Proceedings of the Tenth International AAAI Conference on
Web and Social Media (ICWSM 2016)

Towards an Open-Domain Framework for Distilling the Outcomes of Personal Experiences from Social Media Timelines

Alexandra Olteanu
Ecole Polytechnique
Federale de Lausanne
alexandra@aolteanu.com

Onur Varol
Indiana University
ovarol@indiana.edu

Emre Kıcıman
Microsoft Research
emrek@microsoft.com

Abstract

Millions of people share details about their real-world experiences on social media. This provides an opportunity to observe the outcomes of common and critical situations and actions for individual and societal benefit. In this paper, we discuss our efforts to design and build an open-domain framework for mining the outcomes of any given experience from social media timelines. Through a number of example situations and actions across multiple domains, we discuss the kinds of outcomes we are able to extract and their relevance.

Introduction

Through social media today, 10-100s of millions of individuals regularly and publicly report on their experiences, including the actions they take, the things that happen to them, and the experiences they have afterwards. They talk about work or relations (Ehrlich and Shami 2010; Garimella, Weber, and Dal Cin 2014), health and dietary practices (Teodoro and Naaman 2013; Abbar, Mejova, and Weber 2015), and even log information about their illnesses and coping strategies (Eschler, Dehlawi, and Pratt 2015; Chou et al. 2011). They do this for many reasons: keeping in touch with friends, gaining social capital, or even helping others. Regardless of why people share this information, such *social media posts can be leveraged to better understand common and critical situations and their outcomes.*

Individuals may learn from the experiences of others to better understand their own situations, and help them decide what actions may help them to reach their personal, high-order goals (Manski 1993). For example, someone diagnosed with a medical condition might be interested in learning about the likelihood of specific symptoms, or how long a painkiller's effects normally last. Furthermore, policy makers and others may also benefit from gleaning insights into the perceptions or outcomes of important phenomena (Diener 2006), such as public health, social or financial issues. They might, for instance, be interested in quantifying such outcomes to help assess various risk factors as a step towards designing early interventions (De Choudhury et al. 2016).

Problem Definition: Here, we explore the problem of distilling likely outcomes—broadly defined as results or effects of an action or situation that are more likely to be mentioned post-factum—of *any* given experience (e.g., being in a situation or taking an action) based on multiple users reports of this experience on social media.

Our Approach: We devised an *open-domain* framework for mining the outcomes of a given experience from social media that organizes social media posts into per-user timelines of events (Kıcıman and Richardson 2015). To isolate likely outcomes of a given experience from observed confounding factors, the framework first splits the timelines into two groups based on whether or not they include the experience. Then, it employs propensity score matching—a popular technique for reducing the confounding bias (Rosenbaum and Rubin 1984)—to stratify the groups of timelines such that the variation in the occurrence likelihood of an outcome is measured only across comparable subsets of users.

The Problem Setting

Self-Disclosure on Social Media: Despite the multipurpose nature of social media, messages about users' own experiences account for a notable fraction: (Naaman, Boase, and Lai 2010) found that "me now" messages about personal state and experiences constitute 40% of all messages, while (Kıcıman and Richardson 2015) found 26% of tweets to be experiential. This tendency to disclose information about oneself on social media is part of a broader phenomena, research estimating that self-disclosure represents 30–40% of human speech output (Tamir and Mitchell 2012). This is a key feature of social media that enables our work, although, social media data may not capture all users' experiences or all aspects surrounding their experiences, since users may be likely to share certain type of experiences such as positive and extreme experiences (Guerra et al. 2014). In addition, users might also mention their experiences or related aspects out of order. Such idiosyncrasies of social media influence the kind of insights (outcomes of a given experience in our case) that one is able to draw from it.

Observational Studies of Social Behavior: However, by leveraging this kind of data, prior work examined from how dietary habits vary across locations (Abbar et al. 2015) to what are the links between diseases, drugs, and side-effects (Paul and Dredze 2011; Myslín et al. 2013), or how opinions shift as events unfold (Guerra, Meira Jr, and Cardie 2014).

Controlling for Confounding Bias: Yet, although controlling for confounding bias in observational studies on social media is important (Gayo-Avello 2011), this is rare: many analyses are only co-occurrence based, assuming that co-occurring items share some true relation. A key challenge is discriminating attributes that are merely correlated with a variable of interest from those causally related to it. This is due to subjects not being randomly assigned to the groups under study, making them prone to the correlation-causation fallacy. Recent social media studies looking at shifts in suicidal ideation (De Choudhury et al. 2016), at the effect of exercise on mental health (Dos Reis and Culotta 2015), or of community feedback on individual user behavior (Cheng et al. 2014) try to fill this gap by applying techniques that have come into extensive use in medicine, economics, and other sciences, including the potential outcomes framework of the Rubin causal model (Sekhon 2007) and the structural equation model (Robins and Wasserman 1999; Pearl 2000). To complement these efforts, we aim to develop *generalizable techniques that separate the domain-agnostic mechanics of such analyses from the semantic interpretation of results that often requires domain knowledge*.

Most similar to our work is the study of (Landeiro and Culotta 2016) that focuses on adjusting for confounding factors in social media text classification. While their study considers a socioeconomic confounding variable (gender), we take into account all terms used in the past by users. While these terms are unlikely to capture all variables correlated with the confounding variables (as it is hard to argue that all relevant aspects of users' lives are captured in their social media streams), word use is known, for instance, to correlate with various psyho-socio-economic factors including gender, age or personality (Schwartz et al. 2013). In addition, rather than classifying users based on the prevalence of outcomes, our goal is to uncover likely outcomes of any experience.

A Framework for Outcomes Extraction

In this section, we briefly cover the key aspects of our pipeline for extracting likely outcomes of a target experience from a large corpus of social media messages:

Identifying Treated Users: The first challenge is to identify social media users that had a given experience. To do so, we search for a set of queries that include a conjunction of phrases written to avoid possible ambiguity in the experience term, and that match with reasonable accuracy posts of users that have had the experience, as opposed to using the word in different contexts. To take a concrete example, for prescribe drugs, we wish to avoid including advertisement tweets like "*buy Xanax online*", and instead search for specific phrases such as "*I took Xanax*".

Timeliness Construction: Next, for each user we identify this way, we retrieve their social media messages, which we represent as lists of arbitrarily sorted unigrams and bigrams (that we refer to as *events*), and use their timestamp to organize them in a per-user timeline. We then summarize the timeline using two partial timelines tracking the first, resp. the last, occurrence of each event. We do so to efficiently locate the set of events that have occurred before, resp. after, the experience across user timelines, which we use to compute the propensity to mention the experience, respectively to locate potential outcomes after the experience took place.

Outcomes Extraction & Treatment Effect: The goal of our analysis is to understand whether a user mention of the experience makes her more or less likely to mention certain events in the future, for which we employ propensity score matching (Rosenbaum and Rubin 1984). To isolate the outcomes of the experience from observed confounds—to the degree this is possible—we wish to compare the group of users identified to have had the experience (the *treatment group*) with another set of users that share a similar distribution of features, but did not had the experience (the *control group*). One way to identify such comparable groups is through matching based on users' propensity to mention the experience, which we estimate using the past events in the users' timelines. In our study, the control group is drawn from users identified to have had a different experience, but in the same domain (e.g. we compare users *suffering from gout*, with users suffering from other diseases such as *having kidney stones*). Then, the two groups (treatment and control) are subdivided based on the users propensity to mention the experience into comparable sub-groups. To determine the effects of the target experience we iteratively examining each post-hoc event (possible outcome) reported by users for systematic differences in the event likelihood among comparable treatment and control subgroups. Finally, we report the average treatment effect for each event across subgroups.

Exploratory Analysis

To assess our framework, we examined a variety of experiences about which people actively seek information to make decisions (including health ailments, business and financial decisions, relationship issues), and gathered 3 months of tweets for users that mentioned them within the time frame of one month (varying from 10s to 50k users per experience). We first review the raw results generated by our framework. Table 1 shows the top-most statistically significant outcomes extracted for a sample of experiences. We can see that the extracted outcomes are topically related to the target experience: users discussing gout are significantly more likely to mention *flare ups* of their symptoms, *uric acid*, and the physical locations of their symptoms, while users mentioning high triglyceride levels later discuss *statins*, *cardiovascular* issues, and *dietary* changes. Similarly, we see topically relevant outcomes for other scenarios (e.g., investment-related words such as the stock market, or varieties of anxieties and manifestations such as panic attacks).

Then, we looked at the temporal evolution of outcomes which provides additional context for characterizing the outcomes of an experience. Figure 1 highlights several of examples. We observe, for instance, that the likelihood of someone mentioning their *depression* increases around when she admits to *tak[ing] Prozac*, with other outcomes following a similar pattern such as *get paid* when *increasing the gross income*. In other cases, the outcomes become prominent several days later, like discussing *medication* when *taking Lorazepam* or *joint* when *suffering from gout* (known to lead to swelling joints). Other outcomes are more likely to occur

Outcome	Count	Effect%	Z-Score	Outcome	Count	Effect%	Z-Score	Outcome	Count	Effect%	Z-Score
Health/Diseases: Gout				**Society/Issues: Belly fat**				**Health/Diseases: Triglycerides**			
flare_up	35	4.1	12.33	burn	156	62.2	8.96	your_risk	46	24.8	18.12
uric_acid	27	2.9	10.36	ab_workout	13	8.5	5.82	statin	48	23.1	17.69
uric	28	2.9	10.11	workout_lose	13	8.5	5.82	lower	120	35.9	17.18
flare	81	4.9	9.92	help_burn	8	11.1	5.82	cardiovascular	54	23.0	16.72
big_toe	38	2.9	9.86	add_video	26	14.0	5.75	healthy_diet	55	19.3	16.54
joint	301	7.2	7.22	url_playlist	26	14.0	5.75	fatty_acid	29	18.3	16.37
aged	32	1.7	6.51	fitness	39	18.6	5.51	help_prevent	73	26.9	16.01
correlation	45	2.8	6.11	ab	43	19.1	5.51	risk_factor	33	18.3	15.55
bollock	53	2.5	5.96	playlist_mention	30	15.3	5.39	fish_oil	48	24.4	15.42
shite	108	3.4	5.93	biceps	7	4.7	4.74	inflammation	78	25.1	15.30

Table 1: Most significant 10 outcomes following selected events. We see that individuals mentioning gout are more likely to later mention symptoms of the disease and its cause; those mentioning high triglycerides are more likely to mention treatments; and individuals mentioning a desire to lose "belly fat" later mention workouts, videos and music.

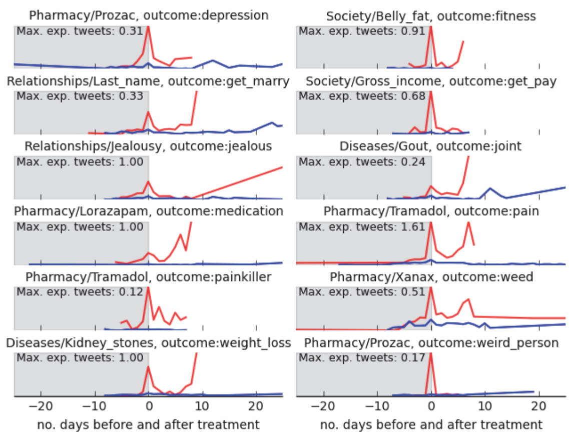

Figure 1: Comparison of temporal evolution of outcomes in treatment (red) and control groups (blue). The volume indicates the expected number of tweets per user (max value highlighted). Best seen in color.

both with the target event, as well as a few days later, e.g., the use of *weed* after *taking Xanax* (also taken for recreational use). We also notice interesting interplays among outcomes of the same experience: we see the mention of *painkiller* peaking around the time users mention *taking Tramadol*, while *pain* seems to recur again after a week.

What Kinds of Outcome Are Found?

To better understand the kinds of outcomes extracted by our framework, and whether these outcomes are conceptually or causally related to the target experience, we contrast our frameworks results with available concepts and relations from a large knowledge base, ConceptNet5 (Speer and Havasi 2012). In ConceptNet5, the relations between concepts are categorized across a variety of types, capturing both conceptual and descriptive relations, e.g., *IsA, DerivedFrom* or *SimilarTo*, as well as more causal like relations, e.g., *Causes, HasSubEvent* or *MotivatedByGoal*. Example relations include [*xanax IsA prescription drug*], [*xanax UsedFor anxiety*] and [*divorce CausesDesire drink*].

Across all domains, we find that our results are generally more likely to cover causal relations, including implementation steps, motivations and prerequisites, and implications. In contrast, our results do not cover as well more conceptual and descriptive relationships, including things that cannot be done, and alternate names and similar actions. These results indicate that our framework can distill outcomes that share a mixture of relations with the target experience. Note that different relation types might be of interest to different stakeholders. For instance, an individual diagnosed with anxiety might be interested in learning about likely, yet *not desired*, symptoms (e.g., *panic attacks* or a *nervous breakdown*), while someone diagnosed with gout might be want to know that this is *related to* high levels of *uric acid* and his *joints* may be affected as a result. On the other hand, a policy-maker might rather be interested in learning about real-world use cases for various drugs—e.g., our results for Xanax indicate that while this drug is typically *used for* medicinal treatment of *anxiety*, others mention *smoking weed* and *getting drunk* around the time they take Xanax, indicating recreational usage.

Future Directions

Our analysis aims to distill the outcomes that are more likely to be mentioned on social media following personal experiences. However, it remains to be seen how relevant these outcomes are across domains and for various tasks. Thus, our future work would naturally include a more comprehensive assessment of our framework, including an evaluation procedure to judge the relevancy and the quality of our results, an analysis of the relationship between the data volume (e.g., number of users, of tweets, timeline length) and the results quality, but also a more detailed characterization of the kinds of outcomes we extract.

Another important area that needs to be explored in the future, is the question of how the information about the outcomes can best be used and presented to aid people in specific application scenarios, as well as, the implications of these application patterns for the analysis framework itself.

Concluding Remarks

Social media, sensors and other computing services capture increasing amounts of data about the behaviors and experiences of millions of people. This provides an opportunity to better understand common and critical situations people are in, the actions they take, and their implications. Our aim is to develop generalizable techniques that separate the domain-agnostic mechanics of such analysis from the semantic interpretation of results that often requires domain knowledge. Driven by the type of experiences and domains about which people actively seek information online, our evaluation demonstrates that our framework may support people interested in gleaning insights into phenomena across a wide variety of semantic domains.

Applications for Individuals: We believe that individuals may benefit from the kind of outcomes we uncover. People planning their retirement might be interested in knowing that others having a *pension* tend to also be concerned with *taxes*, *benefits* or *health care*. Even when the outcomes of an action or situation are known, aggregated statistics about their likelihood can prove informative for those seeking information about them: someone *taking Prozac* to treat a depressive episode might feel relieved to know that while the likelihood of mentioning *depression* is high among others *taking Prozac*, the incidence of these mentions tends to quickly fade away after the treatment starts.

Applications for Policy-makers & Others: Further, while our work is motivated primarily by the desire to help individuals understand their situations and possible implications of their actions, there is also an opportunity to use this kind of analysis to better understand behavioral phenomena of societal importance, third-party interventions and other policy questions. For instance, learning about the concerns (and their likelihood) of people *having a pension* within a given time period is not only informative for individuals, but it is also an important source of information for policy-makers (Diener 2006). Other example, for pharmaceutical and public health research, such a source of information can help understanding drug uses that fall beyond the drug prescription (as we have found about *Xanax*).

Limitations: Our analysis aims to distill the outcomes that are more likely to be mentioned following personal experiences. However, it is important to note that while we borrow propensity score analysis from the causal inference literature, this application is not causal due to key assumptions that fail to hold. Furthermore, our analysis currently ignores population and reporting biases of social media. Future work includes scalability and performance of the analysis, as well as visualization and applications.

References

Abbar, S.; Mejova, Y.; and Weber, I. 2015. You tweet what you eat: Studying food consumption through Twitter. In *CHI*.

Cheng, J.; Danescu-Niculescu-Mizil, C.; and Leskovec, J. 2014. How community feedback shapes user behavior. In *ICWSM*.

Chou, W.-Y. S.; Hunt, Y.; Folkers, A.; and Augustson, E. 2011. Cancer survivorship in the age of YouTube and social media: a narrative analysis. *Journal of medical Internet research* 13(1).

De Choudhury, M.; Kiciman, E.; Dredze, M.; Coppersmith, G.; and Kumar, M. 2016. Discovering shifts to suicidal ideation from mental health content in social media. In *CHI*.

Diener, E. 2006. Guidelines for national indicators of subjective well-being and ill-being. *Applied Research in Quality of Life*.

Dos Reis, V. L., and Culotta, A. 2015. Using matched samples to estimate the effects of exercise on mental health via twitter. In *AAAI*.

Ehrlich, K., and Shami, N. S. 2010. Microblogging inside and outside the workplace. In *ICWSM*.

Eschler, J.; Dehlawi, Z.; and Pratt, W. 2015. Self-characterized illness phase and information needs of participants in an online cancer forum. In *ICWSM*.

Garimella, V. R. K.; Weber, I.; and Dal Cin, S. 2014. From ”i love you babe” to ”leave me alone”-romantic relationship breakups on Twitter. In *SocInfo*.

Gayo-Avello, D. 2011. Don’t turn social media into another ’Literary Digest’ poll. *Communications of the ACM*.

Guerra, P. C.; Meira Jr, W.; and Cardie, C. 2014. Sentiment analysis on evolving social streams: How self-report imbalances can help. In *WSDM*.

Kıcıman, E., and Richardson, M. 2015. Towards decision support and goal achievement: Identifying action-outcome relationships from social media. In *KDD*.

Landeiro, V., and Culotta, A. 2016. Robust text classification in the presence of confounding bias. In *AAAI*.

Manski, C. F. 1993. Dynamic choice in social settings: Learning from the experiences of others. *Journal of Econometrics*.

Myslín, M.; Zhu, S.-H.; Chapman, W.; and Conway, M. 2013. Using twitter to examine smoking behavior and perceptions of emerging tobacco products. *Journal of medical Internet research* 15(8).

Naaman, M.; Boase, J.; and Lai, C.-H. 2010. Is it really about me?: message content in social awareness streams. In *CSCW*. ACM.

Paul, M. J., and Dredze, M. 2011. You are what you tweet: Analyzing twitter for public health. In *ICWSM*, 265–272.

Pearl, J. 2000. *Causality: models, reasoning and inference*, volume 29. Cambridge Univ Press.

Robins, J. M., and Wasserman, L. 1999. On the impossibility of inferring causation from association without background knowledge. *Computation, causation, and discovery* 305–321.

Rosenbaum, P. R., and Rubin, D. B. 1984. Reducing bias in observational studies using subclassification on the propensity score. *Journal of the American Stat. Assoc.* 79(387):516–524.

Schwartz, H. A.; Eichstaedt, J. C.; Kern, M. L.; Dziurzynski, L.; Ramones, S. M.; Agrawal, M.; Shah, A.; Kosinski, M.; Stillwell, D.; Seligman, M. E.; et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one*.

Sekhon, J. S. 2007. The neyman-rubin model of causal inference and estimation via matching methods. In *Oxford handbook of political methodology*.

Speer, R., and Havasi, C. 2012. Representing general relational knowledge in ConceptNet 5. In *LREC*, 3679–3686.

Tamir, D. I., and Mitchell, J. P. 2012. Disclosing information about the self is intrinsically rewarding. *National Academy of Sciences*.

Teodoro, R., and Naaman, M. 2013. Fitter with Twitter: Understanding personal health and fitness activity in social media. In *ICWSM*.