

Social Media-based User Embedding: A Literature Review

Shimei Pan and Tao Ding

Department of Information Systems, University of Maryland, Baltimore County
 {shimei, taoding01}@umbc.edu

Abstract

Automated representation learning is behind many recent success stories in machine learning. It is often used to transfer knowledge learned from a large dataset (e.g., raw text) to tasks for which only a small number of training examples are available. In this paper, we review recent advance in learning to represent social media users in low-dimensional embeddings. The technology is critical for creating high performance social media-based human traits and behavior models since the ground truth for assessing latent human traits and behavior is often expensive to acquire at a large scale. In this survey, we review typical methods for learning a unified user embeddings from heterogeneous user data (e.g., combines social media texts with images to learn a unified user representation). Finally we point out some current issues and future directions.

1 Introduction

People currently spend a significant amount of time on social media to express opinions, interact with friends and share ideas. As a result, social media data contain rich information that is indicative of who we are and predictive of our online or real world behavior. With the recent advent of big data analytics, social media-based human trait and behavioral analytics has increasingly been used to better understand human minds and predict human behavior. Prior research has demonstrated that by analyzing the information in a user's social media account, we can infer many latent user characteristics such as political leaning [Pennacchiotti and Popescu, 2011; Kosinski *et al.*, 2013; Benton *et al.*, 2016], brand preferences [Pennacchiotti and Popescu, 2011; Yang *et al.*, 2015], emotions [Kosinski *et al.*, 2013], mental disorders [De Choudhury *et al.*, 2013], personality [Kosinski *et al.*, 2013; Schwartz *et al.*, 2013; Liu *et al.*, 2016; Golbeck *et al.*, 2011], substance use [Kosinski *et al.*, 2013; Ding *et al.*, 2017] and sexual orientation [Kosinski *et al.*, 2013].

Although social media allow us to easily record a large amount of user data, the characteristics of social media data also bring significant challenges to automated data analytics. For example, the texts and images are unstructured data.

Making sense of unstructured data is always a big challenge. It is also hard to efficiently search and analyze a large social graph. Moreover, social media analytics can easily suffer from the *curse of dimensionality problem*. If we use the basic text features such as unigrams or TF*IDF scores as the features to represent text, we can easily have hundreds of thousands of text features. Moreover, assessing human traits and behavior often requires psychometric evaluations or medical diagnosis, which are expensive to perform at a large scale (e.g., only trained professionals can provide an accurate assessment on whether someone has substance use disorders or not). Without proper user feature learning, a machine learning model can easily overfit the training data and will not generalize well to new data.

Recent years have seen a surge in methods that automatically encode features in low-dimensional embeddings using techniques such as dimension reduction and deep learning [Mikolov *et al.*, 2013; Le and Mikolov, 2014; Bengio *et al.*, 2013; Grover and Leskovec, 2016; Perozzi *et al.*, 2014]. *Representation learning* has increasingly become a critical tool to boost the performance of complex machine learning applications. In this paper, we review recent work on automatically learning user representations from social media data. Since automated user embedding simultaneously performs latent feature learning and dimension reduction, it can help downstream tasks to avoid overfitting and boost performance.

2 Overview

Here we define social media-based user embedding as the function that maps raw user features in a high dimensional space to dense vectors in a low dimensional embedding space. The learned user embeddings often capture the essential characteristics of individuals on social media. Since they are quite general, the learned user embeddings can be used to support diverse downstream user analysis tasks such as user preference prediction [Pennacchiotti and Popescu, 2011], personality modeling [Kosinski *et al.*, 2013] and depression detection [Amir *et al.*, 2017].

Automated user embedding is different from traditional user feature extraction where a pre-defined set of features is extracted from data. For example, based on the Linguistic Inquiry and Word Count (LIWC) dictionary [Pennebaker *et al.*, 2015], a set of psycholinguistic features can be extracted from

text. Similarly, a set of egocentric network features such as degree, size and betweenness centrality can be extracted from one’s social network. The main difference between user embedding and traditional user feature extraction is that in user embedding, the user features are not pre-defined. They are latent features automatically learned from data.

Figure 1 shows the typical architecture of a system that employs automated user embedding for personal traits and behavior analysis. One or more types of user data are first extracted from a social media account. For each type of user data such as text or image, a set of latent user features is learned automatically via *single-view user embedding* (e.g., text-based user embedding and image-based user embedding). The embeddings learned from different types of user data (e.g., text embeddings and image embeddings) are combined to form a single unified user representation via *multi-view user embedding*. The output of multi-view user embedding is then used in subsequent applications to predict human traits and behavior.

Given the page limit, we define the scope of this survey quite narrowly to include only embedding methods published within the last 10 years that have been used to learn user representations from social media data. Although very relevant, We exclude embedding methods that do not learn a representation of social media users. For example, we exclude the papers on learning word embeddings from social media data [Zeng *et al.*, 2018]. Table 1 lists the papers included in our survey. We summarize each paper along six dimensions: Data Type, Single-view Embedding Method, Auxiliary Training Task, Multi-view Embedding Method, Target Task and Supervised Tuning.

Among them, *Data Type* is used to indicate the types of social media data used in each study. Here, *text* refers to user-generated text data (e.g., tweets or status update on Facebook); *like* refers to things/people a social media user likes such as books, celebrities, music, movies, TV shows, photos and products; *user profile* includes demographic information (gender, age, occupation, relationship status etc.) and aggregated statistics (the number of friends, followers, followings etc.); *image* includes the profile and background photos as well as the images shared on social media; *social network* refers to social connections between different user accounts such as the friendship network on Facebook and the follower/retweet network on Twitter.

We also list the main methods used in *Single-view* and *Multi-view user embedding*. They typically employ unsupervised or self-supervised learning to automatically uncover the latent structures/relations in the raw data. To employ self-supervised user embedding, frequently an *Auxiliary Training Task* is employed for which the system can easily construct a large number of training examples. We will explain the details of these methods later. *Target task* describes the downstream applications that make use of the learned embeddings. We also indicate whether the learned embeddings are further tuned/adjusted so that they are optimized for the target tasks.

In the following, We first present the typical Single-view User Embedding methods. Then we summarize the methods that combine multiple types of user information together to form a unified user representation.

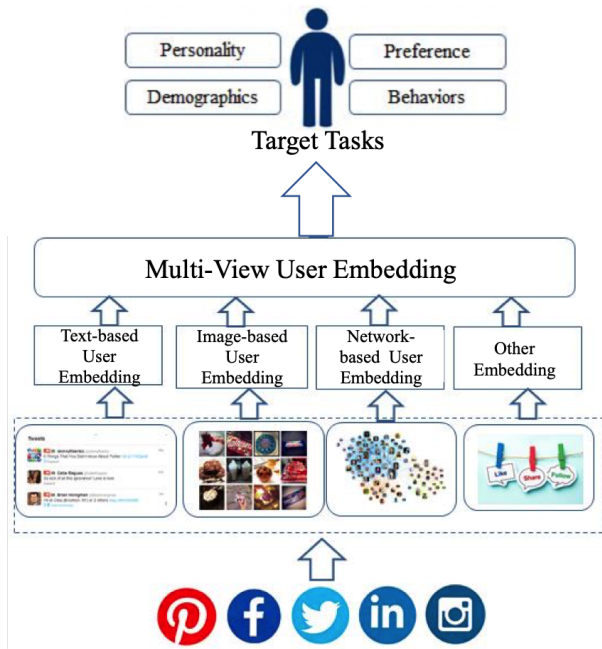


Figure 1: A typical system architecture

3 Single-View User Embedding

Since most papers in table 1 learn user embedding from text, we focus on text-based user embeddings. We will also discuss the typical methods used in learning user embeddings from social networks. Finally, we briefly describe how to learn user embeddings from other types of social media data such as likes and images.

3.1 Text-based User Embedding

The goal of text-based user embedding is to map a sequence of social media posts by the same user into a vector representation which captures the essential content and linguistic style expressed in the text. Here, we focus on methods that do not require any human annotations such as traditional unsupervised dimension reduction methods (e.g., *Latent Dirichlet Allocation* and *Single Value Decomposition*) or the more recent neural network-based prediction methods.

Latent Dirichlet Allocation (LDA)

LDA is a generative graphical model that allows sets of observations to be explained by unobserved latent groups. In natural language processing, LDA is frequently used to learn a set of topics (latent themes) from a large number of documents. Several methods can be used to derive a user representation based on LDA results:(1)User-LDA which treats all the posts from each user as a single document and trains an LDA model to drive the topic distribution for this document. The per-document topic distribution is then used as the representation for this user. (2) Post-LDA which treats each post as a separate document and trains an LDA model to derive a topic distribution for each post. All the per-post topic distribution vectors by the same user are aggregated (e.g., by averaging them) to derive the representation of each user. According to [Ding *et al.*, 2017], Post-LDA often learns better

| Paper | Data Type | Single-view embedding Method | Auxiliary Training Task | Multi-view embedding Task | Target Task | Target-task Tuning |
|--|---------------------------------|------------------------------|--|---------------------------|---|--------------------|
| [Pennacchiotti and Popescu, 2011] | text,user profile | LDA | NA | concatenation | political learning | No |
| [Kosinski <i>et al.</i> , 2013] | like | SVD | NA | NA | ethnicity user preference age,gender, personality | No |
| [Schwartz <i>et al.</i> , 2013] | text | LDA | NA | NA | relationship substance use religion political learning age,gender | No |
| [Gao <i>et al.</i> , 2014] | text | SVD | NA | NA | personality attribute | No |
| [Perozzi <i>et al.</i> , 2014] | network | DeepWalk | node prediction | NA | user interests | No |
| [Preoțiuc-Pietro <i>et al.</i> , 2015] | text, user profile | SVD | NA | concatenation | occupation | No |
| | | Word2Vec | word prediction | | | |
| [Song <i>et al.</i> , 2015] | text | LDA | NA | NA | interest | No |
| [Hu <i>et al.</i> , 2016] | text | LDA | NA | NA | occupation | No |
| [Song <i>et al.</i> , 2016] | text network user profile | LDA | NA | | volunteerism | Yes |
| [Benton <i>et al.</i> , 2016] | text | PCA | NA | CCA | topic engagement | No |
| | network | Word2vec PCA | word prediction NA | | friend recommendation age, gender political learning | |
| [Wallace <i>et al.</i> , 2016] | text | NA | word prediction | NA | sarcasm detection | Yes |
| [Ding <i>et al.</i> , 2017] | text like | Doc2vec Doc2vec | word prediction like prediction | CCA | substance use | No |
| [Preoțiuc-Pietro <i>et al.</i> , 2017] | text | Word2vec | word prediction | NA | political learning | No |
| | | LDA | NA | | | |
| [Amir <i>et al.</i> , 2017] | text | Word2vec | word prediction | NA | depression | No |
| [Zhang <i>et al.</i> , 2017] | network user profile | NA | NA | UPPSNE | gender | No |
| [Wang <i>et al.</i> , 2017] | network | NMF | NA | NA | education political learning community | NA |
| [Ding <i>et al.</i> , 2018a] | like | SVD, LDA Doc2vec | NA | NA | delay discounting | No |
| [Liao <i>et al.</i> , 2018] | network,user profile | NA | NA | SNE | group classification | No |
| [Zhang <i>et al.</i> , 2018a] | image | VGGNet | NA | NA | user image popularity | Yes |
| | text | LSTM | NA | | | |
| [Ribeiro <i>et al.</i> , 2018] | text | GloVe | NA | GraphSage | hateful user | No |
| | network | | | | | |
| [Zhang <i>et al.</i> , 2018b] | network user profile | NA | NA | ANRL | group classification | No |
| [Do <i>et al.</i> , 2018] | text network timestamp | Doc2vec Node2Vec | word prediction node prediction NA | concatenation | location | Yes |

Table 1: Summary of user embedding methods

user representations than User-LDA in downstream applications. This may be due to the fact that social media posts are often short and thus each post may only have a single topic,

which makes it easier for LDA to uncover meaningful topics than from one big document containing all the user posts.

Matrix Representation

Since we can also use a matrix to represent user-word and word-word co-occurrences, matrix optimization techniques are often used in learning user embeddings. If we use a sparse matrix to represent the relations between users and words where each row represents a user and each column represents a unique term in the dataset, we can use matrix decomposition techniques such as Singular Value Decomposition (SVD) and Principle Component Analysis (PCA) to yield a set of more manageable and compact matrices that reveal hidden relations and structures in the data (e.g., correlation, orthogonality and sub-space relations).

Recently, there is a surge of new text embedding methods that are designed to capture the semantics of words and documents. Except for **GloVe (Global Vectors for Word Representation)**, which uses matrix optimization to learn a general representation of words, most text embedding methods employ neural network-based methods. Since neural network-based methods are supervised learning methods, to learn user embeddings, we often need an auxiliary training task for which a large number of training examples can be automatically constructed from raw social media data. We called these methods *self-supervised machine learning*.

Word Embedding

Word2Vec is a popular neural network-based method designed to learn dense vector representations for individual words. The intuition behind the model is the Distributional Hypothesis, which states words that appear in the same context have similar meanings. There are two models for training a representation of word: Continuous Bag of Word (CBOW) and Skip Gram (SG) model. CBOW predicts a target word from one or more context words, while SG predicts one or more context words from a target word. Thus, predicting words in the neighborhood is the auxiliary task used to train word embeddings. The models are frequently trained using either a hierarchical softmax function (HS) or negative sampling (NS) for efficiency. To learn user embeddings from social media posts, the word2vec model is first applied to learn a vector representation for each word. Then a simple average of all the word vectors by the same user is used to represent a user [Benton *et al.*, 2016; Ding *et al.*, 2017].

GloVe is an unsupervised learning algorithm designed to learn vector representations of words based on aggregated global word-word co-occurrence statistics from a text corpus. GloVe employs a global log bi-linear regression model that combines the advantages of global matrix factorization with that of local context window-based methods. GloVe has been used in [Ding *et al.*, 2017] to learn a dense vector for each word. To summarize all the words authored by a user, we can use a vector aggregation function such as *average* to combine the vectors of all the words in a user's posts.

Document Embedding

Doc2Vec is an extension of Word2Vec, which produces a dense low dimensional feature vector for a document. There are two Doc2Vec models: Distributed Memory (DM) and Distributed Bag-of-Words (DBOW). Given a sequence of tokens in a document, DM can simultaneously learn a vector representation for each individual word token and a vector

for the entire document. In DM, each sequence of words (e.g. a document) is mapped to a sequence vector (e.g., document vector) and each word is mapped to a unique word vector. The document vector and one or more word vectors are aggregated to predict a target word in the context. DBOW learns a global vector to predict tokens randomly sampled from a document. Unlike DM, DBOW only learns a vector for the entire document. It does not use a local context window since the words for prediction are randomly sampled from the entire document.

There are two typical methods for learning a user embedding from doc2vec results: (1) User-D2V which combines all the posts by the same user in one document and trains a document vector to represent the user. (2) Post-D2V which treats each post as a document and train a doc2vec model to learn a vector for each post. To derive a user embedding, all the post vectors from the same person can be aggregated using "average".

Recurrent Neural Networks (RNN)

The text embedding methods described above ignore the temporal order of the words in a post and of the posts in a user account. Since the order of text contains important information, to capture the sequential relations between words and posts, Recurrent Neural Network (RNN) models such as *Long Short-Term Memory (LSTM)* can be used [Zhang *et al.*, 2018a]. The input to an LSTM is a sequence of word embeddings and the output of an LSTM is a sequence of hidden states, which are the input to downstream applications. Similar to word2vec, a language model-based auxiliary task is used to train a LSTM model on raw texts.

Among all the text embedding methods we discussed, some employ prediction-based technologies (e.g., Word2Vec, Doc2Vec and LSTM), others use count-based methods (e.g., PCA, SVD, LDA and GloVe). There are some empirical evidence indicating that prediction-based methods may have some advantage over count-based methods in feature learning [Baroni *et al.*, 2014]. Among all the text embedding methods we discussed, only LDA generates embeddings that are somewhat interpretable.

3.2 Social Network-based User Embedding

The objective of social network-based user embedding is to map very large social networks into low-dimensional embeddings that preserve local and global topological similarity. These methods focus primarily on learning a user representation that captures essential social structures and relations of a user. The three most widely used network embedding methods are *DeepWalk* [Perozzi *et al.*, 2014], *Node2vec* [Grover and Leskovec, 2016] and Matrix Factorization.

DeepWalk learns latent representations of vertices in a network from truncated random walks. It first generates short random walks. Each random walk $S = v_1, v_2, \dots, v_l$ is treated as a sequence of words in a sentence. DeepWalk then employs the SkipGram model (SG) in word2vec to learn the latent representation of a vertex. The learned embeddings can be used in many applications such as predicting user interests and anomaly detection [Perozzi *et al.*, 2014].

Node2Vec is a modification of DeepWalk which employs

a biased random walk to interpolate between Breadth-first Sampling (BFS) and Depth-first Sampling (DFS). With biased random walks, Node2vec can better preserve both the second-order and high-order proximity [Grover and Leskovec, 2016]. Given the set of neighboring vertices generated by a biased random walk, Node2Vec learns the vertex representation using the SkipGram model (SG). The learned embedding has been used to characterize a Twitter user’s social network structure [Do *et al.*, 2018] and predict user interests [Grover and Leskovec, 2016].

Non-Negative Matrix Factorization (NMF) is a matrix decomposition method with the additional constraint that all the entries in all the matrices have only positive values. The connections between network vertices are represented in an adjacency matrix. Non-negative matrix factorization is used to obtain a low-dimensional embedding of the original matrix. NMF was used in [Wang *et al.*, 2017] to learn a network-based user embedding that preserves both the first- and second-order proximity.

3.3 Other Single-View User Embedding Methods

In addition to texts and social networks, it is also possible to learn user embeddings from other types of social media data such as likes and images. For example, *User Like Embedding* was used in [Kosinski *et al.*, 2013; Ding *et al.*, 2018a] for personality and delay discounting prediction. Many text-based user embedding methods are also applicable here. For example, SVD was used in [Kosinski *et al.*, 2013]; LDA, GloVe, Word2Vec, Doc2vec were used in [Ding *et al.*, 2017; 2018a]. In addition, *AutoEncoder (AE)* can be used in learning like embeddings. AE is a neural network-based feature learning method [Hinton and Salakhutdinov, 2006]. It learns an identity function so that the output is as close to the input as possible. Although an identity function seems a trivial function to learn, by placing additional constraints (e.g., to make the number of neurons in the hidden layer much smaller than that of the input), we can still force the system to uncover latent structures in the data. Finally, *Image-based User Embedding* can be obtained by extracting low-level latent image features from pre-trained deep neural network models such as VGGNet [Simonyan and Zisserman, 2014].

4 Multi-View User Embedding

To obtain a comprehensive and unified user representation based on all the social media data available, we need to combine user features from different views together. In addition to simply concatenating features extracted from different views, we can also apply machine learning algorithms to systematically fuse them. We categorize these fusion methods into two types: (a) general fusion methods (b) customized fusion methods. General fusion methods can be applied to diverse types of embedding vectors such as text and image embedding or text and like embedding. In contrast, customized fusion methods are specifically designed to combine certain types of user data together. For example, ANRL is a method specifically designed to fuse user attributes and network topology together [Zhang *et al.*, 2018b].

4.1 General Fusion Methods

First, we introduce two widely used general fusion methods.

Canonical Correlation Analysis (CCA) CCA is a statistical method that explores the relationships between two multivariate sets of variables (vectors) [Hardoon *et al.*, 2004]. Given two feature vectors, CCA tries to find a linear transformation of each feature vector so that they are maximally correlated. CCA has been used in [Sharma *et al.*, 2012; Ding *et al.*, 2017] for multi-view fusion.

Deep Canonical Correlation Analysis (DCCA) DCCA is a non-linear extension of CCA, aiming to learn highly correlated deep architectures [Andrew *et al.*, 2013]. The intuition is to find a maximally correlated representation of two feature vectors by passing them through multiple stacked layers of nonlinear transformation. Typically, there are three steps in training DCCA: (1) using a denoising autoencoder to pre-train each single view; (2) computing the gradient of the correlation of top-level representation; (3) tuning parameters using back propagation to optimize the total correlation.

The features learned from multiple views are often more informative than those from a single view. Comparing with single-view user embedding, multi-view embedding achieved significantly better performance in predicting demographics [Benton *et al.*, 2016], politic leaning [Benton *et al.*, 2016] and substance use [Ding *et al.*, 2017].

4.2 Customized Fusion Methods

Several studies in our survey employ algorithms that are specifically designed to combine certain types of data. For example, [Zhang *et al.*, 2017] proposed a new algorithm called User Profile Preserving Social Network Embedding (UPPSNE), which combines user profiles and social network structures to learn a joint vector representation of a user. Similarly, Attributed Network Representation Learning (ANRL) [Zhang *et al.*, 2018a] employed a deep neural network to incorporate information from both network structure and node attributes. It learns a single user representation that jointly optimizes AutoEncoder loss, SkipGram loss and Neighbour Prediction Loss. [Liao *et al.*, 2018] proposed a Social Network Embedding framework (SNE), which learns a combined representations for social media users by preserving both structural proximity and attribute proximity. [Ribeiro *et al.*, 2018] creates embeddings for each node with word embeddings learn from text using GloVe and the activity/network-centrality attributes associated with each user. So far, most of the customized fusion methods are designed to fuse network topology with additional node information (e.g., user profiles).

5 Embedding Fine Tuning Using Target Tasks

In many cases, the learned user embeddings are simply used as the input to a target task. It is also possible that the learned user embeddings can be further refined to better support the target tasks with supervised learning. For example, in [Miura *et al.*, 2017], the authors propose an attention-based neural network model to predict geo-location. It simultaneously learns text, network and metadata embeddings in supervised fine turning. In [Song *et al.*, 2016], the authors collected

multi-view user data from different platforms (e.g., Twitter, Facebook, LinkedIn of the same user) and predicted volunteerism based on user attributes and network features. Then, they combine these two sets of features in supervised fine tuning to enhance the final prediction. [Farnadi *et al.*, 2018] learned a hybrid user profile which is a shared user representation learned from three data sources. They were combined at decision level to predict multiple user attributes (e.g., age, gender and personality).

6 Discussion

Large-scale social media-based user trait and behavior analysis is an emerging multidisciplinary field with the potential to transform human trait and behavior analysis from controlled small scale experiments to large scale studies of natural human behavior in an open environment. Although raw social media data are relatively easy to obtain, it is expensive to acquire the ground truth data at a large scale. The proposed unsupervised/self-supervised user embedding methods can alleviate the “small data” problem by transferring knowledge learned from raw social media data to a new target task. This step is very important for many human trait and behavior analysis applications. According to [Benton *et al.*, 2016; Preoțiuc-Pietro *et al.*, 2015], machine learning models that incorporate unsupervised/self-supervised user embedding significantly outperform the baselines that do not employ user embedding. Based on the survey, we have also identified a few major issues in the current social media analysis research.

6.1 Interpretability

Although systems employing user embeddings significantly outperform baselines in terms of prediction accuracy, these systems also suffer from one significant drawback: low interpretability. Since user embedding features are latent features automatically uncovered by the system, it is often difficult for humans to understand the meaning of these features. This may significantly impact our ability to gain insight into these behavioral models. So far, there have not been much work focusing on learning user representations that are both effective and interpretable.

6.2 Ethical Issues

Due to the privacy concerns in accessing user data on social media and the sensitive nature of the inferred user characteristics, if not careful, there could be significant privacy consequences and ethical implications. So far, most of the studies in our survey focused primarily on optimizing system performance. There have not been sufficient discussion on ethical concerns when conducting research in this field.

7 Future Directions

Each of the main issues we identified above also presents a good opportunity for future work.

7.1 Interpretable User Representation Learning

we need more research on learning high-performance user representations that are also interpretable. Some preliminary work has conducted in this area. In [Ding *et al.*, 2018b], a

knowledge distillation framework was proposed to train behavior models that are both highly accurate and interpretable. Developing causal models for both inference and interpretation is another potential new direction.

7.2 Ethical Research on Data-driven Behavior Analysis

Ethical issues are complex, multifaceted and resist simple solutions. In addition to privacy concerns in data collection, researchers working on social media-based human trait and behavior analysis also face other ethical challenges including informed consent, traceability and working with children and young people. There is an urgent need for the research community to decide an ethical framework to guide researchers to navigate obstacles, gain trust and still allow them to capture salient behavioral and social phenomena. Recently there is a surge of interests and research on fair data-driven decision making. As a researcher, we also need to be aware of the potential impact of social media analytics on the well-being of individuals and our society.

We have also identified a few possible research directions to improve the state of the art user embedding techniques.

7.3 Temporal User Embedding

since most social media data are collected over a long period of time and associated with time stamps, it is an ideal data source for longitudinal data analysis. Also, for many medical and public health applications, analyzing behavioral changes over time is critical to understanding one’s decision making process. Although Recurrent Neural Networks such as LSTM can capture some sequential patterns, they totally ignore the time stamp associated with each event. More work on learning user embedding from time is needed.

7.4 User Embedding with Multi-task Learning

Since individual traits and behavior are highly correlated, building a prediction model that simultaneously infer multiple correlated traits and behavior should yield better performance than predicting each trait/behavior separately. Most existing studies only predict one user attribute/behavior at a time. More research is needed to jointly train and predict multiple user attributes together for better performance.

7.5 Cross-platform Fusion

It is also common for a user to have multiple accounts on different social media platforms. Recently, new technologies have been developed to link different social media accounts of the same user together [Abel *et al.*, 2013]. With this linked data, it is possible to perform novel cross-platform user trait and behavior analysis such as (1) domain bias analysis that focuses on studying the impact of domain or social media platform on user trait and behavior analysis, (2) domain adaptation that addresses how to adjust prediction models trained on one platform (e.g., Twitter) to predict the traits and behavior on another platform (e.g., Facebook). So far, there is some initial work on domain bias analysis and correction [Kiliç and Pan, 2016]. More research is needed in order to develop more robust tools for human trait and behavior analysis.

References

- [Abel *et al.*, 2013] F. Abel, E. Herder, G. Houben, N. Henze, and D. Krause. Cross-system user modeling and personalization on the social web. *UMUAI*, 2013.
- [Amir *et al.*, 2017] Silvio Amir, Glen Coppersmith, Paula Carvalho, Mario J Silva, and Bryon C Wallace. Quantifying mental health from social media with neural user embeddings. In *Machine Learning for Healthcare Conference*, 2017.
- [Andrew *et al.*, 2013] G. Andrew, R. Arora, J. Bilmes, and K. Livescu. Deep canonical correlation analysis. In *ICML*, 2013.
- [Baroni *et al.*, 2014] M. Baroni, G. Dinu, and G. Kruszewski. Don't count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *ACL*, 2014.
- [Bengio *et al.*, 2013] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(8), 2013.
- [Benton *et al.*, 2016] A. Benton, R. Arora, and M. Dredze. Learning multiview embeddings of twitter users. In *ACL*, 2016.
- [De Choudhury *et al.*, 2013] Munmun De Choudhury, Michael Gamon, Scott Counts, and Eric Horvitz. Predicting depression via social media. In *ICWSM*, 2013.
- [Ding *et al.*, 2017] T. Ding, W. Bickel, and S. Pan. Multi-view unsupervised user feature embedding for social media-based substance use prediction. In *EMNLP*, 2017.
- [Ding *et al.*, 2018a] Tao Ding, Warren K Bickel, and Shimei Pan. Predicting delay discounting from social media likes with unsupervised feature learning. In *ASONAM*, 2018.
- [Ding *et al.*, 2018b] Tao Ding, Fatema Hasan, Warren K Bickel, and Shimei Pan. Interpreting social media-based substance use prediction models with knowledge distillation. In *ICTAI*, 2018.
- [Do *et al.*, 2018] Tien Huu Do, Duc Minh Nguyen, Evaggelia Tsiligianni, Bruno Cornelis, and Nikos Deligiannis. Twitter user geolocation using deep multiview learning. In *ICASSP*, 2018.
- [Farnadi *et al.*, 2018] Golnoosh Farnadi, Jie Tang, Martine De Cock, and Marie-Francine Moens. User profiling through deep multimodal fusion. In *WSDM*, 2018.
- [Gao *et al.*, 2014] H. Gao, J. Mahmud, J. Chen, J. Nichols, and Michelle X. Zhou. Modeling user attitude toward controversial topics in online social media. In *ICWSM*, 2014.
- [Golbeck *et al.*, 2011] Jennifer Golbeck, Cristina Robles, and Karen Turner. Predicting personality with social media. In *CHI*, 2011.
- [Grover and Leskovec, 2016] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *KDD*, 2016.
- [Hardoon *et al.*, 2004] D. Hardoon, S. Szedmak, and J. Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural computation*, 16(12), 2004.
- [Hinton and Salakhutdinov, 2006] G. Hinton and R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *science*, 313(5786), 2006.
- [Hu *et al.*, 2016] T. Hu, H. Xiao, J. Luo, and T. Nguyen. What the language you tweet says about your occupation. In *ICWSM*, 2016.
- [Kiliç and Pan, 2016] D. Kiliç and S. Pan. Analyzing and preventing bias in text-based personal trait prediction algorithms. In *ICTAI*, 2016.
- [Kosinski *et al.*, 2013] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *PNAS*, 110(15), 2013.
- [Le and Mikolov, 2014] Quoc Le and Tomas Mikolov. Distributed representations of sentences and documents. In *ICML*, 2014.
- [Liao *et al.*, 2018] Lizi Liao, Xiangnan He, Hanwang Zhang, and Tat-Seng Chua. Attributed social network embedding. *TKDE*, 2018.
- [Liu *et al.*, 2016] L. Liu, D. Preotiuc-Pietro, Z. Samani, M. Moghaddam, and L. Ungar. Analyzing personality through social media profile picture choice. In *ICWSM*, 2016.
- [Mikolov *et al.*, 2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.
- [Miura *et al.*, 2017] Yasuhide Miura, Motoki Taniguchi, Tomoki Taniguchi, and Tomoko Ohkuma. Unifying text, metadata, and user network representations with a neural network for geolocation prediction. In *ACL*, 2017.
- [Pennacchiotti and Popescu, 2011] M. Pennacchiotti and A. Popescu. A machine learning approach to twitter user classification. *ICWSM*, 2011.
- [Pennebaker *et al.*, 2015] James W Pennebaker, Ryan L Boyd, Kayla Jordan, and Kate Blackburn. The development and psychometric properties of liwc2015. Technical report, 2015.
- [Perozzi *et al.*, 2014] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *KDD*, 2014.
- [Preotiuc-Pietro *et al.*, 2015] D. Preotiuc-Pietro, V. Lampos, and N. Aletras. An analysis of the user occupational class through twitter content. In *ACL*, 2015.
- [Preotiuc-Pietro *et al.*, 2017] D. Preotiuc-Pietro, Y. Liu, D. Hopkins, and L. Ungar. Beyond binary labels: political ideology prediction of twitter users. In *ACL*, 2017.
- [Ribeiro *et al.*, 2018] Manoel Horta Ribeiro, Pedro H Calais, Yuri A Santos, Virgílio AF Almeida, and Wagner Meira Jr. Characterizing and detecting hateful users on twitter. In *AAAI*, 2018.
- [Schwartz *et al.*, 2013] A. Schwartz, J. Eichstaedt, M. Kern, L. Dzurzynski, S. Ramones, M. Agrawal, A. Shah, M. Kosinski, D. Stillwell, M. Seligman, et al. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PLoS one*, 8(9), 2013.
- [Sharma *et al.*, 2012] A. Sharma, A. Kumar, H. Daume, and D. Jacobs. Generalized multiview analysis: A discriminative latent space. In *CVPR*, 2012.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014.
- [Song *et al.*, 2015] X. Song, L. Nie, L. Zhang, M. Liu, and T. Chua. Interest inference via structure-constrained multi-source multi-task learning. In *IJCAI*, 2015.
- [Song *et al.*, 2016] X. Song, Z. Ming, L. Nie, Y. Zhao, and T. Chua. Volunteerism tendency prediction via harvesting multiple social networks. *TOIS*, 2016.
- [Wallace *et al.*, 2016] Silvio Amir Byron C Wallace, Hao Lyu, and Paula Carvalho Mário J Silva. Modelling context with user embeddings for sarcasm detection in social media. *CoNLL*, 2016.
- [Wang *et al.*, 2017] Xiao Wang, Peng Cui, Jing Wang, Jian Pei, Wenwu Zhu, and Shiqiang Yang. Community preserving network embedding. In *AAAI*, 2017.
- [Yang *et al.*, 2015] C. Yang, S. Pan, J. Mahmud, H. Yang, and P. Srinivasan. Using personal traits for brand preference prediction. In *EMNLP*, 2015.
- [Zeng *et al.*, 2018] Ziqian Zeng, Xin Liu, and Yangqiu Song. Biased random walk based social regularization for word embeddings. In *IJCAI*, 2018.
- [Zhang *et al.*, 2017] Daokun Zhang, Jie Yin, Xingquan Zhu, and Chengqi Zhang. User profile preserving social network embedding. In *IJCAI*, 2017.
- [Zhang *et al.*, 2018a] Wei Zhang, Wen Wang, Jun Wang, and Hongyuan Zha. User-guided hierarchical attention network for multi-modal social image popularity prediction. In *WWW*, 2018.
- [Zhang *et al.*, 2018b] Zhen Zhang, Hongxia Yang, Jiajun Bu, Sheng Zhou, Pinggang Yu, Jianwei Zhang, Martin Ester, and Can Wang. Anrl: Attributed network representation learning via deep neural networks. In *IJCAI*, 2018.