

Abstractive Tabular Dataset Summarization via Knowledge Base Semantic Embeddings

Paul Azunre
 Craig Corcoran
 David Sullivan
 Garrett Honke
 Rebecca Ruppel
 Sandeep Verma
 Jonathon Morgan

PAUL@NEWKNOWLEDGE.IO
 CRAIG@NEWKNOWLEDGE.IO
 DAVID@NEWKNOWLEDGE.IO
 GARRETT@NEWKNOWLEDGE.IO
 REBECCA@NEWKNOWLEDGE.IO
 SANDEEP@SAMPLE.COM
 JONATHON@SAMPLE.COM

New Knowledge, 3810 Medical Parkway, Austin, Texas

Editor: Editor’s name

Abstract

This paper describes an abstractive summarization method¹ for tabular data which employs a knowledge base semantic embedding to generate the summary. Assuming the dataset contains descriptive text in headers, columns and/or some augmenting metadata, the system employs the embedding to recommend a subject/type for each text segment. Recommendations are aggregated into a small collection of *super types* considered to be descriptive of the dataset by exploiting the hierarchy of types in a prespecified ontology. Using February 2015 Wikipedia as the knowledge base, and a corresponding DBpedia ontology as types, we present experimental results on open data taken from several sources – OpenML, CKAN and data.world – to illustrate the effectiveness of the approach.

Keywords: Dataset Summarization, Type Recommendation, Semantic Embeddings

1. Introduction

The motivation of this work is to develop a method for summarizing the content of tabular datasets. One can imagine the potential utility of automatically assigning a set of tags to each member of a large collection of datasets that would indicate the potential subject being addressed by the dataset. This can allow for semantic querying over the dataset collection to extract all available data pertinent to some specific task subject at scale.

We make the assumption that the dataset contains some text that is semantically descriptive of the dataset subject, whether appearing in columns, headers or some augmenting metadata. As opposed to an *extractive* approach that would merely select some exact words and phrases from the available text, we propose an *abstractive* approach that builds an internal semantic representation and produces subject tags that may not be explicitly present in the text augmenting the dataset.

The result of this work is *DUKE*—**D**ataset **U**nderstanding via **K**nowledge-base **E**mbeddings—a method that employs a pretrained Knowledge Base (KB) semantic embedding to perform *type recommendation* within a prespecified *ontology*. This is achieved by aggregating the

1. Our code is available for download at <https://github.com/NewKnowledge/duke>

recommended types into a small collection of *super types* predicted to be descriptive of the dataset by exploiting the hierarchical structure of the various types in the ontology. Effectively, the method represents employing an existing KB embedding to extensionally generate a *dataset2vec* embedding. Using a February 2015 Wikipedia knowledge base and a corresponding DBpedia ontology to specify types, we present experimental results on open data taken from several sources—OpenML, CKAN, and data.world—to illustrate the effectiveness of the approach.

2. Related Work

The *distributional semantics* (Sahlgren, 2008) concept has been recently widely employed as a natural language processing (NLP) tool to embed various NLP concepts into vector spaces. This rather intuitive hypothesis states that the meaning of a word is determined by its context. By far the most pervasive application of the hypothesis has been the word2vec model (Mikolov et al., 2013)(Pennington et al., 2014) which employs neural networks on large corpora to embed words that are contextually similar to be close to each other in a high-dimensional vector space. Arithmetic operations on the elements of the vector space produce semantically meaningful results, e.g., *King-Man+Woman=Queen*.

Since the original word2vec model, various incremental incarnations of it have been employed to embed sentences, paragraphs and even knowledge graphs into vector spaces via sent2vec(Pagliardini et al., 2017), paragraph2vec(Le and Mikolov, 2014), and RDF2Vec(Ristoski and Paulheim, 2016) respectively.

A topic *domain* is typically expressed as a manually curated ontology. A basic element of an ontology is a *type*, and a *type assertion* statement links specific entities of the knowledge graph to specific types. These statements can be used to augment a semantic embedding space with type information in order to add high level context of the graph to the embedding space. For instance, it was recently shown that one can extend a pretrained Knowledge Graph Embedding (KGE) to contain types of a specific ontology if those were not already present as entities, given a list of assertion statements(Kejriwal and Szekely, 2017). Thus, it can be assumed that a semantic embedding is *typed* for our purposes.

We note that the abstractive tabular dataset summarization problem is closely related to the well-studied problem of *type recommendation*, where the type is a super tag for all text segments in the dataset within a prespecified ontology that needs to be predicted. Systems for type recommendation using both manually curated features(Ma et al., 2013) and automated features (van Erp and Vossen, 2017), e.g., via typed KGEs(Kejriwal and Szekely, 2017), for *individual entities*, have been previously explored. To the best of our knowledge, this is the first application of typed semantic embeddings to abstractive tabular dataset summarization.

3. Approach

3.1. Framework

In this subsection, we present a pair of definitions to aid orientation.

Definition (word2vec) Word2vec models utilize a large corpus of documents to build a vector space mapping words to points in a space, where proximity implies semantic sim-

ilarity(Mikolov et al., 2013). This allows us to calculate distances between words in the dataset and the set of types in our ontology.

Definition (wiki2vec) When discussed in this paper, a *wiki2vec* model is a form of word2vec model trained on a corpus of Wikipedia KB documents². Training on this data ensures that the list of types in the DBpedia ontology are included in the vocabulary of the model, and increases the likelihood that topics are discussed in context with their super-types.

Note that *wiki2vec* is different from a KGE, which is typically trained on relationship triples between entities in a knowledge graph (such as DBpedia)(Ristoski and Paulheim, 2016).

3.2. Generating Type Recommendations

The method for summarizing a tabular dataset can be broken down into three distinct steps:

1. Collect a set of *types* and an *ontology* to use for abstraction
2. Extract any text data from the tabular dataset and embed it into a vector space to calculate the distance to all the types in our ontology
3. Aggregate the distance vectors for every keyword in the dataset into a single vector of distances

3.2.1. TYPE ONTOLOGY

In order to generate an abstract term to describe the dataset, we must first collect an ontology of types to select a descriptive term from. We use an ontology provided by DBpedia³ which contains approximately 400 defined types, including everything from *sound* to *video game* and *historic place*. DBpedia also contains defined parent-child relationships for the types⁴ that we use to build a complete hierarchy of types e.g. that *tree* is a sub-type of *plant* which is a sub-type of *eukaryote*.

3.2.2. WORD EMBEDDING

With the set of topics collected, extract each word from the dataset, embed it in a *wiki2vec* vector space and calculate the distance between that word and every type in the ontology. If a single cell in a column contains more than one word, take the average of the corresponding embedded vectors. This results in a collection of distance vectors representing all text in the dataset. Collect the vectors according to their source within the dataset, i.e. words in the same column are collected into a matrix of distances *for each column*. If column headers are provided, treat them as an additional column in the dataset.

2. See <https://github.com/idio/wiki2vec>

3. Downloaded from http://downloads.dbpedia.org/2015-10/dbpedia_2015-10.nt

4. Defined parent-child type relationships can be found at <http://dbpedia.org/ontology/>

3.2.3. DISTANCE AGGREGATION

The previous steps produce a set of matrices containing distances between every text segment in the dataset and the set of types. The goal of this step is to reduce them to a single vector of distances.

We utilize three successive aggregations in order to compute this final vector. The first aggregation is computed across the rows of each column matrix in order to produce a single vector of distances between the column and all types. Potential functions to use are discussed below. The second aggregation is what we call the *tree* aggregation, where we take this vector of distances for a column and utilize the hierarchy of types described by DBpedia in order to update the scores for each type. For instance, we need to update the score for *means of transportation* based on the scores for *airplane*, *train*, and *automobile*. The third aggregation is performed over the set of distance vectors computed for each column, producing a single vector of distances to every defined type. We tested two simple functions for each aggregation step: *mean* and *max*, as well as a variety of more complex aggregations for the tree aggregation step. Tree aggregation allows for additional complexity because the updated distance for each type was dependent on the original distance for that type and the vector of scores for all the children. We found that the most successful tree aggregation functions were those that utilized different functions for processing the child scores and the original type score, e.g.,

$$type' = mean(type, max(child_1, child_2, \dots, child_n)) \quad (1)$$

3.2.4. AGGREGATION FUNCTION SELECTION

To select the best function for each aggregation, we hand-labelled a collection of datasets with types from our ontology to use as a sort of ‘training set’. Then, for each labelled dataset and each combination of aggregation functions, we computed the percentage of true labels found in the top three labels predicted by DUKE, with results shown in Figure 1. This figure clearly shows that using *mean* for column aggregation, *meanmax* tree aggregation described in equation 1, and then *mean* for the final dataset aggregation step produces the best results.

4. Results and Discussion

The goal of this section is to illustrate the effectiveness of the proposed approach to the tabular dataset summarization problem, in the context of some widely available open data sets for which manually curated summary (i.e., types/tags) are available to facilitate comparison and evaluation. Links to every dataset used is provided to facilitate verification by the reader. For each dataset, we generated one subject tag using the DUKE program, as described in the previous sections, and grade it manually using ‘low’ for low accuracy, ‘medium’ for medium accuracy (where the automatically generated tag is “related to”, but is not exactly one of the manual tags) and ‘high’ for high accuracy (where the automatically generated tag is *exact* in the sense that it is one of manually generated tags). Also, please note that each prediction took roughly 20 seconds to perform (approximately 17 seconds of which was spent loading the wiki2vec model) on a 16 CPU 64 GB D16s v3 Azure Cloud VM executing serially.

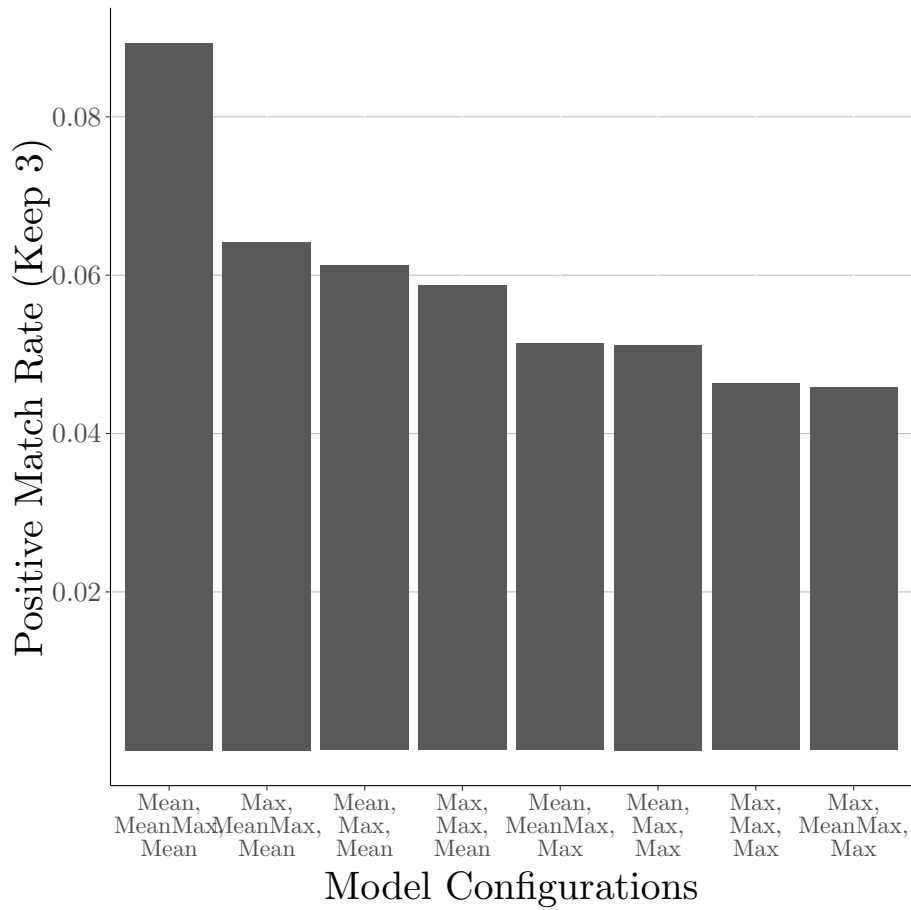


Figure 1: Match rate between true labels and top 3 predicted labels for the best performing aggregation function combinations. The labels for each bar describe the three tested aggregation functions in the order: column, tree, dataset.

Table 1: CKAN tabular dataset summarization results

Dataset	Manual Tags	Predicted Tag	Score
Class Size 2016-2017	class size, public, school, students in classes	school	high
2016 Annual Survey Questions	annual survey, library, public library, public library	library	high
BC Liquor Store Product Price List Oct 2017	BC Liquor Stores, alcohol, beer, price, beverage, wine, spirits	wine region	medium
Coalfile Report	assessment reports, coal, data, maps	river	medium

4.1. Example 1 - CKAN Datasets

Four randomly selected CKAN datasets were used: Class Size 2016-2017⁵, 2016 Annual Survey Questions⁶, BC Liquor Store Product Price List Oct 2017⁷, and Coalfile Report⁸. Manually curated subject tags were available for each dataset (see Table 1). The match between the predicted tags and the manual tags for each dataset is depicted in Table 1.

For the first two datasets, DUKE predicts an exact tag. For the next two datasets, the accuracy is medium, with *wine region* being very close to *wine* and *river* being a common semantic theme in coal field names (examples include *Elk River*, *Hat Creek* and *Peace River*). Specifically, the top 5 tags returned by DUKE in decreasing order for the fourth example were *river*, *stream*, *body of water*, *natural place* and *natural region*, words that are semantically descriptive of the kind of names typically possessed by coal fields. Moreover, we plot the top 5 DUKE-predicted tags and the manual tags for the third example in Figure 2, demonstrating an exact match.

4.2. Example 2 - OpenML Datasets

Four simple OpenML datasets were obtained through the D3M DARPA program: the 185 baseball⁹, 196 autoMpg¹⁰, 30 personae¹¹, and 313 spectrometer¹² datasets. The results for these datasets are shown in Table 2.

5. Available at <https://catalogue.data.gov.bc.ca/dataset/bc-schools-class-size>
 6. See <https://catalogue.data.gov.bc.ca/dataset/bc-public-libraries-statistics-2002-2016>
 7. Available at <https://catalogue.data.gov.bc.ca/dataset/bc-liquor-store-product-price-list-historical-prices>
 8. Available at <https://catalogue.data.gov.bc.ca/dataset/coalfile-database>
 9. Available for download at <http://www.openml.org/d/185>
 10. Available for download at <http://www.openml.org/d/196>
 11. Available for download at <http://www.clips.ua.ac.be/datasets/personae-corpus>
 12. Available for download at <http://www.openml.org/d/313>

Table 2: OpenML tabular dataset summarization results

Dataset	Manual Tags	Predicted Tag	Score
185 baseball	baseball player, play statistics, database	baseball player	high
196 autoMpg	city-cycle, miles per gallon, fuel consumption	engine	high
30 personae	personality, prediction, from text	person	medium
313 spectrometer	measurement, sky, red band, blue band, spectrum, database, flux	band	medium

For the first two datasets, DUKE predicts an exact tag. Note that for the second dataset, we consider *engine* to be an exact tag, since the manual tags are essentially attributes of engines. For the next two datasets, the accuracy is medium, with *person* being very close to *personality* and *band* being descriptive of *red band* and *blue band* manual tags. To verify that bands here referred to the right context, we looked at the top 5 tags returned by DUKE, which in decreasing order are *band*, *brown dwarf*, *inhabitants per square kilometer*, *star* and *celestial body*, words that are fairly consistent with the context suggested by the manual tags. Moreover, we plot these 5 tags and the manual tags for this dataset in Figure 2, demonstrating that while an exact match is not attained, nontrivial subsets of both tag types are ‘very close’ to each other in the wiki2vec embedding space.

4.3. Example 3 - data.world Datasets

The names of some randomly-selected data.world datasets are as follows: US terrorist origins¹³, Occupational Employment Growth¹⁴, CAFOD activity file for Haiti¹⁵ and Queensland Gambling Data¹⁶. The results for this representative set of four data.world datasets are shown in Table 3.

For the first two datasets, DUKE achieves medium accuracy. To see the justification for this, note that the top 5 tags returned by DUKE for the first dataset, in decreasing order, are *person*, *still image*, *legal case*, *supreme court of the USA* and *military person*,

13. Available for download at <https://data.world/tommyblanchard/u-s-terrorist-origins>

14. Available for download at <https://data.world/tommyblanchard/u-s-terrorist-origins>

15. Available for download at <https://data.world/cafod/cafod-activity-file-for-haiti>

16. Downloaded at <https://data.world/queenslandgov/all-gambling-data-queensland>

Table 3: Tabular dataset summarization data.world results

Dataset	Manual Tags	Predicted Tag	Score
US terrorist origins	terrorism, usa politics	person	medium
Occupational Employment Growth	employment, economics	site of scientific interest	medium
CAFOD activity file for Haiti	funding, Haiti, grants, donors, aid transparency	human development index	high
Queensland Gambling Data	expenditure, gambling, Queensland	casino	high

words fairly descriptive of the dataset, which is a list of terrorists, availability of a headshot and details of their legal charges. Moreover, we plot these 5 tags and the manual tags for this dataset in Figure 2, demonstrating that while an exact match is not attained, nontrivial subsets of both tag types are ‘very close’ to each other in the wiki2vec embedding space. The second dataset provides a list of occupations, many of which are scientific, and corresponding wages at various locations, which leads us to believe that *site of scientific interest* is fairly descriptive of the semantics represented in the dataset. For the next two datasets, the accuracy is high, which should be self-explanatory to the reader from Table 3.

5. Conclusion and Future Work

A method for abstractive summarization of tabular datasets, under the assumption that it contains some descriptive text, was presented. Results of numerical experiments on OpenML, CKAN and data.world datasets show good agreement between manual and automatically generated tags by our system, DUKE. These results can be significantly improved by more extensive ontologies included in the model (in place of the 2015 DBpedia ontology). Additionally, retraining wiki2vec on a more complete version of DBpedia (potentially augmented using an Automatic Knowledge Base Completion, or AKBC, Algorithm (Wang et al., 2015) (Groth et al., 2016)) will help improve the accuracy of our system. More sophisticated handling of multi-word phrases also needs to be explored.

Acknowledgments

Work was supported by the Defense Advanced Research Projects Agency (DARPA) under Contract Number D3M (FA8750-17-C-0094). Views, opinions, and findings contained in




Figure 2: Concept embedding space for three of the examined datasets. Point shape depicts DUKE predictions and manual tags. t-SNE dimension reduction was used to project the 1000 dimension concept embeddings into a 2D space for presentation.

this report are those of the authors and should not be construed as an official Department of Defense position, policy, or decision.

References

- Paul T. Groth, Sujit Pal, Darin McBeath, Brad Allen, and Ron Daniel. Applying universal schemas for domain specific ontology expansion. In *Proceedings of the 5th Workshop on Automated Knowledge Base Construction, AKBC@NAACL-HLT 2016, San Diego, CA, USA, June 17, 2016*, pages 81–85, 2016. URL <http://aclweb.org/anthology/W/W16/W16-1315.pdf>.
- Mayank Kejriwal and Pedro Szekely. Supervised typing of big graphs using semantic embeddings. In *Proceedings of The International Workshop on Semantic Big Data, SBD '17*, pages 3:1–3:6, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4987-1. doi: 10.1145/3066911.3066918. URL <http://doi.acm.org/10.1145/3066911.3066918>.
- Quoc V. Le and Tomas Mikolov. Distributed representations of sentences and documents. *CoRR*, abs/1405.4053, 2014. URL <http://arxiv.org/abs/1405.4053>.
- Y. Ma, T. Tran, and V. Bicer. Typifier: Inferring the type semantics of structured data. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*, pages 206–217, April 2013. doi: 10.1109/ICDE.2013.6544826.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119. Curran Associates, Inc., 2013. URL <http://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf>.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. Unsupervised learning of sentence embeddings using compositional n-gram features. *CoRR*, abs/1703.02507, 2017. URL <http://arxiv.org/abs/1703.02507>.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–1543, 2014.
- Petar Ristoski and Heiko Paulheim. Rdf2vec: Rdf graph embeddings for data mining. In Paul Groth, Elena Simperl, Alasdair Gray, Marta Sabou, Markus Krötzsch, Freddy Lecue, Fabian Flöck, and Yolanda Gil, editors, *The Semantic Web – ISWC 2016*, pages 498–514, Cham, 2016. Springer International Publishing. ISBN 978-3-319-46523-4.
- Magnus Sahlgren. The distributional hypothesis. *Italian Journal of Linguistics*, 20(1): 33–54, 2008.
- Marieke van Erp and Piek Vossen. Entity typing using distributional semantics and dbpedia. In Marieke van Erp, Sebastian Hellmann, John P. McCrae, Christian Chiarcos, Key-Sun Choi, Jorge Gracia, Yoshihiko Hayashi, Seiji Koide, Pablo Mendes, Heiko Paulheim, and Hideaki Takeda, editors, *Knowledge Graphs and Language Technology*, pages 102–118, Cham, 2017. Springer International Publishing. ISBN 978-3-319-68723-0.

Quan Wang, Bin Wang, and Li Guo. Knowledge base completion using embeddings and rules. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 1859–1866, 2015. URL <http://ijcai.org/Abstract/15/264>.