Viewpoint

Ethical Challenges and Opportunities Associated With the Ability to Perform Medical Screening From Interactions With Search Engines: Viewpoint

Elad Yom-Tov[1,2], PhD; Yuval Cherlow[3], Rabbi

[1]Microsoft Research, Herzeliya, Israel

[2]Faculty of Industrial Engineering and Management, Technion, Haifa, Israel

[3]The Tzohar Rabbinical Organization, Lod, Israel

Corresponding Author:
Elad Yom-Tov, PhD
Microsoft Research
Alan Turing 3
Herzeliya, 4672415
Israel
Phone: 972 747111359
Email: eladyt@microsoft.com

Abstract

Recent research has shown the efficacy of screening for serious medical conditions from data collected while people interact with online services. In particular, queries to search engines and the interactions with them were shown to be advantageous for screening a range of conditions including diabetes, several forms of cancer, eating disorders, and depression. These screening abilities offer unique advantages in that they can serve a broad strata of the society, including people in underserved populations and in countries with poor access to medical services. However, these advantages need to be balanced against the potential harm to privacy, autonomy, and nonmaleficence, which are recognized as the cornerstones of ethical medical care. Here, we discuss these opportunities and challenges, both when collecting data to develop online screening services and when deploying them. We offer several solutions that balance the advantages of these services with the ethical challenges they pose.

KEYWORDS

search engines; diagnosis; screening

Introduction

Recent work has demonstrated the ability to screen for serious medical conditions using search engine logs [1-5]. The development and deployment of these abilities can open new opportunities for earlier diagnosis and more equitable care but require careful consideration of the associated ethical challenges. The goal of this paper is to discuss the ethical pros and cons of these capabilities and to set the stage for a broader discussion of these issues.

Search engines are used by the vast majority of internet users to obtain information on a variety of topics, including medicine [6]. Search engine operators collect information on the interaction of users with their services to improve the operation of their search engines, for example, by measuring user satisfaction from specific answers given to them [7]. It is important to stress that the data collected by search engine operators are not collected to improve medical research or improve people's health but to enhance search engine operation.

The data collected by search engine operators include, for example, query text, links shown to the user, time of clicking on the links, duration of reading each link, and mouse movements, which serve as a proxy for eye gaze tracking [8]. The data collected by search engines are usually anonymous, in the sense that specific individuals cannot be easily linked to their data, but, unless specified by the user, multiple searches can be attributed, with high likelihood, to the same user.

As noted above, these data have been shown to be effective for screening people for a variety of medical conditions, such as diabetes [4], several forms of cancer [1-3], eating disorders, and depression [5]. Interactions with search engines are useful for such screening because of a combination of factors, including people's limited knowledge of the association between symptoms and conditions [9]; the fact that many conditions (eg,

ovarian cancer) have benign symptoms, of which only the confluence indicates disease, but psychological biases lead people to focus on only the latest symptom [2]; and people's natural tendency to defer treatment but ask about it online.

These screening capabilities offer unique advantages in that they can serve a broad strata of the society, including people in underserved populations and in countries with poor access to medical services [7]. However, these advantages come at a possible cost to privacy, autonomy, and nonmaleficence, which are recognized as the cornerstones of ethical medical care [10]. Note that the legal aspects of providing (and of not providing) these capabilities are not discussed in this work.

We note that other services, including content providers (eg, Wikipedia and patient groups [11]) and social media platforms (eg, Facebook [12] and Twitter [13]) collect similar data. However, for the reasons described above, we focus on search engines. Moreover, for a broader discussion on the ethics of internet research, readers can refer to the article by Buchanan and Zimmer [14].

In our opinion, the ethical questions that arise from the ability to screen search engine logs should be divided into questions that appear during the development of screening capabilities and questions that should be resolved before medical interventions are provided to people as part of the use of a resulting product. Here, we discuss both these areas.

Development

Incidental Finding

"Incidental finding" [15] refers to the case where, during research on one medical topic, data indicates that a person under study has another medical condition of which he/she is (possibly) unaware. Consider a person who contributes their genetic information to build a new screening test for a specific hereditary condition. Upon examination, researchers realize that this person's genetic information reveals that he/she has another, perhaps common, mutation, which indicates that the person has a serious medical condition that he/she may not know about.

The commonly accepted solution to this challenge in genetic research is to screen for mutations that are common, life-saving, and do not require the person to have a deep understanding of genetics in order to decide whether he/she would like to be treated. If such a mutation is identified, the researcher informs the person that he/she should consult with a genetic counsellor but does not provide advice, as this is not the researcher's specialty. This route is taken also because, if the burden of treatment (or advice) is placed on the researcher, medical research will, in practice, be severely restricted. This is also the reason that, in many cases, ethics committees recommend opting for completely anonymous research, which reduces the ethical burden on the researcher.

We claim that there is similarity between the question of incidental finding in the medical domain and the case where researchers use data collected during people's interactions with search engines to later determine that a user may have a medical condition. This can arise from a simple interaction, such as a

query suggesting suicidal ideation, to a more elaborate insight obtained from a predictive model based on interactions with the internet service. However, the analogy is not perfect. People who donate their data for medical research know that their data will be examined for medical purposes, whereas people who use a search engine do not expect their data to be used for medical research. In fact, in many cases, people who use search engines may not realize that their interaction data are being collected. We note in passing that routine experiments, such as Facebook's modification of the order of postings by friends, caused an uproar when they were described in an academic paper [16].

Nevertheless, we argue that these differences should not prevent us from using the insight medical ethics has garnered on the question of incidental findings because people who contribute their data to medical research may not realize that other findings are possible, and on the other hand, as public awareness of search engine data grows, people will realize that these data can provide them with benefits.

Informed Consent and Autonomy

Search engines, as other internet services, have a system of consent that often includes the use of data collected by the search engine for research purposes. People who use a search engine implicitly consent to its use and further can click on the link at the bottom of the search page where its "Terms of Use" are specified. However, it is difficult to refer to this as informed consent in the medical sense. For example, the authors found that in a sample of approximately 116 million users, only 0.05% clicked on the Terms of Use page during a 1-month period. Experience from other web services suggests that even if a pop-up window would require people to consent to their data being used for research, most people would click on the window without considering what they are consenting to [17].

Additionally, Terms of Use are necessarily broad in their description and, we assume, are often broader than informed consent forms signed by people participating in medical research.

Thus, it is still a challenge to develop a form of consent that both satisfies the ethical requirements for data use and does not overburden users in their interactions with the search engine.

Willingness to Provide Search Data for Medical Research

There is often an implicit assumption that people would not want their data, collected for other purposes, to be used for medical research without their specific consent. Gefen et al [18] tried to quantify the value that people assign to their data and found that, in a sample of people from around the world, 99% were willing to provide their search engine data in exchange for monetary compensation lower than US $1500 and 53% were willing to pay to have their data analyzed, even if the value of the analysis would be to the society at large rather than to them directly.

Thus, while a minority of users would not agree to the use of their data regardless of compensation, many would agree, and a relevant portion of the population even sees value in the

availability of these services, which exceeds that of the data itself.

Anonymity

As noted above, most of the search data used for medical research has, to date, been anonymous as far as researchers are concerned. This anonymization is provided through the provision of a random user identifier and by not including information that could easily compromise anonymity (eg, location). However, as shown in the AOL leak [19], a malicious researcher may be able to identify a small number of users when such anonymization is used. Therefore, it may be necessary to assume that data are not fully anonymized to a malicious researcher and perhaps even sometimes to a benevolent researcher.

Companies collecting data may, on the other hand, be able to identify a user. This can happen, for example, if users register to their services with their real name. Thus, even if data are anonymous to a researcher, it could, conceivably, be deidentified by the organization collecting the data. In such a case, the problem of incidental findings can arise, as described above.

Finally, an advantage to having data linked to an individual (either anonymous or identifiable) is the ability of users to control the use of their data, as offered, for example, in the European Union's recent General Data Protection Regulation (GDPR).

Representativeness

The question of representation in internet data appears in several forms. First, there are the questions of who uses the services from which data are collected and whether they faithfully characterize the entire population. Second, not all people use the internet in similar ways to acquire information, which causes another form of bias in the data.

The first source of representation bias could greatly impact populations, especially in financially disadvantaged parts of the society and in countries with lower access to the internet. Although many efforts have been devoted to closing this gap, it still exists. For example, the percentage of people with access to the internet in different countries ranges from almost 0% to 100% [20]. Thus, it is important to account for such representation biases when using the data to build a model that can be useful to all people.

The second source of representation bias is less well known but is no less important. As shown in past work [21], the use of search engines for medical queries, for example, is highly dependent on people's age and gender. Moreover, only around 16% of people use search engines to query for medical information [9], adding to the representation bias.

Summary

Taking the above-mentioned points into consideration, we suggest that in the case of research on medical conditions using search engine data collected for operational purposes, it may be preferable to use anonymized data rather than to obtain consent for the use of identifiable information. If the former route is taken, it is important that ethical committees approve

the research, serving in their capacity as representatives of society. We recognize that this is an imperfect solution because of both the inability to seek informed consent and the difficulty for ethics committees to represent search engine users who come from a range of countries and societies, each with its own norms and expectations. However, we view this as a balance between the competing challenges outlined above.

Production

Approaches for Providing Search-Based Screening Information

Once a screening method is developed, it may be put into regular use. This could be done in several ways, which are described below.

Suppose an anonymous search engine user is predicted to have a medical condition (eg, screened positive according to interactions with the search engine for the medical condition). The first and most intrusive way to provide the user with this information is to display a notice at a prominent location on the screen. This is currently done only to people who search for information on how to kill themselves [22] or for related topics. In such cases, a banner notice is displayed with the telephone number of a local helpline.

Another way that could be used is to bias (modify) search results toward suggesting the suspected condition. For example, if a user searches for "constant thirst," instead of showing the regular set of results, users who are predicted to have diabetes will be shown more results that suggest diabetes. A similar "personalized search" is currently part of the service of all major search engine providers (eg, when results are served such that they are relevant to the user's current location). Therefore, such a solution might not be perceived as a major change by users.

The third way we envision to display information is through the use of search advertisements [3]. Advertisements are not part of the main search results ("organic results") and are assessed differently by users [23]. People who search for diagnostic information ("do I have diabetes?") will be shown advertisements that would suggest help in diagnosis ("Worried you have diabetes? Click here to obtain more information"). People who click on the advertisements will be given diagnostic assistance, for example, in the form of clinical questionnaires. As shown recently [3], it is possible to train advertising systems to focus on people who are the most at risk.

We note that advertising in the health domain is currently limited by the policy of advertising systems. On one hand, this prevents abuse of the system by purveyors of unapproved medical services, but also means that any use of this method will often require approval by advertising system managers.

A fourth method to inform people of a possible medical concern is through the normal use of a search engine after first obtaining informed consent to provide these insights. If this method is adopted, users will be shown an informed consent form whenever they are identified as new users by the search engine. The form will offer the users to receive screening information but will default to not receiving the information unless the users

positively indicate their interest in receiving this information. Users who consent will then be given alerts whenever a possible medical condition is predicted, based on their queries and behaviors.

Finally, a system might be built where users register and agree to provide their search data on a continuous basis in exchange for alerts when a medical concern is identified in these searches. The data collection, storage methods, and data use would be clearly described to the user. This is similar to services that analyze people's genetic material, where their data will be the search data (or browsing data, in general) and the analysis will be conducted on an ongoing basis, rather than a single transfer of data. Such a system could be offered by search engine providers or, perhaps preferably, by medical providers or dedicated companies.

Unsolicited Diagnosis

Unsolicited diagnosis [24] or unsolicited medical opinion [25] refers to the case where people may be provided with medical information when they do not expect it. For example [26], consider the case of a dermatologist who is standing at the back of an elevator at the mall and notices that the person standing in front has a mole that the dermatologist thinks is likely cancerous. In this case, the person who has the mole is not expecting to receive a diagnosis from a random person at the mall (though a specialist in this case), and thus, this is a case of unsolicited diagnosis. Medical ethicists have considered the question of whether the medical specialist has a duty to inform the person to seek medical attention and whether the specialist has a right to do so. On one hand, the person is not expecting a diagnosis and there is no doctor-patient relationship between the two. On the other hand, not informing the person may lead to serious and irreparable damage. The conclusion reached by some ethicists [26] is that medical doctors have a duty to offer their unsolicited medical opinion, especially when the medical condition requires urgent attention for treatment. However, doctors need to consider the possible harm of such an intervention. Note that a legal duty to act is very much country-specific, often defined through legislation (ie, "good Samaritan" laws [27]) to protect people who take such action.

We note in passing that the balance between benefit and harm for the individual may differ from that for the society. For example, some conditions currently have no treatment because they cannot be identified early enough and so many people would prefer not to know that they have such conditions. However, suppose search engine data could provide such an early alert [28]. In such a case, if enough people knew they have these conditions, pharmaceutical companies might be compelled to develop treatments. However, as this is a secondary effect, we have not focused on it.

Risk Compensation

Risk compensation (also referred to as moral hazard [29]) describes increased risk taking caused by the perceived usefulness of safety measures. For example, it has been suggested that condom distribution fosters inhibition among HIV-positive people [30].

If internet platforms disclose offering screening services, users may choose to modify their behaviors in ways that could harm them. For example, as noted above, only around 16% of users queried about medical symptoms prior to diagnosis [9]. It is difficult to predict illness for people who do not query, but they may assume that a screening model is examining their queries and will alert them when it is necessary to visit a health provider, thus preferring not to access medical care even when they think they should. This is especially likely in the fifth solution described above, because users who register with a dedicated service expect it to provide such alerts. Therefore, it may be important for such a service to alert users about its inability to screen when it predicts that they will not ask relevant questions.

Cost of Errors

No system is perfect, including those discussed in this paper. The cost of errors is an important factor in whether and how information should be provided to users. A false positive error means that a person is informed (depending on the method of provision described above) of a medical condition when he/she does not have one. This can cause undue stress and result in unnecessary medical procedures [31]. A false negative error means that a person who should have been provided with an alert does not receive one, possibly causing late diagnosis (as described above).

Summary: Advantages and Disadvantages of Different Notification Methods

The first method described above, whereby a notice is shown to the user, is advantageous in that it provides people with immediate, clear, and actionable information. However, we advocate its use in only the most extreme situations (eg, expressed intention of suicide) because it is intrusive and may cause more harm than good in the form of breaching privacy and impinging on people's autonomy.

The second (biasing results), third (advertisements), and fourth (explicit prior informed consent) methods are advantageous in that they do not force information upon users and allow users to decide if they would like to use the offered information. These methods (especially biasing results and advertisements), however, somewhat impinge on privacy and autonomy. Additionally, not all users will recognize the help offered to them, and only some will make use of it even when they recognize it. We note that in the case of advertising, the act of choosing to click on the advertisement should be considered informed consent (assuming that the advertisement is appropriately phrased). We also recognize that obtaining explicit informed consent (eg, the fourth method) can be difficult from a design perspective and can burden users who, for example, use private browsing. For these reasons, we believe that these methods, especially the one using advertisements, correctly balance benefit and harm.

Finally, the fifth method (dedicated system) is clearly superior in terms of autonomy and consent, but based on prior experience, we assume that only a minority of users, probably skewed toward the more affluent parts of society, will use this method. Thus, while extremely beneficial for individuals, it should be considered less useful at the societal level.

Authors' Contributions

EYT conceived this paper. YC and EYT discussed the questions raised herein, suggested relevant past work, and wrote the paper.

Conflicts of Interest

YC declares no conflicts of interest. EYT is an employee of Microsoft, owner of the Bing search engine. The positions described in the paper are those of the authors and not necessarily those of their respective employers.

References

1. White RW, Horvitz E. Evaluation of the Feasibility of Screening Patients for Early Signs of Lung Carcinoma in Web Search Logs. JAMA Oncol 2017 Mar 01;3(3):398-401. [doi: 10.1001/jamaoncol.2016.4911] [Medline: 27832243]

2. Soldaini L, Yom-Tov E. Inferring Individual Attributes from Search Engine Queries and Auxiliary Information. In: WWW '17: Proceedings of the 26th International Conference on World Wide Web. 2017 Presented at: 26th International Conference on World Wide Web; April 2017; Perth, Australia. [doi: 10.1145/3038912.3052629]

3. Yom-Tov E. Screening for Cancer Using a Learning Internet Advertising System. ACM Trans. Comput. Healthcare 2020 Apr 15;1(2):1-13. [doi: 10.1145/3373720]

4. Hochberg I, Daoud D, Shehadeh N, Yom-Tov E. Can internet search engine queries be used to diagnose diabetes? Analysis of archival search data. Acta Diabetol 2019 Oct;56(10):1149-1154. [doi: 10.1007/s00592-019-01350-5] [Medline: 31093762]

5. Schueller SM, Steakley-Freeman DM, Mohr DC, Yom-Tov E. Understanding perceived barriers to treatment from web browsing behavior. J Affect Disord 2020 Apr 15;267:63-66. [doi: 10.1016/j.jad.2020.01.131] [Medline: 32063574]

6. Fox S, Duggan M. Health Online 2013. Pew Research Center. 2013. URL: https://www.pewinternet.org/wp-content/uploads/sites/9/media/Files/Reports/PIP_HealthOnline.pdf [accessed 2020-03-01]

7. Youngho K, Hassan A, White RW, Zitouni I. Modeling dwell time to predict click-level satisfaction. In: WSDM '14: Proceedings of the 7th ACM international conference on Web search and data mining. 2014 Presented at: 7th ACM international conference on Web search and data mining; February 2014; New York, NY, USA. [doi: 10.1145/2556195.2556220]

8. Youngmann B, Yom-Tov E. Anxiety and Information Seeking: Evidence From Large-Scale Mouse Tracking. In: Proceedings of the World Wide Web Conference on World Wide Web (WWW 2018). 2018 Presented at: World Wide Web Conference on World Wide Web (WWW 2018); 2018; Lyon, France. [doi: 10.1145/3178876.3186156]

9. Hochberg I, Allon R, Yom-Tov E. Assessment of the Frequency of Online Searches for Symptoms Before Diagnosis: Analysis of Archival Data. J Med Internet Res 2020 Mar 06;22(3):e15065 [FREE Full text] [doi: 10.2196/15065] [Medline: 32141835]

10. Gillon R. Medical ethics: four principles plus attention to scope. BMJ 1994 Jul 16;309(6948):184-188 [FREE Full text] [doi: 10.1136/bmj.309.6948.184] [Medline: 8044100]

11. Yom-Tov E. Crowdsourced Health: How What You Do on the Internet Will Improve Medicine. Cambridge, Massachusetts, United States: MIT Press; 2016.

12. De Choundhury M, Counts S, Horvitz E, Hoff A. Characterizing and predicting postpartum depression from shared facebook data. In: CSCW '14: Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing. 2014 Presented at: 17th ACM conference on Computer supported cooperative work & social computing; February 2014; Vancouver, BC, Canada p. 626-639. [doi: 10.1145/2531602.2531675]

13. Reece AG, Reagan AJ, Lix KL, Dodds PS, Danforth CM, Langer EJ. Forecasting the onset and course of mental illness with Twitter data. Sci Rep 2017 Oct 11;7(1):13006 [FREE Full text] [doi: 10.1038/s41598-017-12961-9] [Medline: 29021528]

14. Buchanan EA, Zimmer M. Internet Research Ethics. 2012. URL: https://plato.stanford.edu/entries/ethics-internet-research/ [accessed 2020-05-19]

15. Christenhusz GM, Devriendt K, Dierickx K. To tell or not to tell? A systematic review of ethical reflections on incidental findings arising in genetics contexts. Eur J Hum Genet 2013 Mar;21(3):248-255 [FREE Full text] [doi: 10.1038/ejhg.2012.130] [Medline: 22739341]

16. Jouhki J, Lauk E, Penttinen M, Sormanen N, Uskali T. Facebook's Emotional Contagion Experiment as a Challenge to Research Ethics. MaC 2016 Oct 10;4(4):75-85. [doi: 10.17645/mac.v4i4.579]

17. Acquisti A, Adjerid I, Balebako R, Brandimarte L, Cranor LF, Komanduri S, et al. Nudges for Privacy and Security. ACM Comput. Surv 2017 Oct 09;50(3):1-41. [doi: 10.1145/3054926]

18. Gefen G, Ben-Porat O, Tennenholtz M, Yom-Tov E. Privacy, Altruism, and Experience: Estimating the Perceived Value of Internet Data for Medical Uses. In: WWW '20: Companion Proceedings of the Web Conference 2020. 2020 Presented at: Web Conference 2020; April 2020; Taipei, Taiwan p. 552-556. [doi: 10.1145/3366424.3383414]

19. Hong Y, Vaidya J, Lu H, Karras P, Goel S. Collaborative Search Log Sanitization: Toward Differential Privacy and Boosted Utility. IEEE Trans. Dependable and Secure Comput 2015 Sep 1;12(5):504-518. [doi: 10.1109/tdsc.2014.2369034]

XSL·FO

RenderX

20.   Individuals using the Internet (% of population). The World Bank. 2018. URL: https://data.worldbank.org/indicator/IT.
       NET.USER.ZS [accessed 2020-05-19]

21.   Yom-Tov E. Demographic differences in search engine use with implications for cohort selection. Inf Retrieval J 2019 Jan
       1;22(6):570-580. [doi: 10.1007/s10791-018-09349-2]

22.   Cheng Q, Yom-Tov E. Do Search Engine Helpline Notices Aid in Preventing Suicide? Analysis of Archival Data. J Med
       Internet Res 2019 Mar 26;21(3):e12235 [FREE Full text] [doi: 10.2196/12235] [Medline: 30912753]

23.   Danescu-Niculescu-Mizil C, Broder A, Gabrilovich E, Josifovski V, Pang B. Competing for users' attention: on the interplay
       between organic and sponsored search results. In: WWW '10: Proceedings of the 19th International Conference on World
       Wide Web. 2010 Presented at: 19th International Conference on World Wide Web; 2010; Raleigh, North Carolina p.
       291-300. [doi: 10.1145/1772690.1772721]

24.   Mitchell EW. The ethics of unsolicited diagnosis of mental disorder in acquaintances: benefits and dangers. Psychiatrist
       2018 Jan 02;35(8):297-301. [doi: 10.1192/pb.bp.110.032953]

25.   Ratzan RM. Unsolicited medical opinion. J Med Philos 1985 May;10(2):147-162. [doi: 10.1093/jmp/10.2.147] [Medline:
       4009080]

26.   Preller G, Salloch S. Melanoma in the shopping mall: A utilitarian argument for offering unsolicited medical opinions in
       informal settings. Bioethics 2018 Mar;32(3):193-198. [doi: 10.1111/bioe.12426] [Medline: 29369379]

27.   Pardun JT. Good Samaritan Laws: A Global Perspective. Loyola of Los Angeles International and Comparative Law Journal
       1998;20(3):591-613.

28.   Youngmann B, Allerhand L, Paltiel O, Yom-Tov E, Arkadir D. A machine learning algorithm successfully screens for
       Parkinson's in web users. Ann Clin Transl Neurol 2019 Dec;6(12):2503-2509 [FREE Full text] [doi: 10.1002/acn3.50945]
       [Medline: 31714022]

29.   Reynolds J. A critical examination of the climate engineering moral hazard and risk compensation concern. The Anthropocene
       Review 2014 Oct 08;2(2):174-191. [doi: 10.1177/2053019614554304]

30.   Shelton JD. Ten myths and one truth about generalised HIV epidemics. The Lancet 2007 Dec;370(9602):1809-1811. [doi:
       10.1016/s0140-6736(07)61755-3]

31.   Hatch S. Snowball in a Blizzard: A Physician's Notes on Uncertainty in Medicine. New York, NY, USA: Basic Books;
       2016.