Temporal-Attribute Inference Using Dynamic Bayesian Networks

Lihi Idan[1]

Yale University, New Haven CT 06511, USA
lihi.idan@yale.edu

Abstract. As social networks continue to grow in popularity, it is essential to understand what can be learned about private attributes of social-network users by mining social-network data. Previous work focused on the inference of time-invariant attributes such as personality traits. By contrast, in this paper we focus on the inference of dynamic, time-varying attributes. We present a new approach to modeling social-network users and mining time-varying attributes using dynamic bayesian networks (DBNs). We then explore the extent to which such temporal models can improve the inference results of various dynamic attributes. This work is the first to take a DBN-based approach to the task of private-attribute inference in social networks.

1 Introduction

Knowledge of social-network users' intentions has immense potential to improve the design of recommendation systems, ad-targeting mechanisms, public-health campaigns, and other social and commercial endeavors. At the same time, such knowledge can have a detrimental effect on users' privacy. In this paper, we are interested in inferring intentions of social-network users using public data extracted from their social-network profiles.

Problem description Let u be a social-network user and S_u be the set of social networks on which u has accounts. We use $\xi_{(u,s)}$ to denote user u's account on network s. Each account has a private portion $\xi^{pr}_{(u,s)}$ and a public portion $\xi^{pu}_{(u,s)}$. The private portion contains data that only u's ties and the social-network provider can see, while the public portion contains data that can be seen by everyone. In addition to data that u publishes, $\xi^{pu}_{(u,s)}$ contains metadata information about $\xi_{(u,s)}$ such as the mere existence of $\xi_{(u,s)}$ and the visibility levels of different attributes in $\xi_{(u,s)}$. The goal of this work is to infer *offline* behavioral intentions of a social-network user u using only the public portions, $\{\xi^{pu}_{(u,s)}\}_{s\in S_u}$, of u's *online* social-network accounts. We focus on present or near-future behavioral intentions, *i.e.*, on decisions to perform certain actions within short periods of time after the decisions are made.

The paper makes the following contributions:

A new approach to modeling social-network users using Dynamic Bayesian Networks We present a new approach to modeling social-network

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

users and mining time-varying attributes using DBNs. We evaluate our models when used for the inference of different dynamic attributes given temporal, real-world social-network data. This work is the first to take a DBN-based approach to the task of attribute inference in social networks and the first to offer a DBN-based representation of social-network users.

A unique focus on offline, time-varying behavioral intentions Unlike other existing works that tackle the task of attribute inference, ours is the first work that aims at inferring *offline and dynamic*, non-politically-related behavioral intentions of social-network users (*i.e.*, a *user-centric* approach) solely based on *public social-network data*. Other works either focus on online intentions or time-invariant preferences; use private data or data that is not obtained from social networks; or take an "object-centric" approach by trying to infer the intention associated with a single, standalone and contextless social-network "object" such as a post or a tweet. Furthermore, some of the behavioral intentions that we consider in this paper have never been studied in any prior machine learning (ML) or social-network-related work.

A new multidisciplinary methodology for the inference of behavioral attributes We introduce a novel, multidisciplinary methodology for the inference of behavioral attributes such as decisions and intentions. We design modular bayesian-network (BN) models that are able to capture the evolving nature of the human decision-making process by combining data and priors from multiple domains. Our methodology handles common challenges in social-network mining such as incomplete datasets, unlabeled data and bidirectional influence between features and the target variable.

2 Related work

Inference of personal attributes using social-network data has been extensively researched. Inferring users' personality type was investigated in [8] using regression models and Twitter/Facebook data, respectively. Youyou *et al.* [29] showed that automatic inference methods that rely on Facebook likes achieve better prediction accuracy than those achieved by asking the users' friends. Staiano *et al.* [27] used data gathered through smartphones such as calls and texts; their results significantly vary across different personality dimensions.

Demographic attributes' inference is another well-studied topic, with age and gender being the most researched attributes [14, 15]. A related stream of research focuses on psychological and mental conditions. Depression is the most researched condition, followed by anxiety and stress [9, 18].

The common denominator of all the above works is that they focus on attributes that are either static (their values rarely change), non-self controlled, or both.

Inference of self-controlled attributes has also been extensively studied. However, such works focus on the inference of opinions and attitudes rather than behavioral attributes [4, 25]. While a substantial amount of work does study different types of behavioral attributes, their goals are different than ours. Such

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

works study general correlations between network or linguistic features and a given behavior, identify the prevalence of a certain behavior among the general population, or classify social-network textual objects such as tweets or posts. For example, while there exists a considerable amount of work about the use of social networks for monitoring public health, none of those works aims at inferring vaccination intent of a given social-network user at a given point in time. Rather, existing works analyze collective sentiment towards vaccinations [19], track the spread of infectious diseases and monitor online discussions concerning widespread diseases [23], or perform classification of stand-alone social-network objects according to vaccination attitudes of the object's creator [1].

Inference of time-varying, behavioral attributes using public social-network data has therefore been hardly researched, with two exceptions: voting intentions and *online* purchase intentions. There are several key differences between this work and prior ML work on PI. First, the majority of existing works examine general buying preferences rather than time-varying PIs [30]. Other works try to infer PI of stand-alone social-network objects (content-centric) rather than PI of social-network users (user-centric), an approach which is inherently biased [2,10]. The remaining works that do try to infer a user-centric, time-varying PI use data derived solely from E-commerce platforms. Such data is both platform-specific and oftentimes considered private, unlike our use of public social-network data [21]. The closest work to ours is [17] which infers PI of Pinterest users using temporal features and a logistic regression model. However, they only consider online purchases and do not differentiate between different product categories.

3 Methodology

"Intentions are people's decisions to perform particular actions" [24]. In this paper, we aim at understanding to what extent we can infer behavioral intentions of social-network users. In order to do that, we build a BN model that leans on intentions' most influential factors as shown in behavioral psychology literature [7, 11, 24]. We split those factors into two groups: static factors, such as personality, demographic attributes and self-efficacy, and dynamic factors such as emotions, interest and opinion. A significant challenge, however, is the fact that the values of some of those determinants (*e.g.* personality) can not be directly obtained from the user's social-network profiles ("latent variables"). Therefore, we enrich the model with various observed network features which may assist in both inferring the target intention and inferring its latent determinants.

Though different intentions are influenced by the same high-level factors, their associated BNs still differ in their qualitative, quantitative and temporal specifications. To reflect those differences, we build on our general intention-inference BN and create, for each behavioral intention, an intention-specific DBN. This is achieved using a multistage process: first, we identify the best set of determinants of general behavioral intentions using existing behavioral-psychology literature. Second, for each intention, we identify its unique determinants using existing literature which investigates that specific intention. Third, we identify

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

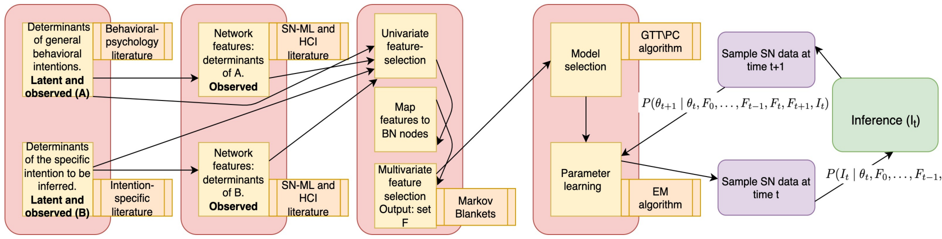

Fig. 1: An illustration of our behavioral-intention-inference methodology

the set of network features that are known to have a strong relation to the set of general and intention-specific determinants described above; the priors used for the third step, collected from existing literature, are not as strong as the priors used for stages 1 and 2 but are still informative — especially those collected from prior attribute-inference works and human-computer interaction (HCI) literature. Fourth, the final feature set of each intention is determined using priors, feature-selection methods, or both. Fifth, the set of selected features is mapped into network nodes; this includes aggregation, state-elicitation and discretization strategy. Sixth, the temporal structure of each intention-specific DBN is specified using priors, structure-learning methods, or both. Lastly, The DBN's temporal parameters are quantified using a combination of prior information and data. The diagram in Figure 1 illustrates our approach as detailed above.

In Section 6 we use the above methodology to infer the values of five dynamic attributes — behavioral intentions, using real-world, social-network datasets. The behavioral intentions that we consider are weight-loss intentions (WI), vaccination intentions (VI), travel-purchase intentions (PI), borrowing intentions (BI) and job-searching intentions (JI).

4 Features

A high-level diagram of our intention-inference model is shown in Figure 2. Note that Figure 2 is brought for illustration purposes and thus only features edges between layers; edges between specific nodes must be determined separately for each intention according to its own unique priors, intention-specific features and results of feature selection methods applied to it. To avoid a large conditional probability table (cpt), we used a layering-divorcing technique and created a layered network model: The first layer contains the target intention nodes that we aim at inferring. The second layer contains either latent or partially-observed nodes which represent external and internal factors that are highly influential on the formation of behavioral intentions. The third layer contains observable network features. They serve two purposes: assisting in inferring the behavioral intentions, and serving as observed predictors for second layer's latent variables.

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

4.1 Second-layer features

Values of second-layer features were obtained using our surveys and included in our training sets. In order to simulate a real-world inference task (which only considers the public portion of online profiles), values of *latent* second-layer features were omitted from our test sets (treated as missing values); instead, we tried to infer them using *publicly available* network features.

Personality This variable represents five broad dimensions of personality obtained from the "Big Five" model of personality dimensions. The big five model distills personality into five traits: neuroticism, extraversion, agreeableness, conscientiousness, and openness to experience. To measure the Big Five personality traits among survey participants we used a short version of the Big Five Inventory based on BFI-10 [22].

Demographic attributes We considered the following demographic attributes: age, gender, ethnicity (and country of origin), marital status, occupation group, income (latent variable). Only a subset of those attributes was used in each model.

Situational variables Events that might trigger a behavioral intention. Those events include personal-life transitions, professional-life transitions, external events (such as a holiday or an election), *etc.* Priors were obtained for some intentions. For instance, life transitions were shown to have an important impact on weight-loss intentions [3].

Emotions Different emotions may serve as either the cause of a behavioral intention or as its effect. Therefore, we went beyond the binary emotion-representation (positive-negative) and also considered fine-grained emotions. The most studied model of discrete emotions is the Ekman model [6] which posits the existence of six basic emotions: anger, disgust, fear, joy, sadness and surprise. Since momentary emotion ratings are not particularly indicative of the behavioral intentions explored in this work, survey participants were presented with eight emotion categories (six basic emotions and two positive-negative emotion categories) and were asked to rate their feelings over the past week/month/three months *in general*.

Interest, Opinion Those variables represent the user's level of interest and opinion regarding topics related to a given intention.

4.2 Network features

The value of a given network feature was included in our datasets only if it was part of the public portion of one of the user's social-network profiles.

Numeric features (NUMERIC) We considered statistics about the user's activity (number of posts, status updates, number of uploaded photos *etc*), reactions to the user's content (number of tagged photos, for instance) and the user's reactions to others' content. The latter measure was sparse, as both Facebook and Instagram limit the visibility of such reactions. We also considered basic statistics about the users' network, but we limit ourselves to statistics that are

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

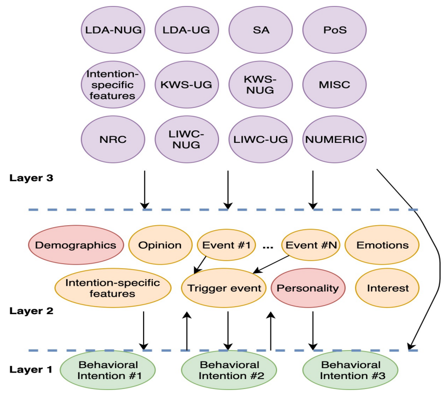

Fig. 2: A diagram of our static network model

both publicly visible and can be directly extracted from the user's own social-network profile/s (number of friends, followers-following ratio, public-figures-non-public-figures following ratio, *etc*).

Raw Textual features (TEXT) Textual features were classified as either user-generated (UG) features (including textual content that was written by the user), or non-user-generated (NUG) features (textual features that were not written by the user such as likes (Facebook) or hashtags (Instagram)). We limit ourselves to textual content that is both publicly visible and was either produced by the target user, or can be directly extracted from the user's own social-network profile/s. Raw textual features were not directly fed to our models. Instead, each textual feature was analyzed using various linguistic methods as described below; Textual-features-related-nodes in our network represent averaged score (frequency) of a given category of a given linguistic feature among the user's raw textual features. Such nodes represent the prevalence of a specific linguistic category among the entire set of raw textual features.

Miscellaneous features (MISC) Miscellaneous features include features that are neither numeric nor textual, such as the mere existence of various social-network accounts, visibility level/s that the user has chosen to apply to her social-network accounts, profile attributes from which demographic attributes can be extracted, *etc*. MISC features can be seen as social-network accounts' metadata rather than data itself (NUMERIC, TEXT).

Linguistic features We use a broad range of linguistic features, created based on our raw textual features.

Keyword-search (KWS-UG, KWS-NUG) For a given intention or an event, $\mathcal{A}$, we manually identified the most prominent keywords related to $\mathcal{A}$. We then performed a keyword search on our textual features. This resulted in two groups of features, KWS-UG and KWS-NUG (keyword search applied to user-generated/non-user-generated content).

LIWC (LIWC-UG, LIWC-NUG) LIWC is a text analysis tool that is widely used in psychological studies [28]. Each list of words is associated with

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

a semantic or syntactic category, such as negations, adverbs or tone. LIWC analysis was applied to UG and NUG textual features.

Sentiment analysis, part-of-speech tagging (SA, PoS) These were only applied to KWS-UG (SA and PoS) and KWS-NUG (SA), *i.e.*, items that were found to contain at least one relevant keyword. SA was applied to items that were found to contain keywords that relate to the behavioral intention to be inferred, in order to assess the user's opinion on related topics. The use of PoS tagging was more implicit; it was applied to items that were found to contain keywords that are related to events in order to assess whether an event is relevant to each inference task (use of first-person writing, tensed verbs *etc*).

Topic modeling (LDA-UG, LDA-NUG) Topics were extracted using Latent Dirichlet Allocation (LDA). Shorter features (such as likes) and longer features (such as posts) were considered separately using different parameters.

Emotions (NRC, LIWC) We automatically quantify emotions from our UG textual features using LIWC and NRC. NRC is a publicly available lexicon of words associated with different emotions, as well as general positive and negative sentiment [20]. We assign a predicted emotion to each UG textual feature and then average across all users' features.

5 Temporal modeling using Dynamic Bayesian Networks

A Dynamic Bayesian Network is a sequence of T static bayesian networks. Each BN represents a time slice ("slice") of the DBN, $i \in T$, corresponding to one instance of time. A DBN adds three components to a static BN: temporal variables, temporal edges and temporal evidence. For instance, if a static BN contains the variables $\{X_j\}_{j\in D}$, a DBN contains variables that can take different values in different time slices, *e.g.* $\{X_{j,i}\}_{j\in D, i\in T}$, as well as temporal edges between them. Formally, a DBN is defined as a pair (B_0, B_t) where B_0 defines the prior $P(X_1)$ and B_t is a two-slice temporal BN that defines $P(X_i|X_{i-1})$ by means of a directed acyclic graph ($PA(X_{j,i})$ represents $X_{j,i}$'s parents):

$$P(X_i|X_{i-1}) = \prod_{j\in D} P(X_{j,i} \mid PA(X_{j,i})) \tag{1}$$

A DBN-based approach to modeling social-network users: Each social-network user is modeled using a *set of Dynamic Bayesian Networks*. Specifically, let u be a social-network user, and let $|\mathcal{K}| = K$ be the set of u's dynamic attributes we aim at inferring. u is represented by the set $\{(D_u^{X_k}, T(X_k))|k \in [\mathcal{K}]\}$. $D_u^{X_j}$ corresponds to a DBN which aims at inferring an attribute X_j^u, the attribute X_j of a user u, and $D_u^{X_j}[i]$ corresponds to the i'th slice of the DBN. $T(X_j)$ refers to X_j's unique "sampling rate": the rate in which data for each attribute is sampled from $\{\xi_{(u,s)}^{pu}\}_{s\in S_u}$. The sampling rate $T(X_k)$ associated with a DBN $D_u^{X_k}$ should be determined according to the unique attribute to be inferred. For instance, if our target attributes are various behavioral intentions, the sampling rate of each intention's DBN should be determined according to the

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

intention-behavior (IB) interval [24] of intention X_k; the shorter the IB interval of an intention X_k is, the higher $T(X_k)$ should be.

After determining its feature set and its structure, as we describe in the following subsection, each of the user's DBNs can be used to perform temporal inference of its associated attribute at any point in time. For an inference performed at time $i = 0$, before any training data has been collected, inference of $X^u_{j,0}$ will be done solely based on $D^{X_j}_u[0]$ as a "prior network" — where cpts are solely determined according to prior information. At time point $i = t$ s.t. $t > 0$ inference of $X^u_{j,t}$ is done by training a new slice of DBN_j, $D^{X_j}_u[t]$ using an up-to-date set of training records where each record is composed of $t-1$ sets of historical features $\{F_i|i < t\}$, a set of current features $\{F_t\}$, and a set of historical labels, $\{I_i|i < t\}$.

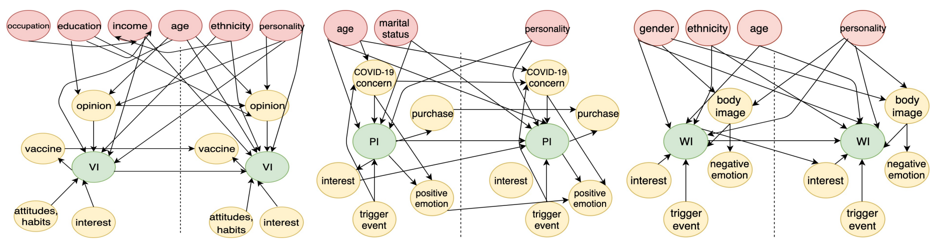

(a) VI (b) PI (c) WI

Fig. 3: A DBN representation of various intentions.

In both cases, inference of attribute $X^u_{j,i}$ will be done using $\{F^u_t\}$, the sampled feature sets of u at time t; $\{F^u_i|i < t\}$, sampled feature sets of the user from prior points in time; and historical labels (if exist), $\{I^u_i|i < t\}$. In addition, when the target attribute to be inferred is a behavioral intention, we can input the inference algorithm with a set of historical *behaviors*, $\{B^u_i|i < t\}$. A behavior at time i may suggest on an associated intention at time $i-1$ or $i-2$ thus allowing us to retroactively update the network's parameters to reflect the new insights.

5.1 Feature selection and model selection

We designed a two-level, hybrid feature-selection method. Due to the high number of correlations between features, we opted for a multivariate feature-selection method based on bayesian networks. However, solely relying on a DBN-based feature selection method may lead to overfitting. Hence, we employed a hybrid feature-selection approach. First, a simple, univariate feature selection method was applied to a subset of the features on which we didn't have strong priors. For that purpose, we used a mutual information-based feature-selection method and removed all the features that received a score below a certain threshold. The resulting features, as well as the set of latent/high-prior features were the

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

input for the second phase which uses two structure-learning algorithms: Greedy Thick-Thinning and the PC algorithm [26]. This phase aimed at identifying the best features using Markov Blankets.

A Markov Blanket of a variable t is a minimal variable subset conditioned on which all other variables are probabilistically independent of t. The Markov Blanket of a DBN node, $MB(t)$ is the set of its parents, $P(t)$; children, $C(t)$; and spouses, $U(t)$ as encoded by the structure of the DBN. As shown in [13], the Markov Blanket of a given target variable is the theoretically optimal set of variables to predict its value. However, simply considering all the features in the Markov Blanket of the behavioral intention node is unsatisfactory in our case, due to the existence of latent variables. Thus, a better strategy would be to first find an "approximated" Markov Blanket of the target node, $MB'(t)$ which includes the variables in the sets $P(t)$, $C(t)$ and $U(t)$ as discussed above. Then, identify the Markov Blanket of *each* latent variable that is also a member of the target's approximated Markov Blanket and include the features in the union of those blankets in our feature set (in addition, of course, to features in $MB'(t)$). That is, our feature set is:

$$\{MB'(t)\} \cup \{MB(I) \mid I \in S \cap MB'(t)\}$$

Where S represents the set of latent variables in our model. The above strategy would have probably been sufficient if our datasets were complete. However, our datasets contain missing values which had to be imputed before running the structure-learning algorithm. Hence, for some variables we consider an "extended" notion of a Markov Blanket which also includes certain variables that belong to the variable's second-degree Markov Blanket. Specifically, if a given variable, v, represents an observed attribute with more than 50% missing values ($m()$) and for which we do not have a strong prior ($p()$), we consider a restricted notion of v's second degree Markov Blanket, and add both its direct parents, $P(v)$, and its direct children, $C(v)$, to our feature set. Let F be our variable-set before applying feature selection, and O the set $F \setminus S$. Our final feature set is:

$$\begin{gathered}\{MB'(t)\} \cup \{MB(I) \mid I \in S \cap MB'(t)\} \cup \\ \{P(I) \mid I \in O \cap MB'(t) \wedge m(I) > 50\% \wedge p(I) = false\} \cup \\ \{C(I) \mid I \in O \cap MB'(t) \wedge m(I) > 50\% \wedge p(I) = false\}\end{gathered}$$

Model selection and parameter learning: The approach described above not only yields a feature set but also a network structure, comprised of the nodes in the feature set and the edges connecting features in the feature set. Some edges were corrected in order to reflect strong prior information. The balance between automatic structure-learning algorithms and the use of priors for structure elicitation, as well as the initial parameters for the structure-learning algorithms (when applicable) were validated using cross-validation. Note that while information gathered from prior literature would have probably been sufficient to model most of the meaningful dependency relations between an intention and second-layer features, relations between third-layer features and other features,

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

as well as between third-layer features and the target intentions can not be captured solely using priors, as those types of relations are not as extensively studied as behavioral intentions-second-layer features relations. Parameter learning was performed using the Expectation-Maximization (EM) algorithm [5]. Hence, we were able to combine both labeled and unlabeled data in our training sets as explained in Section 6 as well as use the *original* training datasets which contain missing values. We believed that since our test datasets include a large number of missing values, training the DBN on incomplete datasets will allow the DBN to learn relations between missing and observed values of different features. Prior information was combined in the model using a Dirichlet prior.

5.2 Intention-specific models

Figure 3 presents a high-level overview of three intention-specific DBNs (DBNs for JI and BI, as well as third-level features are omitted due to lack of space). As can be seen, a temporal link is created between variables that represent our target intentions in consecutive time slices. $P(intention_{i+1} \mid intention_i, U)$ represents the intention's evolution over time, given changes in other temporal variables in the network (U).

Interest-intention is an interesting relation. First, we see that interest may serve as either a cause or an effect of different intentions. Second, interest seems to be a cyclic process as can be concluded from $P(WI_i \mid interest_i, U)$ and $P(interest_{i+1} \mid WI_i)$, for example. Such a temporal relation might be attributed to the fact that interest in a certain topic assists in forming a behavioral intention related to that topic. After the intention has been formed, a new level of interest is formed, aimed at understanding how to fulfill that intention. In addition, $P(PI_{i+1} \mid interest_i, U)$ and $P(interest_{i+1} \mid PI_{i+1})$ show that both prior interest-level and current interest-level are important determinants of some intentions. Such historical data can assist in identifying a sudden increase in the user's interest level.

"Opinion" is another interesting variable; it is influenced by multiple factors such as personality traits and demographics as demonstrated by VI's $P(opinion_{i+1} \mid opinion_i, education, age, personality)$. Note that this cpt also contains $opinion_i$. This represents the fact that oftentimes, opinion is a self-propelling process: opinion at a given point in time, in addition to other factors, influences opinion at future points in time. A similar cpt is seen in "COVID-19 concern".

Fine-grained emotions were not used in any model. Furthermore, we weren't able to extract from the data meaningful inter-slice relations between different emotions and the target intentions. We attribute that difficulty to the fact that unlike other features, emotions change quickly. Thus understanding emotions' temporal evolvement mechanism for each intention requires the use of finer-grained sampling rates

In Section 7 we show our inference results (Lauritzen-Spiegelhalter algorithm [16]) when using a two-slice DBN and social-network data sampled twice.

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

Table 1: Datasets' statistics

Intention	VI	WI	BI	PI	JI
% Intending, first-wave dataset	.58	38	.17	.24	.19
% Intending, second-wave dataset	.66	.4	.13	.19	.23

6 Data collection

We designed and distributed a comprehensive survey, created and hosted using Qualtrics survey platform. The first part of our survey contained questions about the participants' personal attributes, as discussed in Section 4. The second part contained the following statements, which users were asked to rank (as well as dummy statements about unrelated intentions): "I am planning to start a weight-loss regime within the next 1-4 weeks" and "I am currently trying to lose weight" (weight-loss intentions); "I am planning to look for a new job within the next 1-4 weeks" and "I am currently looking for a new job" (job-searching intentions); "I am planning to apply for a loan within the next 1-4 weeks" (borrowing intentions); "I received a flu vaccine this season" — depending on the answer to that statement the following statement was presented for either the upcoming (2020-2021) flu season or the next season (2021-2022): "I am planning to get vaccinated against influenza this upcoming fall-winter/next year" (vaccination intentions); "I am planning to make a travel-related purchase within the next 1-4 weeks" (travel-purchase intentions). All survey data was anonymized after collection. We informed participants that their responses would be used for academic research. We implemented several methods for identifying and excluding data from participants who answered unreliably, as extensively discussed in [12].

Datasets Survey data was collected in two waves with a three-month lag. Training and test datasets include data obtained from Amazon Mechanical Turk (MTurk), Facebook, Instagram and Linkedin. Our datasets include both labeled and unlabeled data; unlabeled data is specifically important when using multi-wave data, as a considerable number of participants dropped out after the first wave: from 1300 respondents who participated in our first-wave survey, only 803 respondents participated in our second-wave survey (0.617 response rate). In order to both reduce non-response bias and create a bigger training dataset, we chose a subset of our partially-labeled data records which belong to participants who dropped out (missing attributes were treated as missing values) and added it to our training set. Our training datasets, $\mathcal{D}_j^1$ (first-wave data for intention j) and $\mathcal{D}_j^2$ (second-wave data for intention j) consist of 780 and 592 labeled and unlabeled data records, respectively. Our test datasets, $\mathcal{D}_j^3$ (first-wave data for intention j) and $\mathcal{D}_j^4$ (second-wave data for intention j) consist of 520 and 361 labeled data records, respectively.

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

Table 2: Results of the DBN models presented in this paper

Intention	Micro F1, (1)	Macro F1, (1)	Micro F1, (2)	Macro F1, (2)
Vaccination intentions	.732	.73	.75	.741
Weight-loss intentions	.832	.815	.831	.82
Borrowing intentions	.663	.54	.662	.526
Travel-purchase intentions	.763	.691	.812	.732
Job-searching intentions	.704	.623	.747	.699

7 Experimental Results

For a given intention, j, we tested its DBN, DBN_j, using our datasets as follows: in the first stage, (1), we trained DBN_j using $\mathcal{D}_j^1$ and tested it on $\mathcal{D}_j^3$. Only the first DBN's slice was affected in this stage. In the second stage, (2), we trained DBN_j using $\mathcal{D}_j^2$ (implicitly using $\mathcal{D}_j^1$ as well due to the use of priors) and tested it on $\mathcal{D}_j^4$, using evidence data from $\mathcal{D}_j^3$ as well. Hence, inference results in (2) were obtained based on data and parameters from two slices of the DBN. Note that "evidence data" contains only historical values of publicly available features, and *does not include historical labels* as in most cases exact information on historical labels for all prior test sets will not be available in real time. As seen in Table 1 (% intending), some of our datasets are *highly imbalanced*. Moreover, each dataset contains a large number of missing values — attributes that the user has not publicly revealed on her social-network accounts. Those facts make the inference task highly challenging.

Table 2 provides a detailed summary of our results. We report Micro F1 and Macro F1 scores for each intention-specific DBN. In addition, we compare our ROC AUC scores to those achieved by a Support Vector Machine (SVM) and a Decision-Tree Ensemble (boosted decision trees, BDT).

As can be seen from Table 2, different intentions achieved significantly different Micro F1 and Macro F1 scores. BI's score is the lowest, whereas WI's score is the highest. A possible explanation for BI's performance is that applying for a loan is an intention that is oftentimes not publicly shared on social networks. However, other non-publicly shared intentions such as JI scored significantly better than BI. This can be attributed to the fact that we were able to find other strong predictors for JI which don't depend on user-generated content, whereas for BI we failed to do so.

Figure 4 compares our ROC AUC scores to those achieved by two different types of classifiers: BDT and SVM (RBF kernel). Hyperparameters were tuned using a grid search over a large grid covering at least 7 options for each numeric

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

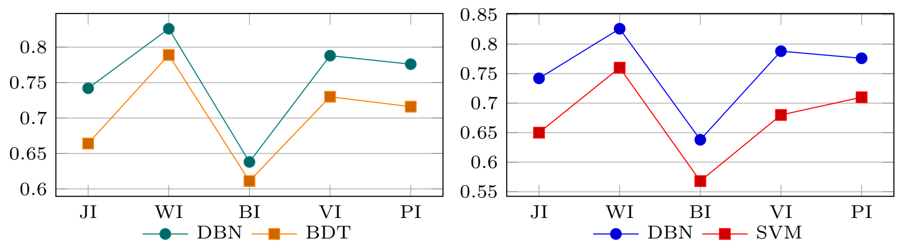

Fig. 4: Average ROC AUC scores

hyperparameter. Imputation of missing values was done using Scikit-Learn's IterativeImputer (using a random-forest regressor), a multivariate imputation method. As seen in Figure 4, our models outperform both SVM and BDT on all five intentions, though the differences in results vary between intentions. A possible explanation is the varying number of missing values within the unique set of features of each intention, or the varying number of latent variables in each DBN. Another possible explanation is the varying number of temporal dependencies *between* features of each target intention.

When comparing Micro F1 and Macro F1 scores achieved in different stages ((1) and (2)) using the same DBN, we can see that the differences are more pronounced for PI and JI. This can be attributed to the underlying differences between different intentions. As evidenced by our data, intentions such as WI and VI can be seen as "continuous intentions" in the sense that their intention-behavior interval is longer than for other intentions; the persistence rate of such intentions is significantly higher than rates reported for PI or JI. Another explanation for the varying differences is the different set of determinants of each intention. While the importance of some of those determinants stems from their intra-slice values (that is, their values at a given point in time), the importance of others is derived from a combination of intra-slice values *and* inter-slice change patterns between slices. For instance, various features related to *non* user-generated content serve as excellent predictors of PI in (2), but only as solid predictors in (1).

8 Acknowledgments

The author was supported by US National Science Foundation (CNS-1407454 and CNS-1409599) and William and Flora Hewlett Foundation (2016-3834).

References

1. Aramaki, E., Maskawa, S., Morita, M.: Twitter catches the flu: detecting influenza epidemics using twitter. In: EMNLP 2011 (2011)
2. Atouati, S., Lu, X., Sozio, M.: Negative purchase intent identification in twitter. In: WWW 2020 (2020)

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67

3. Brink, P.J., Ferguson, K.: The decision to lose weight. Western Journal of Nursing Research **20**(1), 84–102 (1998)
4. Conover, M., et al.: Predicting the political alignment of twitter users. In: SocialCom 2011 (2011)
5. Dempster, A., Rubin, D.: Maximum likelihood from incomplete data via the em algorithm. Journal of the Royal Statistical Society **39**(1), 1–22 (1977)
6. Ekman, P.: An argument for basic emotions. Cognition & emotion **6**(3-4) (1992)
7. Fishbein, M.: The role of theory in HIV prevention. AIDS care **12**(3) (2000)
8. Golbeck, J., et al.: Predicting personality with social media. In: CHI (2011)
9. Guntuku, S.C., et al.: Understanding and measuring psychological stress using social media. In: ICWSM 2019 (2019)
10. Gupta, V., et al.: Identifying purchase intent from social posts. In: ICWSM (2014)
11. Hampshire, S., Hart, H.: Decision, intention and certainty. Mind **67**, 1–12 (1958)
12. Idan, L., Feigenbaum, J.: Show me your friends, and I will tell you whom you vote for: predicting voting behavior in social networks. In: ASONAM 2019 (2019)
13. Koller, D., Sahami, M.: Toward optimal feature selection. In: ICML 1996 (1996)
14. Kulshrestha, J., et al.: Web routineness and limits of predictability: Investigating demographic differences using web tracking data. In: ICWSM 2021 (2021)
15. Lampos, V., Aletras, N., Geyti, J.K., Zou, B.: Inferring the socioeconomic status of social media users based on behaviour and language. In: ECIR 2016 (2016)
16. Lauritzen, S., Spiegelhalter, D.: Local computations with probabilities on graphical structures and their application to expert systems. Journal of the Royal Statistical Society **50**, 157–194 (1988)
17. Lo, C., Frankowski, D., Leskovec, J.: Understanding behaviors that lead to purchasing: A case study of pinterest. In: KDD 2016 (2016)
18. Mann, P., et al.: See and read: Detecting depression symptoms in higher education students using multimodal social media data. In: ICWSM 2020 (2020)
19. Mitra, T., Counts, S., Pennebaker, J.W.: Understanding anti-vaccination attitudes in social media. In: ICWSM 2016 (2016)
20. Mohammad, S.M., Turney, P.D.: Crowdsourcing a word-emotion association lexicon. Computational Intelligence pp. 436–465 (2013)
21. Mokryn, et al.: Will this session end with a purchase? inferring purchase intent of anonymous visitors. Electronic Commerce Research and Applications **34** (2019)
22. Rammstedt, B., John, O.: Measuring personality in one minute: A 10-item version of the big five inventory. Journal of Research in Personality **41**, 203–212 (2007)
23. Šćepanović, S., et al.: The healthy states of america: Creating a health taxonomy with social media. In: ICWSM 2021 (2021)
24. Sheeran, P.: Intention—behavior relations: a conceptual and empirical review. European review of social psychology **12**, 1–36 (2002)
25. Silva, A., et al.: On predicting personal values of social media users using community-specific features and personal value correlation. In: ICWSM (2021)
26. Spirtes, P., et al.: Causation, prediction, and search. Lecture Notes in Statistics (1993)
27. Staiano, J., et al.: Friends don't lie: inferring personality traits from social network structure. In: ACM Conference on Ubiquitous Computing (2012)
28. Tausczik, Y., Pennebaker, J.W.: The psychological meaning of words: Liwc and computerized text analysis methods. Journal of language and social psychology **29**(1), 24–54 (2010)
29. Youyou, W., et al.: Computer-based personality judgments are more accurate than those made by humans. PNAS **112**(4), 1036–1040 (2015)
30. Zhang, Y., Pennacchiotti, M.: Predicting purchase behaviors from social media. In: WWW 2010 (2010)

ICCS Camera Ready Version 2022
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-08754-7_67