

Canonicalization of Database Records using Adaptive Similarity Measures

Aron Culotta, Michael Wick, Robert Hall, Matthew Marzilli, Andrew McCallum
Department of Computer Science
University of Massachusetts
Amherst, MA 01003

ABSTRACT

It is becoming increasingly common to construct databases from information automatically culled from many heterogeneous sources. For example, a research publication database can be constructed by automatically extracting titles, authors, and conference information from online papers. A common difficulty in consolidating data from multiple sources is that records are referenced in a variety of ways (e.g. abbreviations, aliases, and misspellings). Therefore, it can be difficult to construct a single, standard representation to present to the user. We refer to the task of constructing this representation as *canonicalization*. Despite its importance, there is little existing work on canonicalization.

In this paper, we explore the use of edit distance measures to construct a canonical representation that is “central” in the sense that it is most similar to each of the disparate records. This approach reduces the impact of noisy records on the canonical representation. Furthermore, because the user may prefer different *styles* of canonicalization, we show how different edit distance costs can result in different forms of canonicalization. For example, reducing the cost of character deletions can result in representations that favor abbreviated forms over expanded forms (e.g. *KDD* versus *Conference on Knowledge Discovery and Data Mining*). We describe how to learn these costs from a small amount of manually annotated data using stochastic hill-climbing. Additionally, we investigate feature-based methods to learn ranking preferences over canonicalizations. These approaches can incorporate arbitrary textual evidence to select a canonical record. We evaluate our approach on a real-world publications database and show that our learning method results in a canonicalization solution that is robust to errors and easily customizable to user preferences.

Categories and Subject Descriptors

H.2 [Information Systems]: Database Management; H.2.8 [Information Systems]: Database Applications—*data mining*

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.

KDD’07, August 12–15, 2007, San Jose, California, USA.
Copyright 2007 ACM 978-1-59593-609-7/07/0008 ...\$5.00.

General Terms

Algorithms

Keywords

Data mining, information extraction, data cleaning

1. INTRODUCTION

Consider a research publication database such as Citeseer¹ or Rexa² that contains records gathered from a variety of sources using automated extraction techniques. Because the data comes from multiple sources, it is inevitable that an attribute such as a conference name will be referenced in multiple ways. Since the data is also the result of extraction, it may also contain errors. In the presence of this noise and variability, the system must generate a single, canonical record to display to the user.

Record canonicalization is the problem of constructing one standard record representation from a set of duplicate records. In many databases, canonicalization is enforced with a set of rules that place limitations or guidelines for data entry. However, obeying these constraints is often tedious and error-prone. Additionally, such rules are not applicable when the database contains records extracted automatically from unstructured sources.

Simple solutions to the canonicalization problem are often insufficient. For example, one can simply return the most common string for each field value. However, incomplete records are often more common than complete records. For instance, this approach may canonicalize a record as “J. Smith” when in fact the full name (John Smith) is much more desirable.

In addition to being robust to noise, the system must also be able to adapt to user preferences. For example, some users may prefer abbreviated forms (e.g., *KDD*) instead of expanded forms (e.g., *Conference on Knowledge Discovery and Data Mining*). The system must be able to detect and react to such preferences.

In this paper, we first formalize the canonicalization problem and then propose three solutions. The first uses string edit distance to determine which record is *most central* in a set of records. This approach can mitigate the noise contained in outlying records. To enable the system to adapt to user preferences, the second solution optimizes the edit distance parameters using human-labeled data. Finally, we

¹www.citeseer.ist.psu.edu

²www.rexa.info

describe a feature-based solution that can flexibly incorporate textual evidence to select a canonical record. We again estimate the parameters of this method from labeled data. We show that this problem can be more naturally formulated as a *ranking* task, rather than a *classification* task, and modify the learning methods accordingly.

We perform several empirical comparisons of these approaches on publication records culled from the Web. The results indicate that the feature-based approach significantly outperforms competing approaches as well as a number of simpler baselines. Furthermore, we show that the parameters of the feature-based approach can be estimated from just a few training examples, and is quite robust to noise that is common in automatically generated databases.

2. RELATED WORK

While there has been little work explicitly addressing canonicalization, the idea has been present in many application and research areas. In this section we review several of these applications as well as related work in the area of learning string edit distance costs.

Tejada et al. [13] devise a system to automatically extract and consolidate information from multiple sources into a unified database. When a user queries this database, multiple representations of an attribute are inevitable due to naming inconsistencies across the various sources from which they were drawn. Although object deduplication is the primary goal of the that research, canonicalization arises when the system presents results to the user. Tejada et al. propose ranking the strings for each attribute based on the user’s confidence in the data source from which the string was extracted.

One difficulty in this approach is that if the data is extracted from a large number of sources, a non-trivial burden is placed on the users, who may not have the expertise to express preferences about each data source. Additionally, the database must store source-specific meta information for each string. Our canonicalization methods are adaptable to any database, regardless of whether the source of the information is available. Additionally, we enable the user to express preferences independent of the data source.

Other work has focused on learning the parameters of string edit distance with encouraging results. Zhu and Unger [15] apply string edit distances to the task of merging database records. They observe that parameters cannot be optimized individually due to the complex interaction between various edit cost weights on the outcome. Additionally they note that greedy methods are too likely to converge prematurely in local optima and that random restarts are unnecessarily expensive. Instead they propose a genetic algorithm to learn the weights of each cost and find that it stabilizes after 100 generations. In lieu of genetic approaches we propose learning the edit costs using either stochastic search, or an exhaustive search over a relatively small discrete space of possible parameter settings.

McCallum et al. [9] also use a discriminatively learned edit distance to perform record deduplication. They extend Zhu and Unger’s work by using a conditional random field to learn the costs of a variety of flexible edit distance operations. However, they do not explore canonicalization.

Ristad and Yianilos [12] learn a probability distribution over atomic string edit operations (insertion, deletion, substitution) and define a stochastic transducer that defines the

probability of a string as either the Viterbi sequence of edit operations or the sum of all possible sequences of edit operations required to produce that string. The parameters of this generative model are learned using the expectation maximization (EM) algorithm.

Bilenko and Mooney [1] present a method to learn edit distance based similarity measures of each attribute between records in order to perform deduplication. They extend the work of Ristad and Yianilos [12] by accommodating affine gaps. In similar fashion, the weights are learned iteratively with EM.

Our learning methods differ from those outlined in Ristad and Yianilos [12] and Bilenko and Mooney [1] in that we are not concerned with learning a generative model. We propose two methods to learn edit distance parameter settings using stochastic or exhaustive search. Additionally, we propose two feature-based methods that combine the outputs of multiple parameter settings (i.e., multiple edit distance models) and other sources of textual evidence into a discriminative model of canonicalization.

Canonicalization has also been implicitly considered in deduplication research. For example, Milch et al. [11] and McCallum and Wellner [10] propose deduplication models containing variables for canonical attributes of a record. The variables are used to help deduplication, although the accuracy of the resulting canonical records is not optimized or evaluated.

Recently, frameworks have been proposed to handle uncertainty in databases, particularly those containing automatically extracted records. One approach when consolidating extractions from various sources is to perform some type of weighted voting to determine which facts should be inserted in the database [8]. Another approach is to store the uncertainty of these extractions directly in the database. This can be accomplished by storing the top n most confident extraction results (along with corresponding probabilities or confidence measures) for each desired record. Gupta and Sarawagi [5] leverage confidence value outputs from the extraction models to improve query results on databases containing uncertainty. Fundamentally the problem is canonicalization because the system is faced with a choice when presenting multiple query results with various confidence values to the user. It is analogous to our canonicalization task except that we do not have the luxury of confidence values. While the inclusion of such values is clearly beneficial, we propose methods that achieve canonicalization in absence of such information (and often this information is strictly unavailable).

3. PROBLEM DEFINITION

Let a record R be a set of fields, $R = \{F_1 \dots F_p\}$. Let field F_i be an attribute-value pair $\langle a, v \rangle$. Table 1 shows three example records.

Databases constructed from multiple sources often accumulate multiple versions of the same record. *Record deduplication* is the task of detecting these different versions. Table 1 shows three records that have been predicted to be duplicates. In fact, record (c) refers to a book chapter version, whereas (a) and (b) refer to conference proceedings.

Record deduplication is a difficult problem that has been well-studied. However, in this paper we are interested in a subsequent step: how to present the user one *canonical* representation of a record.

author	Brian Milch et al.	author	B. Milch et al.	author	Brian Milch et al.
title	BLOG: Probabilistic Models with Unknown Objects	title	BLOG: Probabilistic Models with Unknown Objects	title	BLOG: Probabilistic Models with Unknown Objects
venue	IJCAI	venue	Intl. Conf. on AI	venue	
editor		editor		editor	L. Getoor and B. Taskar
pages		pages	1352-1359	pages	

(a)

(b)

(c)

Table 1: Three publication records predicted to be duplicates. Note that a de-duplication error has erroneously merged record (c) (a book chapter) with the other two conference papers. De-duplication errors, as well as misspellings, abbreviations, and aliases, can make canonicalization difficult.

We define the *canonicalization* problem as follows: Given a set of duplicate records $\mathbf{R} = \{R_1 \dots R_k\}$, create a *canonical* record R^* that summarizes the information in \mathbf{R} . We refer to the canonicalization operation as $C(\mathbf{R})$.

Note that it is not always clear what the optimal canonicalized record should be. Indeed, different users may prefer different forms of canonicalization. For example, some users may prefer the abbreviated conference string *IJCAI*, while others may prefer the expanded string *International Joint Conference on Artificial Intelligence*. However, there are a few desiderata of a good canonicalization:

- **Error-free:** The canonical record should not contain errors, such as misspellings or incorrect field values. This is especially a concern when the data has been automatically extracted from noisy sources (e.g. when text must be extracted from a PDF file and field assignments are automated). In these cases, there may exist *outlying* records that contain erroneous data. The canonicalizer should attempt to minimize the effect of these errors.
- **Complete:** The canonical record should contain all the accurate information contained in the duplicate records. Thus, even if not all records contain a *date* field, the field should be included in the canonical record.
- **Representative:** The canonical record should reflect the commonality among the duplicate records. Thus, the canonical record should in some sense be *similar* to all of the duplicate records.

4. THREE CLASSES OF CANONICALIZATION SOLUTIONS

We now outline three classes of canonicalization solutions, in increasing order of ambition. We will then explore the first solution class in more depth.

4.1 Record Selection

The *record selection* approach to canonicalization selects an existing record as its output. For example, $C(\mathbf{R})$ must select from the three records in Table 1. Record selection algorithms must ensure that the selected record contains no errors, and that it is representative of other records. Note that this approach is most prone to errors of incompleteness, since one record may not contain all the fields present in the duplicates. For example, selecting record (a) in Table 1 will omit the page numbers, but selecting record (b) will omit the full first name of the author.

4.2 Record Merging

The *record merging* approach to canonicalization constructs a canonical record by piecing together fields from different records.

While this approach can increase the *completeness* of canonicalization, it does so at the risk of introducing errors. In the worst case, an error in record deduplication may merge together records that in fact refer to different objects. Constructing one record containing fields from these non-duplicate records can result in a canonical record containing invalid information. For example, a record merging approach may return the record in Table 2:

author	Brian Milch et al.
title	BLOG: Probabilistic Models with Unknown Objects
venue	Intl. Conf. on AI
editor	L. Getoor and B. Taskar
pages	1352-1359

Table 2: Possible record merging canonicalization for the records in Table 1.

While this result is complete, it erroneously includes the editor field from record (c), which is not truly a duplicate. To address this issue, it may be useful to consider measures of field compatibility, as in Wick et al. [14].

4.3 Record Generation

The *record generation* approach to canonicalization is an extension of the record merging approach that may also propose field values that do not explicitly exist in any of the record duplicates. This allows the system to predict field values based on pattern analysis. For example, a record generation approach may return the following record:

author	Brian Milch et al.
title	BLOG: Probabilistic Models with Unknown Objects
venue	International Joint Conference on Artificial Intelligence
editor	
pages	1352-1359

Table 3: Possible record generation canonicalization for the records in Table 1.

Here, the system has generated an expanded venue value from the abbreviated form, even though this expanded form does not appear among the duplicate records. This predictive operation can be accomplished either by learning statistical patterns in the database, or by a pattern-matching approach.

While in this case record generation succeeded, in general positing field values that do not exist in any of the records can be quite dangerous and lead to unacceptable errors.

These three solution classes motivate a number of implementations and experiments. In this paper, we describe three record selection methods and perform experiments to measure their effectiveness.

5. THREE PROPOSALS FOR RECORD SELECTION CANONICALIZATION

5.1 Edit distance with fixed costs

The motivation for our approach is to minimize the effect of pre-processing errors on canonicalization. As we have described, errors from PDF-to-text conversions, misspellings, and incorrect deduplication can lead to poor canonicalization choices.

We make two assumptions about the behavior of pre-processing errors:

- Correct records are more common than incorrect records. That is, most records are error-free.
- Errors have high variance. For example, it is unlikely for many records to have the same exact spelling mistake.

With these assumptions in mind, we propose selecting the record that has the greatest average similarity to every other document. We define the distance between two records as the string edit distance between them.

Let $D : R_i \times R_j \mapsto \mathcal{R}^+$ be the edit distance between two records. Given a set of duplicate records $\mathbf{R} = \{R_1 \dots R_k\}$, we define the average edit distance of record R_i as

$$A(R_i) = \frac{\sum_{R_j \in \mathbf{R}} D(R_i, R_j)}{k} \quad (1)$$

The canonical record we return is the one with minimum average distance to every other string:

$$C^d(\mathbf{R}) = \operatorname{argmin}_{R_i \in \mathbf{R}} A(R_i)$$

We refer to $C^d(\mathbf{R})$ as the *edit distance canonicalizer*.

We now must decide on the form of D , the metric defining the distance between two strings. A natural choice is the Levenshtein distance: the number of character insertions, deletions, and replacements required to transform one string into another [6]. The recursive definition of the Levenshtein distance for strings s^n and t^m with length n and m is the following:

$$D(s^n, t^m) = \min \begin{cases} c_r(s_n, t_m) + D(s^{n-1}, t^{m-1}) \\ c_i + D(s^{n-1}, t^m) \\ c_d + D(s^n, t^{m-1}) \end{cases} \quad (2)$$

where $c_r(s_n, t_m)$ is the *replacement cost* for swapping character s_n with character t_m , c_i is the *insertion cost*, and c_d

Algorithm 1 Exhaustive cost search

```

1: Input:
   Training set  $\mathcal{S}$ 
   Initial costs  $\mathbf{c} = \{c_i, c_d, c_r^-, c_r^\neq\}$ 
   max, min, step
2: while More Costs do
3:    $\mathbf{c} \leftarrow \text{NextCosts}(\mathbf{c}, \text{max}, \text{min}, \text{step})$ 
4:   if  $\mathcal{L}(\mathbf{c}, \mathcal{S}) < \text{bestLoss}$  then
5:      $\text{bestLoss} \leftarrow \mathcal{L}(\mathbf{c}, \mathcal{S})$ 
6:      $\mathbf{c}^* \leftarrow \mathbf{c}$ 
7:   end if
8: end while

```

is the *deletion cost*. We can further define the replacement cost as

$$c_r(s_n, t_m) = \begin{cases} c_r^\neq & \text{if } s_n \neq t_m \\ c_r^- & \text{if } s_n = t_m \end{cases} \quad (3)$$

That is, c_r^\neq is the cost of replacing one character with another, and c_r^- is the cost of copying a character from one string to the next. We refer to c_r^\neq as the *substitution cost*, and c_r^- as the *copy cost*.

The value of the edit distance costs greatly effects the output of the system. For example, if c_i is small, then abbreviated strings will have a small distance to their expanded version. Abbreviated strings will therefore have lower values of $A(R_i)$.

Rather than requiring the user to manually tune these costs, in the next section we propose ways of learning these costs automatically given labeled examples.

5.2 Edit distance with learned costs

Suppose the user provides a labeled training set

$$\mathcal{S} = \{\langle \mathbf{R}^1, l_1 \rangle \dots \langle \mathbf{R}^n, l_n \rangle\}$$

where each set of duplicate records $\mathbf{R}^i = \{R_1 \dots R_k\}$ is annotated with label $l_i \in \{1 \dots k\}$, indicating which of the duplicates should be selected as the canonical record (i.e., $R_{l_i} \in \mathbf{R}$ is the true canonical record). We wish to use \mathcal{S} to learn the weights of D .

There has been a fair amount of work on methods to automatically learn edit distance costs, mostly applied to record de-duplication; (see Section 2). However, we are not aware of any work that learns edit distance costs for canonicalization.

We propose two simple but effective methods to learn edit distance costs from training data: *exhaustive search* and *stochastic hill climbing*.

5.2.1 Exhaustive search

The simplest method is to exhaustively enumerate settings of each cost and maximize the canonicalization performance on the training set.

Let $\mathcal{L}(\mathbf{c}, \mathcal{S})$ be the loss function for an assignment to \mathbf{c} . For example, \mathcal{L} may be the proportion of records in \mathcal{S} for which $C^d(\mathbf{R}^i)$ returns a non-canonical record; i.e. $C^d(\mathbf{R}^i) \neq R_{l_i}$.

We wish to optimize \mathbf{c} as follows:

$$\mathbf{c}^* = \operatorname{argmin}_{\mathbf{c}} \mathcal{L}(\mathbf{c}, \mathcal{S}) \quad (4)$$

Since we must discretize the cost settings to perform exhaustive search, the input is the following:

- **min:** The minimum cost value

Algorithm 2 Stochastic cost search

```

1: Input:
   Training set  $\mathcal{S}$ 
   Initial costs  $\mathbf{c} = \{c_i, c_d, c_r^-, c_r^+\}$ 
   max, min, step
2: for  $i < \text{NumIterations}$  do
3:    $c \leftarrow \text{SampleCostElement}(\mathbf{c})$ 
4:    $c \leftarrow \text{RandomUpdate}(c, \text{step}, \text{max}, \text{min})$ 
5:    $\mathbf{c} \leftarrow \text{NextCosts}(\mathbf{c}, \text{max}, \text{min}, \text{step})$ 
6:   if  $\mathcal{L}(\mathbf{c}, \mathcal{S}) < \text{bestLoss}$  then
7:      $\text{bestLoss} \leftarrow \mathcal{L}(\mathbf{c}, \mathcal{S})$ 
8:      $\mathbf{c}^* \leftarrow \mathbf{c}$ 
9:   end if
10:   $i \leftarrow i + 1$ 
11: end for

```

- **max:** The maximum cost value
- **step:** The amount to perturb each cost to obtain a new setting.

Search proceeds by cycling through each setting of \mathbf{c} and returning the best found setting \mathbf{c}^* . The details are given in Algorithm 1. The method *NextCosts* generates the next cost setting as determined by the step size.

5.2.2 Stochastic hill-climbing

Computing $\mathcal{L}(\mathbf{c}, \mathcal{S})$ requires computing $C^d(\mathbf{R})$ for all $\mathbf{R}^i \in \mathcal{S}$. This computational cost limits the number of settings we can enumerate using exhaustive search. Instead, we propose a simple stochastic hill-climbing algorithm to optimize Equation 4. Given an initial setting for \mathbf{c} , the algorithm proposes a modification to \mathbf{c} and accepts the change if $\mathcal{L}(\mathbf{c}, \mathcal{S})$ decreases. This can be understood as simulated annealing without the *temperature* parameter. The details of this method are given in Algorithm 2.

The method *SampleCostElement* samples a cost uniformly from the cost vector. The method *RandomUpdate* uniformly chooses between incrementing or decrementing \mathbf{c} by *step*.

5.3 Feature-based Ranking Models

While the adaptive edit distance approach to record selection can be simple and effective, it is limited by the small number of tunable parameters (the four costs), which limits the expressivity of the model.

In this section, we propose two feature-based learning approaches that enable the use of arbitrary features over the records, including the output of various edit distance measures.

Consider a set of duplicate records $\mathbf{R} = \{R_1 \dots R_k\}$. Let $\Phi_i = \{\phi_1(R_i) \dots \phi_m(R_i)\}$ be a vector of *binary feature functions* $\phi : R \mapsto \{0, 1\}$ that compute evidence indicating whether R_i should be selected as the canonical record. For example, $\phi_j(R_i)$ may be 1 if record R_i is the longest record in \mathbf{R} . Let $\Lambda = \{\lambda_1 \dots \lambda_m\}$ be a vector of real-valued weights associated with each feature.

We can compute a score for the event that R_i is chosen as the canonical record by taking the dot product of the features and weights:

$$\tau(R_i, \Lambda) = \Phi_i \cdot \Lambda$$

Below we describe two methods to estimate Λ from the training set \mathcal{S} .

5.3.1 Ranking Logistic Regression

The first method is based on logistic regression (sometimes called *maximum entropy classification*). Here, we modify the traditional logistic regression loss function to *rank*, rather than *classify*, instances.

Let the binary random variable C_i be 1 if and only if record R_i is the true canonical record of \mathbf{R} . Given Λ and F , we can compute the probability of C_i as follows:

$$p(C_i = 1 | \mathbf{R}, \Lambda) = \frac{e^{\tau(R_i, \Lambda)}}{\sum_{R_j \in \mathbf{R}} e^{\tau(R_j, \Lambda)}}$$

where the score for record R_i is normalized by the scores for every other duplicate record.

Note that this formulation differs from traditional classification, which computes an independent binary classification decision for each record:

$$p(C_i = 1 | \mathbf{R}, \Lambda) = \frac{e^{\tau(R_i, \Lambda_1)}}{e^{\tau(R_i, \Lambda_1)} + e^{\tau(R_i, \Lambda_0)}}$$

where the parameters Λ_1 represent the positive class, and Λ_0 represent the negative class.

By formulating canonicalization as a ranking task rather than a classification task, we can compute a loss function that is sensitive to competing records. This can be beneficial for at least two reasons. First, if no record is error-free, a classification loss function is forced to place a positive label on a partially incorrect example. Second, if a non-canonical record shares many features with the canonical record, a classification loss function will erroneously penalize those features. By focusing on the *differences* between examples, a ranking loss function overcomes these deficiencies.

We can estimate Λ from the training set \mathcal{S} by minimizing the negative log-likelihood of the data given Λ :

$$\mathcal{L}(\Lambda, \mathcal{S}) = - \sum_{\mathbf{R}^i \in \mathcal{S}} \log p(C_i | \mathbf{R}, \Lambda) \quad (5)$$

Note that this is the sum of probabilities for each of the *correct* canonical records for the current setting of Λ . We also add a Gaussian prior over Λ with fixed mean and variance to mitigate over-fitting. We find the setting of Λ that minimizes Equation 5 using limited-memory BFGS, a gradient ascent method with a second-order approximation [7].

5.3.2 MIRA

MIRA (Margin Infused Relaxed Algorithm) is a relaxed, online maximum margin training algorithm [4]. It iteratively cycles through the training set and updates the parameter vector with two constraints: (1) the true canonical record must have a higher score than any other record by a given margin, and (2) the change to Λ should be minimal. This second constraint is to reduce fluctuations in Λ . Using the same scoring function τ as in the previous section, this optimization is solved through the following quadratic program:

$$\begin{aligned} \Lambda^{t+1} &= \underset{\Lambda}{\operatorname{argmin}} \|\Lambda^t - \Lambda\|^2 \\ &\text{s.t.} \\ \tau(R_{i_t}, \Lambda) - \tau(R_{j_t}, \Lambda) &\geq 1 \quad \forall R_j \neq R_{i_t} \end{aligned}$$

In this case, the MIRA parameter update is a quadratic program with constraint size equal to the number of non-canonical records in the training example. This QP can be

solved efficiently using the Hildreth and D’esopo method [2]. To improve the stability of this online method, we average the parameter vectors from each update at the end of training, as in voted perceptron [3].

6. EXPERIMENTS

6.1 Data

We collect 3,683 citations to 100 distinct papers from REXA, an online publications search engine.³ These citations were automatically extracted from the headers of research papers as well as from the reference section, and record deduplication was performed automatically by the REXA system. The data therefore contains misspellings, PDF-to-text errors, abbreviations, and possibly deduplication errors. To construct a labeled data set, we collect the corresponding citations to these papers from the Digital Bibliography and Library Project (DBLP).⁴ The DBLP citations are manually curated to ensure accuracy, so they provide a good source of canonical examples. In fact, as part of its pipeline, REXA crawls the DBLP repository and performs record deduplication to merge citations together.

For these experiments, we focus on constructing the canonical representation of the conference string for each paper. This is arguably the most difficult field to canonicalize because of the prevalence of acronyms, abbreviations, and misspellings.

Using the DBLP data, we construct two versions of the dataset. In the first, the true canonicalization is the conference title acronym. This simulates the use case in which the user desires abbreviated canonical forms. In the second version, the true canonicalization is the expanded conference title. This simulates the case when the user does not desire any abbreviations in the canonical form. We refer to the former version as the *acronym* dataset, and the latter as the *expanded* dataset.

Table 4 shows an example with labels from each of the datasets. We can see that the duplicate records contain a variety of abbreviated forms, as well as PDF-to-text conversion errors in the first and fourth duplicate records (*Artificial, Conference on artificial*) that make canonicalization difficult.

Figure 1 shows the distribution of the number of duplicates for each record in the dataset. As we can see, most records have around 20 duplicates, but a few records have over 200 duplicates.

We perform 5-fold cross validation on the data, with each split containing 80 training examples and 20 testing examples. To evaluate performance, we consider two measures:

- **Mean Reciprocal Rank (MRR):** The average rank among the duplicates given to the true canonical record. The ranking is computed by sorting the records according to their canonicalization score. This measure is commonly used in information retrieval to evaluate search results.
- **Accuracy (Acc):** The proportion of predicted canonical records that are truly canonical. This can also be understood as MRR at rank 1.

³Data is available at <http://www.cs.umass.edu/~culotta/data/canonicalization.html>

⁴<http://dblp.uni-trier.de>




Figure 1: Distribution of number of duplicates per record.

6.2 Systems

We evaluate eight different systems:

- **Edit-Distance, Fixed Costs (ED-F)** - Levenshtein string distance canonicalizer with the default costs $c_i = 1, c_d = 1, c_r^{\neq} = 1, c_r^{\bar{}} = 0$. (See Section 5.1.)
- **Edit-Distance, Exhaustive Cost Search (ED-E)** - Levenshtein string distance canonicalizer with costs set by exhaustive search (See Section 5.2.1.) We set $max = 1.0, min = -1.0$, and $step = 0.5$, resulting in 81 different settings.
- **Edit-Distance, Stochastic Cost Search (ED-S)** - Levenshtein string distance canonicalizer with costs set by stochastic search (See Section 5.2.2.) We set $step = 0.2$ and perform 20 iterations.
- **Logistic Regression (LR)** - Feature-based ranking model with exponential loss function. (See Section 5.3.1.)
- **MIRA (M)** - Feature-based ranking model with large-margin loss function. (See Section 5.3.2.)
- **Shortest (S)** - A baseline method that ranks records in increasing order of length.
- **Longest (L)** - A baseline method that ranks records in decreasing order of length;
- **Most Common (C)** - Assign each record a count equal to the number of exact string duplicates contained in the record set. Rank records in decreasing order of count.

The features for the feature-based canonicalizers are as follows:

- **Edit-distance features:** We construct four different edit-distance canonicalizers from four different setting of the Levenshtein costs.⁵ For each cost setting, we

⁵ $(c_d = c_r^{\neq} = c_i = 1, c_r^{\bar{}} = 0); (c_d = 0, c_r^{\neq} = c_i = 1, c_r^{\bar{}} = 0); (c_d = 1, c_r^{\neq} = 0, c_i = 1, c_r^{\bar{}} = 0); (c_d = c_r^{\neq} = 1, c_i = c_r^{\bar{}} = 0)$

acronym canonical record
In AAAI
expanded canonical record
In proceedings of the Ninth National Conference on Artificial Intelligence
duplicate records
In proceedings of the Ninth National Conference on <i>Artifici al</i> Intelligence
proc the 9th National Conf on AI
in proc AAAI
in proceedings of the Ninth National <i>Conferenceonarti cial</i> Intelligence

Table 4: An example of a set of conference strings to be canonicalized annotated with two versions of preferred canonicalization: acronym and expanded forms.

canonicalizer	MRR	Acc	Time (s)
LR	.708 (.017)	.6 (.015)	75 (5)
M	.661 (.017)	.55 (.035)	80 (4)
ED-S	.597 (.053)	.5 (.061)	278 (26)
ED-E	.578 (.053)	.5 (.061)	925 (65)
C	.551 (.023)	.53 (.043)	.05 (.02)
ED-F	.438 (.034)	.37 (.04)	6 (1)
L	.426 (.033)	.28 (.03)	.06 (.05)
S	.087 (.007)	0 (0)	.07 (.06)

Table 5: Mean reciprocal rank, accuracy, and running time on *expanded* dataset. The numbers in parentheses are the standard error over five cross-validation trials.

canonicalizer	MRR	Acc	Time (s)
M	.94 (.014)	.92 (.012)	103 (5)
LR	.935 (.02)	.92 (.025)	63 (5)
ED-E	.868 (.027)	.82 (.04)	866 (99)
ED-S	.865 (.028)	.82 (.04)	254(33)
S	.767 (.038)	.64 (.043)	.02 (.004)
C	.126 (.033)	.06 (.024)	.05 (.03)
ED-F	.059 (.004)	0 (0)	6 (2)
L	.049 (.003)	0 (0)	.011 (.001)

Table 6: Mean reciprocal rank, accuracy, and running time on *acronym* dataset. The numbers in parentheses are the standard error over five cross-validation trials.

compute $C^d(\mathbf{R})$ as in Equation 1. The binary features for a record indicate if it is the first, second, or third highest ranked record according to $C^d(\mathbf{R})$ (for example, one feature is *ranked-second-by-canonicalizer3*). These features allow the classifier to serve as a “meta-canonicalizer” by aggregating the output of many different canonicalizers.

- **Text features:** We compute several features that examine the properties of the strings themselves.
 - *Acronyms:* This feature is true if the record contains a token on a list of known acronyms (e.g., *ICML*).
 - *Abbreviations:* This feature is true if the record contains a token on a list of known abbreviations (e.g., *conf* for *conference* and *proc* for *proceedings*).
 - *Relative length:* We compute the character length of each record and create features that indicate if a record is the first, second, or third longest or shortest record.

6.3 Results

Tables 5 and 6 display results for the eight different methods on the *acronym* and *expanded* datasets.

From these results, we can conclude that the feature-based canonicalizers consistently outperform the edit-distance canonicalizers. We can see that for the *expanded* data logistic regression (**LR**) outperforms the fixed cost edit-distance (**ED-F**) by 27% MRR, and further outperforms the stochastic search edit distance (**ED-S**) by 11% MRR. The difference

between the two feature-based canonicalizers is small: logistic regression outperforms MIRA (**M**) by nearly 5% MRR on the *expanded* data, but MIRA outperforms logistic regression by .5% MRR on the *acronym* data. Similarly, the difference between the two cost learning methods is small (**ED-S** versus **ED-E**).

Furthermore, we can conclude that cost-learning greatly improves the performance of the edit-distance canonicalizers, increasing MRR by nearly 16% on the *expanded* data and by 80% on the *acronym* data. The pronounced difference on the *acronym* data can be attributed to the fact that the default setting used in **ED-F** has unit cost for inserting characters. This gives acronym records a large distance from non-acronym records, making it unlikely they will have the lowest average distance. However, the cost learning methods can discover settings that do not penalize insertions, thereby reducing the average edit-distance of acronyms.

None of the simpler baseline methods perform consistently well across the two datasets. Simply choosing the shortest, longest, or most common record is significantly worse than using one of the more complex record selection algorithms we propose.

6.4 Impact of features

We investigate the impact of features on the performance of the feature-based canonicalizers. Table 7 displays performance of the logistic regression method (**LR**) with and without the textual features described in Section 6.2. These results show that using edit-distance features alone still outperforms the fixed-cost edit-distance canonicalizer **ED-F** by nearly 6% (.496 versus .438 from Table 5). However, the majority of the improvement of the feature-based classifiers appears to come from textual features. This result suggests

features	MRR	Acc
edit-distance + text	.708 (.017)	.6 (.015)
edit-distance	.496 (.027)	.29 (.033)

Table 7: Mean reciprocal rank and accuracy of LR on the expanded dataset with and without textual features. The numbers in parentheses are the standard error over five cross-validation trials.

a simple approximation to improve the scalability of this approach, which we discuss in Section 6.7.

6.5 Learning rates

In a real-world application, it may be difficult to obtain labeled data from the user. We therefore perform experiments to evaluate how many labeled examples are needed to obtain accurate results. Figures 2 and 3 plot performance as the proportion of training data used increases. As we can see, using only 10% of the data (8 examples), performance is already quickly approaching its maximum.

6.6 Robustness to noise

We perform additional experiments to measure how robust the methods are to the introduction of non-canonical records. For each training example \mathbf{R} , we introduce noise as follows:

- Select an incorrect record uniformly at random $R_i \in \mathbf{R}$ s.t. $R_i \neq R_{i^*}$.
- Add n duplicates of R_i to \mathbf{R} .

Figures 4 and 5 show results as n varies from 0 to 20. We compare the four learning methods, as well as the *Most Common* baseline (**C**). These figures show that the feature-based methods are quite robust to noise, as their accuracy drops only slightly as n increases. The *Most Common* baseline degrades significantly, which is unsurprising since as n increases it is very likely that it chooses the incorrect record. The exhaustive cost-learning method also appears relatively robust; however the stochastic cost-learning method degrades significantly.

6.7 Scalability

In Table 5 and 6 we report the wall-clock running time of each method. Note that this includes the time to train each method. The logistic regression requires about one minute to train and evaluate on 80 training examples and 20 testing examples. Note that the long running times of the cost-learning methods is high because for each setting of costs, the average edit-distance for each record must be recomputed to calculate the loss function. This is in contrast to the feature-based methods, which uses fixed costs for the edit-distance features.

For databases containing many records with many duplicates, the computation of the average edit-distance may become burdensome. The edit distance computation has time complexity $O(m^2)$ where $|\mathbf{R}| = m$. $A(R_i)$ requires iterating over all records, and we must compute this m times. Since the edit distance canonicalizer is used as input to the feature-based canonicalizers, these have time complexity $\Omega(m^2)$. Thus, canonicalizing m records requires $\Omega(m^2k)$ time.




Figure 2: Learning curves for expanded dataset.




Figure 3: Learning curves for acronym dataset.

We can alleviate the quadratic dependence on m by pruning elements of \mathbf{R}^i that are unlikely to be chosen as the canonical record. We propose the following method to prune records for the feature-based canonicalizers:

- Build a feature-based canonicalizer C' that only uses text features, not edit-distance features. Thus, this canonicalizer does not require the m^2 computation to compute edit-distance.
- Score each element of \mathbf{R} using C' .
- Remove all R_i with scores less than threshold δ .
- Run the original canonicalizer using edit-distance features.

This method therefore prunes records from consideration prior to computing the edit-distance. We leave empirical evaluation of this approximation for future work.




Figure 4: Noise experiments for *expanded dataset*.




Figure 5: Noise experiments for *acronym dataset*.

7. CONCLUSION AND FUTURE WORK

Record canonicalization is an important and under-studied problem in databases populated with heterogeneous, imperfect data. In this paper, we have formalized the canonicalization problem and proposed three broad classes of solutions. We have implemented three instantiations of one solution class and empirically evaluated them on manually annotated data. These experiments show that it is possible to build a system to accurately learn canonicalization preferences with only a few examples. Furthermore, this approach appears to be quite robust to the types of noise common in automatically extracted records.

In future research, we plan to explore record merging and record generation approaches to canonicalization. We also plan to explore the interaction of canonicalization with deduplication. These two tasks are inter-related, and it is likely that performing joint inference across them can improve performance.

8. ACKNOWLEDGEMENTS

This work was supported in part by the Defense Advanced Research Projects Agency (DARPA), through the Department of the Interior, NBC, Acquisition Services Division, under contract #NBCHD030010, in part by U.S. Government contract #NBCH040171 through a subcontract with BBNT Solutions LLC, in part by The Central Intelligence Agency, the National Security Agency and National Science Foundation under NSF grant #IIS-0326249, in part by Microsoft Live Labs, and in part by the Defense Advanced Research Projects Agency (DARPA) under contract #HR0011-06-C-0023.0

9. REFERENCES

- [1] M. Bilenko and R. J. Mooney. Learning to combine trained distance metrics for duplicate detection in databases. Technical Report AI-02-296, University of Texas at Austin, 2002.
- [2] Y. Censor and S. Zenios. *Parallel optimization: theory, algorithms, and applications*. Oxford University Press, 1997.
- [3] M. Collins. Discriminative training methods for hidden markov models: Theory and experiments with perceptron algorithms. In *EMNLP*, 2002.
- [4] K. Crammer, O. Dekel, J. Keshet, S. Shalev-Shwartz, and Y. Singer. Online passive-aggressive algorithms. *J. Mach. Learn. Res.*, 7:551–585, 2006.
- [5] R. Gupta and S. Sarawagi. Creating probabilistic databases from information extraction models. In *VLDB*, 2006.
- [6] V. Levenshtein. Binary codes capable of correcting deletions, insertions and reversals. *Doklady Akademii Nauk SSR*, 163(4):845–848, 1965.
- [7] D. C. Liu and J. Nocedal. On the limited memory BFGS method for large scale optimization. *Math. Programming*, 45(3, (Ser. B)):503–528, 1989.
- [8] G. Mann and D. Yarowsky. Multi-field information extraction and cross-document fusion. In *ACL*, 2005.
- [9] A. McCallum, K. Bellare, and F. Pereira. A conditional random field for discriminatively-trained finite-state string edit distance. In *Conference on Uncertainty in AI*, 2005.
- [10] A. McCallum and B. Wellner. Conditional models of identity uncertainty with application to noun coreference. In L. K. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems 17*. MIT Press, Cambridge, MA, 2005.
- [11] B. Milch, B. Marthi, S. Russell, D. Sontag, D. L. Ong, and A. Kolobov. BLOG: Probabilistic models with unknown objects. In *IJCAI*, 2005.
- [12] E. S. Ristad and P. N. Yianilos. Learning string edit distance. Technical Report CS-TR-532-96, Princeton University, 1997.
- [13] S. Tejada, C. A. Knoblock, and S. Minton. Learning object identification rules for information integration. *Information Systems*, 26(8):607–633, 2001.
- [14] M. Wick, A. Culotta, and A. McCallum. Learning field compatibilities to extract database records from unstructured text. In *EMNLP*, 2006.
- [15] J. J. Zhu and L. H. Unger. String edit analysis for merging databases. In *KDD*, 2000.