



Tracing State-Level Obesity Prevalence from Sentence Embeddings of Tweets: A Feasibility Study

Xiaoyi Zhang¹ , Rodoniki Athanasiadou² , and Narges Razavian²  

¹ Department of Biomedical Informatics and Medical Education, School of Medicine, University of Washington, Seattle, WA 98195, USA
xyzhang7@uw.edu

² Department of Population Health, School of Medicine, New York University, New York, NY 10012, USA
{rodoniki.athanasiadou,narges.razavian}@nyulangone.org

Abstract. Twitter data has been shown broadly applicable for public health surveillance. Previous public health studies based on Twitter data have largely relied on keyword-matching or topic models for clustering relevant tweets. However, both methods suffer from the short-length of texts and unpredictable noise that naturally occurs in user-generated contexts. In response, we introduce a deep learning approach that uses hashtags as a form of supervision and learns tweet embeddings for extracting informative textual features. In this case study, we address the specific task of estimating state-level obesity from dietary-related textual features. Our approach yields an estimation that strongly correlates the textual features to government data and outperforms the keyword-matching baseline. The results also demonstrate the potential of discovering risk factors using the textual features. This method is general-purpose and can be applied to a wide range of Twitter-based public health studies.

Keywords: Public health informatics · Deep learning · Natural language processing · Social media

1 Introduction

Twitter, or social media in general, is a vast space for users to express opinions and sentiments. The proliferation of social media networks accelerates the generation of public health data at an unprecedented rate, allowing big data computing approaches to achieve innovative and impactful research in health sciences. Since Paul and Dredze [10] proposed the use of Twitter data for public health informatics, several studies [2, 8, 9, 13, 17] discovered strong correlations

This work was conducted during the first author's research intern at NYU Center for Data Science.

between government statistics for specific diseases and tweets on specific topics. These studies suggest that a large number of relevant tweets can provide insight into the general health of a population.

In a recent literature review, Jordan et al. [7] showed the majority of Twitter studies in the last decade on public health informatics has relied on keywords for classification or clustering. However, since tweets contain unpredictable noises such as slang, emoji, and misspellings, the tweets retrieved from keyword-matching models often exclude relevant or include irrelevant messages. These uncertain semantics of retrieved tweets lead to unreliable estimations for public health metrics [5, 7]. Meanwhile, other studies use Latent Dirichlet allocation to learn latent topics [11, 12]. These models depend on reliable word co-occurrence statistics and typically suffer from data sparsity when applied to short documents like tweets [14].

We propose using sentence embeddings by supervised deep learning methods to overcome this shortcoming. Hashtags, the user-annotated label that clusters tweets with shared topics regardless of the diverse textual patterns, provide a natural supervision for training distributed representations of tweets. In this work, we adapt *TagSpace* [15], a convolutional neural network (CNN) that learns word and tweet embeddings in the same vector space using hashtags as supervised signals. In order to demonstrate the feasibility of this method, we address the specific task of estimating state-level obesity from tweets characterizing actual dietary habits. We use both embeddings to cluster and to extract relevant textual features that correspond to population-level dietary habits from over two hundred million tweets. The regression on these textual features strongly correlates to state-level obesity prevalence surveyed by Centers for Disease Control and Prevention [1]. Since our method is not specifically tailored to obesity research, our approach is applicable to a wide range of public health studies that involve Twitter data.

2 Data Acquisition and Pre-processing

We retrieve 272.8 million tweet records posted in 2014 using the Twitter API ¹, and we assign one state among the contiguous United States (48 states plus the District of Columbia) to each of the 261 million records based on user geolocation metadata. Non-English posts are removed from our dataset. We use regex to perform the following steps:

1. Convert all alphabetical characters to lower case
2. Remove all URLs, user mentions, and special characters except the hashtag symbol #
3. Remove numerical characters except those in hashtags
4. Add white space between consecutive emojis, and limit repeating mentions of words or emojis in a post

¹ <https://developer.twitter.com/>.

After initial preprocessing, we find 9.5 million unique vocabularies (including hashtags and emojis) heavily tailed at scarce mentions - 6.3 million are mentioned only once, and 8.9 million less than ten times. Similar distribution is also discovered in hashtags. Including scarcely-mentioned words not only results in a memory-demanding lookup table, but also puts our model under the risk of overfitting, as the model may memorize the rare textual patterns found only in the training set. Hence we select 500 k most mentioned words (excluding stop words) and 50k hashtags for our model, and all out-of-vocabulary words are tokenized as <UNKNOWN>. The data pre-processing pipeline can be visualized in Fig. 1.




Fig. 1. Twitter data pre-processing and keyword acquisition pipeline.

3 Methods

We address the specific task of using the twitter data within a state to estimate the obesity prevalence in that state. Inspired by recent Twitter-derived public health studies [2, 8, 9, 17], we first compile a set of keywords related to dietary habits to form a feature space in a regression scenario. We then adapt two

deep learning models to retrieve food-related tweets by the scoring between embeddings of tweets and keywords. Embeddings of food-related tweets within a state will be aggregated for extracting features later used in regression.

3.1 Constructing Feature Space from Keywords

Following prior works by Nguyen et al. [9], we generate the keyword list from two sources: 1) the U.S. Department of Agriculture’s National Nutrient Database [1] - from over 7000 food records found in the USDA database, we extract only the first-level information (e.g. “strawberry yogurt” and “nonfat yogurt” are both recorded as “yogurt”), which gives us 371 terms; and 2) popular food-related mentions in the press ² - we add food (e.g. “sashimi” and “kimchi”), food-related slangs (e.g. “blt”), and chain restaurants (e.g. KFC and Starbucks) that are not included by the USDA database but frequently appear in user-generated contexts, which results in 131 additional terms. All of the 502 keywords in the list are reduced to their singular forms by NLTK lemmatizer [3], and words in the tweets are also lemmatized for keyword matching. We show the keyword acquisition process in Fig. 1.

3.2 Retrieving Relevant Tweets via Deep Learning




Fig. 2. The model architecture of *TagSpace* [15]. Given an input tweet w and its hashtag t , the forward pass outputs the scoring $f(w, t)$, where N denotes vocabulary size, d denotes embedding dimension, l denotes number of words (i.e. max sequence length) of input tweet, K denotes convolution window size, H denotes hidden dimension.

Keyword Matching. The simple baseline regards tweets explicitly mentioning at the one of the keywords as relevant. We find the relative frequency of food-related tweets ranges from 3.0% to 6.2% across all states with the mean value of 4.7%, which is close to statistics reported by Nguyen et al. [9] and Ghosh and Guha [6]. This implies that our pre-processing pipeline, which set a higher bar for word frequency than [9] and [6], does not drastically change the distribution of food-related tweets.

² App Spring Inc. List Challenges: Food, <https://www.listchallenges.com/lists/food>.

TagSpace. Simple keyword matching results in questionable semantic relevancy of retrieved tweets (e.g. “that problem is a hard *nut* to crack”, “taylor swift is the *cream* of the crop”). In contrast, hashtags provide the user’s labeling of a tweet’s themes. We adapt *TagSpace*, a CNN model that learns the distributed representations of both words and tweets using hashtags as a supervised signal [15]. Given a tweet, the model convolves its unigrams’ embeddings as input, and ranks the scoring (e.g. inner product) between the learnt embeddings of the tweet and candidate hashtags. The ranking is optimized by a pairwise hinge objective function as given by Algorithm 1, which is optimized for retrieving top-ranked hashtags according to Weston et al. [15]. We show the model architecture in Fig. 2. In our study, we consider a tweet with top ranked hashtags containing the keywords as food-related. Among the 50k hashtag candidates, 373 are included in our food keyword list. For this training task, we require the input tweets to mention hashtags in the candidate pool for prediction, and this gives us 14.4 million tweets for training, 1.8 million held out for hyperparameter tuning and 1.8 million for testing. To prevent data leakage, hashtags in the input texts are substituted with the corresponding plain words.

Algorithm 1: WARP ranking loss of *TagSpace*

Data: $e_{conv}(w) \in \mathbb{R}^{N \times d}$: sentence embedding of tweets w ; $e(t) \in \mathbb{R}^d$: word embedding of hashtag t ; T : set of all hashtags in corpus; T^+ : set of hashtags found in w ; $m \in \mathbb{R}$: margin; M : max sample iterations

Result: Optimized CNN weights, $e_{conv}(w)$, and embedding of words in w

Sample $t^+ \in \mathbb{R}^n$ from T^+

Compute $f(w, t^+) = e_{conv}(w) \cdot e(t^+)$

Sample t^- from $T \setminus T^+$

Compute $f(w, t^-) = e_{conv}(w) \cdot e(t^-)$

Initialize $i \leftarrow 0$

while $f(t^-, w) \leq m + f(t^+, w)$ **and** $i \leq M$ **do**

Resample t^- from $T \setminus T^+$;

Compute $f(w, t^-) = e_{conv}(w) \cdot e(t^-)$;

$i \leftarrow i + 1$

end

Compute $Loss = \max(0, m - f(w, t^+) + f(w, t^-))$

Backward propagation on $Loss$

Binary TagSpace. While Weston et al. [15] optimizes the prediction of $p(\text{hashtag} \mid \text{tweet})$ over all hashtag candidates, we are only interested in the tweets’ semantic relevancy with food (i.e. $p(\text{hashtags about food} \mid \text{tweet})$). Based on whether a tweet contains hashtags found in our food keyword list, we label all tweets with hashtags as either food-related or not. The word and tweet embeddings learnt from the CNN discussed in the previous method are optimized for

a binary classification objective function instead. We use the same training and testing sets as those of the previous method.

3.3 Feature Engineering and Obesity Estimation by Elastic Net

For a given state, features are calculated from the scoring (e.g. inner product or cosine similarity) between the keyword embeddings and the average sentence-level embedding of food-related tweets within that state. The scoring function is the same for both CNN models. Both CNNs’ objective functions internally train tweet embeddings in the word vector space [15, 16], and hence the scoring provides information about the semantic relevance between the tweets and the keywords. By aggregating food-related tweets (i.e. tweets with sentence embeddings that have a high score with keyword vectors) within a state, we represent the dietary characteristics of that state in the word vector space. For obesity prevalence estimation, we apply the elastic net, a regression method that combines L1 and L2 regularization and has been shown to surpass ridge or lasso regressions in text regression tasks [17]. In particular, given the regression task

$$y_s = \mathbf{w}^T \mathbf{x}_s + \beta + \varepsilon$$

where $y_s \in \mathbb{R}$ denotes the obesity prevalence of a given state s , $x_s \in \mathbb{R}^{373}$ the vector of extracted textual features of state s , $\beta \in \mathbb{R}$ the intercept, and $\varepsilon \in \mathbb{R}$ an independent and zero-centered noise, the weight vector w is learnt by optimizing the objective function

$$\arg \min_{\mathbf{w}, \beta} \left(\sum_{s \in S} (\mathbf{w}^T \mathbf{x}_s + \beta + \varepsilon - y_s)^2 + \lambda_1 \sum_{k=1}^m |w_k| + \lambda_2 \sum_{k=1}^m |w_k|^2 \right)$$

where λ_1 and λ_2 are the regularization coefficients, and in practice they are chosen by random search in the range $[1e-5, 1e2]$. We randomly hold out four states for validation and eight states for testing, and apply cross-validations for training.

4 Results and Discussion

4.1 Deep Learning Model Performance

We show the performance of two deep learning models in Table 1 based on their objective functions. Table 1a evaluates the ranking performance of our adaptation of *TagSpace*, and the result is comparable to the implementation by Weston et al. [15] on less noisy text data, which yield 37.42% P@1 and 43.01% R@10. This implies that *TagSpace* maintains its ability to predict hashtags on short and noisy documents and hence applicable to Twitter texts in general. As for the binary version of *TagSpace* shown in Table 1b, there is no prior studies for comparison. The low recall can be explained by the unbalanced labels, as in average only 9.4% of tweets in the test set contain food-related hashtags. The

Table 1. Performance of CNNs on the test set

(a) Ranking by <i>TagSpace</i>			(b) Classification by <i>Binary TagSpace</i>		
Embedding dim	Precision@1	Recall@10	Embedding dim	Precision	Recall
64	15.37%	40.13%	64	73.05%	53.24%
128	28.39%	43.97%	128	84.35%	62.14%
256	32.72%	45.65%	256	87.48%	66.37%

precision of binary *TagSpace* is high, and hence we suspect if the model optimizes objective function by over-generalizing hashtag predictions (i.e. tagging tweets with only general and frequent hashtags such as *#restaurant* and *#diet*). As the model internally learns tweet embeddings, we use them to rank hashtags and find that the most frequent 100 food-related hashtags in the prediction account for 11.3% of the food-related tweets. This implies that binary *TagSpace* gives more granular information about a tweet than whether food-related or not.

4.2 Estimating Obesity Prevalence by Tweet Embeddings

Table 2. Regressions on obesity prevalence by extracting features from word and tweet embeddings

Model	dim	MAE	Pearson Corr.
Bag-of-Words	-	2.596	0.607
<i>TagSpace</i>	64	1.653	0.795
	128	1.571	0.813
	256	1.452	0.836
<i>Binary TagSpace</i>	64	1.239	0.871
	128	1.018	0.904
	256	0.839	0.927

We evaluate the regression results using mean absolute error (MAE) and Pearson correlation with government obesity data. Since no prior study has used our dataset, we handcraft a *Bag-of-Words* baseline that uses tweets filtered by *keyword matching* method and extracts features by frequencies of keywords mentioned within a state. The BOW approach is used in previous Twitter-derived obesity research [8, 9]. The regression results by our baseline moderately correlates to government data, which agrees with prior works that that dietary characteristics mined from Twitter data is informative in actual obesity estimation [2, 8, 9]. Both CNNs generate features resulting in more accurate estimation of state-level obesity compared to our baseline, and binary *TagSpace* outperforms

all other methods. Hence we are optimistic that word and tweet embeddings trained from *TagSpace* models optimizing for selective topics results in better indicators of specific diseases (Table 2).

4.3 Discovering Dietary Risk Factors with Obesity




Fig. 3. Spearman correlations between selected features and obesity prevalence

We are interested in features correlating to higher obesity prevalence, and we obtain such features using Spearman correlation, which quantifies monotonic relationship between two variables. The highest positive correlation to obesity prevalence is given by ‘‘#macncheese’’ ($corr = 0.4910$), ‘‘#wendys’’ (0.4853), ‘‘#doughnut’’ (0.4796), ‘‘#blt’’ (0.4359), and ‘‘#dominospizza’’ (0.4307). We also observe that more general and frequently-mentioned features (such as #dinner, #food) usually have weaker monotonic relationship with our target variable as shown in Fig. 3. While the correlation values are modest, it points to a possibility to learn behavior risk factors from Twitter data using sentence-level embeddings of tweets.

5 Conclusion

In conclusion, we propose a deep learning approach to extract textual features using sentence-level embeddings of tweets for public health monitoring. In the case study, our adaptation of the two CNNs performs reliably on Twitter data

and provides informative textual features for obesity prevalence estimation. We have also shown that features constructed via word and tweet embeddings can potentially learn risk factors for specific diseases, which is useful for monitoring acute public health incidents such as influenza tracking, allergy ailments, and infectious diseases [10, 11, 17]. Our data acquisition and deep learning methods do not include any obesity-related settings, which implies that our approach can be applied to a wide range of Twitter-based public health studies and for various purposes. One limitation of our study is that the demographics of Twitter users over-represent younger age-groups, and one remedy is to standardize tweets based on user ages inferred from probabilistic models [4] for future study. This work serves to inspire future studies to explore the potential of using sentence-level embeddings of social media texts for a wide scope of public health surveillance.

References

1. United states department of agriculture. national nutrient database (2014). <http://ndb.nal.usda.gov/ndb/search/list?format=&count=&max=25&sort=&fg=&man=&lfacet=&qlookup=&offset=50>
2. Abbar, S., Mejova, Y., Weber, I.: You tweet what you eat: studying food consumption through twitter. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, pp. 3197–3206. ACM (2015)
3. Bird, S., Klein, E., Loper, E.: Natural language processing with python, July 2009
4. Chamberlain, B.P., Humby, C., Deisenroth, M.P.: Probabilistic inference of Twitter users’ age based on what they follow. In: Altun, Y., et al. (eds.) ECML PKDD 2017, Part III. LNCS (LNAI), vol. 10536, pp. 191–203. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71273-4_16
5. Culotta, A.: Towards detecting influenza epidemics by analyzing Twitter messages. In: Proceedings of the First Workshop on Social Media Analytics, pp. 115–122. ACM (2010)
6. Ghosh, D., Guha, R.: What are we ‘tweeting’ about obesity? Mapping tweets with topic modeling and geographic information system. *Cartogr. Geogr. Inf. Sci.* **40**(2), 90–102 (2013)
7. Jordan, S., Hovet, S., Fung, I., Liang, H., King-Wa, F., Tse, Z.: Using Twitter for public health surveillance from monitoring and prediction to public response. *Data* **4**(1), 6 (2019)
8. Nguyen, Q.C., et al.: Building a national neighborhood dataset from geotagged twitter data for indicators of happiness, diet, and physical activity. *JMIR Public Health Surveill.* **2**(2), e158 (2016)
9. Nguyen, Q.C., et al.: Twitter-derived neighborhood characteristics associated with obesity and diabetes. *Sci. Rep.* **7**(1), 16425 (2017)
10. Paul, M.J., Dredze, M.: You are what you tweet: analyzing twitter for public health. In: Fifth International AAAI Conference on Weblogs and Social Media (2011)
11. Paul, M.J., Dredze, M.: Discovering health topics in social media using topic models. *PLoS ONE* **9**(8), e103408 (2014)
12. Prier, K.W., Smith, M.S., Giraud-Carrier, C., Hanson, C.L.: Identifying health-related topics on Twitter. In: Salerno, J., Yang, S.J., Nau, D., Chai, S.-K. (eds.) SBP 2011. LNCS, vol. 6589, pp. 18–25. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19656-0_4

13. Sarma, K.V., Spiegel, B.M.R., Reid, M.W., Chen, S., Merchant, R.M., Seltzer, E., Arnold, C.W.: Estimating the health-related quality of life of twitter users using semantic processing. *Stud. Health Technol. Inf.* **264**, 1065–1069 (2019)
14. Sridhar, V.K.R.: Unsupervised topic modeling for short texts using distributed representations of words. In: *Proceedings of the 1st workshop on vector space modeling for natural language processing*, pp. 192–200 (2015)
15. Weston, J., Chopra, S., Adams, K.: # tagspace: semantic embeddings from hashtags. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1822–1827 (2014)
16. Wu, L.Y., Fisch, A., Chopra, S., Adams, K., Bordes, A., Weston, J.: Starspace: embed all the things! In: *Thirty-Second AAAI Conference on Artificial Intelligence* (2018)
17. Zou, B., Lampos, V., Gorton, R., Cox, I.J.: On infectious intestinal disease surveillance using social media content. In: *Proceedings of the 6th International Conference on Digital Health Conference*, pp. 157–161. ACM (2016)