

Data Cleaning for XML Electronic Dictionaries via Statistical Anomaly Detection

Michael Bloodgood

Center for Advanced Study of Language
University of Maryland
College Park, MD 20742
meb@umd.edu

Benjamin Strauss†

Department of Computer Science and Engineering
The Ohio State University
Columbus, OH 43210
strauss.105@osu.edu

Abstract—Many important forms of data are stored digitally in XML format. Errors can occur in the textual content of the data in the fields of the XML. Fixing these errors manually is time-consuming and expensive, especially for large amounts of data. There is increasing interest in the research, development, and use of automated techniques for assisting with data cleaning. Electronic dictionaries are an important form of data frequently stored in XML format that frequently have errors introduced through a mixture of manual typographical entry errors and optical character recognition errors. In this paper we describe methods for flagging statistical anomalies as likely errors in electronic dictionaries stored in XML format. We describe six systems based on different sources of information. The systems detect errors using various signals in the data including uncommon characters, text length, character-based language models, word-based language models, tied-field length ratios, and tied-field transliteration models. Four of the systems detect errors based on expectations automatically inferred from content within elements of a single field type. We call these *single-field* systems. Two of the systems detect errors based on correspondence expectations automatically inferred from content within elements of multiple related field types. We call these *tied-field* systems. For each system, we provide an intuitive analysis of the type of error that it is successful at detecting. Finally, we describe two larger-scale evaluations using crowdsourcing with Amazon’s Mechanical Turk platform and using the annotations of a domain expert. The evaluations consistently show that the systems are useful for improving the efficiency with which errors in XML electronic dictionaries can be detected.

I. INTRODUCTION

There is increasing interest in the research, development, and use of automated techniques for assisting with *data cleaning*, also called *data cleansing* or *scrubbing*, which deals with detecting and removing errors and inconsistencies from data in order to improve the quality of data [1]. In this paper we deal with data cleaning of electronic dictionaries stored in Extensible Markup Language (XML) format. XML is a markup language that defines a set of rules for encoding documents in a format that is both human-readable and machine-readable. Defined by free open standards, XML is a textual data format with strong support via Unicode for different human languages. It is widely used for the representation of

electronic dictionaries and other forms of structured data [2], [3].

Electronic dictionaries are a fundamentally important semantic computing resource. They are a core resource consumed by downstream processes in the provision of various human language technologies as well as consumed directly by human language learners and as reference materials more generally. When dictionaries are digitized, whether via manual entry, Optical Character Recognition (OCR), or a mixture of these methods, it is inevitable that errors are introduced into the digitized version that is produced.

Prior work has focused on providing editing tools to assist with manual curation of the data and on providing tools for automatically detecting structural errors using only structural information. Although these systems have had success in finding dictionary errors, there are many errors that cannot be detected without analyzing the text content of the dictionary. For example, suppose that in some dictionary a lexical entry requires a headword, a part of speech, and a definition. Suppose there is an error in which the definition for some entry appears in the part of speech field, and the definition field contains the headword of the next entry. These errors will not be found by examining only the structure. Unlike previous work, in this paper we present methods for flagging statistical anomalies as likely errors in the textual content itself in electronic dictionaries in XML format using information from the textual content of the data.

We present six systems that detect errors using various signals in the data. The types of data quality problems that our systems are designed to detect are single-source instance level data quality problems [1]. Four of our systems detect errors based on information automatically inferred from content within elements of a single field type. We call these systems *single-field* systems. Two of the systems detect errors based on correspondence information automatically inferred from content within elements of multiple related field types. We call these latter systems *tied-field* systems. For each system, we provide an intuitive analysis of the type of error that it is successful at detecting. Finally, we describe two larger-scale evaluations using crowdsourcing with Amazon’s Mechanical Turk platform and using domain expert annotations. The evaluations consistently show that the systems are useful for

†This research was conducted while this author was a Faculty Research Assistant at the University of Maryland Center for Advanced Study of Language.

improving the efficiency with which errors in XML electronic dictionaries can be detected.

In the next section we situate our work with respect to previous related work. In section III we describe the error detection systems in detail and provide intuitive examples of the sorts of errors that each error detection system finds. In section IV we describe our experimental tests and provide experimental results and discussion of results. Finally, in section V we conclude.

II. RELATED WORK

A categorization of data quality problems addressed by data cleaning and an overview of data cleaning methods is provided in [1]. Many data cleaning methods are based on identifying discrepancies for user auditing, e.g., [4], [5]. The system in [4] is highly interactive; discrepancy detection is not their sole main focus. The discrepancy detection approach they outline is for users to define domains and then for discrepancies to be located through checking against the user-defined domains for constraint violation. In contrast, the methods in the current paper do not require user specification of domains and the methods in the current paper operate on the basis of different sources and indicators of discrepancies. In particular, the methods in the current paper are fundamentally different in that they operate on the basis of statistically anomalous events instead of constraint violations. In [5], a similar overall workflow is presented whereby the most suspicious examples are flagged for a human operator to inspect and annotate as clean or “garbage.” In contrast to the current paper, there is only one method in [5], which is to flag the examples with the highest information gain for inspection by a human operator. The work in [5] assumes the context of construction of an automated classifier in the computation of the information gain. The method in [5] is evaluated on the task of handwritten digit recognition by seeing how well classifier performance improves with the addition of the data cleaning approach. In contrast, the current paper uses different methods for detecting suspicious examples and does not assume the context of construction of automated classifiers. Also in contrast, the methods in the current paper are evaluated on XML electronic dictionaries by measuring how many and what percentages of the detected anomalies are annotated as data errors by domain experts. The methods in the current paper may be able to be used in a complementary fashion with the methods from [4] and [5] in future work.

Past work presented a method for repairing a digital dictionary in an XML format using a dictionary markup language called DML [6]. It remains time-consuming and error-prone however to have a human exhaustively read through and manually correct a digital version of a dictionary, even with languages such as DML available for making corrections once errors are detected. The methods we present in the current paper automatically scan through and detect errors in dictionaries. The methods in the current paper can be used in concert with error correction techniques such as dictionary markup languages.

Previous approaches have been presented for detecting structural errors in digitized dictionaries [7], [8]. The method in [8] works by linearizing the lexicon structure, converting the opening tags in XML into tokens and then considering the likelihoods of various strings of tokens using a language modeling approach. Anomalous branches of XML tags are flagged as structural errors. The method ignores the underlying text data within the dictionary and only detects structural errors. The methods in [7] outperform the method from [8]. The methods in [7] use a mixture of unsupervised methods, supervised machine learning methods, and system combination approaches. The highest-performing method uses a random forest system combination approach. The methods in [7] only detect structural errors. Errors in the textual data content within the XML elements are not detected. In contrast, the current paper presents methods that detect errors in the text (data) content of XML elements. The methods in the current paper also use different approaches and different sources of information than were used in [7], [8] and do not require training data, which is often not available. The error detection methods in the current paper can be used in concert with structural error detection methods.

III. METHODS

This section describes how our methods work. We categorize our methods into two types: *single-field* methods and *tied-field* methods. Single-field methods work by utilizing information within the data content of a single XML field type in order to detect errors. Tied-field methods work by utilizing information within the data content of multiple XML field types, exploiting various relationships between the data in the different fields, in order to detect errors.

For each candidate error, all of our methods return a numeric score indicating the system’s confidence that the candidate is an error. A threshold can be set for each method to control which candidates are detected as errors. The threshold for each method can be adjusted according to recall-precision¹ preferences and dictionary characteristics. In general, higher thresholds will return results with higher precision and lower recall whereas lower thresholds will return results with lower precision and higher recall. The optimal threshold for each method depends on data characteristics and user preferences. We are not aware of a method for determining optimal thresholds. We set our thresholds to a level that yielded a reasonable number of error candidates for human review. The exact thresholds for each experiment are given in the following subsections.

Subsection III-A describes how the single-field methods work and subsection III-B describes how the tied-field methods work. Subsection III-C provides examples of anomalies detected by the various methods.

¹Recall and precision are standard measures for systems that perform search. For the case of detecting dictionary errors, recall is the percentage of true errors that are found by the system. Precision is the percentage of system-proposed errors that are in fact true errors in the dictionary.

A. Single-Field Methods

The single-field error detection systems do not require any advance knowledge about the structure of the electronic dictionary. The only structural context information they use is the tag name of the elements containing the text data to be checked. Single-field methods can be used to check for errors in elements of any individual tag name. All single-field methods work according to the following high-level description: all entries of an individual tag name are processed and then any entries that are anomalous are flagged as errors. Each single-field method processes the entries and flags anomalies in different ways based on different aspects of the data. The rest of this subsection describes four single-field methods in detail.

1) *Uncommon Characters Method*: Uncommon characters can be a frequent source of errors in electronic dictionaries. They can arise due to OCR errors, author typographical errors, and mislabeled and/or incorrectly merged fields. Texts that contain uncommon characters are reported as potential errors. For each element in the dictionary with a particular tag, we consider the texts inside those elements to be a collection of documents D , and the characters in the texts as the tokens. We calculate the inverse document frequency of each character c observed in D as follows:

$$idf(c, D) = \log_{10} \frac{N}{|\{d \in D : c \in d\}|}, \quad (1)$$

where N is the number of documents in the collection.

When $idf(c, D) > threshold$, we consider c to be an uncommon character. The threshold is configurable; we use a default value of four. Users can adjust the threshold according to their recall-precision preferences and according to dictionary characteristics. The elements containing uncommon characters are flagged by the system as potential errors.

2) *Text Length Method*: When two text fields are inappropriately combined into one field, the result can be text that is unusually long. When a single text unit is inappropriately truncated, or split across two fields, the result can be text that is unusually short. Texts with unusually long or unusually short length are reported as potential errors. For each element in the dictionary with a given tag, we treat the lengths of the texts inside those elements as a sample of a normally distributed population. We calculate the mean and standard deviation of the sample, and for each value, we calculate the z-score, i.e., the signed number of standard deviations the value is above or below the mean. If the absolute value of the z-score of the length of a text is above a threshold, we flag the text as a possible error. The threshold is configurable; we use a default value of four. As the threshold is raised, only the most unusually long, or short, fields will be returned as errors.

3) *Word Sequence Method*: Language modeling can capture the probability of a sequence of words occurring in textual content. This gives us the capability to flag unlikely sequences of words. These unlikely sequences can often be indicative of typographical errors, OCR errors, incorrect field joining and splitting, etc. Texts that are unlikely given a word-level

language model of texts in the same context are reported as potential errors. For each element in the dictionary with a given tag, we build a language model of the text content of all the elements using n-grams of words. We calculate the entropy of the model with respect to each individual element's text, and treat the entropies as a sample of a normally distributed population. We calculate the z-score of each entropy score. High z-scores indicate texts that are unlikely in the language model. The texts with entropy z-scores above a threshold are flagged as errors. The size of the n-grams in the language model and the threshold are configurable. We use 4-grams and threshold five by default. Using a larger n-gram size in the language model allows one to potentially capture more nuanced sequence characteristics, however, it would require a much larger amount of data to estimate properly and avoid introducing spurious estimates due to data sparsity. Also, larger n-gram sizes are more computationally intensive. The entropy z-score thresholds can be adjusted according to recall-precision preferences and dictionary characteristics.

4) *Character Sequence Method*: This method is similar to the Word Sequence Method, except that instead of n-grams of words, we build language models using n-grams of characters. The size of the n-grams in the language model and the threshold are configurable. We use 4-grams and a threshold of five by default. Larger n-gram sizes could potentially capture more nuanced character sequence models, however, they would require a much larger amount of data to estimate properly and avoid introducing spurious estimates due to data sparsity. The n-gram size can also be adjusted based on the language of the field's textual content. The entropy z-score thresholds can be adjusted according to recall-precision preferences and dictionary characteristics. Note that the error detection system based on character sequences will in some cases find errors that the Uncommon Character system also detects, but the Character Sequence Method is also capable of finding some errors that the Uncommon Character system is not able to find. This is because the Character Sequence Method can find errors in which none of the characters in the textual content is particularly uncommon, but in which the ordering of those characters is incorrect.

B. Tied-Field Methods

In many structured data sets, there are pairs of fields that are related to each other in predictable ways. For example, in dictionaries a word in a language's native orthography and a phonetic transcription of the word are related because in many languages there are predictable relationships between spelling and pronunciation. Another related pair of fields in bilingual dictionaries is an example sentence demonstrating usage of a word and its translation. We call these sorts of related fields *tied fields* and we call error detection methods that exploit relationships between content in different field types *tied-field* methods. It is possible for there to be errors in a single field where the data value in that field is not anomalous in the context of only other values in that field type. However, the value might be anomalous in the context of data

TABLE I
EXAMPLES OF ORTHOGRAPHY-PRONUNCIATION PAIRS. SHORT STRINGS
HAVE A DIFFERENT DISTRIBUTION OF LENGTH RATIOS THAN LONG
STRINGS.

Short			Long		
Orth	Pron	Ratio	Orth	Pron	Ratio
t	tē	0.50	groundwork	ground''wûrk'	0.83
ease	ēz	2.00	lithargyrum	li''thâr''ji''rûm	0.79
v	vē	0.50	haidingerite	hi''ding''ēr''it	0.92

values in related fields. The single-field methods presented in subsection III-A will be unable to detect these sorts of errors. The rest of this subsection describes two tied-field methods that can detect these sorts of errors.

1) *Tied-Field Length Ratio Method*: This method determines the ratio of length in characters of tied fields; pairs of data values with unusual length ratios are then reported as potential errors. We treat the ratio of the length in characters of the tied-field data values to be a sample of a normally distributed population. We calculate the absolute value of the z-score of each length ratio. High values indicate tied-field instances where the data values have an unusual length ratio. Tied-field instances with scores above a threshold are flagged as potential errors. The threshold is configurable; we use a threshold of two by default. The threshold can be adjusted according to recall-precision preferences and dictionary characteristics.

To handle situations in which the distribution of length ratios is significantly different for short and long strings, we have an option to partition the tied-field instances by the length of the data in the first field, and treat each partition as its own population. Table I illustrates why it could be beneficial to partition tied-field instances by length. In these examples of correct pairs of data, the length ratios of the short tied-field pairs of data can be seen to vary more than the length ratios of the long tied-field pairs of data.

2) *Tied-Field Transliteration Method*: For some tied fields, the data in the two fields have a more specific relationship than a length relationship. For example, in dictionaries some fields are transliterations of each other. Such transliterations will usually have character-level correspondences (not necessarily a one-to-one correspondence). If the correspondence can be modeled, then pairs of texts that do not correspond well can be reported as errors. We use Phonetisaurus² to learn transliteration models across tied-fields. The Phonetisaurus system has been described in detail and has been shown to perform well in [9]. The resulting transliteration model represents how the first field can be transliterated into the second field, and can be used to generate scored transliteration candidates of the first field. For each pair of tied-fields, we use the transliteration model to transliterate the first field into the n-best candidates for the second field. Each candidate is given a transliteration cost by Phonetisaurus. By taking the inverse of the cost, we obtain a score indicating the model’s confidence in that transliteration candidate. We calculate the normalized

²<https://code.google.com/p/phonetisaurus/>

TABLE II
EXAMPLES OF ANOMALIES DETECTED BY SINGLE-FIELD SYSTEMS.

ExampleID	FieldName	Value
Example 1	GENDER	/F.
Example 2	NUMBER	PLU.
Example 3	PART-OF-SPEECH	PARTICLE

TABLE III
EXAMPLE ANOMALY DETECTED BY TIED-FIELD SYSTEMS.

ExampleID	FieldName	Value
Example 4	ORTHOGRAPHY	مگ ہف مگ مہو حسہ لس حسہ مگ ا
	PRONUNCIATION	rā

edit distance (NED) of each candidate to the observed data that actually is present in the second field, and calculate the mean NED weighted by transliteration score. We treat the weighted means as a sample of a normally distributed population. We calculate the z-score of each weighted mean. High values indicate that the data occurring in the second field is an unusually large NED away from what our learned model would have expected based on the data that occurred in the first field. Instances with weighted mean z-scores above a threshold are flagged as errors. The threshold is configurable; we use a default of two. The weighted mean z-score threshold can be adjusted according to recall-precision preferences and dictionary characteristics.³

C. Examples

In this subsection we provide examples of anomalies detected by the various systems. To obtain these examples, we ran the systems over a digitized sample of an Urdu-English dictionary [10] and selected illustrative examples that can be understood by most readers without having to know too many details about specific dictionary representations in XML. Table II shows examples detected by the various single-field systems and Table III shows an example detected by the tied-field systems.

Example 1 in Table II was detected by all four of our single-field systems. This is an error in the data since the value should have been just “F.” without the “/”. Since the GENDER field almost exclusively contains values of “M.” or “F.”, the Word Sequence Method found any other words such as “/F.” to have unusually high entropy. The character-based language model detected this error similarly. The Text Length Method detected this error since it is three characters long, which is unusually long given the predominance of two-character-long values in this field. The Uncommon Characters Method detected this error since the “/” character is uncommon in this field. Example 1 shows how the different methods can sometimes all detect the same error albeit through different views of the data.

Example 2 in Table II shows an example of an error that was detected by some of the systems and not others. This is an error in the data since the value should have been just “PL.”

³We also offer an option to normalize the tied-field text data to all lowercase letters.

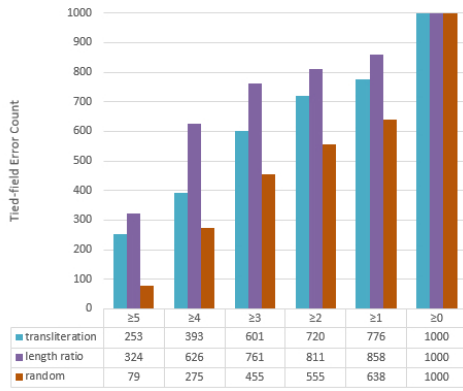


Fig. 2. Counts of proposed errors judged to be real errors by at least 5, 4, 3, 2, 1, and 0 Turkers.

the time. In contrast, for randomly selected proposed errors, three or more Turkers agree they are really errors only about 45.5% of the time. These results are evidence that the error detection systems could substantially increase the efficiency with which domain experts can clean XML data.

Diving a little deeper, Figure 3 and Figure 4 show the average number of Turkers that agree a proposed error is really an error for proposed errors at varying score cutoffs. Figure 3 shows the information for the tied-field length ratio error detection system. In Figure 3, the score cutoff on the x-axis is the absolute value of the z-score. The absolute value is used because this error detection system finds errors with both unusually high and unusually low z-scores. The y-axis is the average number of Turkers (out of 5) who marked the proposed errors with scores above the corresponding cutoff as real errors. Figure 4 shows similar information for the tied-field transliteration error detection system. In Figure 4, the score cutoff on the x-axis is the z-score. The z-score is used because this error detection system finds errors with unusually high z-scores. The y-axis is the average number of Turkers (out of 5) who marked the proposed errors with scores above the corresponding score cutoff as real errors.

For systems that will be used for purposes of ranking proposed errors in order from most likely to least likely, it is desirable that they predict errors with higher accuracy above a particular score threshold. This has positive implications for application settings where users will go through errors in a sorted order from most likely to least likely. Figure 3 and Figure 4 show that from this perspective the length ratio system produces better results. The length ratio system correctly predicts consistently with a cutoff over 5.5 standard deviations from the average. In contrast, the transliteration system does not perform as well because it is unable to predict errors with as high a degree of precision as the length ratio system at any score cutoff. A perhaps surprising result in Figure 4 is that the transliteration system decreases in precision when the score cutoff increases past about 2.5 standard deviations from the average. This result can be explained by the low number of proposed errors with very high scores with the transliteration

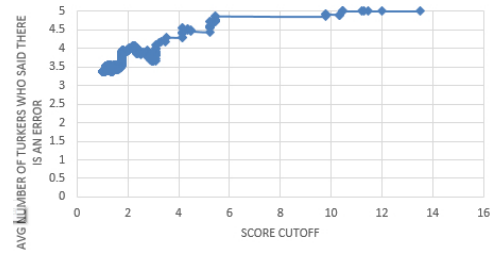


Fig. 3. The average number of Turkers (out of 5) who believe the orthography-pronunciation pair is an error for varying score cutoffs for the length ratio system. The score is the absolute value of the z-score.

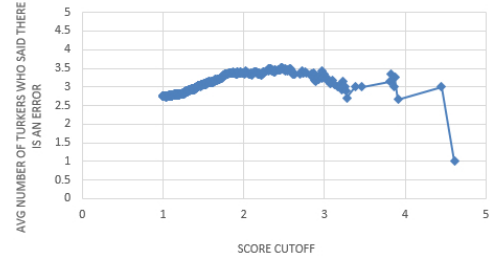


Fig. 4. The average number of Turkers (out of 5) who believe the orthography-pronunciation pair is an error for varying score cutoffs for the transliteration system. The score is the z-score.

system. This can create the situation where there are too few data points to compute a precision score reliably.

Table IV displays the overlap of the errors most strongly proposed by the length ratio error detection system and the transliteration error detection system for varying cutoff sizes. This table allows us to evaluate the similarity between the two systems. The two systems do not have a high degree of similarity; for example, only 4 of the top 100 anomalies detected by the length ratio system also appear in the top 100 anomalies detected by the transliteration system. The low levels of overlap indicate that the systems have complementary behavior. This complementary behavior can be leveraged to build improved hybrid systems.

The four single-field methods can also be combined with each other and with the tied-field methods. By combining methods a hybrid system could be built that takes into account different perspectives on the data. There are many possible ways of combining methods to build a hybrid system, e.g., see [7]. In future work, it would be worthwhile to explore how to optimally combine the error detection methods to create an improved system.

We conducted an additional evaluation of our systems using the annotations of a language expert in the process of correcting errors in a Tamasheq dictionary containing 5968 lexical entries. The language expert evaluated 175 anomalies proposed by the single-field text length system and the single-field uncommon character system for data in fields POS (part of speech) and MAIN (headword). The language expert marked each proposed anomaly as either a real error or not a real error. The results are in Table V. These results are further evidence that the error detection systems could substantially

TABLE IV
OVERLAP OF THE ANOMALIES MOST STRONGLY PROPOSED BY THE LENGTH RATIO AND THE TRANSLITERATION SYSTEMS.

Number of Proposed Errors	Number of Common Results	Percent
10	0	0%
25	1	4%
50	3	6%
100	4	4%
200	16	8%
300	31	10%
400	52	13%
500	81	16%
600	110	18%
700	143	20%
800	173	22%
900	186	21%
1000	227	23%
1500	396	26%
2000	563	28%

TABLE V
LANGUAGE EXPERT ANNOTATIONS ON 175 ANOMALIES PROPOSED BY THE TEXT LENGTH SYSTEM AND THE UNCOMMON CHARACTER SYSTEM FOR DATA IN FIELDS POS AND MAIN.

System	Field	Real Error	No Error	Total
Uncommon Character System	POS	16	1	17
	MAIN	4	10	14
Text Length System	POS	110	3	113
	MAIN	30	1	31

increase the efficiency with which domain experts can clean XML data.

V. CONCLUSION

There is increasing interest in methods for computer-aided rapid data cleaning. We presented multiple methods for data cleaning of XML electronic dictionaries. The methods detect errors in the data content of the XML, unlike previous work that detected errors in the structure of the XML electronic dictionaries and ignored the content. The methods are based on different underlying sources and indicators of errors and have complementary behavior with each other and with previously developed methods. The methods can be classified into single-field methods and tied-field methods. Single-field methods detect anomalies on the basis of expectations inferred from content in the same single field type. Tied-field methods detect anomalies on the basis of expectations of content correspondences inferred from content in multiple related fields. Four single-field methods for error detection were presented that work by using expectations of word sequence information, expectations of character sequence information, expectations of length, and expectations of individual characters in particular fields. Two tied-field methods for error detection were presented that work by using expectations of length ratios and expectations of string correspondences via transliteration models.

The precision of the error detection systems tends to correlate with the internal scores of the systems. This desirable behavior supports the scenario where a domain expert would go down a sorted list of proposed errors from most likely to

least likely. The performance of the different error detection systems varies based on the XML fields and the dictionaries on which they are applied. Domain experts can choose to invoke error detection system-field combinations with score cutoffs to suit their needs.

We evaluated these methods using the crowdsourcing platform Mechanical Turk and using expert annotations on multiple datasets. Sometimes the systems have overlapping behavior, detecting the same errors albeit through different views of the data. Often the systems have complementary behavior, which is promising for hybrid system construction in the future. We provided intuitive examples of the sorts of errors each of the systems can detect. In the evaluations, the systems are consistently helpful in increasing the efficiency with which data errors can be identified.

ACKNOWLEDGMENT

This material is based upon work supported, in whole or in part, with funding from the United States Government. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the University of Maryland, College Park and/or any agency or entity of the United States Government.

REFERENCES

- [1] E. Rahm and H. H. Do, "Data cleaning: Problems and current approaches," *IEEE Data Eng. Bull.*, vol. 23, no. 4, pp. 3–13, 2000.
- [2] N. Ide and J. Veronis, *Text encoding initiative: Background and contexts*. Springer Science & Business Media, 1995, vol. 29.
- [3] G. Francopoulo, M. George, N. Calzolari, M. Monachini, N. Bel, M. Pet, and C. Soria, "Lexical markup framework (LMF)," in *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC 2006)*. Genoa, Italy: European Language Resources Association, May 2006, pp. 233–236.
- [4] V. Raman and J. M. Hellerstein, "Potter's wheel: An interactive data cleaning system," in *VLDB*, vol. 1, 2001, pp. 381–390.
- [5] I. Guyon, N. Matic, and V. Vapnik, "Discovering informative patterns and data cleaning," in *Advances in knowledge discovery and data mining*. American Association for Artificial Intelligence, 1996, pp. 181–203.
- [6] D. Zajic, M. Maxwell, D. Doermann, P. Rodrigues, and M. Bloodgood, "Correcting errors in digital lexicographic resources using a dictionary manipulation language," in *Proceedings of the Conference on Electronic Lexicography in the 21st Century*. Bled, Slovenia: Trojina, Institute for Applied Slovene Studies, November 2011, pp. 297–301. [Online]. Available: http://www.trojina.si/elex2011/elex2011_proceedings.pdf
- [7] M. Bloodgood, P. Ye, P. Rodrigues, D. Zajic, and D. Doermann, "A random forest system combination approach for error detection in digital dictionaries," in *Proceedings of the Workshop on Innovative Hybrid Approaches to the Processing of Textual Data*. Avignon, France: Association for Computational Linguistics, April 2012, pp. 78–86. [Online]. Available: <http://www.aclweb.org/anthology/W12-0511>
- [8] P. Rodrigues, D. Zajic, D. Doermann, M. Bloodgood, and P. Ye, "Detecting structural irregularity in electronic dictionaries using language modeling," in *Proceedings of the Conference on Electronic Lexicography in the 21st Century*. Bled, Slovenia: Trojina, Institute for Applied Slovene Studies, November 2011, pp. 227–232. [Online]. Available: http://www.trojina.si/elex2011/elex2011_proceedings.pdf
- [9] J. R. Novak, N. Minematsu, and K. Hirose, "WFST-based grapheme-to-phoneme conversion: Open source tools for alignment, model-building and decoding," in *Proceedings of the 10th International Workshop on Finite State Methods and Natural Language Processing*. Donostia-San Sebastián: Association for Computational Linguistics, July 2012, pp. 45–49. [Online]. Available: <http://www.aclweb.org/anthology/W12-6208>
- [10] B. A. Qureshi and A. Haq, *Standard Twenty First Century Urdu-English Dictionary*. Delhi: Educational Publishing House, 1991.

- [11] R. Snow, B. O'Connor, D. Jurafsky, and A. Ng, "Cheap and fast – but is it good? evaluating non-expert annotations for natural language tasks," in *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*. Honolulu, Hawaii: Association for Computational Linguistics, October 2008, pp. 254–263. [Online]. Available: <http://www.aclweb.org/anthology/D08-1027>
- [12] M. Bloodgood and C. Callison-Burch, "Bucking the trend: Large-scale cost-focused active learning for statistical machine translation," in *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. Uppsala, Sweden: Association for Computational Linguistics, July 2010, pp. 854–864. [Online]. Available: <http://www.aclweb.org/anthology/P10-1088>
- [13] M. Negri, L. Bentivogli, Y. Mehdad, D. Giampiccolo, and A. Marchetti, "Divide and conquer: Crowdsourcing the creation of cross-lingual textual entailment corpora," in *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*. Edinburgh, Scotland, UK.: Association for Computational Linguistics, July 2011, pp. 670–679. [Online]. Available: <http://www.aclweb.org/anthology/D11-1062>
- [14] C. Callison-Burch, "Fast, cheap, and creative: Evaluating translation quality using Amazon's Mechanical Turk," in *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*. Singapore: Association for Computational Linguistics, August 2009, pp. 286–295. [Online]. Available: <http://www.aclweb.org/anthology/D09/D09-1030>
- [15] M. Bloodgood and C. Callison-Burch, "Using mechanical turk to build machine translation evaluation sets," in *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*. Los Angeles, California: Association for Computational Linguistics, June 2010, pp. 208–211. [Online]. Available: <http://www.aclweb.org/anthology/W10-0733>
- [16] E. Amigó, J. Artiles, J. Gonzalo, D. Spina, B. Liu, and A. Corujo, "Weps-3 evaluation campaign: Overview of the online reputation management task," in *CLEF 2010 (Notebook Papers/LABs/Workshops)*, 2010.
- [17] D. Chen and W. Dolan, "Collecting highly parallel data for paraphrase evaluation," in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Portland, Oregon, USA: Association for Computational Linguistics, June 2011, pp. 190–200. [Online]. Available: <http://www.aclweb.org/anthology/P11-1020>
- [18] M. Bloodgood and B. Strauss, "Translation memory retrieval methods," in *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*. Gothenburg, Sweden: Association for Computational Linguistics, April 2014, pp. 202–210. [Online]. Available: <http://www.aclweb.org/anthology/E14-1022>
- [19] G. A. Miller, R. Beckwith, C. Fellbaum, D. Gross, and K. J. Miller, "Introduction to wordnet: An on-line lexical database*," *International Journal of Lexicography*, vol. 3, no. 4, pp. 235–244, 1990.
- [20] G. A. Miller, "Wordnet: a lexical database for english," *Communications of the ACM*, vol. 38, no. 11, pp. 39–41, 1995.
- [21] C. Fellbaum, *WordNet: An Electronic Lexical Database*. MIT Press, 1998.