

affinity: A System for Latent User Similarity Comparison on Texting Data

Tobias Eichinger, Felix Beierle

Service-centric Networking

Telekom Innovation Laboratories

Technische Universität Berlin

Berlin, Germany

{tobias.eichinger,beierle}@tu-berlin.de

Sumsam U. Khan, Robin Middelanis,
Veeraraghavan Sekar, and Sam Tabibzadeh

Technische Universität Berlin

Berlin, Germany

Abstract—In the field of social networking services, finding similar users based on profile data is common practice. Smartphones harbor sensor and personal context data that can be used for user profiling. Yet, one vast source of personal data, that is text messaging data, has hardly been studied for user profiling. We see three reasons for this: First, private text messaging data is not shared due to their intimate character. Second, the definition of an appropriate privacy-preserving similarity measure is non-trivial. Third, assessing the quality of a similarity measure on text messaging data representing a potentially infinite set of topics is non-trivial. In order to overcome these obstacles we propose *affinity*, a system that assesses the similarity between text messaging histories of users reliably and efficiently in a privacy-preserving manner. Private texting data stays on user devices and data for comparison is compared in a latent format that neither allows to reconstruct the comparison words nor any original private plain text. We evaluate our approach by calculating similarities between Twitter histories of 60 US senators. The resulting similarity network reaches an average 85.0% accuracy on a political party classification task.

Index Terms—Latent Features, Social Networking, User Similarity, User Profiling, Document Similarity, Word Mover’s Distance, Word Embedding, tf-idf

I. INTRODUCTION

Existing Social Networking Services (SNSs) are commonly used to connect people one already knows. They often suggest new contacts based on user profile data or the social graph. Especially with the proliferation of smartphones and the popularity of location-based services, there is also a growing interest in more dynamic scenarios. Combining the ideas of ubiquitous computing and social networking, it is possible to connect to similar and close by people who not necessarily are close in the social graph [1], [2]. In the field of psychology, the principle of *homophily* states that individuals tend to bond if they are similar. This tendency also holds true within social networks services employing digital communication [3].

Nowadays, a large portion of digital communication happens on smartphones which harbor highly personalized sensor and personal context data such as visited locations, music listened to, or app usage statistics [4]. Profiling users upon this data allows to estimate user similarities [3], [5]. A very commonly used – yet virtually never mined – resource on

smartphones for user profiling is text messaging¹ or simply *texting*. When texting, users will cover topics of their interest, include personal information on their background and use individual vocabulary such as slang and emoticons. In particular, there is scientific evidence that personality is projected and perceived through digital communication [6]. Hence, people with similar personality traits will tend to write similarly, which makes it eligible as a base for similarity estimation. However, a definition of similarity is not straightforward since most profiling approaches involve a full view on the data to derive similarity. Even if private texting data is promising as a meaningful basis for user profiling and similarity estimation, to the best of our knowledge, it essentially is uncharted terrain.

We see three main problems: First, private text messaging data is not shared due to its intimate character. Second, the definition of an appropriate privacy-preserving similarity measure on texting data is non-trivial. Third, assessing the quality of a similarity measure on text messaging data representing a potentially infinite set of topics is non-trivial. In particular, it is unclear which similarity benchmark to use. Should we evaluate on single benchmarks such as personality traits or user demographics such as gender, age, and political affiliation, or all together? In order to tackle these three problems we propose *affinity*, a system for comparing user similarity on latent, that is in particular incomprehensible for humans, word vectors derived from a word embedding trained on texting data. We estimate similarity with the Word Mover’s Distance (WMD) [7], which allows to measure the distance between two sets of word vectors independently of any benchmark. Even though *affinity* employs a central server harboring minimal amounts of global knowledge between users, similarity comparisons are performed decentrally, that is the central server neither knows the basis for comparison, nor the similarities between users, nor who compared with whom. In order to evaluate *affinity*, we propose an evaluation scheme in the form of a political party classification task in order to measure *affinity*’s performance. We use Twitter data of US politicians for which *affinity* derives pairwise similarities. Finally, the pairwise similarities are translated into a similarity network

¹<http://www.pewinternet.org/2015/04/01/us-smartphone-use-in-2015/>

which can be interpreted as knowledge-based nearest neighbor machine learning classifier. Our main contributions are:

- Design and implementation of *affinity*, a system which leverages private texting data for real-time user similarity comparison, while preserving privacy.
- Evaluation of *affinity*'s performance on Twitter data of 60 US politicians. *affinity* reaches an average accuracy of 85.0% on a political party classification task.

II. RELATED WORK

We see two main areas of related work converge in the field of social networking via user profiling.

Profiling on smartphone data: Previous works have proposed paradigms for finding similar people in one's vicinity leveraging sensor and context data. One of the first works utilizing mobile phones for this purpose follows the idea of exchanging identifiers of existing Online Social Networks (OSN) via Bluetooth in order to view public profiles of nearby users [8]. Other works enable users to broadcast short profiles or manually entered topics of interest over a short distance [2], [9], [10]. There are also automated approaches, which utilize personal context data in the form of user profiles for comparison [1], [5], [11].

Profiling on natural language resources: In the field of Natural Language Processing (NLP), extracting a user's writing style and selection of topics which reflect the author's demographic and psychological properties is referred to as *author profiling*. Author profiling includes classifying an author's gender or demographic information and has been largely exercised on Twitter data. Similarities to our proposed approach involve the usage of term and document weights in combination with Part-of-Speech (POS) n-grams [12], and word embedding features combined with logistic regression [13] or with a multimodal model for gender classification [14]. Other related work aims at constructing thematic user profiles by clustering word embedding features of topic words for full-text item recommendation [15]. Predicting topics of interest or personal opinions on tweets has been covered in [16]. A very recent work proposes to simultaneously and dynamically embed users and words even for streams of tweets [17].

In the present paper, we propose an approach based on NLP techniques in order to profile users on private texting data in a social networking scenario such that private data stays on the user device.

III. FUNDAMENTALS

A. From Documents to Users

Text is sequential data, that is *one-dimensional* and *directed*. Given an unfinished sentence such as

An apple is ...

we may guess the next word in the sentence. The sentence may be continued in a meaningful way with the words *edible*, *green*, or *big*, yet less so with the word *blue*. We have guessed possible subsequent words to the sentence upon the preceding 3 words (3-gram) in the sentence. Since some words are more

likely to follow than others, we see that text production can be interpreted as a (conditional) random process depending on close words.

Furthermore, text documents consist of interdependent sentences. Producing sentences, analogously to producing a subsequent word, is not an arbitrary process. In the related fields of text classification, text categorization, and image retrieval, topic models assume that documents have an underlying conceptual *topic* [18]–[21] that can be described by a set of characteristic words. For example, when writing on the citric acid cycle, the probability of producing the term *adenosine triphosphate* will be more probable than for example when writing a fairy tale. Thus the *topicality* of a text document influences the writing process.

In the following we will view a user's own part of his/her texting history as a distinct document. Most users have a diverse set of people they text with on probably an even more diverse set of topics. Hence texting histories will not harbor a single topic yet rather be a mixed lot of topics. We will call this mixture a person's *state of mind*. We construe a user's state of mind as weighted yet discriminative set of words. We then consider people similar if they have similar states of mind, and else dissimilar. Since texting histories in plain text certainly are of private character, we will use word representations which neither allow to reconstruct the comparison words nor any original private plain text.

B. On Document Similarity

Many search engines and information retrieval techniques revolve around the isolation of characteristic words which yield much information on the topic of a search query or a document. Virtually unlimited possibilities to rephrase any given sentence into a topically and logically equivalent sentence makes it hard to exclusively map words to topics. Many words will be shared between topics. For example the occurrence of the word *adenosine triphosphate* does not imply that the topic of the text document necessarily involves the citric acid cycle. Vice versa its absence does not imply that the topic of the text document excludes the citric acid cycle. In order to tackle this problem of uncertainty, NLP techniques provide means to map single words or entire text documents into semantic space. In semantic space, a word or a text document is represented by a vector wherein commonly the cosine distance is used to measure distances. However, there are many different means of measuring similarity between documents. See for instance [22] for a survey on traditional approaches. More recently, with the advent of powerful word embedding models, a new method of estimating document similarity upon latent word features has emerged, which we deem to be fit for texting histories.

C. On Word Embeddings

Word embeddings, also known as Vector Space Models (VSM) [23], map words to vectors in real space by evaluating distributional context information. The *context* of a word are other words in its proximity such as for example the two

words to the left and right. Originally being high dimensional and sparse by nature, recent low-dimensional and dense word embedding algorithms [24]–[26] have proven to be similar in performance [27]. State-of-the-art word embeddings trained on very large text corpora capture many semantic and syntactic similarities between words [28], yet there does not seem to be a *perfect word embedding* being able to capture similarities in a best way [29]. There are also indications that domain-specific embeddings can perform better on domain-specific tasks [30], which would suggest that word embeddings trained on corpora derived from for instance SMS messages perform better on linguistic tasks on SMS messages.

D. Word Mover’s Distance

The Word Mover’s Distance (WMD) introduced by Kusner et al. [7] is an application of the Earth Mover’s Distance² (EMD) to word embeddings. The EMD measures the cost of moving earth from i points to j points in real space and may thus be reformulated as a transportation problem. Among others, the EMD has been traditionally used to compare grey-scale and color histograms of images for image comparison [31], [32] and image retrieval [33]. Yet Kusner et al. [7] apply the EMD to word embeddings in order to compare sentences such as (A): “*Obama speaks to the media in Illinois.*” and (B): “*The President greets the press in Chicago.*”. They do so by dropping *stop words* (non-bold), which are very common and thus do not convey content information, and afterwards moving the vector representations of the remaining words in sentence (A) to those in sentence (B). Word embeddings allocate word vectors of semantically similar words close to each other, for example the word vectors of the words *Obama* and *President* are close to each other. Thus, when we move one vector to the other, the cost is low / the distance is short. Now, when the words of two sentences are pairwise semantically similar, their pairwise costs are low / distances are short. The sentences are then assumed to convey messages on similar topics³. Instead of sentences, we will apply the WMD to entire texting histories. Using an entire texting history’s vocabulary is computationally intractable. The WMD scales super cubic on the number of points to move [34] and even small texting histories easily exceed a thousand vocabulary words. In order to circumvent this problem, we pick a small set of *relevant* words for each and every user as a condensate of their states of mind, and compare similarity on this limited set of words. Note that the WMD has to be equipped with a designated distance function. Throughout our experiments, we employ the cosine distance⁴ since it is naturally bounded and commonly used with word vectors. For formal information on the underlying transportation problem and efficient solvers we refer to [7] and the references therein.

²more generally: *Wasserstein distance*

³Note that in general we cannot expect that the messages convey the same information.

⁴ $\cos(x, y) = x \cdot y / \|x\| \|y\|$, where \cdot denotes the dot product.

E. tf-idf Features

Term frequency–inverse document frequency (tf–idf) features are commonly used as relevance measures for text retrieval [35], [36]. It follows the rationale that frequently used words *within* a document are important, yet they are only relevant if they are simultaneously less frequent *over all* documents. We define tf–idf features for a term t and a document d by

$$\begin{aligned} \text{tf-idf}(t, d) &= \text{tf}(t, d) \cdot \text{idf}(t, d) \\ &= \frac{f_{t,d}}{\max_{x \in \mathcal{D}} f_{x,d}} \cdot \log \frac{|\mathcal{D}|}{|\{d \in \mathcal{D} : t \in d\}|}, \end{aligned} \quad (1)$$

where $f_{t,d}$ is the raw count of occurrences of the term t in the document d and \mathcal{D} is a collection of documents. Note that there are alternative definitions of tf–idf features [37]. We choose this definition in order to mitigate a bias toward long documents.

IV. CONCEPT

We will describe *affinity’s* similarity comparison along three main blocks, that is text selection, latent feature pre-calculation, and lastly similarity comparison upon request. For a seamless user experience, we designed an extended pre-calculation procedure (see Figure 1) paired with quick similarity comparisons.

A. Text Selection

The user selects a texting history in plain text form upon which he/she wants to be compared with other users. All messages that are not self-written are dropped. We call the resulting file *reference file*. It is stored on the user’s smartphone.

B. Latent Feature Pre-Calculation

The reference file is transformed into a Bag-of-Words (BoW) vector, which is a dictionary with keys of unique vocabulary words and values of word occurrence counts. The BoW vector is used to derive a user document frequency (DF) vector which shares the dictionary keys of the BoW vector, yet only holds ones in the values. The user DF vector is submitted to the backend, hence the backend only sees that the user has used a certain word yet not how often. The user DF vector updates the global DF vector that stores all words ever submitted and the number of their respective users. After updating the global DF vector, a truncated DF vector corresponding to the words submitted by the user, and the total number of users is returned to the smartphone. The user can now calculate tf–idf features for all of his/her used words according to Equation (1), which are then added to his/her BoW vector. The WMD calculates distances between two (word vectors, word weights) pairs, which are commonly referred to as *signature*. Note that the WMD calculates distances on a word level and does only consider structural and distributional information such as word order that is condensed in the central word embedding. We see that the quality of similarity comparison hinges on the quality of the underlying word embedding.

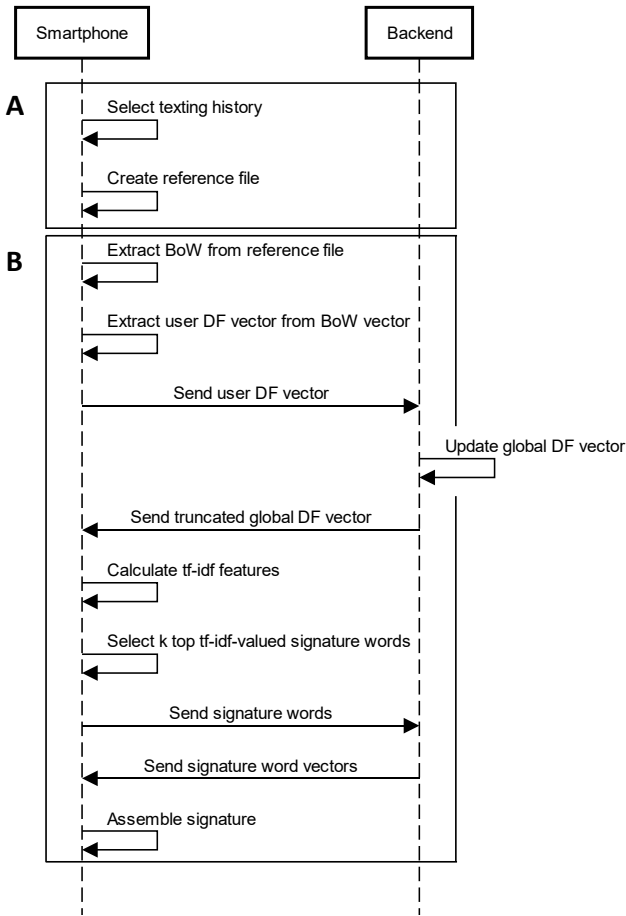


Fig. 1. Message flow diagram of *affinity*'s pre-calculation procedure before similarity comparison. The smartphone holds the plain texting history, and its signature. The backend holds the global DF vector and a reference word embedding model.

In order to create user signatures, the top k tf-idf valued words are submitted to the backend⁵. The backend derives the submitted signature words' word vectors via some secret reference word embedding model and sends them to the user's smartphone⁶. The signature word vectors (as a $k \times D$ matrix) and their corresponding normalized BoW vector (of length k) form the user's signature, where D is the embedding dimension of the reference word embedding model.

Example: Signature Calculation

Let $k = 2$ be the signature size and $D = 3$ the dimension of the word embedding. Let T be a texting history of some user A . In T , we drop all messages that have not been written by A resulting in some reference file R . Let V be the vocabulary of R and $\#V$ its vocabulary size. We can now derive the user BoW vector $(\#v)_{v \in V}$, where $\#v$ is the number of occurrences of the vocabulary word v in R . The user DF vector can then

⁵The backend sees the user's k words to compare similarity on. Submitting randomly selected additional words can sufficiently obfuscate this insight.

⁶Techniques from differential privacy can be used to add jitter to the true word vectors to obfuscate knowledge on the hidden word embedding model.

be formalized as $DF_{\text{user}}(A) = [1]_{v \in V}$. The user DF vector is sent to the backend, where it updates the global DF vector

$$DF_{\text{global}} = \sum_{u \in U} DF_{\text{user}}(u), \quad (2)$$

where U is the set of registered users. We truncate the global DF vector along the vocabulary V , that is we only keep the entries associated to vocabulary words in V :

$$DF_{\text{trunc}}(A) = DF_{\text{global}}|_{v \in V}. \quad (3)$$

The truncated global DF vector together with the number of total users $\#U$ are sent back to the smartphone. We can now calculate the tf-idf values for all vocabulary words in V according to Equation (1). We then pick the top 2 tf-idf valued vocabulary words, for instance *country* and *nation*, and send them to the backend. The backend prompts two associated 3-dimensional vectors $v_{\text{country}} = [-0.25 \ 0.5 \ 0.75]$ and $v_{\text{nation}} = [-0.23 \ 0.51 \ 0.6]$. We find in the user BoW vector that R contains 5 and 6 occurrences of the words *country* and *nation* respectively, which results in the normalized BoW weight vector $[5/11, 6/11]$. User A 's signature is thus:

$$\text{signature}(A) = \left[\begin{array}{c} [-0.25 \ 0.50 \ 0.75] \\ [-0.23 \ 0.51 \ 0.60] \end{array}, \begin{array}{c} [5/11] \\ [6/11] \end{array} \right].$$

□

We point out that user signatures, which are the basis for similarity comparison, are snapshots. As both individual texting histories and the global DF vector evolve over time, signatures will have to be updated once in a while. Note that replacing the central word embedding only requires an update of signature vectors by all users, independent of the global DF vector.

C. Similarity Comparison

Any two users who are ready for comparison can share their signatures locally without the backend knowing, and calculate their pairwise similarity via the WMD on their smartphones. Note that comparison partners cannot view their partner's comparison words since they have been converted into latent feature vectors. Thus, signatures can be shared at will as long as the reference word embedding model remains secret.

V. IMPLEMENTATION

We use the technology bundle *nginx*, *unicorn*, and *Flask* to set up *affinity* as a Web-application⁷ written in Python on an Ubuntu backend. We will restrict ourselves to the *fasttext* word embedding algorithm [24] as reference word embedding model for three reasons. First, *fasttext* takes into account similarities and dissimilarities between morphologically related words (for example *entangle*, *entangled*, *disentangled*). Second, it embeds *out-of-vocabulary* (OOV) words, that is words that have not been seen during training time. Third, it performs language identification in order to – for example – filter out non-English conversations.

⁷Find *affinity*'s code at: <https://github.com/moxplayer/affinity>.

Note that in this prototype setup, in contrast to *affinity*'s concept, all data processing takes place on the backend, whereas the concept also involves processing on the user device. In Section VI, we will prove that *affinity*'s concept is realizable as conceived in Section IV.

A. Cold Start Problem

The proposed concept requires an eco-system of users who have submitted their DF vectors to the backend, because we use tf-idf word relevance measurements. For example, if there were exactly two registered users, then the idf part of Equation (1) would either be zero for words both users submitted or $\log(2)$ else. This contradicts the initial purpose of finding relevant words since the word they have in common is evaluated to have zero relevance. Therefore, we will require a sufficiently large initial user base for the tf-idf features to work properly. For this purpose, we use the NUS SMS chat corpus⁸ [38], which contains about 40,000 English and 30,000 Mandarin Chinese SMS messages with mainly colloquial bits of conversations. We convert unique users' English messages into *seed documents*. We then concatenate all seed documents into a big corpus and train an initial *fasttext* reference model with default settings (in particular embedding dimension $D = 100$), which is then kept secret. The *fasttext* model then reflects linguistic properties of short colloquial messages.

B. Text Selection

Instead of submitting personal texting histories, the backend fetches Facebook Messenger histories excluding messages not written by the user. We use a pre-trained *fasttext* model⁹ for language classification. We drop conversations where the majority of words is non-English. The result is then appended to the user's reference document.

C. Latent Feature Pre-Calculation

The user's BoW vector excluding stop words is updated. Then the global DF vector is updated. In order to avoid comparison on words containing typos or user-related artifacts, we exclude words from comparison if they occur in less than $p_{min} = 0.05$ reference documents which we found to perform well. For the remainder of the words, we update the tf-idf values in the BoW vector. We select the top $k = 50$ tf-idf valued words as signature vectors and generate signature word vector representations via the initially trained *fasttext* reference model. The signature word vectors together with their corresponding normalized BoW vector form the user's signature. Note that the size of a user signature is approximately 10 kB for the signature size $k = 50$ on half-precision floats.

D. Similarity Comparison

When a user scans another user's id encoded as a QR code, both user ids are sent to the backend, their pre-calculated signatures are passed to the WMD. The result distance score

is returned to both users and displayed on their smartphones. Note that the Word Mover's Distance (WMD) equipped with the cosine distance is bound between zero and two, where zero denotes perfect similarity and two perfect dissimilarity. Therefore, we renormalize as follows

$$\text{sim}(id_1, id_2) = 1 - \frac{1}{2} \text{WMD}((M_1, B_1), (M_2, B_2)), \quad (4)$$

where WMD denotes the Word Mover's Distance equipped with the cosine distance, M_i and B_i denote the i -th id's signature word vector matrix and the normalized BoW vector respectively. Now similarity scores are bound between zero and one with one denoting perfect similarity and zero perfect dissimilarity.

VI. EVALUATION

A. Datasets

We used the NUS SMS corpus in order to derive necessary seed users in order to be able to apply tf-idf relevance measures in a meaningful way. Furthermore, we used its concatenated English messages in order to train a *fasttext* word embedding model, which is able to reflect linguistic properties of language used in short and colloquial messages.

Tweets differ from SMS messages in for example their average words per message, the fact that SMS messages are likely to contain only a single word, or that tweets are publicly available and therefore often contain more structured and non-sensitive information [38]. However, their short character limitations to 280 and 160 respectively create a common ground for shortened forms and abbreviations of words, or pictograms in both formats. Furthermore, since we use users' entire histories instead of single (potentially single-word) messages, we deem the impact of structural differences between tweet and SMS histories negligible in our scenario. Hence, we posit that tweet histories contain linguistic information similar to that of SMS histories or histories from any other (short) messaging service such as *Facebook Messenger* or *WhatsApp*.

For evaluation, we use plain text tweets – excluding meta-data such as timestamps, number of retweets, or likes – published until mid September 2018 by 60 US senators (30 Republicans, 30 Democrats) of the 115th US Congress^{10,11}. We drop stop words and web links.

B. Performance

affinity is designed to pre-calculate user signatures, that is a numerical condensate of linguistic information hidden in private texting data. We already know that the WMD scales super cubic in the signature size k . We benchmark calculation times of the *affintiy* prototype on a single CPU core (*i7-3520M*). The results are shown in Figure 2. We observe that the comparison time decreases to less than a second for $k \leq 100$, whereas for $k \geq 400$ comparison times exceed 10 seconds on average. In order to enable similarity comparison in a real-time scenario, comparisons should not take more than

⁸Dataset available at: <https://doi.org/10.25540/WVM0-4RNX>

⁹<https://fasttext.cc/docs/en/language-identification.html#content>

¹⁰<https://www.socialseer.com/resources/us-senator-twitter-accounts/>

¹¹Our dataset is available at: <https://doi.org/10.7910/DVN/NMT4HP>

a second, yet without jeopardizing the quality of the similarity score. We will see that signature sizes of as small as $k = 50$ yield good discriminative power.




Fig. 2. Average comparison times on a single CPU core (*i7-3520M*) between Twitter reference accounts of 60 US senators subject to the signature size k .

C. Similarity Networks

We define *similarity networks* as a fixed set of fully connected users (nodes), where every edge is weighted by its corresponding pairwise distance. We call the corresponding weight matrix *similarity matrix*. Multi-dimensional Scaling (MDS) [39] allows to visualize the relative similarities between all nodes of the graph in two dimensions. MDS approximately preserves relative distances, in particular very similar users will be mapped to points close to each other. Since MDS requires a matrix of pairwise *dissimilarities* instead of a similarity matrix, we transform similarities into dissimilarities via the mapping $x \mapsto 1 - x$. An example MDS output for the signature size $k = 50$ is depicted in Figure 3.




Fig. 3. MDS visualization of pair-wise similarities based on $k = 50$ signature words. Nodes represent 60 Twitter reference user accounts of US senators.

Applying labels for Republican and Democratic party membership reveals that members of the same political party are mapped closer to each other. In order to measure *affinity*'s quality of separating Democrats from Republicans, we propose a knowledge-based classifier following the principle of nearest neighbors. The classifier predicts the political party of a node

by taking the majority vote over all its n nearest neighbors. The results are shown in Table I. Note that a random classifier would predict party membership with an average accuracy of 50.0% (30 Republicans : 30 Democrats).

TABLE I
Accuracy of predicting the political party correctly by majority vote on n nearest neighbors of Twitter reference users of 60 US senators (30 Republicans, 30 Democrats) over several signature sizes k .

n	k		
	10	50	100
1	0.733	0.850	0.817
3	0.733	0.876	0.833
5	0.750	0.850	0.867
7	0.783	0.850	0.850
9	0.783	0.833	0.816
Avg.	0.756	0.850	0.837

We see that the similarities derived from signature sizes as small as $k = 50$ allow to predict the political party with an average accuracy of 85.0%. *affinity* is thus able to separate Democrat from Republican rhetorics based on their tweets without having seen any tweet history explicitly. In particular, the backend does not know in which regard users are similar since signatures are exchanged between users only. A user's signature data is incomprehensible for other users since it is (a) shared in latent form, and (b) the central word embedding model is not published. In that sense, *affinity* is privacy-preserving.

VII. CONCLUSION

In social networks, similar people tend to form links. On smartphone data, similarity is traditionally measured on sensor data, personal context data, or the social graph. We propose *affinity*, a system for latent – in particular privacy-preserving – user similarity comparison based on texting data. Texting data stays on the smartphone and only a list of used words is disclosed in order to find individually relevant words. We select a small set of relevant words that are represented by latent word vectors and their corresponding word weights, which forms the necessary information for comparison via the Word Mover's Distance. Upon a local exchange of comparison data, similarity comparisons can be performed on the smartphone quickly and reliably without the backend knowing. *affinity* does not assess similarity on interpretable text features yet on individual latent word embedding features. Comparison data in this latent format is incomprehensible to other users as long as the reference word embedding model is kept secret. Sharing comparison data is thus uncritical. In this sense, *affinity* is privacy-preserving and allows to safely leverage the information potential hidden in private texting data. We evaluate our approach on a political party classification task and reach an average accuracy of 85.0%.

VIII. FUTURE WORK




We evaluated *affinity* on a political party classification task. Evaluating *affinity*'s performance on for instance gender or age

classification tasks will substantiate *affinity*'s generalizability and applicability. We did not measure *affinity*'s similarity comparison performance on smartphones. The evaluation of a smartphone implementation is due.

REFERENCES

- [1] F. Beierle, "Do You Like What I Like? Similarity Estimation in Proximity-based Mobile Social Networks," in *Proceedings 2018 IEEE International Conference On Trust, Security And Privacy In Computing And Communications (TrustCom)*. IEEE, 2018, pp. 1040–1047.
- [2] J. Shu, S. Kosta, R. Zheng, and P. Hui, "Talk2Me: A Framework for Device-to-Device Augmented Reality Social Network," in *2018 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, mar 2018, pp. 1–10.
- [3] F. Beierle, K. Grunert, S. Göndör, and V. Schlüter, "Towards Psychometrics-based Friend Recommendations in Social Networking Services," in *2017 IEEE 6th International Conference on AI & Mobile Services (AIMS 2017)*. IEEE, 2017, pp. 105–108.
- [4] F. Beierle, V. Tran, M. Allemand, P. Neff, W. Schlee, T. Probst, R. Pryss, and J. Zimmermann, "Context Data Categories and Privacy Model for Mobile Data Collection Apps," *Procedia Computer Science*, vol. 134, pp. 18–25, 2018.
- [5] F. Beierle, S. Göndör, and A. Küpper, "Towards a Three-tiered Social Graph in Decentralized Online Social Networks," in *Proceedings of the 7th International Workshop on Hot Topics in Planet-scale Mobile Computing and Online Social Networking*, ser. HotPOST 2015. ACM, 2015, pp. 1–6.
- [6] A. J. Gill, "Personality and language: The projection and perception of personality in computer-mediated communication," dissertation, University of Edinburgh, 2003.
- [7] M. J. Kusner, Y. Sun, N. I. Kolkin, and K. Q. Weinberger, "From word embeddings to document distances," in *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ser. ICML 2015. JMLR.org, 2015, pp. 957–966.
- [8] A. Beach, M. Gartrell, S. Akkala, J. Elston, J. Kelley, K. Nishimoto, B. Ray, S. Razgulim, K. Sundaresan, B. Surendar, M. Terada, and R. Han, "WhozThat? Evolving an Ecosystem for Context-Aware Mobile Social Networks," *IEEE Network*, vol. 22, no. 4, pp. 50–55, 2008.
- [9] A.-K. Pietiläinen, E. Oliver, J. LeBrun, G. Varghese, and C. Diot, "MobiClique: Middleware for Mobile Social Networking," in *Proc. 2nd ACM Workshop on Online Social Networks (WOSN)*, ser. WOSN 2009. New York, NY, USA: ACM, 2009, pp. 49–54.
- [10] J. Teng, B. Zhang, X. Li, X. Bai, and D. Xuan, "E-Shadow: Lubricating Social Interaction Using Mobile Phones," *IEEE Transactions on Computers*, vol. 63, no. 6, pp. 1422–1433, Jun. 2014.
- [11] F. Beierle, V. Tran, M. Allemand, P. Neff, W. Schlee, T. Probst, R. Pryss, and J. Zimmermann, "TYDR - Track Your Daily Routine. Android App for Tracking Smartphone Sensor and Usage Data," in *2018 ACM/IEEE 5th International Conference on Mobile Software Engineering and Systems (MOBILESoft '18)*. ACM, 2018, pp. 72–75.
- [12] T. R. Reddy, B. V. Vardhan, and P. V. Reddy, "N-gram approach for gender prediction," in *2017 IEEE 7th International Advance Computing Conference (IACC)*, Jan 2017, pp. 860–865.
- [13] L. Akhtyamova, J. Cardiff, and A. Ignatov, "Twitter author profiling using word embeddings and logistic regression," in *Working Notes of CLEF 2017 - Conference and Labs of the Evaluation Forum, Dublin, Ireland, September 11-14, 2017.*, 2017.
- [14] B. G. Patra, K. G. Das, and D. Das, "Multimodal author profiling for twitter: Notebook for PAN at CLEF 2018," in *Working Notes of CLEF 2018 - Conference and Labs of the Evaluation Forum, Avignon, France, September 10-14, 2018.*, 2018.
- [15] A. Alekseev and S. Nikolenko, "Word Embeddings for User Profiling in Online Social Networks," *Computación y Sistemas*, vol. 21, pp. 203–226, 06 2017.
- [16] M. A. Raghuram, K. Akshay, and K. Chandrasekaran, "Efficient user profiling in twitter social network using traditional classifiers," in *Intelligent Systems Technologies and Applications*, S. Berretti, S. M. Thampi, and S. Dasgupta, Eds. Cham: Springer International Publishing, 2016, pp. 399–411.
- [17] S. Liang, X. Zhang, Z. Ren, and E. Kanoulas, "Dynamic embeddings for user profiling in twitter," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, ser. KDD 2018*. New York, NY, USA: ACM, 2018, pp. 1764–1773.
- [18] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman, "Indexing by latent semantic analysis," *Journal of the American Society for Information Science*, vol. 41, no. 6, pp. 391–407, 1990.
- [19] T. K. Landauer and S. T. Dumais, "A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge," *Psychological review*, vol. 104, no. 2, pp. 211–240, 1997.
- [20] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, Mar. 2003.
- [21] D. M. Blei, "Probabilistic topic models," *Commun. ACM*, vol. 55, no. 4, pp. 77–84, Apr. 2012.
- [22] W. H. Gomaa and A. A. Fahmy, "Article: A survey of text similarity approaches," *International Journal of Computer Applications*, vol. 68, no. 13, pp. 13–18, April 2013.
- [23] P. D. Turney and P. Pantel, "From frequency to meaning: Vector space models of semantics," *J. Artif. Int. Res.*, vol. 37, no. 1, pp. 141–188, Jan. 2010.
- [24] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching word vectors with subword information," *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 135–146, 2017.
- [25] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543.
- [26] A. Salle, M. Idiart, and A. Villavicencio, "Matrix factorization using window sampling and negative sampling for improved word representations," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, ser. ACL 2016, vol. abs/1606.00819, 2016.
- [27] O. Levy and Y. Goldberg, "Linguistic regularities in sparse and explicit word representations," in *CoNLL*, ser. ACL 2014, R. Morante and W. Tau Yih, Eds., 2014, pp. 171–180.
- [28] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," in *Proceedings of the International Conference on Learning Representations*, ser. ICLR 2013, 2013.
- [29] T. Schnabel, I. Labutov, D. M. Mimno, and T. Joachims, "Evaluation methods for unsupervised word embeddings," in *Proceedings of the 14th Conference on Empirical Methods in Natural Language Processing*, ser. ACL-EMNLP 2015, L. Márquez, C. Callison-Burch, J. Su, D. Pighin, and Y. Marton, Eds., 2015, pp. 298–307.
- [30] M. Joshi, E. Hart, M. Vogel, and J.-D. Ruvini, "Distributed word representations improve ner for e-commerce," ser. HLT-NAACL 2015, 2015.
- [31] M. A. Ruzon and C. Tomasi, "Edge, junction, and corner detection using color distributions," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 11, pp. 1281–1295, 2001.
- [32] K. Grauman and T. Darrell, "Fast contour matching using approximate earth mover's distance," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.
- [33] Y. Rubner, C. Tomasi, and L. J. Guibas, "The earth mover's distance as a metric for image retrieval," *Int. J. Comput. Vision*, vol. 40, no. 2, pp. 99–121, Nov. 2000.
- [34] J. Huang, R. Zhang, R. Buyya, J. Chen, and Y. Wu, "Heads-join: Efficient earth mover's distance similarity joins on hadoop," *IEEE Transactions on Parallel & Distributed Systems*, vol. 27, no. 6, pp. 1660–1673, June 2016.
- [35] G. Salton and C. Buckley, "Term-weighting approaches in automatic text retrieval," in *INFORMATION PROCESSING AND MANAGEMENT*, 1988, pp. 513–523.
- [36] J. Beel, S. Langer, and B. Gipp, "TF-IDuF: A Novel Term-Weighting Scheme for User Modeling based on Users' Personal Document Collections," in *Proceedings of the iConference 2017*, Wuhan, China, Mar. 22-25 2017.
- [37] C. D. Manning, P. Raghavan, and H. Schütze, *Introduction to Information Retrieval*. New York: Cambridge University Press, 2008.
- [38] T. Chen and M. Kan, "Creating a live, public short message service corpus: The NUS SMS corpus," *Language Resources and Evaluation*, vol. 47, no. 2, pp. 299–355, 2013.
- [39] W. S. Torgerson, "Multidimensional scaling: I. theory and method," *Psychometrika*, vol. 17, pp. 401–419, 1952.

Additional Information

Bibliographic Data	T. Eichinger, F. Beierle, S. U. Khan, R. Middelanis, V. Sekar, and S. Tabibzadeh, "affinity: A System for Latent User Similarity Comparison on Texting Data," in <i>Proc. 2019 IEEE International Conference on Communications (ICC)</i> . IEEE, 2019.
Pre-print from	https://beierle.de
Online at	(in press)
Authors	Felix Beierle   Tobias Eichinger  Sumsam U. Khan Robin Middelanis Veeraraghavan Sekar Sam Tabibzadeh
BibTeX	@inproceedings{Eichinger2019ICC, title = {{affinity: A System for Latent User Similarity Comparison on Texting Data}}, author = {Eichinger, Tobias and Beierle, Felix and Khan, Sumsam U. and Middelanis, Robin and Sekar, Veeraraghavan and Tabibzadeh, Sam}, booktitle = {{Proc. 2019 IEEE International Conference on Communications (ICC)}}, publisher = {IEEE}, year = {2019} }
Copyright Note	IEEE Copyright Notice Copyright (c) 2019 IEEE Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.