251 - \|D_2\| = 255$	$\|I\| = 410$	22.2	77.8
U7	DBLP-ACM	$\|D_1\| = 2419 - \|D_2\| = 2238$	$\|I\| = 2511$	85.5	14.5
U8	DBLP-ACM	$\|D_1\| = 2406 - \|D_2\| = 2220$	$\|I\| = 2507$	84.5	15.5
U9	DBLP-GoogleScholar	$\|D_1\| = 2491 - \|D_2\| = 9877$	$\|I\| = 7959$	29.0	71.0
U10	DBLP-GoogleScholar	$\|D_1\| = 2488 - \|D_2\| = 9286$	$\|I\| = 7865$	29.0	71.0
U11	Walmart-Amazon	$\|D_1\| = 1578 - \|D_2\| = 4297$	$\|I\| = 5080$	14.0	86.0
U12	Walmart-Amazon	$\|D_1\| = 1524 - \|D_2\| = 4014$	$\|I\| = 4784$	14.0	86.0

Table 3: The use cases considered.

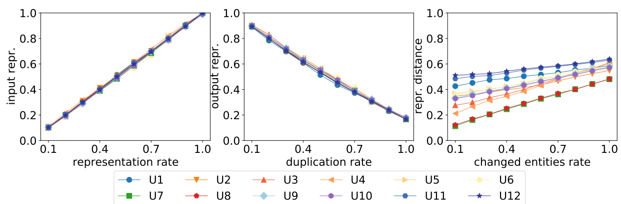

Figure 3: Verification and validation: measures.

a clique. In this case, we adopt a simple merging strategy by randomly selecting one of the entity as the one resulting from the merging process. The third and fourth columns in Table 3 show the cardinalities of the input and integrated datasets. The Table also shows for each use case the ratio of shared entities (i.e., entities in the integrated dataset which are generated by merging more input entities) and unique entities (i.e., entities which come from one of the input sources only). The distribution of these kinds of entities in the ground truth is typically unbalanced: only in U1 shared and unique entities have a similar distribution.

We run all experiments on commodity hardware: a server with 4 virtual cores, 16GB of RAM, 256GB of local (SSD) storage and that runs Ubuntu version 20.04.

3.2 Verification and validation of an integration processes

We evaluate the extent to which representativeness supports the verification and validation of an integration process: input representativeness for the totality, output representativeness for the minimality, and representativeness distance for the validation. The idea of the experiment is to modify the datasets in a controlled way and to check if the representativeness measures vary as expected to reflect the changes.

The first plot on the left of Figure 3 shows how the input representativeness scores measures the totality of the integration process. For each use case, a number of integrated datasets have been created from the ground truth, by selecting an increasing percentage of ground truth entities, as specified in the x-axis. The input representativeness score computed with these reduced datasets is shown on the y-axis. We expect that low input representativeness scores correspond to integrated datasets composed of a reduced numbers of entities. This is due to the existence of entities in the input datasets that do not have any correspondence in the integration. The first plot in Figure 3 shows that the score increases with the number of entities included in the integrated dataset. In a similar way, the second plot shows on the x-axis the percentage of duplicated entities that we have introduced in a "perfect" integrated dataset and, on the y-axis, the corresponding output representativeness score. As expected, the higher the number of duplicates, the lower the value of the output representativeness. Finally, the third plot on the right of Figure 3 evaluates how well the representativeness distance measures the validation of an integration approach. We alter the datasets by removing and by duplicating the same percentage of entities; therefore, for example, a value of 10% on the x-axis means that 5% of the entities are duplicated and 5% are removed; the y-axis shows the corresponding value of the representativeness distance. As expected, the distance grows with the increase of duplicated and missing entities, providing an overall validation of the process. Note that the slope of the curves is less sharp than the previous ones. This is due to the joint contribution of the input and output representativeness in the definition of this measure. Indeed, an entity duplication generates both a reduction of the minimality and an increase of the totality.

Take-away: the input and output representativeness are effective implementations of the totality and minimality properties respectively, while the representativeness distance is a valuable validation measure for an integration process.

3.3 Quality of the representativeness scores

3.3.1 Robustness to randomness in the data. We assess to what extent randomness affects our proposed representativeness scores. To this end, for each representativeness score, we repeat 100 times each of the three experiments reported in the previous Section 3.2 by randomly and uniformly sampling with replacement the data used in each configuration of the experiment. In this way, we can compute mean and standard deviations for each score (i.e., the input, output, and distance representativeness) and verify how often a given score falls in the expected range as defined in Figure 3. Indeed, the more a score falls in the expected range using random and equivalent samples of the same data, the more robust is its predictions, and the less we would change our conclusions due to the observed sample.

Figure 4 shows the results of this experiment for each representativeness score and use case. We considered three ranges: one standard deviation in blue; two standard deviations in orange; and, three standard deviations in green. Each bar in the histograms indicates which ratio of the 100 scores falls in the blue, orange, or green interval. For example, in Figure 4a for use case U1 and a deterioration of 50% of the samples, i.e. 50% of the entities have been removed in this case, we can observe that roughly 70% of the input representation scores fall in the one standard deviation range (blue bar); 20% in the two standard deviations range (orange bar on top of the blue one); 10% (or less) in the three standard deviations range (tiny green bar on top of the orange one).

In the case of the input representativeness in Figure 4a we can observe as the scores fall in the one standard deviation range in 50% to 75% of the cases, indicating a quite stable measure; almost all the other cases fall in the two standard deviations range, and just few of them in the three standard deviations range. We can observe a similar behaviour also for the output representativeness in Figure 4b and for the representativeness distance in Figure 4c.

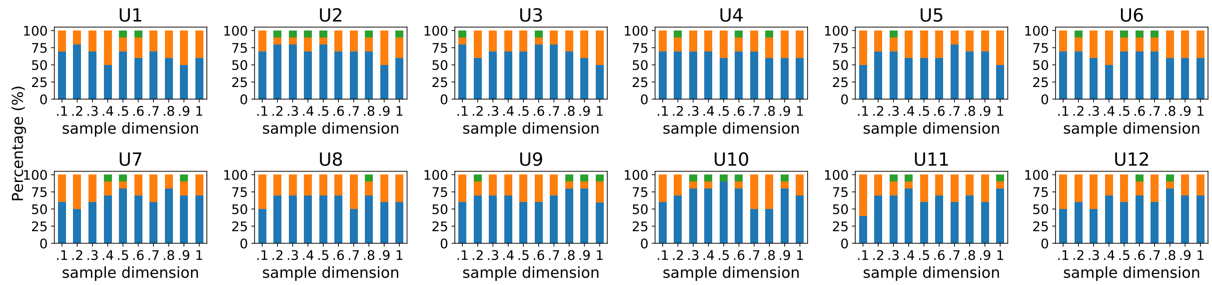

(a) Input representativeness score.

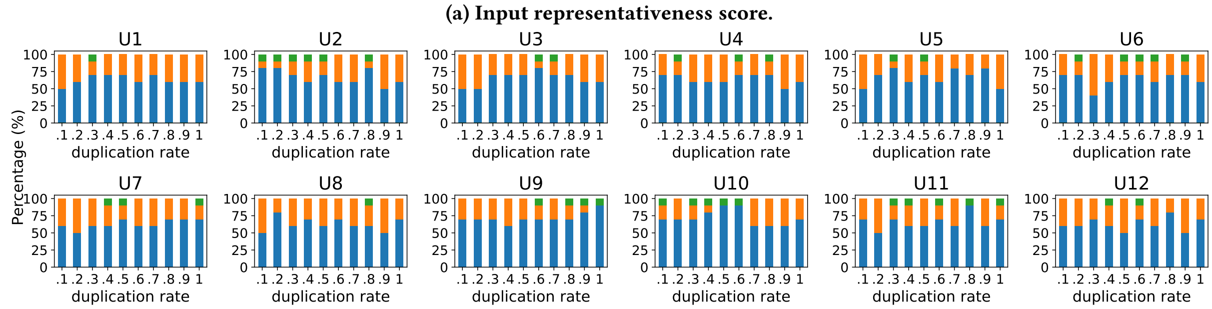

(b) Output representativeness.

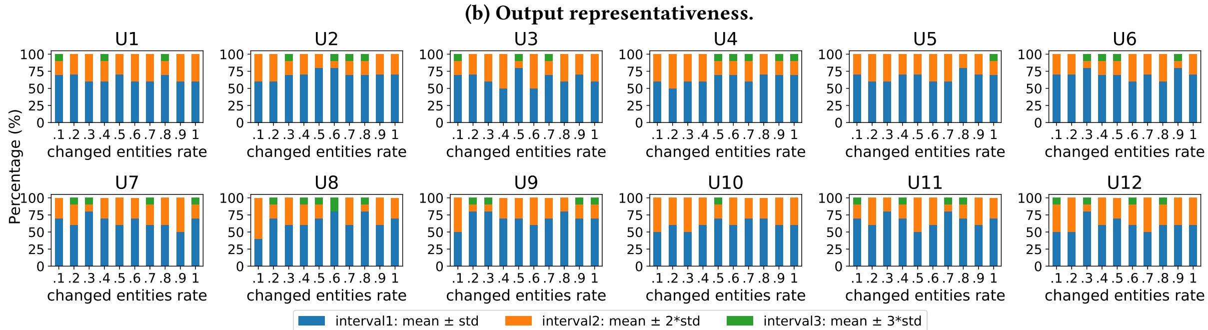

(c) Representativeness distance.

Figure 4: Ratio of representativeness scores.

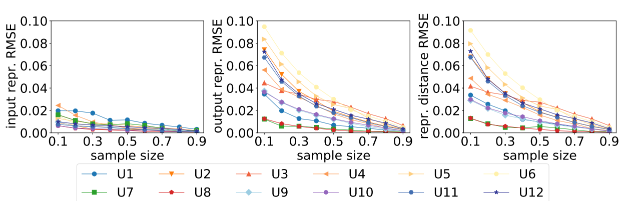

Figure 5: RMSE between the mean representatives scores and the ideal score based on the ground truth.

3.3.2 Robustness of the representativeness scores varying the dataset size. In this experiment we evaluate if the representativeness measures vary as the size of the considered datasets varies. The more the measures are stable, the more the approach is robust to the randomness of the data in the datasets. This experiment provides a complementary assessment compared to previous experiments that focused on the variability of results. We selected samples of increasing size from the ground truth (equal to 10%, 20%, ..., 100%) and we repeated this sampling process 100 times for each target size. For each type of representativeness score, Figure 5 shows the

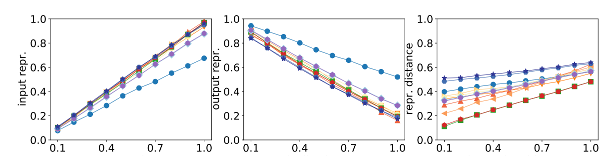

(a) Random Records.

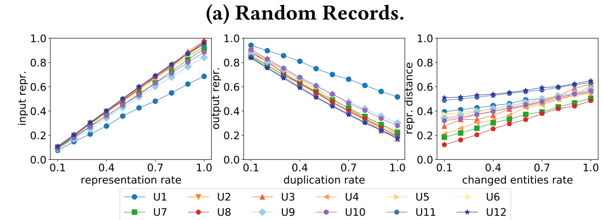

(b) Random Values.

Figure 6: Verification and validation of the representativeness measures against random merging strategies.

Root Mean Square Error (RMSE) between the score computed using the entire ground truth and the mean score computed over the samples related to a target sample dimension. The representativeness metrics do not show significant variations: only for use cases containing small datasets there are higher variations, although never greater than 0.1. This demonstrates their robustness even when significant changes in the size of the involved datasets are applied.

3.3.3 Robustness to the selected merging approach. We evaluate how much the behaviour of our representativeness measures depends on the actual merging strategy adopted to perform the integration of the matching entities. Ideally, we would like to observe some differences in the scores but not drastically different behaviours, otherwise, we could not reliably compare alternative integration processes. To this end, we repeat the experiment of Section 3.2 but we use two different alternatives for merging. Figure 6a shows the results for the first approach which randomly selects which entities to merge. Figure 6b shows the results for the second approach which randomly selects the values of the merged attributes. In both figures, we can observe a trend which is consistent with all the previous experiments.

Take-away: the proposed representativeness scores are quite robust to different types of deterioration and randomness in the data, have a good predictive accuracy, and they are not biased by the considered data fusion techniques.

3.4 Alternative techniques for measuring input representativeness

In Section 2.3 we proposed a specific way of computing representativeness based on word frequency distributions computed on the whole dataset. These distributions can be inaccurate for describing entity similarities, computed at the tuple level.

In this section, we consider the following alternatives for computing the input representativeness score: the jaccard-based similarity as a baseline, for its simplicity; the bleu-score [18] as a reliable unsupervised measure for evaluating the quality of machine-translated text; finally, embeddings largely used in NLP tasks to capture both syntactic and semantic similiarity.

Jaccard and Bleu score-based representativeness. Firstly, we tokenize the entries in the input and the integrated datasets and then we measure the similarity between input and integrated entities. For each input entity, we consider the maximum value computed. The mean of all maximum values is the representativeness measure for the considered input data source.

Embedding-based representativeness. We applied three different techniques (word2vec [11], fasttext [3], and glove [19]) for computing the embeddings of the tokenized entries of input and integrated entities. We measured the similarity between input and integrated entities through the cosine similarity. For each input entity, we consider the maximum value computed and we average the results for all the entities as before.

3.4.1 Alternatives for the representativeness distance. We conduct the experiment described in Section 3.2 comparing the representativeness distance obtained with the alternative measures. Figure 7 shows the results obtained. As expected, the representativeness distance increases as the deterioration of the datasets increases. Nevertheless, the measure defined in Equation 1 assumes the highest values in the majority of the scenarios, indicating that it better recognizes the errors in the integrated dataset.

Take-away: our measure outperforms alternative representativeness metrics based on syntactic and semantic similarities.

3.4.2 Alternatives for input and output representativeness. We compare alternative representativeness measures on the basis of how they react to possible errors in the integration process. We consider two error types: items in the input dataset which are merged even if they represent different entities and items referring to the same real world entities which are not merged. Note that this experiment may resemble the one of Section 3.2 but here we operate directly on the input datasets and on the different categories of entities.

Let us consider the "merge errors". We defined as *unique entities* those entities in the input datasets which are not to be merged with other entities in the integration process. When we erroneously merge unique entities with other entities, the dimension of the integrated dataset decreases as well as its totality, since there are input entities which are not represented in the integrated dataset, i.e. the wrongly merged ones. As a consequence, this kind of error will affect the input representativeness. To evaluate the impact of these errors, we created variations of the use case datasets, where different amounts of errors have been introduced in the ground truth, and we measured the difference of the input representativeness score measured with respect to the ground truth. The results of the experiments are shown in Figure 8, where for each use case, selected percentages of wrong merged entities have been introduced. The input representativeness (independently from the approach used for its computation) decreases when the error increases in all use cases and with all the approaches. Nevertheless, we observe that our measure introduced in Equation 1 better represents these mistakes, by showing the largest variations. Note that Figure 8 shows the results on the overall dataset, not only on the portion of the dataset composed of unique entities. The unbalanced distribution of unique entities (see Table 3) can introduce different amounts of wrong merges in the use cases. Table 4 shows the "real" impact of the perturbations introduced in the ground truth, by showing the percentage of missing unique entities for each experiment. We see that the variation in use cases U7 and U8 are less marked since the reduced number of wrong entities introduced. The plots describing U3, U11, and U12 are those with the largest variations, and this is consistent with the perturbed integrated entities.

	U1	U2	U3	U4	U5	U6	U7	U8	U9	U10	U11	U12
0.2	8.26	13.44	17.23	15.17	16	15.61	2.91	3.07	14.19	14.16	17.22	17.18
0.4	16.52	26.88	34.47	30.09	31.56	30.98	5.81	6.18	28.4	28.35	34.41	34.36
0.6	24.96	40.41	51.46	45.02	47.56	46.83	8.72	9.29	42.58	42.52	51.63	51.55
0.8	33.22	53.85	68.69	59.95	63.11	62.2	11.63	12.41	56.79	56.71	68.82	68.73
1	41.48	67.29	85.92	75.12	79.11	77.8	14.54	15.48	70.98	70.87	86.04	85.91

Table 4: Percentage of unique entities removed from the integrated dataset for each experiment.

We conduct a similar analysis for the second issue, i.e. duplicated entities. We called *shared entities* those entities obtained from merging multiple input entities. In this case, errors in the shared entities result on items in the integrated dataset which are not merged and this will affect the output representativeness. As before, we create a controlled deterioration of the ground truth, where we introduce

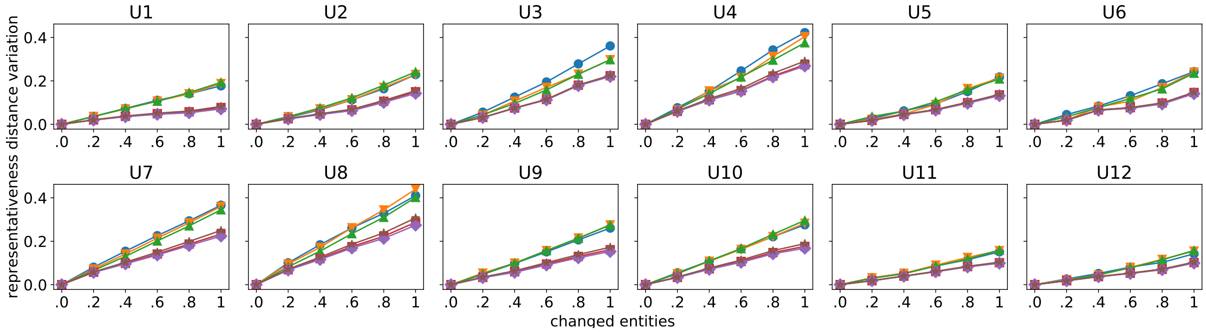

Figure 7: Comparison among the measures introduced for computing the representativeness distance

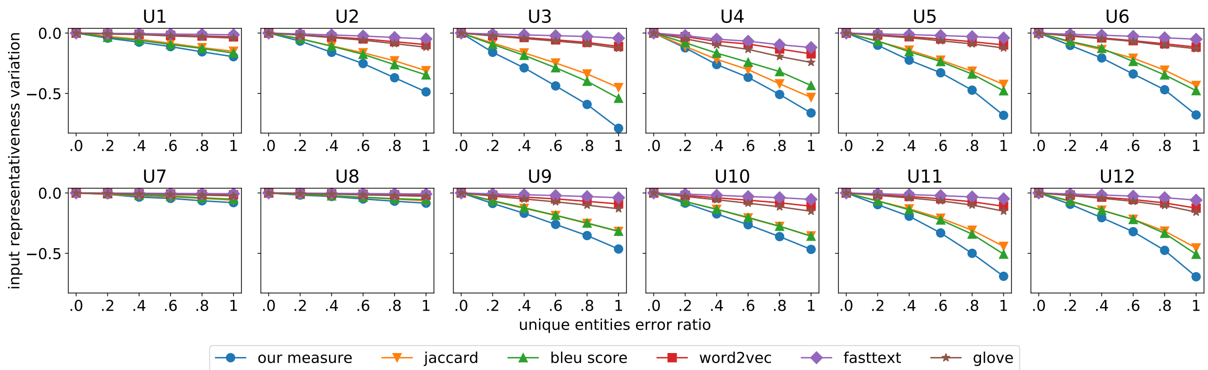

Figure 8: Impact of wrongly merged entities on the input representativeness.

errors on 20%, 40%, ...100% of the shared entities. Figure 9a shows the difference of the output representativeness score with respect to the ground truth: the more the decrease, the more errors in shared entities are detected. Note that, as before, Table 9b is needed to support the analysis. It shows the percentage of new entities introduced with the perturbation: U7 and U8 show the largest amount of entity introduced. This is consistent with the results in the figure that show the largest variation.

Take-away: by examining the variations of input and output representativeness, we understand the nature of the error affecting the integration task. A predominant variation of the input representativeness indicates the presence of errors in the recognition of the no-match class. Errors in the match class produce a more marked variation in the output representativeness.

3.5 Controlled data integration scenarios

Creating the datasets. For each use case in Table 3 we generate four datasets, D_1, D_2, D_3 and D_4. D_1 has a cardinality double than D_2 which has the same cardinality as D_3. D_2 contains a subset of the entities of D_1. D_3 contains entities that are not in D_1. D_4 concatenates D_2 and D_3. We evaluate the datasets in three controlled scenarios. The first column in Table 5 shows the cardinalities of the datasets and the associate vocabularies. The datasets are experimented in three controlled scenarios.

Scenario 1: Datasets describing the same entities. We consider D_1 and D_2, which describe same entities. Since D_1 is a superset of D_2, it can be considered as a possible integration, called $I_M = D_1$ in Figure 10a. I_C is the integration obtained by a concatenation of the

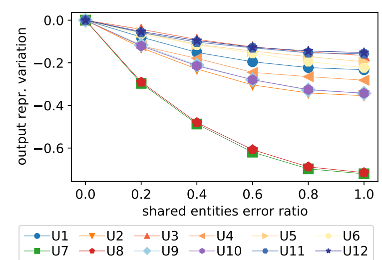

(a) Output representativeness variation in case of non-merged entities.

	U1	U2	U3	U4	U5	U6	U7	U8	U9	U10	U11	U12
0.2	21.55	12.46	4.13	8.77	7.33	7.07	30.63	30.55	16.15	15.04	5.45	5.33
0.4	37.56	22.49	8.01	14.69	12.67	11.46	54.4	53.45	30.66	28.2	9.78	9.62
0.6	49.66	29.48	10.92	20.85	16.89	16.1	71.64	70.68	42.73	38.88	12.93	12.86
0.8	56.64	33.29	12.62	22.99	19.78	20.24	82.04	80.93	51.84	46.68	15.02	15.07
1	59.2	35.04	14.08	24.88	22.22	23.41	85.46	84.52	55.4	49.7	15.65	15.76

(b) Percentage of duplicated entities introduced in the integrated dataset for each use case.

Figure 9: Impact of errors on shared entities on the input representativeness.

tuples in D_1 and D_2. Let us consider for example use case U10: we know the ground-truth and it is thus possible to compute the error rate, which is 0 for I_M, and 0.333 for I_C. Our measure shows that, from a dataset perspective, the concatenation I_C is the best integration scenario, since it does not generate any loss of information. This is clear in Figure 10a, where I_C assumes the maximum value of input representativeness on the x-axis. Nevertheless, concatenation introduces data duplication (D_1 is a superset of D_2) and this is the

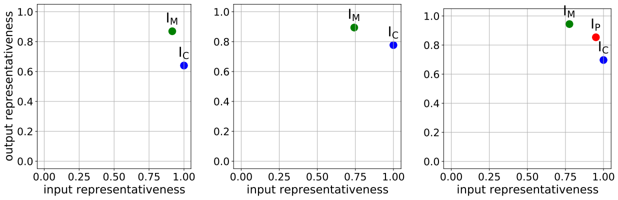

(a) Scenario 1 (b) Scenario 2 (c) Scenario 3

Figure 10: The scenarios applied to use case U10.

reason why in Figure I_C has an output representativeness value on the y-axis lower than I_M. The plot clearly shows that I_M is a better integration than I_C, as we can expect by analyzing the data sources.

Scenario 2: Datasets describing different entities. We consider D_1 and D_3, which describe different entities. As in the previous scenario, we consider D_1 also as integration and we call it I_M in Figure 10b. I_C is the integration obtained by the concatenation of D_1 and D_3, which does not contain duplicates in this case. In this scenario, I_C should be the best integration since all entities are included in this source. This is confirmed by the error rate, 0.5 for I_M and 0 for I_C. This is also clear by our measure applied to U10 (see Figure 10b), comparing the coordinates of I_C and I_M in the Figure. I_C has coordinates (1, 0.79). This means the maximum input representativeness value. I_M has coordinates (0.73,0.9). The output value is due to the low representativeness value for D_3 in I_M (0.46). Note that even if I_M does not contain the entities described in D_3 the representativeness is not zero since there is still a low number of words in D_3 which are contained in I_M anyway. The high level measured from the integration perspective is because I_M completely includes D_1 which has twice the cardinality of D_3.

Scenario 3: Datasets describing common entities. We consider D_1 and D_4 which contain a half common and a half different entities. I_P in Figure 10c is generated by concatenating D_1 and D_3. This is a perfect integration since it includes all entities described by the D_1 and D_4 datasets. I_M, as in the previous scenarios, is D_1 only which, in this case, does not describe half of the entities in D_4. Finally, I_C is obtained by the concatenation of D_1 and D_4. This integration suffers from redundancy, generated by the duplicated entities of D_1 contained in D_4 and included twice in I_C. The error rates of these integrations are 0.5 for I_M and I_C, and no error rate for I_P. Figure 10c shows our measures applied to U10 and correctly reflects the datasets included in the integration, by showing the input representativeness values on the x-axis of I_P and I_M close, but not equal to 1, thus meaning that there is some loss of information in the integration. In I_C, the input representativeness values are equal to 1, since the datasets are completely represented, but the integration suffers from redundancy as shown by the lowest output representativeness value on the y-axis.

Extended evaluation. Table 5 summarizes the results of the experiments performed on all datasets in the benchmark. The second column reports the scenarios, and the other columns outline the measures obtained by considering the I_M, I_C, and I_P integrations. The bold values are the best ones, i.e. the closest to the point (1,1). According to the previous discussion, we expect I_M to be the best integration in Scenario 1, I_C in Scenario 2, and I_P in Scenario 3. The measure performs correctly in almost all evaluations. Wrong best

Use case params	Sc.	I_M	I_C	I_P
U1 (|D1|=600, |D2|=300, |D3|=300, |D4|=600, |V1|=4776, |V2|=1431, |V3|=2258, |V4|=3092)	1	**(0.83, 0.83)**	(1.0, 0.71)	
	2	**(0.80, 0.89)**	(1.0, 0.73)	
	3	(0.7, 0.92)	(1.0, 0.73)	**(0.93, 0.77)**
U2 (|D1|=700, |D2|=350, |D3|=350, |D4|=700, |V1|=1664, |V2|=1139, |V3|=986, |V4|=1699)	1	**(0.91, 0.89)**	(1.0, 0.66)	
	2	(0.75, 0.91)	**(1.0, 0.76)**	
	3	(0.78, 0.95)	(1.0, 0.68)	**(0.92, 0.85)**
U3 (|D1|=50, |D2|=25, |D3|=25, |D4|=50, |V1|=208, |V2|=120, |V3|=136, |V4|=235)	1	**(0.9, 0.94)**	(1.0, 0.70)	
	2	(0.62, 0.97)	**(1.0, 0.89)**	
	3	(0.70, 0.97)	(1.0, 0.76)	**(0.92, 0.92)**
U4 (|D1|=100, |D2|=50, |D3|=50, |D4|=100, |V1|=375, |V2|=192, |V3|=192, |V4|=347)	1	**(0.98, 0.95)**	(1.0, 0.65)	
	2	(0.65, 0.96)	**(1.0, 0.88)**	
	3	(0.78, 0.98)	(1.0, 0.72)	**(0.98, 0.92)**
U5 (|D1|=90, |D2|=45, |D3|=45, |D4|=90, |V1|=697, |V2|=433, |V3|=462, |V4|=736)	1	**(0.92, 0.87)**	(1.0, 0.65)	
	2	(0.72, 0.93)	**(1.0, 0.79)**	
	3	(0.75, 0.94)	(1.0, 0.70)	**(0.95, 0.85)**
U6 (|D1|=90, |D2|=45, |D3|=45, |D4|=90, |V1|=503, |V2|=293, |V3|=335, |V4|=529)	1	**(0.95, 0.89)**	(1.0, 0.61)	
	2	(0.71, 0.93)	**(1.0, 0.79)**	
	3	(0.77, 0.95)	(1.0, 0.67)	**(0.98, 0.85)**
U7 (|D1|=2100, |D2|=1050, |D3|=365, |D4|=1415, |V1|=7359, |V2|=4854, |V3|=1790, |V4|=5460)	1	**(0.96, 0.87)**	(1.0, 0.59)	
	2	**(0.87, 0.79)**	(1.0, 0.7)	
	3	**(0.93, 0.91)**	(1.0, 0.61)	(0.97, 0.86)
U8 (|D1|=2100, |D2|=1050, |D3|=388, |D4|=1438, |V1|=7396, |V2|=4858, |V3|=1863, |V4|=5509)	1	**(0.98, 0.87)**	(1.0, 0.59)	
	2	**(0.87, 0.80)**	(1.0, 0.70)	
	3	**(0.93, 0.91)**	(1.0, 0.61)	(0.97, 0.86)
U9 (|D1|=2300, |D2|=1150, |D3|=1150, |D4|=2300, |V1|=5993, |V2|=4119, |V3|=3979, |V4|=6364)	1	**(0.92, 0.89)**	(1.0, 0.67)	
	2	(0.72, 0.91)	**(1.0, 0.8)**	
	3	(0.76, 0.95)	(1.0, 0.71)	**(0.93, 0.86)**
U10 (|D1|=2200, |D2|=1100, |D3|=1100, |D4|=2200, |V1|=5802, |V2|=3905, |V3|=3632, |V4|=5939)	1	**(0.92, 0.89)**	(1.0, 0.66)	
	2	(0.73, 0.90)	**(1.0, 0.79)**	
	3	(0.77, 0.95)	(1.0, 0.71)	**(0.96, 0.86)**
U11 (|D1|=700, |D2|=350, |D3|=350, |D4|=700, |V1|=2875, |V2|=2096, |V3|=1694, |V4|=3195)	1	**(0.86, 0.91)**	(1.0, 0.71)	
	2	(0.72, 0.93)	**(1.0, 0.80)**	
	3	(0.73, 0.96)	(1.0, 0.73)	**(0.88, 0.89)**
U12 (|D1|=600, |D2|=300, |D3|=300, |D4|=600, |V1|=2152, |V2|=1713, |V3|=1231, |V4|=2453)	1	**(0.88, 0.91)**	(1.0, 0.69)	
	2	(0.74, 0.96)	**(1.0, 0.79)**	
	3	(0.77, 0.95)	(1.0, 0.71)	**(0.88, 0.88)**

Table 5: The evaluation of the scenarios in other datasets.

integrations in U1, U7 and U8 have all a very close distance to the best one. The mistakes are due to the sparse vocabularies (and the low cardinalities in the second dataset).

Take-away: the representativeness scores offer a fine-grained explanation on why an integration strategy can be preferred to another one.

3.6 Efficiency

The time performace has been evaluated on integration processes involving datasets with increasing dimensionality (1K, 10K, 50K, 100K, 500K and 1M). These datasets have been obtained by applying sampling with replacement to the data contained in use case U10 (the largest one). The experiment was repeated 5 times and Figure 11 shows the average times. All embedding-based approaches show the same time performance, since they adopt the same algorithm for mapping the datasets into the vector space of embeddings and for computing the similarity. Moreover, they could not be applied to the largest datasets since they overcame the maximum time (48 hours) we fixed for the duration of the experiment.

Our approach shows the best performance in all configurations: it takes less than 2 minutes to compute the representativeness of the largest dataset. The vectorized implementation of the cosine similarity makes the embedding-based approaches fast, but for running on datasets larger than 100K entities it requires more memory than the one available in our system. The approach based on Jaccard's similarity has a poor performance since it cannot be vectorized for performing our computation. For this reason, the execution time grows quadratically with the size of the datasets.

Take-away: the developed approach is efficient in evaluating high dimensional data integration scenarios.

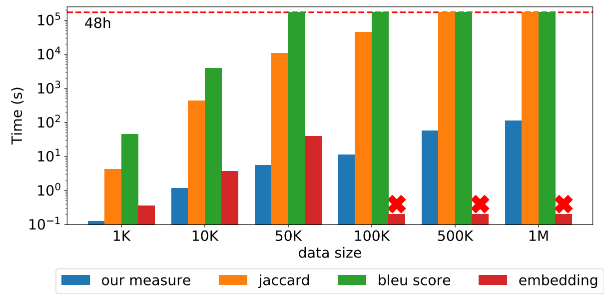

Figure 11: Computing representativeness: efficiency.

4 RELATED WORK

Data Integration and Entity Resolution. Data integration is one of the most challenging and long-lasting issues that the research community is confronted with for the last 30 years. The focus of the research community in the last years was mainly oriented to Entity Resolution (ER), the task concerning the development of techniques for detecting and merging entities. A number of "integration functions" to discover and match the different structures that represent the same real-world entity have been proposed [1, 9, 12, 15, 20]. Among these, rule-based and machine learning (ML) techniques are the most common ones. Regardless of the use of ML or not, ER approaches require either careful manual configuration by domain experts or a large amount of labeled data [13]. To cope with the first issue, methods have been proposed for the fine tuning of parameters such as [16], but all proposals require some human supervision. Regarding to the second problem, many semi-supervised approaches in the field of active learning [2] and crowd-sourcing [22] have been introduced. The fundamental idea behind these techniques is to limit the validation intervention required by domain experts to a minimum or to resort to crowd-workers. However, these methods suffer from a poor quality control mechanism: indeed, the former approach focuses on optimizing recall while ensuring a user-specified precision level [4, 7], while crowd-based solutions are affected by uncertain labels provided by inexperienced workers [5].

Evaluating Data Integration and Entity Resolution. The effectiveness of ER and data integration processes is typically measured against ground truths. The availability of labeled data is a problem in real scenarios, where experts have to manually assess the results obtained. This is also a problem for the evaluation of the approaches proposed by the research community since most of the techniques are evaluated against the same small number of sources (typically the benchmark made available by the Magellan tool[4]) with few hundreds of labeled data. This makes possible the development and promotion of approaches overfitting on those sources (which can have features really different from the ones in sources available in real scenarios). To the best of our knowledge, only recently [10] addressed this issue, by proposing techniques for providing samples on datasets guaranteeing a fair evaluation. Similarly to other techniques [8, 17], our approach is part of this human-machine cooperation framework, but it mainly focuses on supporting analysts in the unsupervised evaluation of the integration process.

[4]https://github.com/anhaidgroup/deepmatcher/blob/master/Datasets.md

5 CONCLUSION AND FUTURE WORK

We introduced the representativeness score, an unsupervised measure to evaluate the quality of an integration process by analyzing the word frequency distributions of the datasets involved. The experimental evaluation showed that the representativeness is able to provide a means for verifying and validating an integration process.

The approach is conceived for textual datasets. Future work will deal with numeric datasets. Our idea is to exploit functional dependencies (FD), by extending our model based on word frequencies to a model based on FD frequencies.

REFERENCES

[1] Andrea Baraldi, Francesco Del Buono, Matteo Paganelli, and Francesco Guerra. 2021. Using Landmarks for Explaining Entity Matching Models. In *EDBT*. OpenProceedings.org, 451–456.

[2] Kedar Bellare, Suresh Iyengar, Aditya G. Parameswaran, and Vibhor Rastogi. 2012. Active sampling for entity matching. In *KDD*. ACM.

[3] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching Word Vectors with Subword Information. *Trans. Assoc. Comput. Linguistics* (2017).

[4] Zhaoqiang Chen, Qun Chen, Fengfeng Fan, Yanyan Wang, Zhuo Wang, Youcef Nafa, Zhanhuai Li, Hailong Liu, and Wei Pan. 2018. Enabling Quality Control for Entity Resolution: A Human and Machine Cooperation Framework. In *ICDE*.

[5] Mohamad Dolatshah, Mathew Teoh, Jiannan Wang, and Jian Pei. 2018. Cleaning Crowdsourced Labels Using Oracles For Statistical Classification. *Proc. VLDB Endow.* 12 (2018).

[6] Xin Luna Dong and Divesh Srivastava. 2015. *Big Data Integration.* Morgan & Claypool Publishers. https://doi.org/10.2200/S00578ED1V01Y201404DTM040

[7] Donatella Firmani, Barna Saha, and Divesh Srivastava. 2016. Online Entity Resolution Using an Oracle. *Proc. VLDB Endow.* 9 (2016).

[8] Boyi Hou, Qun Chen, Zhaoqiang Chen, Youcef Nafa, and Zhanhuai Li. 2020. r-HUMO: A Risk-Aware Human-Machine Cooperation Framework for Entity Resolution with Quality Guarantees. *IEEE Trans. Knowl. Data Eng.* 32 (2020).

[9] Pradap Konda, Sanjay Subramanian Seshadri, Elan Segarra, Brent Hueth, and AnHai Doan. 2019. Executing Entity Matching End to End: A Case Study. In *Proceedings of EDBT 2019, Lisbon, Portugal, March 26-29, 2019*. OpenProceedings.org.

[10] Neil G. Marchant and Benjamin I. P. Rubinstein. 2017. In Search of an Entity Resolution OASIS: Optimal Asymptotic Sequential Importance Sampling. *Proc. VLDB Endow.* 10 (2017).

[11] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *ICLR 2013*.

[12] Sidharth Mudgal, Han Li, Theodoros Rekatsinas, AnHai Doan, Youngchoon Park, Ganesh Krishnan, Rohit Deep, Esteban Arcaute, and Vijay Raghavendra. 2018. Deep Learning for Entity Matching: A Design Space Exploration. In *Proceedings of the 2018 SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*. ACM.

[13] Stefano Ortona, Venkata Vamsikrishna Meduri, and Paolo Papotti. 2018. Robust Discovery of Positive and Negative Rules in Knowledge Bases. In *ICDE*.

[14] Matteo Paganelli, Francesco Del Buono, Francesco Guerra, and Nicola Ferro. 2020. Unsupervised Evaluation of Data Integration Processes. In *iiWAS*. ACM, 77–81.

[15] Matteo Paganelli, Francesco Del Buono, Marco Pevarello, Francesco Guerra, and Maurizio Vincini. 2021. Automated Machine Learning for Entity Matching Tasks. In *EDBT*. OpenProceedings.org, 325–330.

[16] Matteo Paganelli, Paolo Sottovia, Francesco Guerra, and Yannis Velegrakis. 2019. TuneR: Fine Tuning of Rule-based Entity Matchers. In *CIKM*.

[17] Fatemah Panahi, Wentao Wu, AnHai Doan, and Jeffrey F. Naughton. 2017. Towards Interactive Debugging of Rule-based Entity Matching. In *EDBT*.

[18] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *ACL*. ACL.

[19] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the EMNLP Conference, October 25-29, 2014, Doha, Qatar*. ACL.

[20] Rohit Singh, Venkata Vamsikrishna Meduri, Ahmed K. Elmagarmid, Samuel Madden, Paolo Papotti, Jorge-Arnulfo Quiané-Ruiz, Armando Solar-Lezama, and Nan Tang. 2017. Synthesizing Entity Matching Rules by Examples. *Proc. VLDB Endow.* 11 (2017).

[21] Noah A. Smith. 2020. Contextual Word Representations: Putting Words into Computers. *Commun. ACM* 63 (2020).

[22] Norases Vesdapunt, Kedar Bellare, and Nilesh N. Dalvi. 2014. Crowdsourcing Algorithms for Entity Resolution. *Proc. VLDB Endow.* 7 (2014).