

Analysis of Twitter Users' Lifestyle Choices using Joint Embedding Model

Tunazzina Islam, Dan Goldwasser

Department of Computer Science, Purdue University
 West Lafayette, Indiana 47907
 islam32@purdue.edu, dgoldwas@purdue.edu

Abstract

Multiview representation learning of data can help construct coherent and contextualized users' representations on social media. This paper suggests a joint embedding model, incorporating users' social and textual information to learn contextualized user representations used for understanding their lifestyle choices. We apply our model to tweets related to two lifestyle activities, 'Yoga' and 'Keto diet' and use it to analyze users' activity type and motivation. We explain the data collection and annotation process in detail and provide an in-depth analysis of users from different classes based on their Twitter content. Our experiments show that our model results in performance improvements in both domains.

Introduction

Nowadays, people express opinions, interact with friends and share ideas and thoughts via social media platforms. The data collected by these platforms provide a largely untapped resource for understanding lifestyle choices, health, and well-being (Islam 2019; Amir et al. 2017; Schwartz et al. 2016; Yang and Srinivasan 2016; Schwartz et al. 2013a).

In this paper, we use Twitter to study two lifestyle-related activities, Yoga – a popular multi-faceted activity and Keto-genic diet (often abbreviated as Keto) – a low-carbohydrate, high-fat, adequate-protein diet. Various studies show that yoga offers physical and mental health benefits for people of all ages (Ross and Thomas 2010; Smith and Pukall 2009; Yurtkuran, Alp, and Dilek 2007; Khalsa 2004). Ketogenic diet recently discovered benefits include weight loss (Johnstone et al. 2008), reversal/control of type 2 diabetes (McKenzie et al. 2017) as well as therapeutic potential in many pathological conditions, such as polycystic ovary syndrome, acne, neurological diseases, cancer, and the amelioration of respiratory and cardiovascular disease risk factors (Paoli et al. 2013).

The goal of this paper is to analyze the different lifestyle choices of users based on their tweets. These users can correspond to practitioners who share their journey and explain their motivation when taking a specific lifestyle, but also to commercial parties and interest groups that use social media platforms to advance their interests. Because of the short and

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

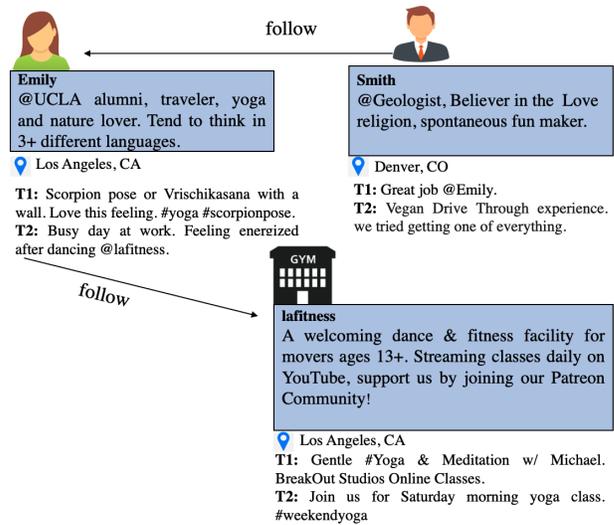


Figure 1: User and Network information on Twitter. Blue rectangle boxes are user profile description and $T1, T2$ are tweets of corresponding users.

often ambiguous nature of tweets, a simple pattern-based analysis using yoga-related keywords often falls short of capturing relevant information. Fig. 1 shows three different Twitter users, the same "#yoga" is used by two different types of users. Our main insight in this paper is that understanding user types and their motivation should be done collectively over the tweets, profile information, and social behavior. In the above example, **Emily**, a practitioner, tweets about specific yoga poses and the emotions the activity evokes. The user **lafitness**, a gym, tweets about online yoga classes in their studios. In addition to the tweets' contents, the profile description of **Emily** indicates that she is a practitioner. And on the other hand, the profile description of **lafitness** indicates that it is a promotional account (Fig. 1). Moreover, social information can help to further disambiguate the text based on the principle of homophily (McPherson, Smith-Lovin, and Cook 2001), the user types and motivations are likely to be reflected by their social circles.

Past work aiming to understand Twitter users' demographic properties has also used a combination of their tweets,

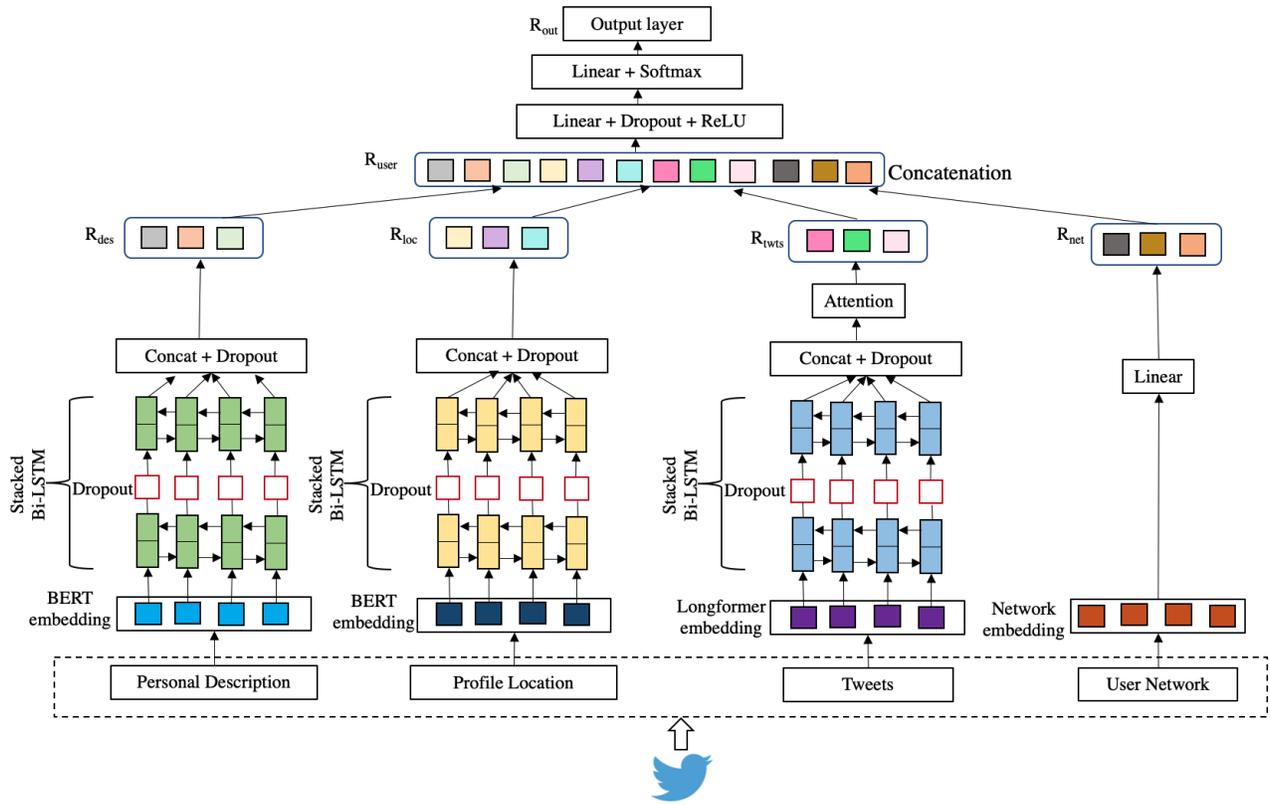


Figure 2: Architecture of our model.

and social information (Li, Ritter, and Jurafsky 2015; Benton, Arora, and Dredze 2016; Yang and Eisenstein 2017; Mishra, Yannakoudakis, and Shutova 2018; Del Tredici et al. 2019), however unlike these works, which look at general demographic properties, our challenge is to construct a user representation relevant for characterizing nuanced, activity and lifestyle specific properties. Recently, pre-trained contextualized language models, such as ELMo (Peters et al. 2018), OpenAI GPT (Radford et al. 2018), and BERT (Devlin et al. 2019), have led to significant improvements in several NLP tasks. However, not much improvement on our task was obtained by directly using the pre-trained BERT model, as it falls short of representing non-linguistic information.

We suggest a method for combining large amounts of Twitter content and social information associated with each user. We concatenate all yoga-related (for keto diet, all keto-related) tweets associated with a given user and use pre-trained Longformer (Beltagy, Peters, and Cohan 2020) – a BERT-like model for long documents, for the contextualized embedding of the user tweets. Many users share their profile description and their location on Twitter; we add this information to our model using a pre-trained BERT model. We refer to our model as BERT based joint embedding model. Finally, we embed the social information associated with the Twitter users and concatenate the user’s embedded representation to their profile description and content representations.

Using this model, we predict (i) the user type, i.e., whether

they are a practitioner, a promotional user, or other (ii) the user’s motivation, e.g., practicing yoga for health benefits, spiritual growth, or other reasons such as commercially motivated users. We compare our model’s performance against several other modeling choices. We describe the data collection and annotation details as well as the results of predicting user type and motivation using our model and in-depth analysis of users from different classes. The main contributions of this work are as follows:

1. Formulating an information extraction type for lifestyle activities, characterizing activity-specific user types, and motivations. We create an annotated dataset related to ‘yoga’ and ‘keto diet’.
2. We suggest a model for aggregating users’ tweets as well as metadata and contextualizing this textual content with social information.
3. We perform extensive experiments to empirically evaluate the contribution of different types of information to the final prediction, and we show that their combination results in the best performing model.
4. We conduct a qualitative analysis aimed at describing the relationship between the output labels and several different indicators, including the tweets, profile descriptions, and location information.

The rest of the paper is organized as follows: we start with the discussion of related work; next, we provide the problem

definition of our work; then, describe the technical approach; next, dataset and annotation details; later, we elaborate details of experimental settings, including discussion of the baseline models, hyperparameter tuning; finally, we show results and analysis containing ablation study, error analysis. Our code and public data are available here.¹

Related Work

Prior research has demonstrated that we can infer many latent user characteristics by analyzing the information in a user’s social media account, i.e., personality (Kosinski, Stillwell, and Graepel 2013; Schwartz et al. 2013b), emotions (Wang and Pal 2015), happiness (Islam and Goldwasser 2020), mental health (Amir et al. 2017), mental disorders (De Choudhury et al. 2013; Reece et al. 2017). (Mishra, Yannakoudakis, and Shutova 2018; Mishra et al. 2019) exploited user’s community information along with textual features for detecting abusive instances. (Ribeiro et al. 2018) characterized hateful users using content as well as user’s activity and connections. (Miura et al. 2017; Ebrahimi et al. 2018; Huang and Carley 2019) used a joint model incorporating different types of available information including tweet text, user network, and metadata for geolocating Twitter users. In contrast, our approach relies on contextualized embeddings for user metadata and tweets. In this paper, we leverage Twitter content and social information associated with each user to learn user representation for predicting the following three tasks: (i) yoga user type, (ii) user motivation related to yoga, (iii) keto user type.

Problem Definition

We formulate our problem as multiview representation neural network based fusion. Suppose we have two learned feature maps X_a and X_b for views a and b , where weights are shared across two views. Concatenation fusion is as follows:

$$h_{cat} = [x_a, x_b] \quad (1)$$

where data from multiple views are integrated into a single representation h , which exploits the knowledge from multiple views to represent the data.

Methodology

We use the following sources of information to train our model: 1) Tweet text; 2) User network; and 3) Metadata, including user location and description. Our model employs those sources and then jointly builds a neural network model to generate a dense vector representation for each field and finally concatenates these representations. Fig. 2 shows the overall architecture of our proposed model.

Metadata Representation

The metadata embedding transforms each metadata into a fix-sized embedding vector. In this paper, we focus on two metadata fields: user description and location. We use pre-trained uncased BERTbase model. Transformers (Vaswani

et al. 2017) in BERT consist of multiple layers, each of which implements a self-attention sub-layer with multiple attention heads. We pass metadata embedding to stacked Bi-LSTM (Hochreiter and Schmidhuber 1997). We get the final hidden representation of metadata by concatenating the forward and backward directions. R_{des} and R_{loc} are the representations of user description and location respectively.

Tweet Representation

For user tweets, we concatenate all yoga-related tweets, which represents a long document. Similarly, for keto, we concatenate all keto-related tweets. Then we use pre-trained longformer-base-4096 model started from the RoBERTa (Liu et al. 2019) checkpoint and pre-trained on long documents. Longformer uses a combination of a sliding window (local) attention and global attention. We forward the tweet embedding to stacked Bi-LSTM. We get the hidden representation of tweets by concatenating the forward and backward directions. To assign important words in the final representation (R_{twts}), we use a context-aware attention mechanism (Bahdanau, Cho, and Bengio 2015).

User Network Representation

To build a dense user network, we consider those users from our dataset if they are @-mentioned (Rahimi et al. 2015) in other users’ (from our data) tweets. We create an undirected and unweighted graph from interactions among users via retweets and/or @-mentions. Nodes are all users in our dataset. An edge is created between two users if either user mentions the other (from our data). In this work, we do not consider edge weights. For yoga, we have 534 nodes with 1831 edges, and for the keto diet, there are 234 nodes with 809 edges. To compute node embedding, we use Node2Vec (Grover and Leskovec 2016). For every node u , Node2Vec’s mapping function maps u to a low dimensional embedding of size d that maximizes the probability of observing nodes belonging to $S(u)$ which is the set of n nodes contained in the graph by taking k random walks starting from u . We generate the embedding of user network $E_{net} = (e_{u_1}, \dots, e_{u_V})$ and forward to a linear layer to compute user network representation, R_{net} .

User Representation

The final user representation, R_{user} is built by concatenation of the four views generated from four sub-networks description, location, tweets, and user network respectively (Fig. 2). We define R_{user} as follows:

$$R_{user} = [R_{des}, R_{loc}, R_{twts}, R_{net}] \quad (2)$$

R_{user} is passed through a fully connected two-layer classifier where the first linear layer with ReLU (Nair and Hinton 2010) activation function. The final prediction R_{out} is passed through a softmax activation function. The risk of overfitting is handled by using dropout (Srivastava et al. 2014) between individual neural network layers. We use stochastic gradient descent over shuffled mini-batches with Adam (Kingma and Ba 2014) and cross-entropy loss as the objective function for classification.

¹<https://github.com/tunazislam/Joint-Embedding-Model>

Model	lr	opt	reg	batch	hd	lstm	attn	ls	eut	eum
Description	$1e^{-3}$	Adam	0	32	300	2	-	-	6	5
Location	$1e^{-3}$	Adam	0	32	300	2	-	-	5	5
Tweets	$1e^{-3}$	Adam	0	32	300	2	300	-	8	6
Network	$1e^{-3}$	Adam	0	32	150	-	-	-	4	8
Des_BF	$2e^{-5}$	AdamW	.01	32	-	-	-	-	4	4
Loc_BF	$2e^{-5}$	AdamW	.01	32	-	-	-	-	2	2
Twts_BF	$2e^{-5}$	AdamW	.01	32	-	-	-	-	2	4
Des + Loc	$1e^{-3}$	Adam	0	32	300	2	-	600	5	6
Des + Net	$1e^{-3}$	Adam	0	32	300	2	-	600	5	5
Des + Loc + Twt	$1e^{-3}$	Adam	0	32	300	2	300	600	5	7
Des + Loc + Net	$1e^{-3}$	Adam	0	32	300	2	-	600	6	7
Our model	$1e^{-3}$	Adam	0	32	300	2	300	600	7	6

lr	Learning rate.
opt	Optimizer.
reg	Weight decay (L^2 regularization).
batch	Batch size.
hd	Hidden dimension.
lstm	Number of LSTM layer as we use stacked Bi-LSTM.
attn	Attention vector size.
ls	Size of the first layer of two-layer classifier.
eut	Best result achieved at epochs for user type classification.
eum	Best result achieved at epochs for user motivation classification.

Table 1: Hyperparameter details of the models.

Dataset

We download tweets using Tweepy by Twitter streaming API sub-sequentially from May to November, 2019. For yoga, we collect 419608 tweets related to yoga containing specific keywords : ‘yoga’, ‘yogi’, ‘yogalife’, ‘yogalove’, ‘yogainspiration’, ‘yogachallenge’, ‘yogaeverywhere’, ‘yogaeveryday’, ‘yogadaily’, ‘yogaeverydamnday’, ‘yogapractice’, ‘yogapose’, ‘yogalover’, ‘yogajourney’. There are 297350 different users among them 13589 users have at least five yoga-related tweets in their timelines. For this work, we randomly pick 1298 users and collect their timeline tweets. We have 3097678 timeline tweets in total.

For ketogenic diet, we focus on several keywords i.e., ‘keto’, ‘ketodiet’, ‘ketogenic’, ‘ketosis’, ‘ketogenicdiet’, ‘ketolife’, ‘ketolifestyle’, ‘ketogenicfood’, ‘ketogenicfoodporn’, ‘ketone’, ‘ketogeniclifestyle’, ‘ketogeniccommunity’, ‘keto-community’, ‘ketojourney’ to extract 75048 tweets from 38597 different users. Among them 16446 users have at least two keto-related tweets in their timelines. In this paper, we randomly pick 1300 users and we have total 3253833 timeline tweets.

To pre-process the text, we first convert them into lower case, remove URLs, smiley, emoji. To prepare the data for input to BERT and Longformer, we tokenize the text using BERT and RoBERTa’s wordpiece tokenizer.

Data Annotation

To annotate the data, for each user, we check both their profile description and timeline tweets. For the tweets, we consider only yoga/keto-related tweets from their timeline. We first

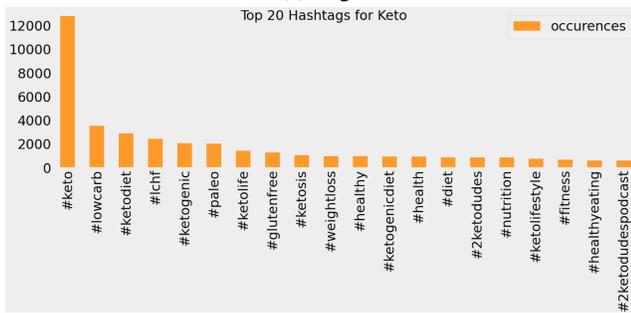
look at the user profile description for user type, whether they explicitly mention practicing a specific lifestyle (i.e., yogi, ketosis); then, we look for the user timeline tweets. If the user tweets about the first-hand experience of practicing yoga (i.e., *love practicing yoga early in the morning*)/keto diet (i.e., *lost 5lb in 1st week of keto*), we annotate the user as a ‘practitioner’. After looking at the description and tweets, if we observe that they are promoting a gym/studio (i.e., *offering free online yoga class*), online shop (i.e., *selling yoga mat*), app (i.e., *sharing keto food recipe*), restaurant, community etc., rather than sharing their first-hand experience about a particular lifestyle, we annotate them as a ‘promotional’ user. If we notice a user has all retweets in their timeline tweets related to yoga/keto (they might have an interest in a particular lifestyle), we annotate them as ‘others’.

For user motivation, we check for the user timeline tweets. If the user tweets about practicing yoga for health benefit (i.e., *yoga heals my back pain*), we annotate the user motivation as ‘health’. If the user tweets about practicing yoga for spiritual help (i.e., *yoga gives me a spiritual wisdom path*), we annotate the user motivation as ‘spiritual’. Otherwise, we annotate the motivation as ‘others’.

To calculate inter-annotator agreement, two graduate students manually annotate a subset of tweets. This subset has an inter-annotator agreement of 64.7% (substantial agreement) using Cohen’s Kappa coefficient (Cohen 1960). After checking annotators’ disagreement for yoga data, we notice that some yoga teacher/instructor has their yoga studio. They tweet about promoting their yoga studio and sharing their first-hand yoga practice experience. One annotator labels them as practitioners, and another annotates them as pro-



(a) Yoga



(b) Keto

Figure 3: Top 20 hashtags related to yoga (green bars) and keto diet (orange bars).

motional. Multi-label classification would be a good option, but we consider only a single label in this paper. Annotators face challenges with user type ‘others’ when they indicate being practitioners (i.e., *yoga lover*) in their Twitter profile description. Still, they retweet about yoga most of the time rather than sharing their experiences in the particular timeline tweets. The annotators then come to a fair agreement to label the remaining tweets of the dataset.

Data Distribution

In our user type annotated yoga data, we have 42% practitioner, 21% promotional, and 37% other users who love to tweet/retweet about yoga but do not practice yoga. In the user motivation annotated data, we have 51% users who tweet about yoga regarding health benefit, 5% spiritual, and 41% other motivation e.g., business. After annotating 1300 keto users, we have 50.8% practitioners, 19% promotional, and 30.2% other users.

Experimental Settings

To run the experiment, we shuffle our dataset and then randomly split it into train (60%), validation (20%), and test (20%).

Baseline Models

In our experiments, we evaluate our model under twelve different settings: (i) Description; (ii) Location; (iii) Tweets; (iv) Network; (v) BERT fine-tuned with Description (Des_BF); (vi) BERT fine-tuned with Location (Loc_BF); (vii) BERT fine-tuned with Tweets (Twts_BF); (viii) joint embedding on

description and location (Des + Loc); (ix) joint embedding on description and network (Des + Net); (x) joint embedding on description, location, and tweets (Des + Loc + Twt); (xi) joint embedding on description, location, and network (Des + Loc + Net), (xii) Word2Vec based joint embedding on description, location, tweets, and network (Islam and Goldwasser 2021).

Description For user description embedding, we use pre-trained uncased BERTbase model using a masked language modeling (MLM) objective. Transformers in BERT consist of multiple layers, each of which implements a self-attention sub-layer with multiple attention heads. We pass the embedding to stacked Bi-LSTM with a dropout value 0.5. We get the hidden representation of description by concatenating the forward and backward directions with dropout (0.5). We forward the description representation to one-layer classifier activated by softmax.

Location To represent location, we follow the same approach as user description.

Tweets For user tweets, we concatenate all yoga-related tweets. Then we use pre-trained longformer-base-4096 model for tweet embedding. We forward the tweet embedding to stacked Bi-LSTM with dropout layer (0.5). We get the hidden representation of tweets by concatenating the forward and backward directions with dropout (0.5). To assign important words in the final representation, we use a context-aware attention mechanism. We forward the tweet representation to one-layer classifier activated by softmax.

Network We use Node2Vec for network embedding after the construction of the user network. We forward the network embedding to a linear layer for obtaining network representation. We pass the network representation to a dropout layer of value 0.5 with ReLU activation and then forward to one-layer classifier activated by softmax.

Des + Loc We concatenate profile description and location representations and pass through a two-layer classifier activated by ReLU and softmax, respectively.

Des + Net We concatenate user’s profile description and user network representations and forward the joint representation to a two-layer classifier activated by ReLU and then softmax.

Des + Loc + Twt We concatenate description, location, tweet representations and pass the joint representation to a two-layer classifier activated by ReLU and softmax correspondingly.

Des + Loc + Net In this case, concatenation of description, location, and user network representations are fed to the two-layer classifier with ReLU and softmax activation function.

Fine-tuning pre-trained BERT BERT uses bidirectional transformers to pre-train a large corpus and fine-tunes the pre-trained model on other tasks. We use description, location, and tweets separately and fine-tune the pre-trained BERT (base-uncased) for Des_BF, Loc_BF, and Twts_BF baselines.

Des_BF For this setting, we use a pre-trained BERT (*Bert-ForSequenceClassification*) model and fine-tune it with an

Model	User type		User motivation	
	Accuracy	Macro-avg F1	Accuracy	Macro-avg F1
Description	0.694	0.611	0.707	0.523
Location	0.639	0.520	0.694	0.517
Tweets	0.795	0.704	0.786	0.595
Network	0.726	0.561	0.798	0.590
Des_BF	0.718	0.681	0.771	0.528
Loc_BF	0.679	0.606	0.695	0.476
Twts_BF	0.760	0.669	0.805	0.551
Des + Loc	0.734	0.653	0.806	0.661
Des + Net	0.808	0.702	0.823	0.653
Des + Loc + Twt	0.778	0.705	0.808	0.603
Des + Loc + Net	0.774	0.725	0.806	0.663
Word2Vec based joint embedding	0.790	0.742	0.844	0.610
Our Model	0.802	0.757	0.853	0.708

Table 2: Performance comparisons on yoga data.

added single linear layer on top. In this case, BERT’s input is the user’s profile description constructed by the summation of the corresponding token, segment, and position embeddings for a given token.

Loc_BF We use a pre-trained and fine-tuned BERT model with user location to classify user type and user motivation.

Twts_BF For Twts_BF model, we use a pre-trained and fine-tuned *BertForSequenceClassification* model with user tweets to classify our tasks.

Word2Vec based joint embedding Instead of using pre-trained BERT, we use pre-trained Word2Vec (Mikolov et al. 2013) for tweets, location, and description embedding. Then concatenate the four sub-networks description, location, tweets, and user networks and pass them to the two-layer classifier with ReLU and softmax activation function.

Hyperparameter Details

We set the hyperparameters of our final model as follows: batch size = 32, learning rate = 0.001, epochs = 10. The dropout rate between layers is set to 0.5. We perform grid hyperparameter search on the validation set using early stopping for all the models except these three models – Des_BF, Loc_BF, Twts_BF. For learning rate, we investigate values 0.001, 0.01, 0.05, 0.1; for L^2 regularization, we examine 0, 10^{-3} , 10^{-2} values; and for dropout, values 0.2, 0.25, 0.4, 0.5. We run the models total 10 epochs and plot curves for loss, accuracy, and macro-avg F1 score. Our early stopping criterion is based on the validation loss when it starts to increase sharply². For Word2Vec based joint embedding model, we have the same hyperparameters used by (Islam and Goldwasser 2021).

In the Des_BF, Loc_BF, and Twts_BF models, for padding or truncating text, we chose maximum sentence length = 160, 50, 500 correspondingly. We use learning rate = $2e - 5$, optimizer= AdamW (Loshchilov and Hutter 2018), epochs = 4, epsilon parameter = $1e - 8$, batch size = 32,

²The plots are in the supplementary material.

Network embeddings are trained using Node2Vec with following parameters: dimension = 300, number of walks per source = 10, length of walk per source = 80, minimum count = 1, window size = 10 and then forwarded to a linear layer of size 150. If users do not appear in the @-mentioned network, we set their network embedding vectors 0. Also, for the users without having location and/or description, we set the embedding vectors as 0 correspondingly. Table 1 summarizes the hyperparameter settings of all models.

Results and Analysis

As reported in Table 2, our proposed model achieves the best test accuracy and macro-avg F1 score for classifying yoga user type and motivation and outperforms the baseline models. Our model obtains the highest test accuracy (80.2%) and macro-avg F1 score (75.7%) for classifying yoga user type (Table 2). We achieve noticeable performance for classifying yoga user motivation where our model obtains the highest test accuracy (85.3%) and macro-avg F1 score (70.8%). To predict user type in keto data, we get test accuracy (71.9%) and macro-avg F1 score (67.6%).

Ablation Study

We train an individual neural network model for each field – description, location, tweets, and network. The 1st four rows of the Table 2 show the performance breakdown for each model over the test dataset. The results conclude that tweets and profile description of any user is informative fields for our task. However, our experiments show that either excluding user network information (Des + Loc + Twt model) or tweets information (Des + Loc + Net model) declines the final model’s performance in terms of both accuracy and macro-avg F1 for both user type and motivation classification task (10th and 11th rows of Table 2).

Top Hashtags

Fig. 3 shows the top 20 hashtags (#) related to yoga and keto diet and the number of occurrences of those hashtags

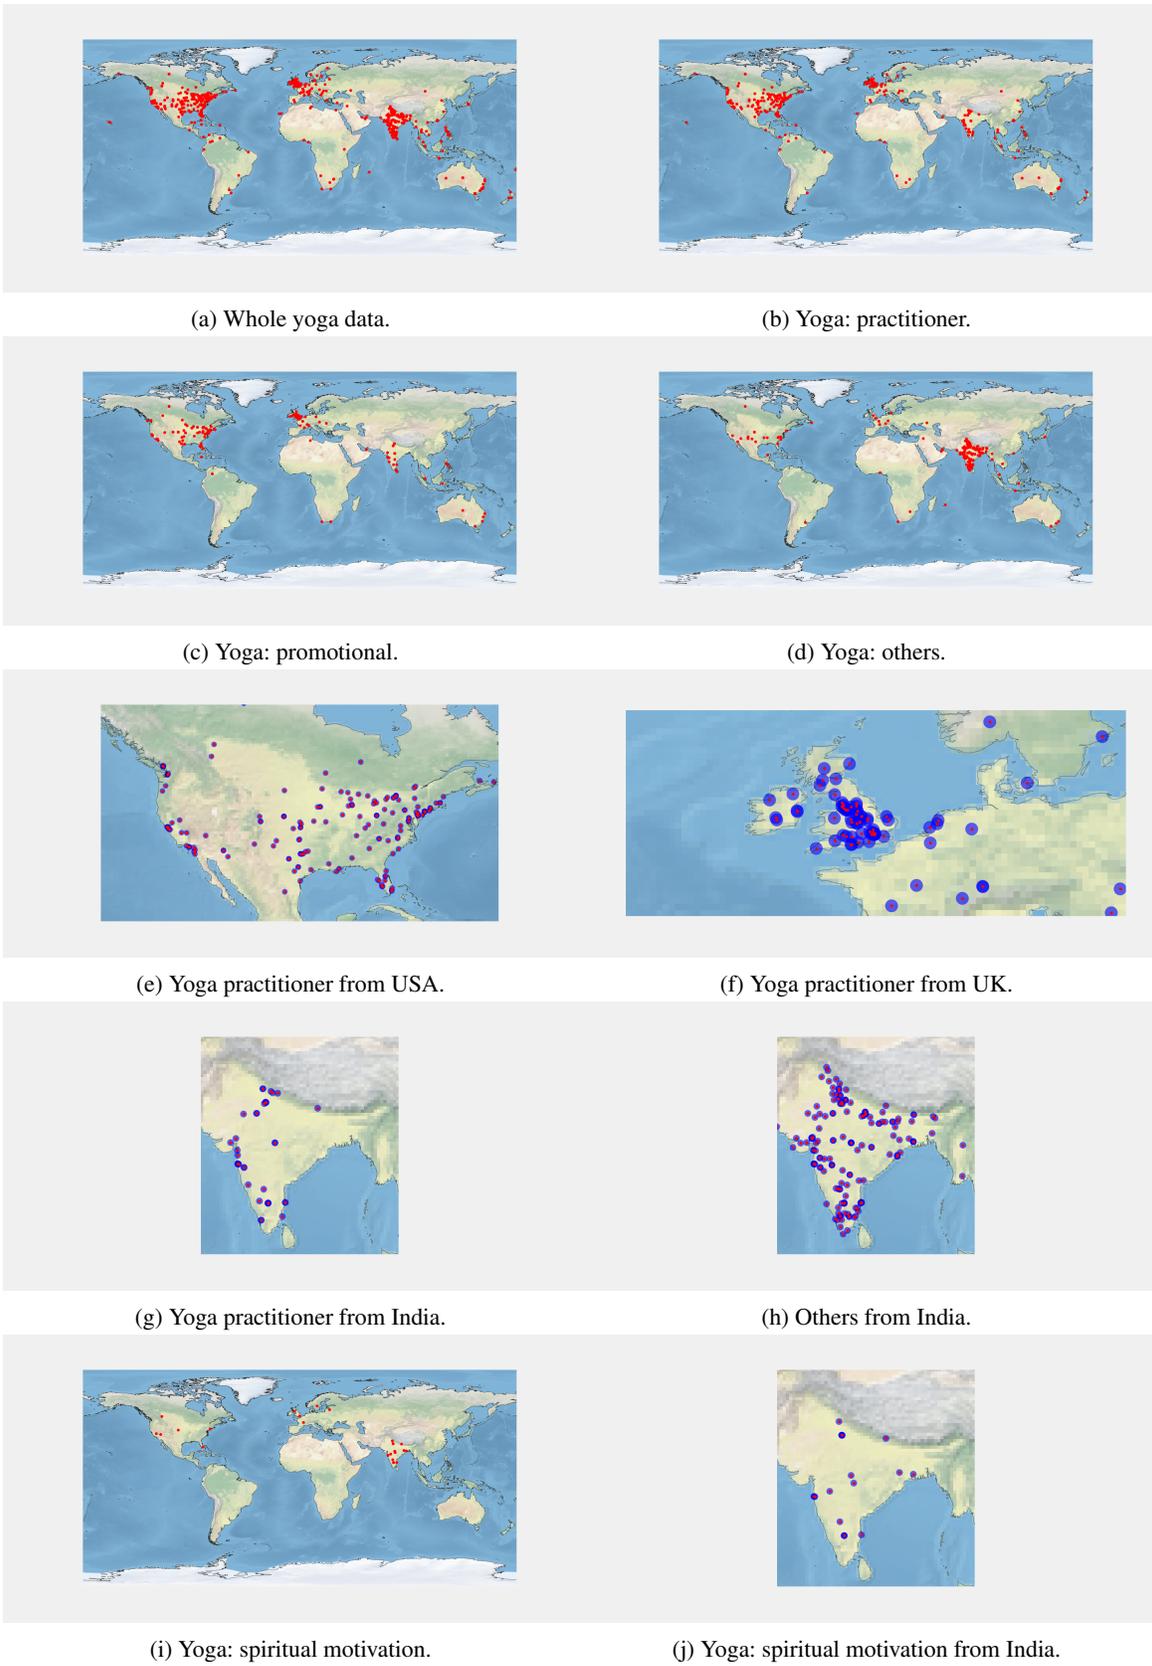


Figure 6: Yoga user distribution over location.

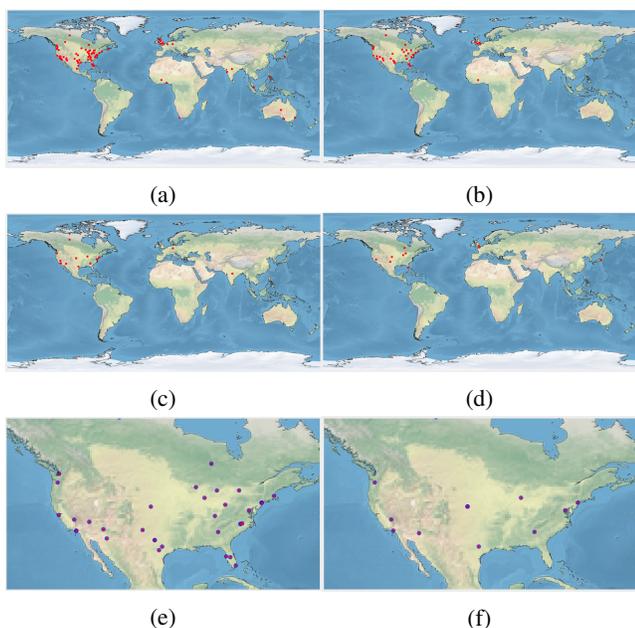


Figure 7: Keto user distribution over location. (a) whole keto data, (b) practitioner, (c) promotional, (d) others, (e) keto practitioner from USA, (f) keto promotional from USA.

community, event, offer, free, product, market, business, program, design (Fig. 5b). Users who practice yoga for health benefits have similar wordcloud to yoga practitioners’ descriptions (Fig. 5c). Fig. 5d shows the wordcloud of the profile description of spiritually motivated yoga user having words like *yoga, spiritual, spirituality, devotee, wisdom, peace, seeker, yogi, meditator, Indian*.

Fig. 5e and 5f show the wordcloud of the profile description of keto practitioner having words *keto, love, life, food, family* and promotional containing *food, health, keto, meal, recipe, product, online, free* words respectively.

Relationship between Location Information and Labels

In this section, we illustrate the relationship between location information and user type. We use Nominatim package from GeoPy³ that given a location (either full address or city name) can identify a real-world location and provide some extra details such as latitude and longitude. To visualize the map, we use Cartopy⁴ to plot individual locations (seen as red dots on the map), as well as blue circles whose radius varies by how many tweets come from that particular place.

In Fig. 6a, we plot yoga data distribution over the user location. We observe that we have more practitioners (Fig. 6b) and promotional users (Fig. 6c) from the USA than the rest of the world. We find South-Asian users mostly retweet about yoga (Fig. 6d). Fig. 6e, 6f, and 6g show yoga practitioners location from USA, UK, and India. We notice more ‘others’ users than practitioners in India (Fig. 6h). In Fig. 6i, we show

yoga motivation for spirituality distribution and most of them are from India. Fig. 6j shows the yoga users from India who are motivated spiritually.

In Fig. 7a, we show whole keto data distribution over the user location. Fig. 7b, 7c, and 7d show keto practitioner, promotional, and other users data distribution over the user location. We notice that our data is skewed towards the USA. Fig. 7e and 7f show keto practitioners and promotional users location from the USA.

Error Analysis

Overall accuracy in detecting yoga user types from our data is 85.2%. Our model correctly predicts 1107 users. We have 191 misclassifications in the yoga dataset, including 67 misclassifications in predicting yoga practitioners, 43 promotional users are misclassified. We notice the highest number of misclassifications (81) in predicting other types of users. 38 users are misclassified as practitioners, and 43 users are misclassified as promotional users.

For yoga user’s motivation, our model correctly predicts 1113 users with an overall accuracy 85.6%. We have 185 misclassifications, including 102 misclassifications in predicting health-related motivation, 28 for spiritual and 55 other motivations are misclassified. We observe the highest number of misclassifications in predicting users’ health motivation for doing yoga where 15 and 87 users’ motivation are misclassified as spiritual and others correspondingly.

For detecting keto users, the overall accuracy is 77.6%. Our model correctly predicts 1009 users. We have 291 misclassifications in keto users, including 94 misclassifications in predicting keto practitioners, 53 promotional. We notice the highest number of misclassifications (144) in predicting other types of users. There are 96 and 48 users who are misclassified as practitioner and promotional, respectively.

Our ablation study demonstrates that the profile description, tweets, and network field contribute mainly to the classification task. However, some prediction errors arise when description fields are absent or misleading. We notice that the user location has relatively low accuracy and macro-avg F1 score from our ablation study. Users sometimes do not provide location information on Twitter. Besides, as Longformer supports sequences of length up to 4096, we might lose some information from tweets if the size of concatenated tweets > 4096. Moreover, we construct @-mentioned network directly from retweets/mentions in tweets, which is less expensive to collect than the following network.

Conclusion and Future Work

We propose a BERT based joint embedding model that explicitly learns contextualized user representations by leveraging users’ social and textual information. We show that our model outperforms multiple baselines. Besides yoga, we demonstrate that our model can effectively predict user type on another lifestyle choice, e.g., ‘keto diet’ and our approach is a general framework that can be adapted to other corpora. In the future, we aim to investigate our work to a broader impact like community detection based on different lifestyle decisions using minimal supervision.

³<https://geopy.readthedocs.io/en/stable/#>

⁴<https://scitools.org.uk/cartopy/docs/v0.16/>

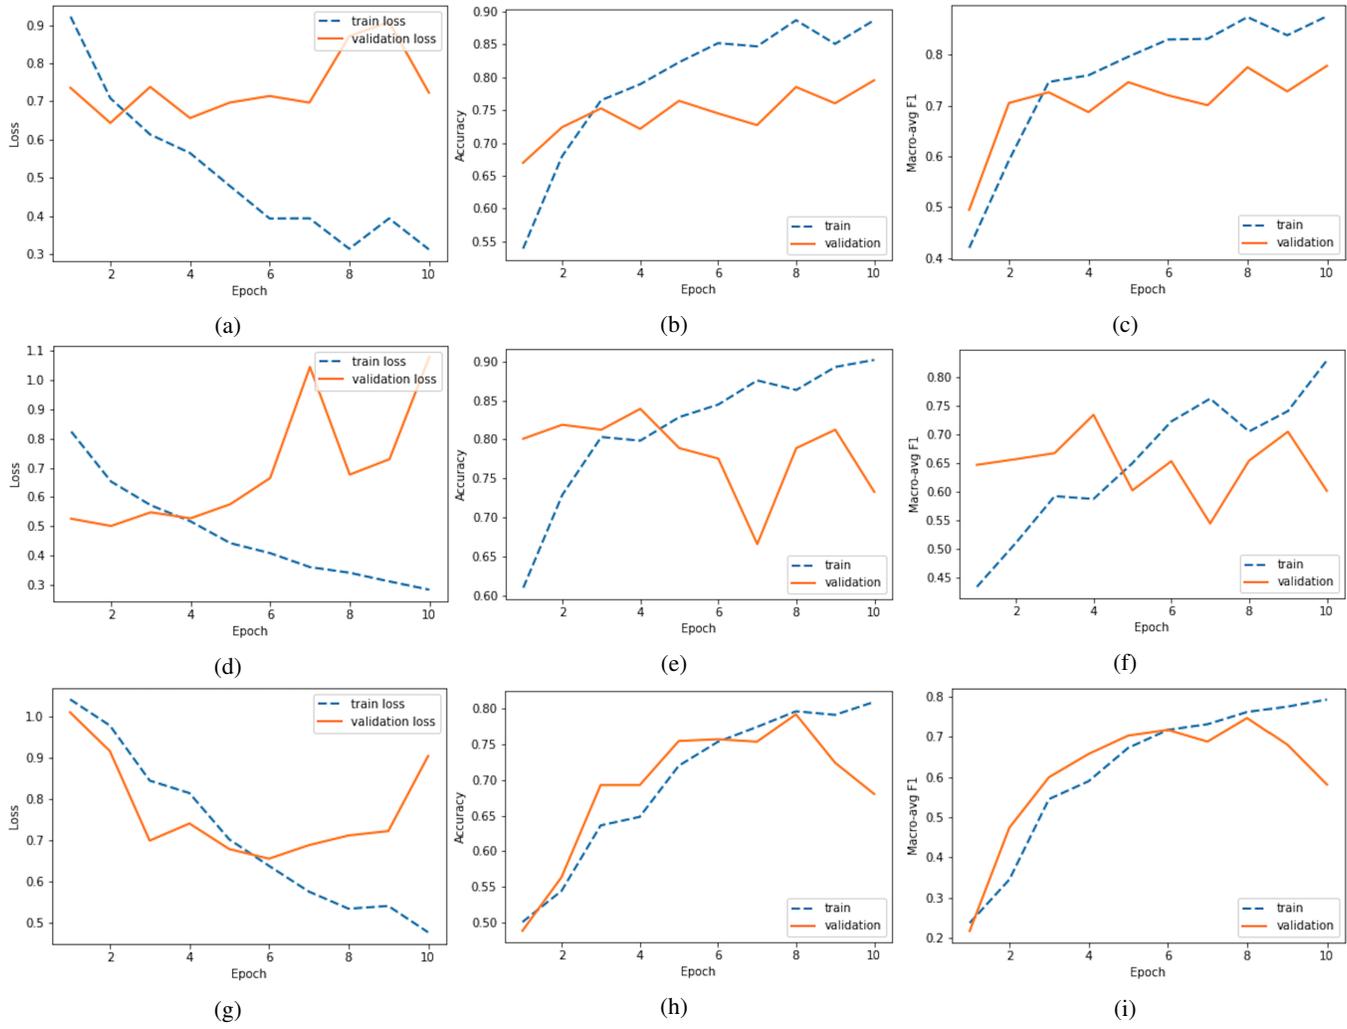


Figure 8: Learning curves of our BERT based joint embedding model for training and validation data based on yoga and keto dataset. The blue dashed line represents train data and the orange solid line represents validation data. The x -axis shows number of epochs and y -axis corresponds to loss, accuracy, and macro-avg F1 score respectively. (a) loss vs. epochs for yoga user type, (b) accuracy vs. epochs for yoga user type, (c) macro-avg F1 score vs. epochs for yoga user type, (d) loss vs. epochs for yoga user motivation, (e) accuracy vs. epochs for yoga user motivation, (f) macro-avg F1 score vs. epochs for yoga user motivation, (g) loss vs. epochs for keto user type, (h) accuracy vs. epochs for keto user type, (i) macro-avg F1 score vs. epochs for keto user type.

Acknowledgments

We are grateful to the anonymous reviewers for providing insightful reviews and Md Masudur Rahman for helping in human evaluation.

Supplementary Material

Fig. 8 shows the learning curves loss (train and validation) vs. epochs, macro-avg F1 score (train and validation) vs. epochs, and accuracy (train and validation) vs. epochs for our BERT based joint embedding model on yoga and keto dataset.

References

- Amir, S.; Coppersmith, G.; Carvalho, P.; Silva, M. J.; and Wallace, B. C. 2017. Quantifying mental health from social media with neural user embeddings. *arXiv preprint arXiv:1705.00335*.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Beltagy, I.; Peters, M. E.; and Cohan, A. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Benton, A.; Arora, R.; and Dredze, M. 2016. Learning multiview embeddings of twitter users. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 14–19.
- Cohen, J. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement* 20(1): 37–46.
- De Choudhury, M.; Gamon, M.; Counts, S.; and Horvitz, E. 2013. Predicting depression via social media. In *Seventh international AAAI conference on weblogs and social media*.
- Del Tredici, M.; Marcheggiani, D.; im Walde, S. S.; and Fernández, R. 2019. You Shall Know a User by the Company It Keeps: Dynamic Representations for Social Media Users in NLP. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4701–4711.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186.
- Ebrahimi, M.; ShafieiBavani, E.; Wong, R.; and Chen, F. 2018. A unified neural network model for geolocating twitter users. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, 42–53.
- Grover, A.; and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 855–864.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8): 1735–1780.
- Huang, B.; and Carley, K. M. 2019. A Hierarchical Location Prediction Neural Network for Twitter User Geolocation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4734–4744.
- Islam, T. 2019. Yoga-Veganism: Correlation Mining of Twitter Health Data. *arXiv preprint arXiv:1906.07668*.
- Islam, T.; and Goldwasser, D. 2020. Does Yoga Make You Happy? Analyzing Twitter User Happiness using Textual and Temporal Information. In *2020 IEEE International Conference on Big Data (Big Data)*, 4241–4249. doi:10.1109/BigData50022.2020.9378461.
- Islam, T.; and Goldwasser, D. 2021. Do You Do Yoga? Understanding Twitter Users’ Types and Motivations using Social and Textual Information. In *2021 IEEE 15th International Conference on Semantic Computing (ICSC)*, 362–365. doi:10.1109/ICSC50631.2021.00067.
- Johnstone, A. M.; Horgan, G. W.; Murison, S. D.; Bremner, D. M.; and Lobley, G. E. 2008. Effects of a high-protein ketogenic diet on hunger, appetite, and weight loss in obese men feeding ad libitum. *The American journal of clinical nutrition* 87(1): 44–55.
- Khalsa, S. B. S. 2004. Treatment of chronic insomnia with yoga: A preliminary study with sleep-wake diaries. *Applied psychophysiology and biofeedback* 29(4): 269–278.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kosinski, M.; Stillwell, D.; and Graepel, T. 2013. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the national academy of sciences* 110(15): 5802–5805.
- Li, J.; Ritter, A.; and Jurafsky, D. 2015. Learning multi-faceted representations of individuals from heterogeneous evidence using neural networks. *arXiv preprint arXiv:1510.05198*.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Loshchilov, I.; and Hutter, F. 2018. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- McKenzie, A. L.; Hallberg, S. J.; Creighton, B. C.; Volk, B. M.; Link, T. M.; Abner, M. K.; Glon, R. M.; McCarter, J. P.; Volek, J. S.; and Phinney, S. D. 2017. A novel intervention including individualized nutritional recommendations reduces hemoglobin A1c level, medication use, and weight in type 2 diabetes. *JMIR diabetes* 2(1): e5.
- McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* 27(1): 415–444.

- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space .
- Mishra, P.; Del Tredici, M.; Yannakoudakis, H.; and Shutova, E. 2019. Abusive Language Detection with Graph Convolutional Networks. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2145–2150.
- Mishra, P.; Yannakoudakis, H.; and Shutova, E. 2018. Neural Character-based Composition Models for Abuse Detection. In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, 1–10.
- Miura, Y.; Taniguchi, M.; Taniguchi, T.; and Ohkuma, T. 2017. Unifying text, metadata, and user network representations with a neural network for geolocation prediction. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1260–1272.
- Nair, V.; and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, 807–814.
- Paoli, A.; Rubini, A.; Volek, J.; and Grimaldi, K. 2013. Beyond weight loss: a review of the therapeutic uses of very-low-carbohydrate (ketogenic) diets. *European journal of clinical nutrition* 67(8): 789–796.
- Peters, M. E.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365* .
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving language understanding by generative pre-training.
- Rahimi, A.; Vu, D.; Cohn, T.; and Baldwin, T. 2015. Exploiting text and network context for geolocation of social media users. *arXiv preprint arXiv:1506.04803* .
- Reece, A. G.; Reagan, A. J.; Lix, K. L.; Dodds, P. S.; Danforth, C. M.; and Langer, E. J. 2017. Forecasting the onset and course of mental illness with Twitter data. *Scientific reports* 7(1): 13006.
- Ribeiro, M. H.; Calais, P. H.; Santos, Y. A.; Almeida, V. A.; and Meira Jr, W. 2018. Characterizing and detecting hateful users on twitter. *arXiv preprint arXiv:1803.08977* .
- Ross, A.; and Thomas, S. 2010. The health benefits of yoga and exercise: a review of comparison studies. *The journal of alternative and complementary medicine* 16(1): 3–12.
- Schwartz, H. A.; Eichstaedt, J. C.; Kern, M. L.; Dziurzynski, L.; Lucas, R. E.; Agrawal, M.; Park, G. J.; Lakshminanth, S. K.; Jha, S.; Seligman, M. E.; et al. 2013a. Characterizing Geographic Variation in Well-Being Using Tweets. *ICWSM* 13: 583–591.
- Schwartz, H. A.; Eichstaedt, J. C.; Kern, M. L.; Dziurzynski, L.; Ramones, S. M.; Agrawal, M.; Shah, A.; Kosinski, M.; Stillwell, D.; Seligman, M. E.; et al. 2013b. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one* 8(9): e73791.
- Schwartz, H. A.; Sap, M.; Kern, M. L.; Eichstaedt, J. C.; Kapelner, A.; Agrawal, M.; Blanco, E.; Dziurzynski, L.; Park, G.; Stillwell, D.; et al. 2016. Predicting individual well-being through the language of social media. In *Biocomputing 2016: Proceedings of the Pacific Symposium*, 516–527. World Scientific.
- Smith, K. B.; and Pukall, C. F. 2009. An evidence-based review of yoga as a complementary intervention for patients with cancer. *Psycho-Oncology: Journal of the Psychological, Social and Behavioral Dimensions of Cancer* 18(5): 465–475.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15(1): 1929–1958.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Wang, Y.; and Pal, A. 2015. Detecting emotions in social media: A constrained optimization approach. In *Twenty-fourth international joint conference on artificial intelligence*.
- Yang, C.; and Srinivasan, P. 2016. Life satisfaction and the pursuit of happiness on Twitter. *PloS one* 11(3): e0150881.
- Yang, Y.; and Eisenstein, J. 2017. Overcoming language variation in sentiment analysis with social attention. *Transactions of the Association for Computational Linguistics* 5: 295–307.
- Yurtkuran, M.; Alp, A.; and Dilek, K. 2007. A modified yoga-based exercise program in hemodialysis patients: a randomized controlled study. *Complementary therapies in medicine* 15(3): 164–171.