

The Werther Effect Revisited: Measuring the Effect of News Items on User Behavior

Elad Yom-Tov
Microsoft Research
Herzeliya, Israel
eladyt@microsoft.com

Shira H. Fischer
RAND Corporation
Boston, MA, USA
sfischer@rand.org

ABSTRACT

People are moved to act following exposure to media coverage of specific events. For example, the “Werther Effect” is the popular term for the observed increase in suicides following media coverage of suicides. Here we develop a fine-grained method for assessing the effect of news stories on the intentions of internet users. Our method assesses the likelihood that a person was exposed to a given news story via the temporal and spatial distances between the location of the person and the location of the news story and/or the website where it was published. This analysis of likelihoods allows us to estimate the contribution of a particular news story to a person’s intent, as manifested in specific, intent-driven, search engine queries.

Data were gathered over a ten-month period and cover both the search engine queries of a large population and the news stories to which this population was exposed. We estimated the contribution of news stories to negative effects (i.e. media coverage of suicides and their effect on queries indicating suicidal intention), and positive effects (e.g. media coverage of disease prompting queries into disease screening). We demonstrate that the contribution of news stories can be assessed at the level of individual users, and we analyzed titles and phrases therein for their effect. Finally, we propose a predictive model to be utilized by media outlets to predict the likely effect of specific stories prior to their publication.

Keywords

News;Query intent;Temporal correlates

1. INTRODUCTION

“Copycat suicide” is colloquial term referring to the occurrence of a person learning of a suicide via the media, and their subsequent attempt to commit suicide because of this. Such behaviors, which sometimes lead to clusters of “copycat suicides” are collectively known as the “Werther Effect”,

named after the first description of these suicide clusters, ostensibly caused by the publication of the novel *The Sorrows of Young Werther*, by von Göthe [18, 21]. Extensive research has demonstrated that the language used to describe suicides (i.e. critical language vs. empowering language), and the volume of attention given to suicides, affect the observed number of suicides [18]. However, this prior research is focused on suicides that received significant media attention [22], i.e., the exaggeration of sensational incidents such as the suicides of well-known political figures [18], those shown on popular TV shows [10], and those linked to famous companies (e.g., the “Foxconn suicides” [4]). Moreover, research into the protective effect media can have, by either not reporting on a suicide, or by changing the quality and content of the reporting (the “Papageno effect”), is still lacking [22]. Moreover, the temporal aspects of such “copycat suicides” has not been investigated in any detail, most likely because of the dearth of relevant, high-resolution data.

Studies in areas unrelated to health have explored the temporal and spatial correlates of user behavior, as well as their interactions. The temporal dynamics of user search behaviors has been a topic of extensive study. For example, Richardson [19] showed that the temporal correlation between queries represents common human behaviors, e.g., updating a resume and subsequently moving to a new location. Much of this research has focused on predicting user behavior for better personalization [2] and more generally to improve user satisfaction [9].

Spatial dynamics of user searches have been investigated in the context of news seeking [14, 24] and of local search [8]. When modeling local interest in news events, Backstrom et al. [1] showed that search engine queries regarding local events can be used to pinpoint user location. It was also shown [24] that information needs and interest levels vary with the user’s physical distance from an event. Both studies demonstrated that interest in a news event decreases exponentially as a function of the physical distance of a person from an event epicenter. These insights are closely aligned with Diaz [5], who found that temporally local variables, such as query volume, were important in detecting the newsworthiness of a search engine query.

The temporal analysis of news events has received some attention. For example, Leskovec et al. [13] showed that media interest in an event peaks within two days and decays somewhat quicker than it arises. Social media outlets follow a similar attention temporal profile, and peak, on average, 2.5 hours later than mainstream media. These findings were reinforced in a later study of news events and their mani-

©2017 International World Wide Web Conference Committee (IW3C2), published under Creative Commons CC BY 4.0 License.
WWW’17 Companion, April 3–7, 2017, Perth, Australia.
ACM 978-1-4503-4913-0/17/04.
<http://dx.doi.org/10.1145/3038912.3038914>



festation in social media, mainstream media, and search engine queries [23], though it also drew attention to the fact that the temporal profile of interest in news events is significantly different among these information sources, and varies according to the characteristics of news events.

The relationships between news and user behavior, on the internet, in general, and on search engine activities in particular, have been studied extensively. Diaz [5] showed how news-specific queries to search engines are indicative of interest in news stories, and how this interest can be used to integrate news into search results. Several researchers examined how news may propel people to conduct specific behaviors. Yom-Tov and boyd [26] showed that news events about specific celebrities cause certain social media activities and search engine queries, and that a fraction of users search using queries highly indicative of an eating disorder. Another indication of how news pushes people to search for relevant behaviors was provided when Google Flu Trends, a service for predicting the number of influenza cases in a location according to the number of influenza-related searches, overestimated the number of cases during the 2009 flu season, ostensibly because of the outsized media attention to the H1N1 influenza virus which caused more people than usual to search for influenza symptoms [17].

Exposure to news items can occur via many routes, including television, radio, and printed newspapers or magazines, as well as the Internet. On the Internet, exposure can happen via news channel websites, search queries, news aggregators, and social media, among other sources. Here we first analyze how likely we are to detect exposure to a given news item using search query data. Having shown that this is a relatively low fraction of exposures, we develop methods for inferring exposure based on temporal and spatial attributes of news sources and users, and use these to infer exposure to news items and then measure their effect.

Our aims are three-fold: first to develop a methodology to evaluate the “Werther Effect” with high spatial and temporal resolution then utilize this methodology to evaluate the “Werther Effect”; second, to use the methodology to measure the effects other health-related behaviors; and finally to develop a tool to predict the likely behaviors and subsequent health effects of a given news article.

2. PROBLEM DEFINITION

A news item is published by a given news source at time $T_{Publish}$. At time $T_{Exposure}$ user U_i reads the news item (partially or entirely). We denote the time between publication and exposure by:

$$\Delta_{P \rightarrow E} = T_{Exposure} - T_{Publish}$$

After reading a news article, the reader might decide to take an action of interest. The indication of the reader’s decision to action is identified through an expression of his or her information need [25] given at time T_{Action} through a query issued to a search engine. It is assumed that T_{Action} is shortly after the time of decision. If the “Werther Effect” is valid, it would cause a fraction of people exposed to some of the news items to take the action of interest within a short time. That is, we expect that some news items, with text that has specific, measurable, attributes, will cause a fraction of the users to make queries indicative of intended

actions, such that:

$$\Delta_{E \rightarrow A} = T_{Action} - T_{Exposure}$$

is short.

Yom-Tov and Diaz [25] showed that people’s information need, $P(T|u)$, defined as the probability that user u will be interested in topic T , is contingent on properties of u . Specifically, in the context of news events, they analyzed the effect of the geographical distance between a topic and the user and the social affinity between u and people affected by T on $P(T|u)$. Here we analyze the another property of u namely, previous exposure to a specific news story. We hypothesize that particular information provided in news stories will drive information needs.

3. METHODS

3.1 Data

Three sources of data were used in this work and for all sources, user identifiers were anonymized, and users were not identified across datasets. The data of all four sources were extracted for the ten-month period between January 1st and October 31st, 2010. Following is a description of each data source:

3.1.1 News data

News items were all those articles displayed on the Bing News aggregator during the data period. For each news item, we obtained the title, publication date and time, and the news source (i.e., website) on which it was published.

3.1.2 Search engine data

We extracted all English language queries submitted to the Bing search engine by users in the United States during the data period. For each query, we extracted the query text, time and date, a list of pages displayed to the user as a result of the query, the zip code from where the user issued the query and an anonymized user identifier. We note two intrinsic limitations of the data used in the study: (1) we cannot distinguish between multiple users with the same login credentials; and (2) if a searcher uses the search service on multiple devices, they would appear with separate identifiers per each device.

3.1.3 Browsing data

To analyze behavior beyond search engine queries, we extracted user browsing information from the logs of consenting users of a browser toolbar distributed by Bing. The logs contained the URL addresses of all web browsing behavior (a significantly smaller sample than that of the search engine data), the time and date, and an anonymous user identifier.

3.2 Topics

Five topics were investigated in this paper. They were:

1. Suicides: As per the Werther Effect, we analyzed news items on suicides, and their effect on queries indicating suicidal intents.
2. HIV-AIDS screening: News items on AIDS and/or HIV and their effect on queries indicating an interest in HIV screening.

Table 1: Terms for identifying relevant search queries and news items. Each item was required to contain at least one of the included terms, and none of the excluded terms.

Topic	News items			Search queries		
	Included	Excluded	Number matching	Included	Excluded	Number matching
Suicides	suicide	bomb, blast, attack, squad	106,440	suicide, kill AND how to, how do I, ways to commit	suicide girl, stop, prevent, help, talk, cope, avoid, spell, hotline	335,976
HIV AIDS	hiv, aids		1,002,983	hiv, aids AND symptoms, screening, risk, test, signs		150,229
Breast cancer	mammogram, mammograph		391,438	mammogram, mammograph AND location, at, screening, recommend, provide, how, where, getting		401,502
Mastectomies	mastectomy		420,287	breast cancer AND symptoms, screening, risk, test, signs		10,636
Cancer	cancer		4,947,622	cancer AND symptoms, screening, risk, test, signs		867,939

- Breast cancer screening: News items on mammography and their effect on queries indicating an interest in breast cancer screening.
- Mastectomies: News items on preventive (or other) mastectomies and their effect on queries indicating an interest in breast cancer screening.
- Cancer screening: News items on cancers and their effect on queries indicating an interest in cancer screening.

While the first topic is of a negative behavior or nature, the latter four are positive, and of interest to public health institutions interested in promoting attention and adherence to health guidelines.

The terms used to identify and extract relevant news items and query terms for each topic are shown in Table 1. We henceforth refer to queries that contained relevant terms as “action queries” (AQs). AQs are queries that suggest that a user is interested in performing a relevant action. For example, in the case of suicide, an AQ might be “How do I commit suicide?”, whereas for cancer screening an AQ can be “Where can I find information on Thyroid cancer screening”. When a user made repeated AQs within a 30-minute time span, only the first query was retained. Table 1 also shows the number of queries and news items for each topic.

4. RESULTS

4.1 Accuracy of terms match

The accuracy of the identified AQs was assessed by manually inspecting the 100 most common queries (accounting for, on average, 40.0% of the volume of AQs) and labeling them according to their relevance to the topic. On average, 95.8% of these queries were correctly labeled as indicating a relevant action. The lowest accuracy was obtained for HIV-AIDS (89.2%) and the highest accuracy was obtained for suicides (99.1%).

Thus, our set of AQs represent a high precision set of queries.

4.2 Search engine queries as indicators of news exposure

People can be exposed to news items in many ways, both online and offline (e.g., TV, radio, etc.). To begin to explore this, we estimated the fraction of people reached by a news item following a search engine query. To accomplish this, the toolbar dataset was utilized, and for each of the five topics, we counted the fraction of instances a news item was accessed following a search query, out of all the instances a given news item was accessed. Note that users might be more inclined to search for specific topics (for example, they might have a long-standing interest in them). However, such prior biases do not detract from the fact that the appearance of a news item is followed by the user being interested in the topic.

As Table 2 shows, there is a large variation in the percentage of news items accessed via search queries and the average indicates that about half of all news items accessions occur subsequent to a search query. This percentage is in close agreement to that found in surveys of Internet users [11]. Therefore, to be able to use search engine data, which is by far the most voluminous of our data sources, and the most popular central repository of such information [11], we infer exposure to news items via the interactions of users and news sources, rather than assume exposure solely when accessed through search queries or other browsing behaviors.

Given user U_i at time T , we infer the probability of exposure to a news item published at time $T_{Publish}$ from source S_j as the probability

$$P(U_i, T | S_j, T_{Publish}) = P_{Exposure}(U_i, P_j, \Delta_{P \rightarrow E})$$

That is, we assume that the probability that the user was exposed at time t is related to the probability of the user being exposed to news items from source j , and the time difference between the publication time and time t .

4.3 Estimating time-to-exposure and time-to-action

In this section we estimated $\Delta_{P \rightarrow E}$, i.e., the time between publication ($T_{Publish}$) and exposure ($T_{Exposure}$) in search queries. Figure 1 shows the distribution of the $\Delta_{P \rightarrow E}$ for all news items and search queries where any of these news items were displayed in their results.

Table 2: Percentage of times that a news item was reached as a result of a search engine query versus all accesses to news items.

Topic	Percentage of times
Suicides	46.9
HIV AIDS	24.8
Breast cancer	79.5
Mastectomies	41.1
Cancer	37.5
Average	44.2

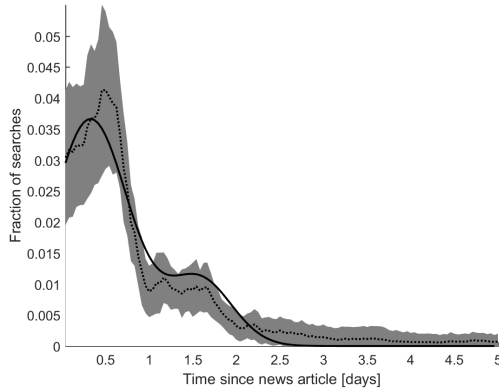


Figure 1: Distribution of the time difference between news publication and exposure by users. The dotted line shows the average for all data. The shaded area represents two standard deviations from the average, computed according to the 100 most common zip codes of users. The bold line shows the best fit of a two Gaussian distribution.

Several observations can be made from this figure. First, the distribution of the time difference is bimodal. The largest fraction of news views occurs in the first twelve hours after publication. A secondary peak is visible 24 hours later. This spiking patterns corresponds to the multiple-spike pattern identified by Kulkarni et al. [12], which can be attributed to the diurnal pattern of human activity.

Additionally, the distribution is within a small range as the two standard deviations (computed according to the distribution of the 100 most common zip codes) are close. This suggests that the time to exposure is independent from the location of exposure. We therefore treat the time to exposure and the relationship between news outlets and user as an independent variable. This product is computed to estimate the probability of viewing a news item given the time difference and the news outlet.

Finally, we parametrized the observed distribution by fitting to it a two Gaussian distribution, using the EM algorithm. The resulting distribution is shown in Figure 1 (in bold). Throughout the remainder of this paper, we use this approximation to compute the probability of viewing a news item given the time since its publication.

The time $\Delta_{E \rightarrow A}$ between viewing a relevant news item ($T_{Exposure}$) and issuing an action query (AQ, T_{Action}) was estimated separately for the different topics, by calculating the average time between news views and AQ for those users

Table 3: Time between exposure and action queries, across topics.

Topic	Average time [sec]	Standard deviation [sec]
Suicides	90.2	506.2
HIV AIDS	127.1	660.2
Breast cancer	88.4	338.4
Mastectomies	162.1	863.3
Cancer	192.7	883.4

that were observed to make both. The average times are shown in Table 3. As shown, the time between exposure and action is relatively short (under five minutes in all cases), especially when compared to the time between publication and exposure.

4.4 Estimating news outlet readership

Given a news item from news outlet S_j , it is possible to compute the probability that user U_i will be exposed to the news item in query data. This is done by counting the number of times that a user was exposed to news items from S_j in query log data (i.e., that items from S_j were shown in search results), divided by the total size of the dataset.

Figure 2 shows two examples for the resulting distributions, one for a local news outlet and the second for news outlet with viewership throughout the United States (US). As the figure shows the viewership of the CNN channel (<http://cnn.com>) is distributed across the US, with larger viewerships in major population centers. This is compared to the distribution of the Florida Times-Union (<http://jacksonville.com>), a local newspaper published in Jacksonville, Florida.

4.5 Effects of individual terms and entire titles

Using the insights obtained in the previous sections, we estimated the contribution of each news item to Aqs. Additionally, the words comprising each news item were similarly analyzed for their estimated impact on future Aqs.

We represent each news item S_j by its words, bigrams, and lexical affinities [3]. Let $P_{Time} = P(\Delta_{P \rightarrow A})$. Since $\Delta_{E \rightarrow A} \ll \Delta_{P \rightarrow E}$, we assume that $P_{Time} = P(\Delta_{P \rightarrow A}) \cong P(\Delta_{P \rightarrow E})$, and compute $P(\Delta_{P \rightarrow E})$ using the two Gaussian approximation detailed above. Let $P_{Location}$ be the probability of viewing a news item from S_j given the location of U_i . For each news item and AQ pair, we compute the probability that the AQ of U_i was the result of viewing news item S_j as the product of $P_{Time} \cdot P_{Location}$, and denoted by $P_{i,j}$.

The total contribution of a news item to Aqs made after it is calculated as the median of all $P_{i,j}$ for which $T_{Action} > T_{Publish}$. We refer to this score, given to news item S_j , as the News Item Action Score (NIAS).

Analysis of the news items with the highest NIAS values shows that they mention specific suicide events as well as coping with the suicide of others. The terms least associated with Aqs were regarding suicide prevention or regarding people rescued after suicide attempts. These findings are in line with those of Niederkrotenthaler et al. [15], who showed that reporting on suicide prevention was associated with a reduction in suicides in Austria (the ‘‘Papageno effect’’). The terms most associated with Aqs are also about

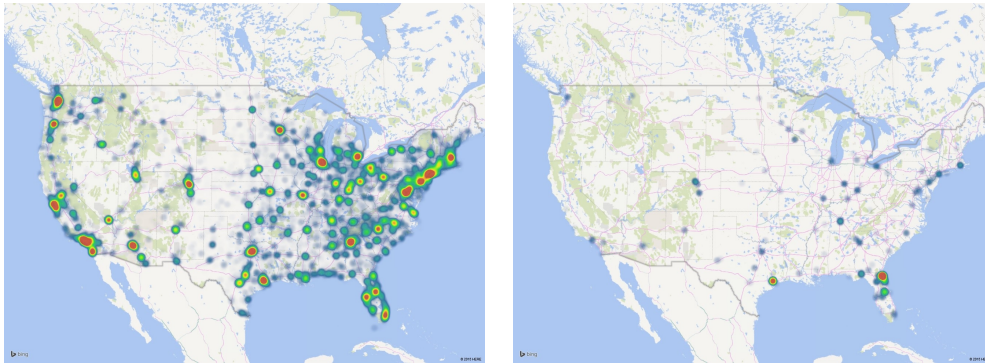


Figure 2: Distribution of the location of users viewing news items from CNN, a national news channel (left), and the Florida Times-Union, a local news outlet published in Jacksonville, Florida (right). Red represents a higher number of viewers at these areas.

specific forms of death (e.g., assisted suicides and murder-suicides), an effect also reported by Sisask and Värnik [22]. Publicity of assisted suicides was correlated with an increase in suicides in Switzerland [7], thus lending credence to our finding. High NIAS terms also mention friends and family who try to understand the suicide.

The other four topics relate to screening for different health conditions. For all conditions, positive news on treatment and negative news on the prevalence of disease had high NIAS values. Low NIAS titles were commonly reports regarding awareness.

4.6 Predicting NIAS

Another indication to the effect of news items on AQs is the ability to predict NIAS values from the terms of news items. The first method to accomplish this is a direct prediction of NIAS values, obtained by a linear regression model which predict NIAS values from news document phrases. The second method is a prediction of whether a news item would have a NIAS value in the top 1% of NIAS values. The goal of this latter task is to be able to provide feedback to journalists on those news items which might pose a risk (in the case of suicides) or a benefit (for other topics). We performed 10-fold cross-validation [6] to evaluate the success in prediction of both methods.

Table 4 shows the Spearman correlation (ρ) between the predicted and actual NIAS values. As the table shows, these values can be predicted with relative accuracy, demonstrating that the “Werther Effect” holds true for a variety of topics. Table 4 also shows the Area under the Receiver Operating Characteristic (ROC) Curve (AUC) to predict which articles have the highest NIAS values. For all cases, the AUC is greater than 0.84, indicating that news items can indeed be filtered for their risks factors or for their beneficial factors.

5. DISCUSSION

Most people in industrialized countries consume information from the media. It is estimated that the average American uses electronic media more than 11 hours per day [20]. This exposure is expected to affect behavior. The area in which this effect was most thoroughly investigated is that of suicides, though even for suicides, less is known about the specific attributes of stories that contribute to action. In-

Table 4: Spearman ρ between the actual NIAS values and those predicted using a regression model from document phrases and the Area Under the ROC curve (AUC) for identifying documents in the top 1% of NIAS values. $P < 10^{-5}$ for all models.

Topic	Spearman ρ	AUC
Suicides	0.477	0.943
HIV AIDS	0.346	0.894
Breast cancer	0.728	0.845
Mastectomies	0.705	0.910
Cancer	0.733	0.916
Average	0.525	0.902

deed, guidelines for reporting on suicides published by the US Centers for Disease Control [16] offer only general information on what aspects of suicides should not be reported. These include, for example, simplistic explanations of suicides, sensational coverage, the presentation of suicides as a means to an end, and the glorification of the person who committed suicide. Here we proposed a new methodology for detecting the effects of news coverage and demonstrate its application to both suicides and more positive health impacts. We further offer a model which can predict the effect of a given news story on future actions of the reader.

One of the limitations of this work is that it relies on action queries (AQs), that is, an information need that signifies a relevant intent (see Section 3). Naturally, many (perhaps even most) people making an AQ will not go on to perform an action. The vast majority of people who make a suicide AQ do not commit suicide, and (perhaps unfortunately) most people who search for cancer screening do not subsequently visit their doctor, even if guidelines do recommend that they undergo screening.

People are exposed to news through many channels. This means that it is difficult to directly identify all the sources of news a person is exposed to, leading us instead to use a probabilistic framework which infers the likelihood that a person was exposed to a news story based on her spatial and temporal distances from a news item. Our results, as well as prior work [24] shows that these parameters are strong predictors of exposure. Future work will investigate whether news items propagate differently via different media, such as social networks, leading to a more refined model

of news exposure. Incorporating user demographics into the exposure model is another area for future investigation.

Interestingly, though the time from publication to exposure takes many hours, the time from exposure to action is extremely short, typically only a few minutes. This has important implications for the topics in our study and more generally. First, there is little time to take mitigating actions once a person is exposed to a harmful news item. Therefore, news items should be evaluated for their effect prior to publication, as we describe below. Second, and more broadly, when for example, firms wishing to test the effectiveness of ads or news items about their company can expect that actions will follow almost immediately after exposure.

Our results show that NIAS values are predictable, based on the terms used in news articles. This implies that the appearance of Aqs can be predicted for a given news item with relatively high accuracy (AUC greater than 0.84 in all cases). Therefore, it is possible to evaluate a news story prior to publication and predict its effect before it is shown to people. For news stories on suicides or other cases where news stories might have negative effects, our models to predict NIAS values offer journalists a useful tool for evaluating (and perhaps thereafter modifying) their publications to reduce the chance of negative effects. For public health officials, our tools offer the ability to evaluate news releases to maximize their effectiveness. Future work will investigate how specific demographic traits interact with terms to form even more accurate models.

6. REFERENCES

- [1] L. Backstrom, J. Kleinberg, R. Kumar, and J. Novak. Spatial variation in search engine queries. In *Proceedings of the 17th international conference on World Wide Web*, pages 357–366. ACM, 2008.
- [2] P. N. Bennett, R. W. White, W. Chu, S. T. Dumais, P. Bailey, F. Borisyuk, and X. Cui. Modeling the impact of short-and long-term behavior on search personalization. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 185–194. ACM, 2012.
- [3] D. Carmel, E. Farchi, Y. Petruschka, and A. Soffer. Automatic query refinement using lexical affinities with maximal information gain. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 283–290. ACM, 2002.
- [4] Q. Cheng, F. Chen, and P. S. Yip. The foxconn suicides and their media prominence: is the werther effect applicable in china? *BMC public health*, 11(1):841, 2011.
- [5] F. Diaz. Integration of news content into web results. In *Proceedings of the Second ACM International Conference on Web Search and Data Mining*, pages 182–191. ACM, 2009.
- [6] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern classification*. John Wiley & Sons, 2012.
- [7] A. Frei, T. Schenker, A. Finzen, V. Dittmann, K. Kraeuchi, and U. Hoffmann-Richter. The werther effect and assisted suicide. *Suicide and Life-Threatening Behavior*, 33(2):192–200, 2003.
- [8] Q. Gan, J. Attenberg, A. Markowetz, and T. Suel. Analysis of geographic queries in a search engine log. In *Proceedings of the first international workshop on Location and the web*, pages 49–56. ACM, 2008.
- [9] A. Hassan, Y. Song, and L.-w. He. A task level metric for measuring web search satisfaction and its application on improving relevance estimation. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, pages 125–134. ACM, 2011.
- [10] K. Hawton, S. Simkin, J. J. Deeks, S. O’Connor, A. Keen, D. G. Altman, G. Philo, C. Bulstrode, et al. Effects of a drug overdose in a television drama on presentations to hospital for self poisoning: time series and questionnaire study. *Bmj*, 318(7189):972–977, 1999.
- [11] A. P. Institute. How americans get their news. <https://www.americanpressinstitute.org/publications/reports/survey-research/how-americans-get-news/>, 2014.
- [12] A. Kulkarni, J. Teevan, K. M. Svore, and S. T. Dumais. Understanding temporal query dynamics. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 167–176. ACM, 2011.
- [13] J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506. ACM, 2009.
- [14] Q. Mei, C. Liu, H. Su, and C. Zhai. A probabilistic approach to spatiotemporal theme pattern mining on weblogs. In *Proceedings of the 15th international conference on World Wide Web*, pages 533–542. ACM, 2006.
- [15] T. Niederkrotenthaler, M. Voracek, A. Herberth, B. Till, M. Strauss, E. Etzersdorfer, B. Eisenwort, and G. Sonneck. Role of media reports in completed and prevented suicide: Werther v. papageno effects. *The British Journal of Psychiatry*, 197(3):234–243, 2010.
- [16] P. W. O’Carroll and L. B. Potter. Suicide contagion and the reporting of suicide: Recommendations from a national workshop. *Morbidity and Mortality Weekly Report: Recommendations and Reports*, pages 9–18, 1994.
- [17] D. R. Olson, K. J. Konty, M. Paladini, C. Viboud, and L. Simonsen. Reassessing google flu trends data for detection of seasonal and pandemic influenza: a comparative epidemiological study at three geographic scales. *PLoS Comput Biol*, 9(10):e1003256, 2013.
- [18] D. P. Phillips. The influence of suggestion on suicide: Substantive and theoretical implications of the werther effect. *American Sociological Review*, pages 340–354, 1974.
- [19] M. Richardson. Learning about the world through long-term query logs. *ACM Transactions on the Web*, 2(4):21, 2008.
- [20] F. Richter. Americans use electronic media 11+ hours a day. <http://www.statista.com/chart/1971/electronic-media-use/>, 2015.
- [21] A. Schmidtke and H. Häfner. The werther effect after television films: new evidence for an old hypothesis. *Psychological medicine*, 18(03):665–676, 1988.
- [22] M. Sisask and A. Värnik. Media roles in suicide prevention: a systematic review. *International journal of environmental research and public health*, 9(1):123–138, 2012.
- [23] E. Yom-Tov and F. Diaz. Location and timeliness of information sources during news events. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, pages 1105–1106. ACM, 2011.
- [24] E. Yom-Tov and F. Diaz. Out of sight, not out of mind: on the effect of social and physical detachment on information need. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, pages 385–394. ACM, 2011.
- [25] E. Yom-Tov and F. Diaz. The effect of social and physical detachment on information need. *ACM Transactions on Information Systems (TOIS)*, 31(1):4, 2013.
- [26] E. Yom-Tov and L. Fernandez-Luque. Information is in the eye of the beholder: Seeking information on the mmr vaccine through an internet search engine. In *AMIA Annual Symposium Proceedings*, volume 2014, page 1238. American Medical Informatics Association, 2014.