Abyss or Shelter? On the Relevance of Web Search Engines' Search Results When People Google for Suicide

Mario Haim, Florian Arendt, and Sebastian Scherr

Department of Communication Studies and Media Research, Ludwig Maximilians University Munich

ABSTRACT

Despite evidence that suicide rates can increase after suicides are widely reported in the media, appropriate depictions of suicide in the media can help people to overcome suicidal crises and can thus elicit preventive effects. We argue on the level of individual media users that a similar ambivalence can be postulated for search results on online suicide-related search queries. Importantly, the filter bubble hypothesis (Pariser, 2011) states that search results are biased by algorithms based on a person's previous search behavior. In this study, we investigated whether suicide-related search queries, including either potentially suicide-preventive or -facilitative terms, influence subsequent search results. This might thus protect or harm suicidal Internet users. We utilized a 3 (search history: suicide-related harmful, suicide-related helpful, and suicide-unrelated) × 2 (reactive: clicking the top-most result link and no clicking) experimental design applying agent-based testing. While findings show no influences either of search histories or of reactivity on search results in a subsequent situation, the presentation of a helpline offer raises concerns about possible detrimental algorithmic decision-making: Algorithms "decided" whether or not to present a helpline, and this automated decision, then, followed the agent throughout the rest of the observation period. Implications for policy-making and search providers are discussed.

Suicide is a substantial public health problem, annually accounting for approximately 800,000 deaths worldwide (World Health Organization [WHO], 2014). In this regard, the mass media can be described as a "double-edged sword": On the one hand, previous research has revealed that suicide rates can increase after suicides are widely reported in the media(Phillips, 1974; Stack, 2005). On the other hand, appropriate depictions of suicide in the media can help people to overcome suicidal crises and can thus elicit preventive effects. For example, it has been shown that suicide rates can decrease after media reports on positive coping strategies (Niederkrotenthaler et al., 2010).

Media effects on suicide-related behavior are mostly explained through Bandura's (1977) social learning theory, stating that individuals are prone to imitating appropriate coping strategies or suicidal behaviors. Factors influencing media effects include a reported person's level of prominence (Ladwig, Kunrath, Lukaschek, & Baumert, 2012), similarity in terms of age and gender of the role model (Yip et al., 2006), as well as explicit or detailed descriptions of suicide methods (Chan, Lee, Lee, & Yip, 2003). Against the backdrop of these consistent findings, media organizations have advised their journalists to adequately report about suicide (e.g., Corbo & Zweifel, 2013; Mann et al., 2005). Most essentially, such reports should include coping strategies and exclude any romanticizing (e.g., "united for eternity") or rationalizing (e.g., "suicide was the logical next step") of a suicidal act (Niederkrotenthaler et al., 2010). Additionally, announcements of suicide telephone helplines is deemed appropriate (Sudak & Sudak, 2005), such as the announcement of telephone hotlines (e.g., Till, Sonneck, Baldauf, Steiner, & Niederkrotenthaler, 2013).

The importance of a person's online search behavior has already been acknowledged by previous research (Kemp & Collings, 2011; Recupero, Harms, & Noble, 2008). For example, Gunn and Lester (2013) found a correlation between the total amount of searches for suicide-related terms and actual suicide rates. Unfortunately, the role that search engines play when individual people search for suicide-related information is not fully understood.

Online search engines are increasingly seen as "personal gatekeepers" for the overwhelming amount of information: Through the personalization of online search results, providers pursue the goal of producing faster and more adequate results for their users. Based on individual users' contexts—such as a user's past search behavior or their geographic location, and aggregate statistical clustering—search engines such as *Google* claim to present different users with different selections and/or rankings for search results on the same search query in order to optimally match with a user's expectations (Feuz, Fuller, & Stalder, 2011), including

CONTACT Mario Haim haim@ifkw.lmu.de Department of Communication Studies and Media Research, Ludwig Maximilian University, Oettingenstr. 67, 80538 Munich, Germany.

All authors are research associates at the Department of Communication Studies and Media Research, LMU Munich.

the risk of partial information blindness: Beneficial information (e.g., preventive content) might be excluded by the algorithm as it is presumed to be of no vital interest to the user—the *filter bubble* hypothesis (Pariser, 2011).

In this study, we tested the filter bubble hypothesis in the suicide domain: Suicide-related search results can include preventive as well as facilitative information, both of which, in turn, have the potential to either protect (i.e., shelter) or harm (i.e., abyss) suicidal Internet users. As algorithms may recognize suicidal individuals' suicide-related search-behavior histories and personalize further web search results, they may have an important impact on suicide-related search results. Following the filter bubble hypothesis, we questioned whether individuals with different suicide-related search-behavior histories—they may have searched for suicide-facilitative (e.g., how to commit suicide) or suicide-preventive (e.g., how to cope with suicidal ideation) information—get biased search results in a subsequent situation.

To the best of our best knowledge, this is the first study that has looked into the assumed personalizing effects of search engine algorithms in the health communication context. Due to *Google*'s influential position, the present study focuses on *Google*.

Evidence for the Filter Bubble Hypothesis

Empirical evidence on the filter bubble phenomenon is rare. One example is a recently published study that provides evidence for small algorithmic effects in line with the filter bubble hypothesis, but stronger effects of individual's clicking behavior on the content presented by *Facebook* (Bakshy, Messing, & Adamic, 2015).

In the context of search engines, Feuz and colleagues (2011) detected a small degree of personalization within *Google*'s search results: Using a qualitative approach, the researchers created three different *Google* accounts, manifesting hypothetical individuals (i.e., similar to our agents) through an adequate usage of the search engine. After this training phase, equal searches were performed for all three accounts, testing their degree of personalization. The researchers found substantial personalization, but only when looking at several hundred search results per user—a rather unrealistic setting.

A study by Hannak and colleagues (2013) asked real human subjects to search for certain terms, logging only *Google*'s top-10 search results. Within these results, the researchers could not confirm filter bubble effects as they found only minor personalization effects (i.e., for logged in users and different geolocations). Moreover, most differences due to personalization referred to the ranking of search results as compared to the selection of the presented results. We try to contribute to these mixed findings by testing the filter bubble in the health communication context.

Research Questions and Hypotheses

Following the filter bubble hypothesis, we assume that if a user searches for query terms that can be assumed to contain suicide-related *harmful* information (i.e., "how to commit suicide"), then the filter bubble hypothesis suggests that results on subsequent search queries will contain more suicide-facilitative information. The same applies to suicide-related *helpful* queries and thus suicide-preventive results. The filter bubble claim reviewed earlier leads to the first hypothesis.

H1: The obtained suicide-related Internet search results are biased toward one's individual search history.

In addition to one's prior search behavior, Dou and colleagues (2007) have shown that clicks depict a major influence on personalization algorithms. A click on a specific search result item can be deemed as a confirmation of the search engine's algorithm. As their findings show, an algorithm accounting for a user's previous click decisions produces the best results. We put this into a further hypothesis.

H2: Clicking the first search result increases the effect that suicide-related Internet search results are biased toward one's individual search history.

As already noted, media outlets were instructed to provide offers of help (e.g., hotlines). *Google* introduced a "suicide prevention result" (Zeiger, 2010). Consequently, for certain search queries, *Google* presents a special helpline result above all other search results. This result does not depict a website but is presented as a box including a helpline that is presented "for certain search queries" (Zeiger, 2010). Unfortunately, *Google* lacks further transparency. Due to the importance of this matter for suicide prevention, we decided to put it into a research question. To the best of our knowledge, no prior research is available on these matters.

RQ1: How often does *Google*'s suicide-prevention result show up and how is its presence influenced by a user's search and click history?

Method

We modeled online behaviors of N = 1,200 virtual agents using agent-based testing. We utilized a three (search history terms: suicide-related harmful, suicide-related helpful, and suicide-unrelated) × 2 (reactive: clicking the top-most result link and not clicking any link) experimental design. Thus, every experimental group consisted of 200 virtual agents. The whole study consisted of two phases: a one-week training phase and a subsequent test phase. Our software is built upon the CasperJS.org framework, which enables a server-based scraping implementation of JavaScript. In addition to the emulation of single participants, our software simulated multiple participants through regular repetition.

In the *training phase*, we simulated search engine user behavior in order to train the agents. That is, the computer randomly chose one agent of one of the six experimental groups, loaded this agent's past cookies, and randomly selected one search query out of the agent's experimental group's word list (i.e., suicide-related harmful, suicide-related helpful, suicide-unrelated). Then the computer navigated to https://www.google.de and searched for the selected search query. The first 10 results and their order were stored in addition to whether *Google*'s suicide-prevention result (i.e., helpline) was displayed. The agent then randomly clicked on the first search result or did not. Moreover, all cookies were individually saved for each agent. This training procedure led to $M = 19.1$ ($SD = 1.4$) training sessions for each of the 1,200 virtual agents. This roughly corresponds to 23,000 search queries and a total of nearly 64,000 cookies.

For the construction of the search terms, we followed Wong et al. (2013), who differentiated three main types of online search behaviors: (1) suicidal searches without any request for help (harmful search behavior), (2) suicidal searches incorporating a demand for help (helpful search behavior), and (3) a control group (neutral search behavior). Each search term list comprised a set of words or short phrases. The harmful search terms were formulated as suggested by Biddle, Donovan, Hawton, Kapur, and Gunnell (2008). Their list of 12 search terms was translated into German using two common expressions for "suicide" (*Selbstmord, Suizid*). As some translations were grammatically not applicable, the final list comprised 22 harmful search terms (e.g., "best method for suicide"). The list of helpful terms (e.g., "overcoming suicidal thoughts") also includes a total of 22 queries, which were also created in accordance with those proposed by Biddle et al. (2008). The full list can be obtained upon request.

For the control group, we used the well-established Berlin Affective Word List (Võ et al., 2009) that consists of 2,900 pretested German words, which had been validly rated in terms of emotional valence, arousal, and imageability. We filtered the list depending on the words' emotional valence —originally ranging from very negative (−3) to very positive (+3)—into a list of more valence-neutral words rated between −2 and +2. The control group's ultimate list consisted of 2,647 words and did not include any suicide-related terms.

In the subsequent *test phase*, the computer randomly chose one agent of one of the six experimental groups. The previously stored individual cookies for each agent were again loaded in order to allow the search engine to identify a revisiting user (i.e., the search history). The dependent variables were measured for every single agent by performing two separate neutral suicide-related terms that allow for disambiguation (i.e., "suicide" and "suicide method" in German).

All simulations took place over the course of one week in February 2015 on one computer. IP addresses varied over time due to provider settings. However, they did not vary systematically from agent to agent. This is beneficial because it equalizes IP address influence (Hannak et al., 2013). The computer was run in a large German-speaking city without any other applications running in order not to reveal any academic background (e.g., through IP subnet identification).

Dependent Variables

Google Search Results

For both queries ("suicide" and "suicide method"), all search result links from the first result page were stored. Any sponsored links as well as results from *Google News* or *Google Images* were ignored. We collected all of the website addresses (URLs) as well as their ordinal rank within the results. The procedure yielded between 8 and 10 result links for each agent.

Suicide-Prevention Result

We collected dichotomous information about the presence of the suicide-prevention result for every agent. We coded whether the helpline was shown (coded as 1) or not (coded as 0).

Result

Search Results

Result pages were almost identical for all agents. As the visual length of *Google*'s first result page is held approximately constant, the number of results per agent depends on the appearance of the suicide-prevention result and varies between 8 and 10. Thus, for all 1,200 agents, a post hoc derived maximum sum of 9,605 different URLs would have been possible for the first search (suicide) and 10,805 for the second search (suicide method). However, only a total of 12 (first search term "suicide") and 15 (second search term "suicide method") unique search results appeared. In other words, for every single result link, one specific URL dominated over almost all agents. Only very few agents were presented with deviating result links. For example, for the first search (suicide), each (100%) of the 1,200 agents was presented with the Wikipedia page on suicide as the top result, whereas the second result was primarily taken up by a publisher on health literature (n = 1,183) and was subordinately occupied by an accompanying website for a book on psychosocial health (n = 17). The latter 17 agents, however, were presented with the publisher's website in the third position, thus indicating a small variation, not in terms of selection, but rather in terms of prioritization.

This lack of variance led us to an alternative (i.e., more descriptive) way of handling the data. We planned to test for group differences by coding single websites in order to assign indices for (1) suicide-preventive, (2) suicide-facilitative, and (3) neutral search results. However, this analysis strategy could not be utilized due to the (surprising) lack of variability regarding the search result: Only 2 out of 12 websites within the first search using "suicide" were presented exclusively to a subgroup of random agents—the other 10 websites were shown to all 1,200 agents. For the second search using "suicide method," this finding applies to 3 out of 15 websites (as visualized in Figure 1).

Hypothesis 1 (H1) predicted that the obtained suicide-related Internet search results would be biased toward one's individual search history (i.e., the filter bubble hypothesis). However, due to the lack of variance, this hypothesis was not supported by the data. There was simply no variability. Similarly, Hypothesis 2 (H2) predicted that regularly clicking the first search result would increase the filter bubble effect specified in H1. Similarly, as with the first hypothesis, clicking could not show an effect due to the lack of variability.

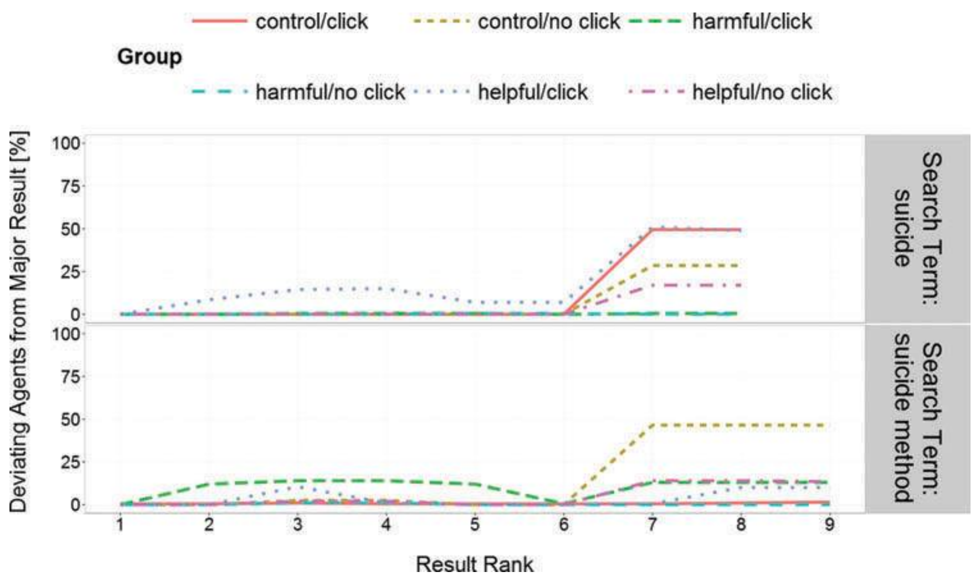

Figure 1. Share of agents with deviating URL from major result (ordinate axis) per ordinal result rank (abscissa axis). Line styles depict experimental groups.

Helpline

Research question 1 asked how often *Google*'s suicide-prevention result shows up and how a user's search and click history influences its presence. As noted, a dichotomous value (whether the helpline was presented or not) was coded per agent and for each training cycle, resulting in 22,962 binary measures (1,200 agents × M = 19.1 training cycles).

We found an interesting pattern (visualized in Figure 2). As expected, the suicide-prevention result was not presented to any agent of the conditions searching for suicide-unrelated terms (control group). Conversely, it was shown in 11% of all cases where helpful search terms were applied. Of interest, harmful search queries resulted in the presence of the suicide-prevention result in 25% of the cases where agents did not regularly follow the top-most result (i.e., nonreactive "no click" condition) and in 31% of all reactive cases (i.e., "click" condition). The percentage values, surprisingly, did not change over time. They were independent of the agent's search history. In fact, the percentage values were constant over time, which is why we could not rely on significance testing.

Post Hoc Analysis

What is the reason for the constant values? We questioned whether it was related to the specific search terms. However, there was also no clear evidence of a relationship between the helpline's appearance and the specific search term entered. Importantly, a follow-up analysis revealed spillover effects as agents either were presented with the suicide-prevention result all the time or never at all: If an agent was presented with the suicide-prevention result for the first search query, the agent was presented with the helpline for all subsequent

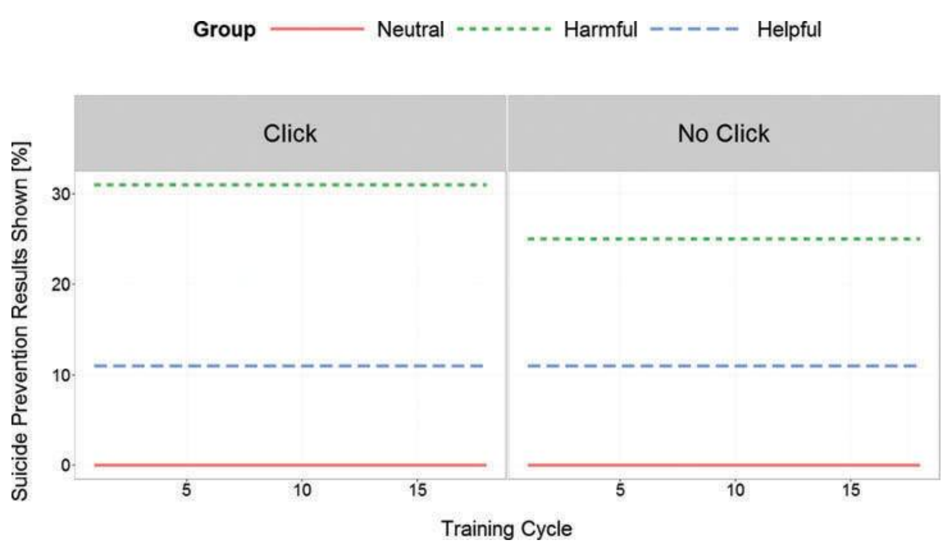

Figure 2. Share of suicide-prevention results over time and across experimental groups. Ordinate axis depicts the share of suicide-prevention results. Abscissa axis depicts the training cycle in question. That is, no matter if an agent queried *Google* for the first or fifth time, the share of queries for which the suicide-prevention result was shown did not vary.

queries as well. A total of 86 out of 1,200 agents were placed in this category, receiving a virtual "suicide-affiliation tag" by *Google*. As this is an important finding for suicide prevention, we named this post hoc finding the *tag hypothesis*. We offer three possible explanations for this phenomenon.

(1) The identified spillover effects called for a post hoc analysis of the *first search query* per agent as this initial term seemed to determine whether the suicide-prevention result tag was applied. Yet, there was no clear pattern on this term–result relationship. Of interest, 32 out of 44 suicide-related search terms led to the helpline result and thus to the tag. Moreover, every suicide-related term that was randomly selected by 1 of the 86 tagged agents for their initial search was also selected a minimum of nine times for initial searches by other agents, which did not lead to the helpline result. An explanation using the dependence on the initially used search term can thus be rejected.

(2) *Time of the day* depicted another varying factor within the first search and thus represents another possible reason for tag variance. Yet, while 86 agents receiving the helpline result conducted their first search equally distributed throughout the first day of training, the 1,114 remaining agents did so as well. That is, the first tagged agent started its training at 6:16 a.m. and the last one ended its initial search at 8:53 p.m., whereas the first untagged agent started at 5:58 a.m. and the last one ended at 9:15 p.m. Thus, the time of the day does not seem to influence the tagging of agents. There was no clear pattern.

(3) Despite these somehow controllable influences, a *random* factor could also be included. Such a random setting would include an algorithm's analysis of an agent's search query in order to apply a rating regarding whether the helpline should be shown or not. Subsequently, this rating would be subject to adaptation in two possible ways: Either a completely random factor is applied or an incrementing counter depicts a ratio of users (agents) that are presented with the helpline (i.e., every *n*-th user with a substantial suicidal rating). This random factor explanation, however, cannot be verified validly but instead needs to be explained by the search provider. Yet, when confronted with these findings, *Google* did not answer our requests. We return to this finding in the following section.

Discussion

The present study looked at filter bubble effects in the context of suicide-related search queries. Our findings showed that search histories did not influence search results for a subsequent search query. Search result personalization took place, if at all, by rearranging identical results. This is in line with some similar findings (e.g., Dou et al., 2007; Hannak et al., 2013). Previous studies documenting filter bubble effects (e.g., Feuz et al., 2011) find rather small effects of personalization, but mostly in highly artificial study settings. Additionally, none of the mentioned studies (including our own) was capable of identifying a general filter effect. That is, no study can support the claim that *Google* blinds out (i.e., censors) specific information. A personalizing filter bubble effect (at least) in the context of suicide-related search queries must be rejected.

We obtained an important result for suicide prevention: *Google*'s suicide-prevention result (helpline box) was shown to a limited number of agents only. This is unfortunate because the helpline box should be presented in more search results for an increase in beneficial preventive effects. *Google*'s algorithms seemed to attach a tag to every agent based on his or her first search query—a tag the agents could not get rid of throughout future search queries. For the presentation of the helpline, we offered three possible explanations—search term dependency, relationship to the time of a day, and randomness following search provider systematics—where only the third explanation could align with empirical findings. Due to the critical nature of our context, where the presentation of a helpline might prevent actual suicides, this finding suggests that more algorithmic transparency is required in the context of search engines (e.g., Diakopoulos, 2014). Future discussions should focus on a higher appearance rate of helplines, more transparency in terms of their search algorithm, or a (transparent) censorship of potentially harmful results.

As with every study, the present one has its limitations. First, we tested the filter bubble hypothesis within the specific context of suicides. Future studies should build upon the gained knowledge and extend it to other health communication topics. Second, we used agent-based testing as a method to test the filter bubble hypothesis. Although the present study provides a methodological surplus value as it combines real-world observation with the benefits of virtual-agent testing (e.g., no harm to real humans), the method may elicit external validity concerns. For example, we used a total of 22 search terms for harmful and helpful searches. However, real media users can use many more search terms.

Yet, despite the limitations, the current study revealed two-fold results on the filter bubble hypothesis in the suicide context. While personalization obviously does not apply in terms of search results, algorithmic effects on partial information blindness were found in terms of helpline presentation. The latter manifests itself in a phenomenon we termed the *tag hypothesis*. That is, algorithms "decided" when an agent submitted the very first search query whether or not to present a helpline (i.e., whether or not to attach a tag to the agent). This automated decision, then, followed the agent throughout the rest of the experiment. Yet, as it cannot be ascribed to a user's prior search behavior, it might not be an effect essentially similar to the filter bubble hypothesis. While this finding raises great concerns about autonomous algorithmic decision-making, it also raises more explicit implications: (1) Due to the changing media environment and the increasing importance of the Internet as a source for health-related information, a targeted research effort on factors influencing online search behavior is important. Following *Google*'s announcements on the suicide-prevention result, the helpline presentation differs slightly from country to country, thus requiring further (country-specific) research. (2) The number of agents receiving the suicide-prevention box was small. We suggest increasing the rate of helpline presentation.

Acknowledgments

The findings reported in the manuscript have not been published previously, the manuscript is not simultaneously under consideration

elsewhere, and the findings have complied with the American Psychological Association's ethical standards in the treatment of any participants in the work being reported.

References

Bakshy, E., Messing, S., & Adamic, L. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, *1160*, 1–2. http://doi.org/10.1126/science.aaa1160

Bandura, A. (1977). *Social learning theory*. Englewood Cliffs, NJ: Prentice-Hall.

Biddle, L., Donovan, J., Hawton, K., Kapur, N., & Gunnell, D. (2008). Suicide and the internet. *British Medical Journal*, *336*, 800–802. doi:10.1136/bmj.39525.442674.AD

Chan, K. P. M., Lee, D. T. S., Lee, S., & Yip, P. S. F. (2003). Media influence on suicide. Media's role is double edged. *British Medical Journal*, *326*, 498–499. doi:10.1136/bmj.326.7387.498

Corbo, A. M., & Zweifel, K. L. (2013). Sensationalism or sensitivity: Reporting suicide cases in the news media. *Studies in Communication Sciences*, *13*, 67–74. doi:10.1016/j.scoms.2013.04.012

Diakopoulos, N. (2014). Algorithmic accountability. Journalistic investigation of computational power structures. *Digital Journalism*, *3*, 398–415. doi:10.1080/21670811.2014.976411

Dou, Z., Song, R., & Wen, J.-R. (2007). *A large-scale evaluation and analysis of personalized search strategies*. In Proceedings of the 16th international conference on World Wide Web (pp. 581–590). New York, NY: ACM. http://doi.org/10.1145/1242572.1242651

Feuz, M., Fuller, M., & Stalder, F. (2011). Personal web searching in the age of semantic capitalism: Diagnosing the mechanisms of personalisation. *First Monday*, *16*(2). doi:10.5210/fm.v16i2.3344

Gunn, III, J. F., & Lester, D. (2013). Using google searches on the internet to monitor suicidal behavior. *Journal of Affective Disorders*, *148*, 411–412. doi:10.1016/j.jad.2012.11.004

Hannak, A., Sapiezynski, P., Molavi Kakhki, A., Krishnamurthy, B., Lazer, D., Mislove, A., & Wilson, C. (2013). *Measuring personalization of web search*. In Proceedings of the 22nd international conference on World Wide Web (pp. 527–538). Rio de Janeiro: International World Wide Web Conferences Steering Committee.

Kemp, C. G., & Collings, S. C. (2011). Hyperlinked suicide: Assessing the prominence and accessibility of suicide websites. *Crisis*, *32*, 143–151. doi:10.1027/0227-5910/a000068

Ladwig, K.-H., Kunrath, S., Lukaschek, K., & Baumert, J. (2012). The railway suicide death of a famous German football player: Impact on the subsequent frequency of railway suicide acts in Germany. *Journal of Affective Disorders*, *136*, 194–198. doi:10.1016/j.jad.2011.09.044

Mann, J. J., Apter, A., Bertolote, J., Beautrais, A., Currier, D., Haas, A. … Hendin, H.; others. (2005). Suicide prevention strategies: A systematic review. *Jama*, *294*, 2064–2074. doi:10.1001/jama.294.16.2064

Niederkrotenthaler, T., Voracek, M., Herberth, A., Till, B., Strauss, M., Etzersdorfer, E. … Sonneck, G. (2010). Role of media reports in completed and prevented suicide: Werther v. Papageno effects. *The British Journal of Psychiatry*, *197*, 234–243. doi:10.1192/bjp.bp.109.074633

Pariser, E. (2011). *The Filter Bubble: How the new personalized web is changing what we read and how we think*. New York, NY: Penguin.

Phillips, D. P. (1974). The influence of suggestion on suicide: Substantive and theoretical implications of the Werther effect. *American Sociological Review*, *39*, 340–354. doi:10.2307/2094294

Recupero, P. R., Harms, S. E., & Noble, J. M. (2008). Googling suicide: Surfing for suicide information on the Internet. *The Journal of Clinical Psychiatry*, *69*, 878–888. doi:10.4088/JCP.v69n0601

Stack, S. (2005). Suicide in the media: A quantitative review of studies based on nonfictional stories. *Suicide & Life-Threatening Behavior*, *35*, 121–133. doi:10.1521/suli.35.2.121.62877

Sudak, H. S., & Sudak, D. M. (2005). The media and suicide. *Academic Psychiatry*, *29*, 495–499. doi:10.1176/appi.ap.29.5.495

Till, B., Sonneck, G., Baldauf, G., Steiner, E., & Niederkrotenthaler, T. (2013). Reasons to love life. *Crisis*, *34*, 382–389. doi:10.1027/0227-5910/a000212

Võ, M. L.-H., Conrad, M., Kuchinke, L., Urton, K., Hofmann, M. J., & Jacobs, A. M. (2009). The Berlin Affective Word List Reloaded (BAWL-R). *Behavior Research Methods*, *41*, 534–538. doi:10.3758/BRM.41.2.534

Wong, P. W.-C., Fu, K.-W., Yau, R. S.-P., Ma, H. H.-M., Law, Y.-W., Chang, -S.-S., & Yip, P. S.-F. (2013). Accessing suicide-related information on the internet: A retrospective observational Study of search behavior. *Journal of Medical Internet Research*, *15*(1). doi:10.2196/jmir.2181

World Health Organization [WHO]. (2014). *Preventing suicide. A global imperative*. Luxemburg: WHO.

Yip, P. S. F., Fu, K.-W., Yang, K. C. T., Ip, B. Y. T., Chan, C. L. W., Chen, E. Y. H., & Hawton, K. (2006). The effects of a celebrity suicide on suicide rates in Hong Kong. *Journal of Affective Disorders*, *93*, 245–252. doi:10.1016/j.jad.2006.03.015

Zeiger, R. (2010, November 11). Helping you find emergency information when you need it. Retrieved from http://blog.google.org/2010/11/helping-you-find-emergency-information.html