

Scalable Approach for Normalizing E-commerce Text Attributes (SANTA)

Ravi Shankar Mishra
India Machine Learning
Amazon
rmsam@amazon.com

Kartik Mehta
India Machine Learning
Amazon
kartim@amazon.com

Nikhil Rasiwasia
India Machine Learning
Amazon
rasiwasi@amazon.com

Abstract

In this paper, we present *SANTA*, a scalable framework to automatically normalize E-commerce attribute values (e.g. “*Win 10 Pro*”) to a fixed set of pre-defined canonical values (e.g. “*Windows 10*”). Earlier works on attribute normalization focused on fuzzy string matching (also referred as syntactic matching in this paper). In this work, we first perform an extensive study of nine syntactic matching algorithms and establish that ‘cosine’ similarity leads to best results, showing 2.7% improvement over commonly used Jaccard index. Next, we argue that string similarity alone is not sufficient for attribute normalization as many surface forms require going beyond syntactic matching (e.g. “*720p*” and “*HD*” are synonyms). While semantic techniques like unsupervised embeddings (e.g. word2vec/fastText) have shown good results in word similarity tasks, we observed that they perform poorly to distinguish between close canonical forms, as these close forms often occur in similar contexts. We propose to learn token embeddings using a twin network with triplet loss. We propose an embedding learning task leveraging raw attribute values and product titles to learn these embeddings in a self-supervised fashion. We show that providing supervision using our proposed task improves over both syntactic and unsupervised embeddings based techniques for attribute normalization. Experiments on a real-world attribute normalization dataset of 50 attributes show that the embeddings trained using our proposed approach obtain 2.3% improvement over best string matching and 19.3% improvement over best unsupervised embeddings.

1 Introduction

E-commerce websites like Amazon are marketplaces where multiple sellers can list and sell their products. At the time of product listing, these sell-

ers often provide product title and structured product information (e.g. *color*), henceforth, termed as product attributes¹. During the listing process, some attribute values have to be chosen from dropdown list (having fixed set of values to choose from) and some attributes are free-form text (allowing any value to be filled). Multiple sellers may express these free-form attribute values in different forms, e.g. “*HD*”, “*1280 X 720*” and “*720p*” represents same TV resolution. Normalizing (or mapping) these raw attribute values (henceforth termed as surface forms) to same canonical form will help improve customer experience and is crucial for multiple underlying applications like search filters, product comparison and detecting duplicates. E-commerce websites provide functionality to refine search results (refer figure 1), where customers can filter based on attribute canonical values. Choosing one of the canonical values restricts results to only those products which have the corresponding attribute value. A good normalization solution will ensure that products having synonym surface form (e.g. ‘*720p*’ vs ‘*HD*’) are not filtered out on applying such filters.

Normalization can be considered as a two step process consisting of - a) identifying list of canonical forms for an attribute, and, b) mapping surface forms to one of these canonical forms. Identification task is relatively easier as most attributes have only few canonical forms (usually less than 10), whereas attributes can have thousands of surface forms. Hence, we focus on the mapping task in this paper, leaving identification of canonical forms as a future task to be explored.

Building an attribute normalization system for thousands² of product attributes poses multiple

¹We use the terms ‘product attributes’ and ‘attributes’ interchangeably in this paper.

²E.g. Xu et al. (2019) have 77K attributes only from ‘Sports & Entertainment category’

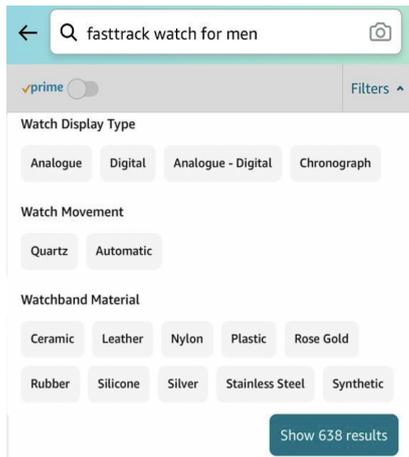


Figure 1: Search filters widget on Amazon.

challenges such as:

- Presence of spelling mistakes (e.g. “grey” vs “gray”, “crom os” vs “chrome os”)
- Requirement of semantic matches (e.g. “linux” vs “ubuntu”, “mac os” vs “ios”)
- Existence of abbreviations (“polyurethane” vs “PU”, “SSD” vs “solid state drive”)
- Presence of multi-token surface forms and canonical forms (e.g. “windows 7 home”)
- Presence of close canonical forms (e.g. “windows 8.1” and “windows 8” can be two separate canonical forms)

Addressing these challenges in automated manner is the primary focus of this work. One can use lexical similarity of raw attribute value (surface form) to a list of canonical values and learn a normalization dictionary (Putthividhya and Hu, 2011). For example, lexical similarity can be used to normalize “windows 7 home” to “windows 7” or “light blue” to “blue”. However, lexical similarity-based approaches won’t be able to handle cases where understanding the meaning of attribute value is important (e.g. matching “ubuntu” to “linux” or “maroon” to “red”). Another alternative is to learn distributed representation (embeddings) of surface forms and canonical forms and use similarity in embedding space for normalization. One can use unsupervised word embeddings (Kenter and De Rijke, 2015) (e.g. word2vec and fastText) for this. However, these approaches are designed to keep embeddings close by for tokens/entities which appear in similar contexts. As we shall see, these unsupervised embeddings do a poor job at distinguishing close canonical attribute forms.

In this paper, we describe SANTA, a scalable framework for normalizing E-commerce text attributes. Our proposed framework uses twin network (Bromley et al., 1994) with triplet loss to learn embeddings of attribute values (canonical and surface forms). We propose a self supervision task for learning these embeddings in automated manner, without requiring any manually created training data. To the best of our knowledge, our work is first successful attempt at creating an automated framework for E-commerce attribute normalization that can be easily extended to thousands of attributes.

Our paper has following contributions : (1) we do a systematic study of nine lexical matching approaches for attribute normalization, (2) we propose a self supervision task for learning embeddings of attribute surface forms and canonical forms in automated manner and describe a fully automated framework for attribute normalization using twin network and triplet loss, and, (3) we curate an attribute normalization test set of 2500 surface forms across 50 attributes and present an extensive evaluation of various approaches on this dataset. We also show an independent analysis on syntactic and semantic portions of this dataset and provide insights into benefits of our approach over string similarity and other unsupervised embeddings. Rest of the paper is organized as follows. We do a literature survey of related fields in Section 2. We describe string matching and embeddings based approaches, including our proposed SANTA framework, in Section 3. We describe our experimental setup in Section 4 and results in Section 5. Lastly, we summarize our work in Section 6.

2 Related Work

2.1 E-commerce Attribute normalization

The problem of normalizing E-commerce attribute values have received limited attention in literature. Researchers have mainly focused on normalizing *brand* attribute, exploring combination of manual mapping curation or lexical similarity-based approaches (More, 2016; Putthividhya and Hu, 2011). More (2016) explored use of manually created key-value pairs for normalizing brand values extracted from product titles. Putthividhya and Hu (2011) explored two fuzzy matching algorithms of Jaccard similarity and Jaro-Winkler distance and found n-gram based Jaccard similarity to be performing better for brand normalization. We use this Jaccard similarity approach as a baseline for comparison.

2.2 Fuzzy String Matching

Fuzzy string matching has been explored for multiple applications, including address matching, names matching (Cohen et al., 2003; Christen, 2006; Recchia and Louwerse, 2013), biomedical abbreviation matching (Yamaguchi et al., 2012) and query spelling correction. Although extensive work has been done for fuzzy string matching, there is no consensus on which technique works best. Christen (2006) explored multiple similarity measures for personal name matching, and reported that best algorithm depends upon the characteristics of the dataset. Cohen et al. (2003) experimented with edit-distance, token-based distance and hybrid methods for matching entity names and reported best performance for a hybrid approach combining TF-IDF weighting with Jaro-Winkler distance. Recchia and Louwerse (2013) did a systematic study of 21 string matching methods for the task of place name matching. While they got relatively better performance with n-gram approaches over commonly used Levenshtein distance, they concluded that best similarity approach is task-dependent. Gali et al. (2016) argued that performance of the similarity measures is affected by characteristics such as text length, spelling accuracy, presence of abbreviations and underlying language. Motivated by these learnings, we do a systematic study of fuzzy matching techniques for the problem of E-commerce attribute normalization. Besides, we use latest work in the field of neural embeddings for attribute normalization.

3 Overview

Attribute normalization can be posed as a matching problem. Given an attribute surface form and a list of possible canonical forms, similarity of surface form with each canonical form is calculated and surface form is mapped to the canonical form with highest similarity or mapped to ‘other’ class if none of the canonical forms is suitable (refer Figure 2 for illustration). Formally, given a surface form s_i ($i \in [1, n]$) and a list of canonical forms c_j ($j \in [0, k]$), where c_0 is the ‘other’ class, n is number of surface forms and k is number of canonical forms. The aim is to find a mapping function M such that:

$$M(s_i) = c_j \text{ where } i \in [1, n], j \in [0, k] \quad (1)$$

In this paper, we explore fuzzy string matching and similarity in embedding space as matching

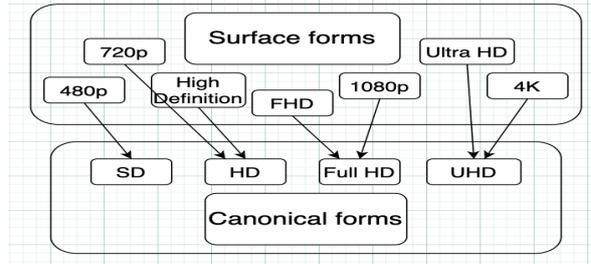


Figure 2: Illustration of Attribute Normalization Task

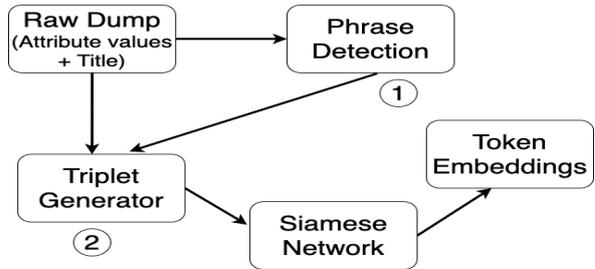


Figure 3: SANTA framework for training embeddings suitable for attribute normalization

techniques. We describe multiple string matching approaches in Section 3.1, followed by unsupervised token embedding approaches in Section 3.2 and our proposed SANTA framework in Section 3.3.

3.1 String Similarity Approach

We study three different categories of string matching algorithms³ and explore three algorithms in each category⁴:

- **Edit distance-based:** These algorithms compute the number of operations needed to transform one string to another, leading to higher similarity score for less operations. We experimented with six algorithms in this category, a) Hamming, b) Levenshtein, and c) Jaro-Winkler.
- **Sequence-based:** These algorithms find common sub-sequence in two strings, leading to higher similarity score for longer common sub-sequence or a greater number of common sub-sequences. We experimented with three algorithms in this category, a) longest common subsequence similarity, b) longest common substring similarity, and c) Ratcliff-Obershelp similarity.

³<https://itnext.io/string-similarity-the-basic-know-your-algorithms-guide-3de3d7346227>

⁴For algorithms which return distance metrics rather than similarity, we use lowest distance as substitute for highest similarity.

- **Token-based:** These algorithms represent string as set of tokens (e.g. ngrams) and compute number of common tokens between them, leading to higher similarity score for higher number of common tokens. We experimented with three algorithms in this category - a) Jaccard index, b) SorensenDice coefficient, and c) Cosine similarity. We converted strings to character ngrams of size 1 to 5 before applying this similarity.

We used python module `textdistance`⁵ for all string similarity experiments. For detailed definition of these approaches, we refer readers to [Gomaa and Fahmy \(2013\)](#) and [Vijaymeena and Kavitha \(2016\)](#).

3.2 Unsupervised Embeddings

[Mikolov et al. \(2013\)](#) introduced word2vec model that uses a shallow neural network to obtain distributed representation (embeddings) of words, ensuring words that appear in similar contexts are closer in the embedding space. To deal with unseen and rare words, [Bojanowski et al. \(2017\)](#) proposed fastText model that improves over word2vec embeddings by considering sub-words and representing word embeddings as average of embeddings of corresponding sub-words. To learn domain-specific nuances, we trained a word2vec and fastText model using a dump consisting of product titles and attribute values (refer Section 4 for details of this dump). We found better results with using concatenation of title with attribute value as compared to using only title, likely due to including surface form from title and attribute canonical form (or vice versa) in a single context.

3.3 Scalable Approach for Normalizing Text Attributes (SANTA)

Figure 3 gives an overview of learning embeddings with our proposed SANTA framework. We define an embedding learning task using twin network with triplet loss to enforce that embeddings of attribute values are closer to corresponding titles as compared to embeddings of a randomly chosen title from the same product category. To deal with multi-word values, we use a simple step of treating each multi-word attribute value as a single phrase. Overall, we observed 40K such phrases, e.g. “back cover”, “android v4.4.2”, “9-12 month” and “wine red”. For both attribute values and product titles, we converted these multi-token phrases to single tokens (e.g. ‘back cover’ is replaced with ‘back_cover’).

⁵<https://pypi.org/project/textdistance/>

We describe details of the embedding learning task and triplet generation in Section 3.3.1, and twin network in Section 3.3.2.

3.3.1 Triplet Generation

There are scenarios when title contains canonical form of attribute value (e.g. “3xl” could be size attribute value for a title ‘Nike running shoes for men xxxl’). We can leverage this information to learn embeddings that not only capture semantic similarity but can also distinguish between close canonical forms. Motivated by work in answer selection ([Kulkarni et al., 2019](#); [Bromley et al., 1994](#)), we define an embedding learning task of keeping surface form closer to corresponding title as compared to a randomly chosen title. We created training data in form of triplets of anchor (q), positive title (a_+) and negative title (a_-), where q is attribute value, a_+ is corresponding product title and a_- is a title selected randomly from product category of a_+ . One way to select negatives is to pick a random product from any product category, but that may provide limited signal for embedding learning task (e.g. choosing an Apparel category product when actual product is from Laptop category). Instead, we select a negative product from same product category, which acts as a hard negative ([Kumar et al., 2019](#); [Schroff et al., 2015](#)) and improves the attribute normalization results. Selecting products from same category may lead to few incorrect negative titles (i.e. negative title may contain the correct attribute value). We screen out incorrect negatives where anchor attribute value (q) is mentioned in title, reducing noise in the training data.

3.3.2 Twin Network and Triplet Loss

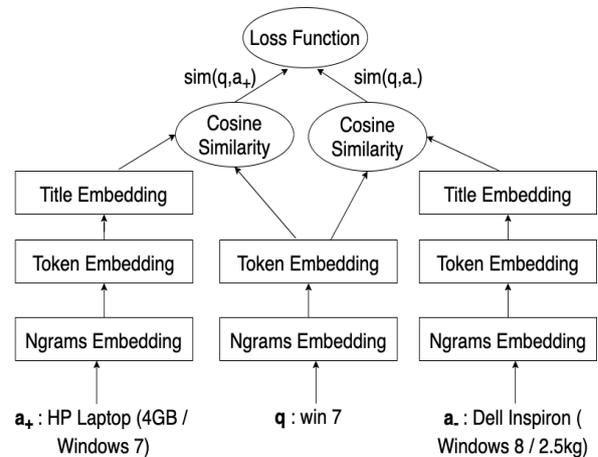


Figure 4: Illustration of Twin network with Triplet loss.

We choose twin network as it projects surface forms and canonical forms in same embedding

space and triplet loss helps to keep surface forms closer to the most appropriate canonical form. Figure 4 describes the architecture of our SANTA framework. Given a (q, a_+, a_-) triplet, the model learns embedding that minimize the triplet loss function (Equation 2). Similar to fastText, we represent each token as consisting of sub-words (n-gram tokens). Embedding for a token is created using a composite function on sub-word embeddings, and similarly, embeddings for title are created using composite function on word embeddings. We use averaging of embeddings as composite function (similar to fastText), though the framework is generic and other composite functions like LSTM, CNN and transformers can also be used.

Let E denote the embedding operator and \cos represent cosine similarity metric, then triplet loss function is given as:

$$Loss = \max \{0, M - \cos(E(q), E(a_+)) + \cos(E(q), E(a_-))\} \quad (2)$$

where M is margin.

The advantage of this formulation over unsupervised embeddings (Section 3.2) is that in addition to learning semantic similarities for attribute values, it also learns to distinguish between close canonical forms, which may appear in similar contexts. For example, the embedding of surface form ‘720p’ will move closer to embedding of ‘HD’ mentioned in a_+ title but away from embedding of ‘Ultra HD’ mentioned in a_- title.

4 Experimental Setup

In this section, we describe our experimental setup, including dataset, metrics and hyperparameters of our model. There is no publicly available data set for attribute normalization problem. More (2016) and Putthividhya and Hu (2011) worked on brand normalization problem but the datasets are not published for reuse. Xu et al. (2019) published a dataset collected from AliExpress ‘Sports & Entertainment’ category for attribute extraction use-case. This dataset belongs to a single category and is restricted to samples where attribute value is present in title, hence limiting its applicability for attribute normalization. To ensure robust learnings, we curate a real-world attribute normalization dataset spread across multiple categories and report all our evaluations on this dataset.

4.1 Training and Test data

We selected 50 attributes across 20 product categories including electronics, apparel and furniture for our study and obtained their canonical forms from business teams. These selected attributes have on average 7.5 canonical values (describing the exact selection process for canonical values is outside the scope of current work). For each of these attributes, we picked top 50 surface forms and manually mapped these values to corresponding canonical forms, using ‘other’ label when none of the existing canonical forms is suitable. We, thus, obtain a labelled dataset of 2500 samples (50 surface forms each for 50 attributes), out of which 38% surface forms are mapped to ‘other’ class. Surface forms mapping to ‘other’ are either junk value (e.g. “5MP” for operating system) or coarser value (e.g. “android” when canonical forms are “android 4.1”, “android 4.2” etc.). It took 20 hours of manual effort for creating this dataset. We split this data into two parts (20% used as dev set and 80% as test set).

For training, we obtain a dump of 100K products corresponding to each attribute, obtaining a dump of 5M records (50 attributes X 100K products per attribute), having title and attribute values. This data (5M records) is used for training unsupervised embeddings (Section 3.2). For each record, we select one negative example for triplet generation (Section 3.3.1) and use this triplet data (5M records) for learning SANTA model. Kindly note that training data creation is fully automated, and does not require any manual effort, making our approach easily scalable.

4.2 Metric

There are no well-established metrics in literature for attribute normalization problem. One simple approach is to consider canonical form with highest similarity as predicted value for evaluation. However, we argue that an algorithm should be penalized for mapping a junk value to any canonical form. Based on this motivation, we define two evaluation metrics that we use in this work.

4.2.1 Accuracy

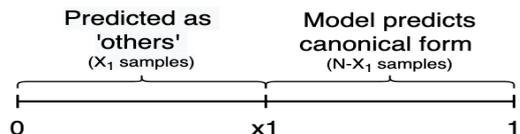


Figure 5: Illustration for Accuracy metric.

We divide predictions on all samples (N) into two sets using a threshold x_1 (see Figure 5). ‘Other’ class is predicted for samples having score less than x_1 (low similarity to any canonical form) and canonical form with highest similarity is considered for samples having score greater than x_1 (confident prediction). We consider prediction as correct for samples in X_1 set if true label is ‘other’ and for samples in $N - X_1$ set, if model prediction matches the true label. We define Accuracy as ratio of correct predictions to the number of cases where prediction is made (N in this case). The threshold x_1 is selected based on performance on dev set.

4.2.2 Accuracy Coverage Curve

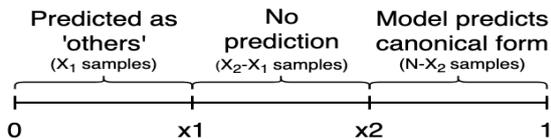


Figure 6: Illustration for Accuracy Coverage metric.

It can be argued that a model is confident about surface forms when prediction score is on either extreme (close to 1 or close to 0). Motivated by this intuition, we define another metric where we divide predictions into three sets using two thresholds x_1 and x_2 (see Figure 6). ‘Other’ class is predicted for samples having score less than x_1 (low similarity to any canonical form), no prediction is made for samples having score between x_1 and x_2 (model is not confidently predicting any canonical form but confidence score is not too low to predict ‘other’ class) and canonical form with highest similarity is considered for samples having score greater than x_2 . We define Coverage as fraction of samples where some prediction is made $((X_1 + N - X_2)/N)$, and Accuracy as ratio of correct predictions to the number of predictions. For samples in X_1 set, we consider prediction correct if true label is ‘other’ and for samples in $N - X_2$ set, we consider prediction correct when model prediction matches the true canonical form. The thresholds are selected based on performance on dev set and based on different choice of thresholds, we create Accuracy-Coverage curve for comparison.

4.3 SANTA Hyperparameters

We set the value of M as 0.4, embedding dimension as 200, minimum n-gram size as 2 and maximum n-gram size as 4. We run the training using Adadelta optimizer for 5 epochs, which took approximately 8 hours on a NVIDIA V100 GPU. The parameters to be learned are ngram embeddings

Table 1: Evaluation of String similarity approaches.

STRING SIMILARITY	ACCURACY
EDIT DISTANCE BASED	
HAMMING	51.6
LEVENSHTEIN	61.1
JARO-WINKLER	62.1
SEQUENCE BASED	
LC SUBSEQUENCE	57.6
LC SUBSTRING	64.7
RATCLIFF-OBERSHELP	64.9
TOKEN BASED	
JACCARD INDEX	74.6
SORENSEN-DICE	74.6
COSINE SIMILARITY	76.6

(0.63M ngrams X 200 embedding dimension = 127M parameters). Ngram embeddings are shared across the twin network.

5 Results

We present systematic study on string similarity approaches in Section 5.1, followed by experiments of unsupervised embeddings in Section 5.2. We compare best results from Section 5.1 and Section 5.2 with our proposed SANTA framework in Section 5.3. We study these algorithms separately on syntactic and semantic portion of test dataset in Section 5.4 and perform qualitative analysis based on t-SNE visualization in Section 5.5.

5.1 Evaluation of String Similarity

Table 1 shows comparison of string similarity approaches for attribute normalization. We observe that token based methods performs best, followed by comparable performance of sequence based and edit distance based methods. We believe that token based approaches outperformed other approaches as they are insensitive to the position where common sub-string occurs in the two strings (e.g. matching “*half sleeve*” to “*sleeve half*” for *sleeve type* attribute). Putthividhya and Hu (2011) evaluated n-gram based ‘Jaccard index’ (token based approach) and ‘Jaro-Winkler distance’ (character based approach) for brand normalization and got similar observations, obtaining best results with ‘Jaccard index’. We observe that ‘Cosine similarity’ obtains 2.7% accuracy improvement over Jaccard index in our experiments.

5.2 Evaluation of Unsupervised Embeddings

Table 2 shows performance of word2vec and fastText approach. We observe that presence of n-grams information in fastText leads to significant improvement over word2vec, as use of n-grams helps with matching of rare attribute values. However, fastText is not able to match string similarity

Table 2: Comparison of normalization approaches

MODEL	ACCURACY
RANDOM	37.8
MAJORITY CLASS PREDICTION	48.5
JACCARD INDEX	74.6
COSINE SIMILARITY	76.6
WORD2VEC	48.4
FASTTEXT	65.7
SANTA (WITHOUT NGRAMS)	47.4
SANTA (WITH NGRAMS)	78.4

baseline (refer Table 1). We believe unsupervised embeddings shows relatively inferior performance for attribute normalization task, as embeddings are learnt based on contexts in product titles, keeping different canonical forms (e.g. “*HD*” and “*Ultra HD*”) close by as they occur in similar context.

5.3 Evaluation of SANTA framework

Table 2 shows comparison of SANTA with multiple normalization approaches, including best solutions from Section 5.1 and Section 5.2. To understand the difficulty of this task, we introduce two baselines of a) randomly mapping surface form to one of the canonical forms (termed as ‘RANDOM’), and b) predicting the most common class based on dev data (termed as ‘MAJORITY CLASS’). We observe 37.8% accuracy with ‘RANDOM’ and 48.5% accuracy with ‘MAJORITY CLASS’, establishing the difficulty of the task. SANTA (with ngrams) shows best performance with 78.4% accuracy, leading to 2.3% accuracy improvement over ‘Cosine Similarity’ (best string similarity approach) and 19.3% over fastText (best unsupervised embeddings). We discuss few qualitative examples for these approaches in appendix.

Figure 7 shows Accuracy-Coverage curve for these algorithms. As observed from this curve, SANTA consistently outperforms string similarity and fastText across all coverages.

5.4 Study on Syntactic and Semantic Dataset

In this section, we do a separate comparison of normalization algorithms on samples requiring semantic and syntactic matching. We filtered test dataset where true label is not ‘Others’, and manually labelled each surface form as requiring syntactic or semantic similarity. Based on this analysis, we observe that 45% of test data requires syntactic matching, 17% requires semantic matching and remaining 38% is mapped to ‘other’ class. For current analysis of syntactic and semantic set, we

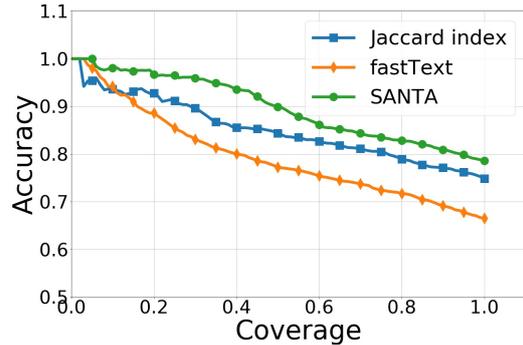


Figure 7: Accuracy-Coverage plot for various Normalization techniques.

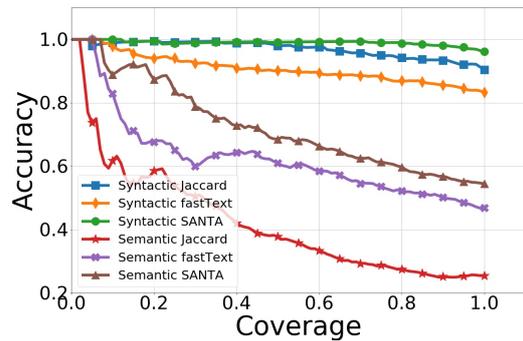


Figure 8: Study of various normalization algorithms on semantic and syntactic dataset

use a special case of metric defined in section 4 (since ‘other’ class is not present). We set $x_1 = 0$, ensuring that ‘other’ class is not predicted for any samples of test data. We show Accuracy-Coverage plot for semantic and syntactic cases in Figure 8.

For semantic set, we observe that fastText performs better than string similarity, due to its ability to learn semantic representation. Our proposed SANTA framework, further improves over fastText for better semantic matching with close canonical forms. For syntactic set, we observe comparable performance of SANTA and string similarity. These results demonstrate that our proposed SANTA framework performs well on both syntactic and semantic set.

5.5 Word Embeddings Visualization

For qualitative comparison of fastText and SANTA embeddings, we project these embeddings into 2-dimensions using t-SNE (van der Maaten and Hinton, 2008). Figure 9 shows t-SNE plots⁶ for 3 attributes (Headphone Color, Jewelry Necklace type and Watch Movement type). For color attribute, we observe that values based on SANTA have ho-

⁶<https://scikit-learn.org/stable/modules/generated/sklearn.manifold.TSNE>

mogenous cohorts of canonical values and corresponding surface forms (e.g. there is a cohort for ‘black’ color on bottom-right and ‘blue’ color on top-left of the plot.). However, with fastText, the color values are scattered across the plot without any specific cohorts. Similar patterns are seen with necklace type where SANTA results show better cohorts than fastText. These results demonstrate that embeddings learnt with SANTA are better suited than fastText embeddings to distinguish between close canonical forms.

6 Conclusion

In this paper, we studied the problem of attribute normalization for E-commerce. We did a systematic study of multiple syntactic matching algorithms and established that use of ‘cosine similarity’ leads to 2.7% improvement over commonly used Jaccard index. Additionally, we argued that attribute normalization requires combination of syntactic and semantic matching. We described our SANTA framework for attribute normalization, including our proposed task to learn embeddings in a self-supervised fashion with twin network and triplet loss. Evaluation on a real-world dataset for 50 attributes, shows that embeddings learnt using our proposed SANTA framework outperforms best string matching algorithm by 2.3% and fastText by 19.3% for attribute normalization task. Our evaluation based on semantic and syntactic examples and t-SNE plots provide useful insights into qualitative behaviour of these embeddings.

References

- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *TACL*, 5:135–146.
- Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1994. Signature verification using a “siamese” time delay neural network. In *NIPS*, pages 737–744.
- Peter Christen. 2006. A comparison of personal name matching: Techniques and practical issues. In *Sixth IEEE ICDM-Workshops (ICDMW’06)*, pages 290–294. IEEE.
- William W Cohen, Pradeep Ravikumar, and Stephen E Fienberg. 2003. A comparison of string distance metrics for name-matching tasks. In *IWeb*, pages 73–78.
- Najlah Gali, Radu Marinescu-Istodor, and Pasi Fränti. 2016. Similarity measures for title matching. In *2016 23rd ICPR*, pages 1548–1553. IEEE.
- Wael H Gomaa and Aly A Fahmy. 2013. A survey of text similarity approaches. *International Journal of Computer Applications*, 68(13):13–18.
- Tom Kenter and Maarten De Rijke. 2015. Short text similarity with word embeddings. In *24th ACM CIKM*, pages 1411–1420.
- Ashish Kulkarni, Kartik Mehta, Shweta Garg, Vidit Bansal, Nikhil Rasiwasia, and Srinivasan Sengamedu. 2019. Productqna: Answering user questions on e-commerce product pages. In *Companion Proceedings of The 2019 WWW Conference*, pages 354–360.
- Sawan Kumar, Shweta Garg, Kartik Mehta, and Nikhil Rasiwasia. 2019. Improving answer selection and answer triggering using hard negatives. In *Proceedings of the 2019 EMNLP-IJCNLP*, pages 5913–5919.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-sne](#). *Journal of Machine Learning Research*, 9(86):2579–2605.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119.
- Ajinkya More. 2016. Attribute extraction from product titles in ecommerce. *arXiv preprint arXiv:1608.04670*.
- Duangmanee Pew Putthividhya and Junling Hu. 2011. Bootstrapped named entity recognition for product attribute extraction. In *EMNLP*, pages 1557–1567. Association for Computational Linguistics.
- Gabriel Recchia and Max M Louwerse. 2013. A comparison of string similarity measures for toponym matching. In *COMP@ SIGSPATIAL*, pages 54–61.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823.
- MK Vijaymeena and K Kavitha. 2016. A survey on similarity measures in text mining. *Machine Learning and Applications: An International Journal*, 3(2):19–28.
- Huimin Xu, Wenting Wang, Xinnian Mao, Xinyu Jiang, and Man Lan. 2019. Scaling up open tagging from tens to thousands: Comprehension empowered attribute value extraction from product title. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5214–5223.
- Atsuko Yamaguchi, Yasunori Yamamoto, Jin-Dong Kim, Toshihisa Takagi, and Akinori Yonezawa. 2012. Discriminative application of string similarity methods to chemical and non-chemical names for biomedical abbreviation clustering. In *BMC genomics*, volume 13, page S8. Springer.

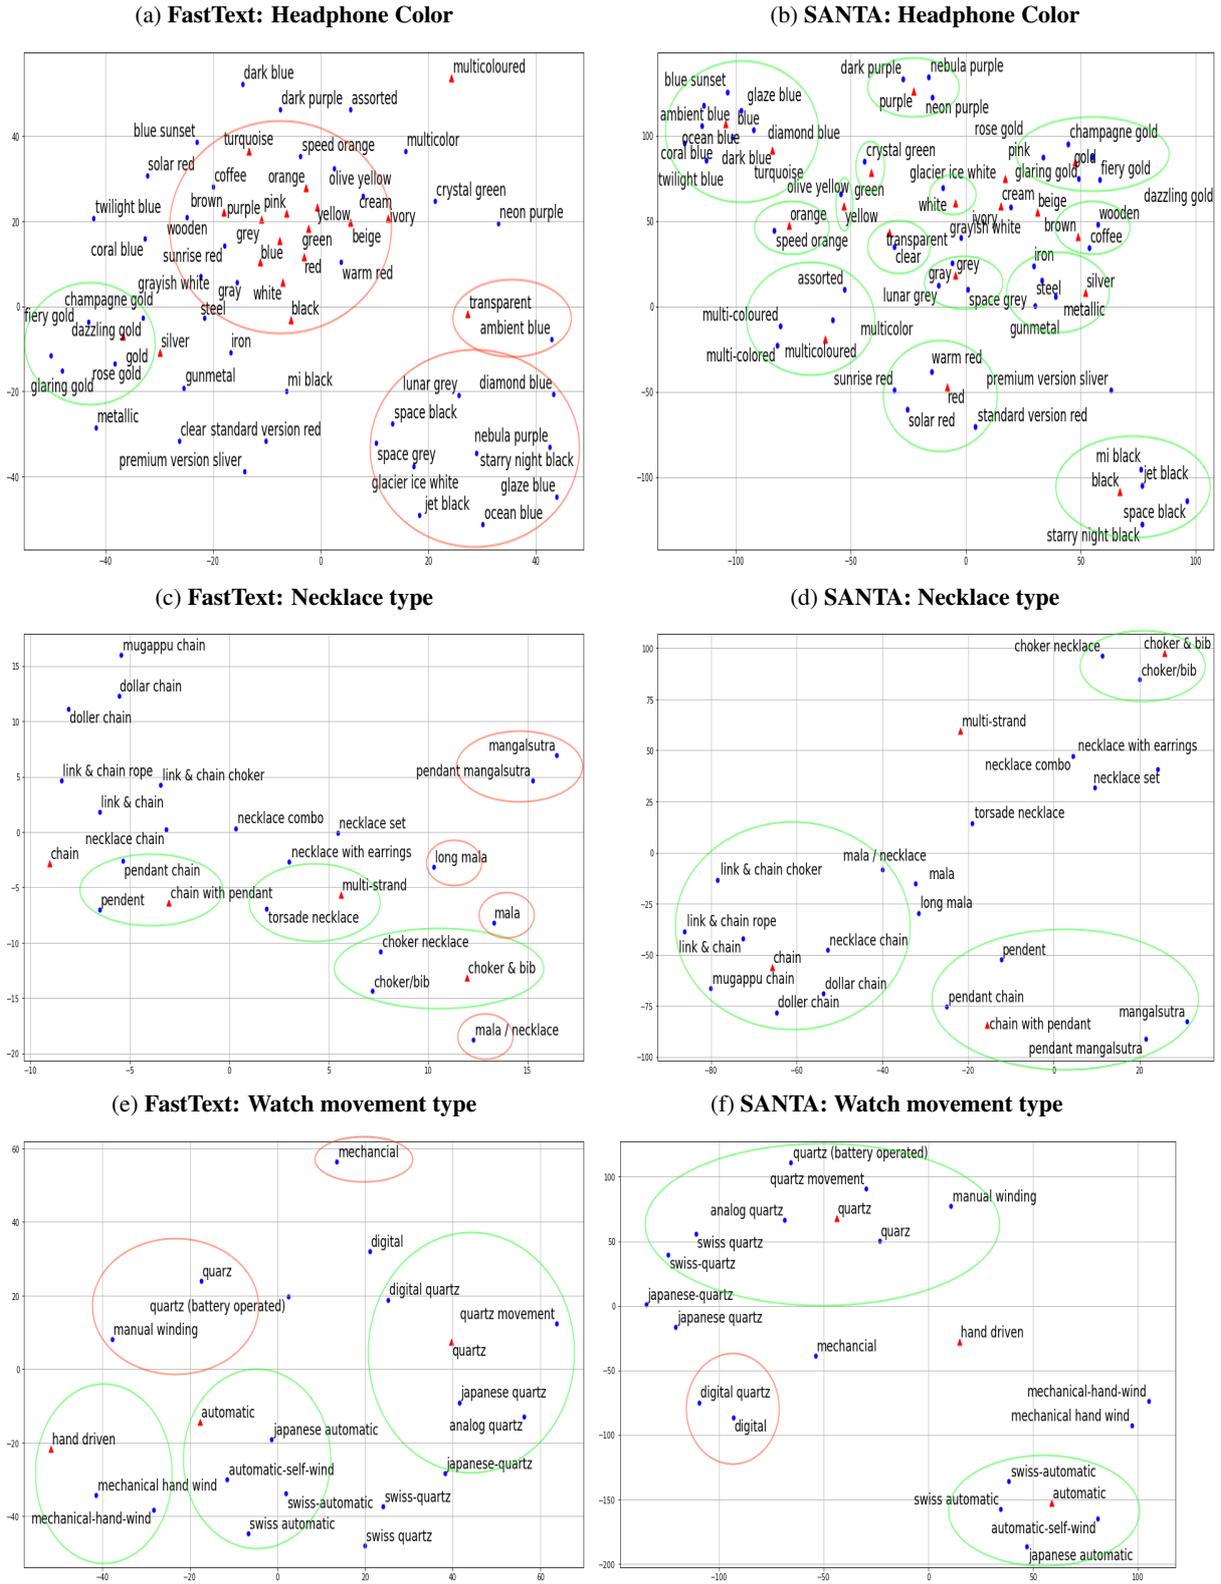


Figure 9: Figure showing t-SNE plot of fastText and SANTA embeddings for three attributes. Surface forms are shown with green dots and canonical forms with red triangles. For better understanding of results, we use green oval selection to show correct homogenous cohorts and red oval selection for incorrect cohorts. This figure is best seen in colors.

7 Appendix

We list few interesting examples in Table 3. It can be observed that string similarity makes correct predictions for cases requiring fuzzy matching (e.g. matching “*multi*” with “*multicoloured*”), but, makes incorrect predictions for examples requiring semantic matching (e.g. incorrectly matching “*cane*” with “*polyurethane*”). With fastText, we get correct predictions for many semantic examples, however, we get incorrect predictions for close canonical forms (e.g. incorrectly mapping “*3 seater sofa set*” to “*five seat*” as “*three seat*” and “*five seat*” occur in similar contexts in title). Our proposed SANTA model does well for most of these examples, but it fails to make correct predictions for rare surface forms (e.g. “*no assembly required, pre-aseembled*”).

Surface Form	Actual Canonical Form	Cosine Similarity Prediction	FastText Prediction	SANTA prediction	Comment
multi	multicoloured	multicoloured	green	multicoloured	FastText fails
thermoplastic	plastic	plastic	silicone	plastic	
amd radeon r3	ati radeon	ati radeon	nvidia geforce	ati radeon	
free size	one size	one size	small	one size	
2 years	2 - 3 years	11 - 12 years	3 - 4 years	2 - 3 years	Both String Similarity and fastText fails but SANTA gives correct mapping
elbow sleeve	half sleeve	3/4 sleeve	short sleeve	half sleeve	
cane	bamboo	polyurethane	rattan	bamboo	
product will be assembled	requires assembly	already assembled	d-i-y	require assembly	
3 seater sofa set	three seat	four seat	five seat	three seat	String Similarity fails
nokia os	symbian	palm web os	symbian	symbian	
silicone	rubber	silk	rubber	rubber	
coffee	brown	off-white	brown	brown	SANTA fails
no assembly required, pre-aseembled	already assembled	requires assembly	already assembled	requires assembly	
mechancial	hand driven	hand driven	hand driven	automatic	

Table 3: Qualitative Examples for multiple normalization approaches. Correct predictions are highlighted in green color and incorrect predictions are highlighted in red color.