

Web Spam, Social Propaganda and the Evolution of Search Engine Rankings

Panagiotis Takis Metaxas*

Wellesley College, Wellesley, MA 02481, USA,
pmetaxas@wellesley.edu
<http://cs.wellesley.edu/~pmetaxas/>

Summary. Search Engines have greatly influenced the way we experience the web. Since the early days of the web, users have been relying on them to get informed and make decisions. When the web was relatively small, web directories were built and maintained using human experts to screen and categorize pages according to their characteristics. By the mid 1990's, however, it was apparent that the human expert model of categorizing web pages does not scale. The first search engines appeared and they have been evolving ever since, taking over the role that web directories used to play.

But what need makes a search engine evolve? Beyond the financial objectives, there is a need for quality in search results. Search engines know that the quality of their ranking will determine how successful they are. Search results, however, are not simply based on well-designed scientific principles, but they are influenced by web spammers. Web spamming, the practice of introducing artificial text and links into web pages to affect the results of web searches, has been recognized as a major search engine problem. It is also a serious users problem because they are not aware of it and they tend to confuse trusting the search engine with trusting the results of a search.

In this paper, we analyze the influence that web spam has on the evolution of the search engines and we identify the strong relationship of spamming methods on the web to propagandistic techniques in society. Our analysis provides a foundation for understanding why spamming works and offers new insight on how to address it. In particular, it suggests that one could use social anti-propagandistic techniques to recognize web spam.

Key words: Search Engines, Web search, Web graph, link structure, PageRank, HITS, Web Spam, Social Networks

1.1 Introduction

Search Engines have greatly influenced the way we experience the web. Since the early days of the web people have been relying on search engines to find

* Part of this work was supported by a Brachman-Hoffman grant.

useful information. When the web was relatively small, Web directories were built and maintained that were using human experts to screen and categorize pages according to their characteristics. By the mid 1990's, however, it was apparent that the human expert model of categorizing web pages would not scale. The first search engines appeared and they have been evolving ever since.

But what influences their evolution? The way a user interacts with a search engine is through the search results to a query that he or she has issued. Search engines know that the quality of their ranking will determine how successful they are. If users perceive the results as valuable and reliable, they will come again. Otherwise, it is easy for them to switch to another search engine.

Research in Information Retrieval has produced a large body of work that, theoretically, produces high quality search results. Yet, search engines admit that IR theory is but one of their considerations. One of the major issues that influences the quality of ranking is the effect that web spam has on their results. *Web spamming* is defined as the practice of manipulating web pages in order to influence search engines rankings in ways beneficial to the spammers. Spammers aim at search engines, but target the end users. Their motive is usually commercial, but can also be political or religious.

One of the reasons behind the users' difficulty to distinguish trustworthy from untrustworthy information comes from the success that both search engines and spammers have enjoyed in the last decade. Users have come to trust search engines as a means of finding information, and spammers have successfully managed to exploit this trust.

From their side, the search engines have put considerable effort in delivering spam-free query results and have developed sophisticated ranking strategies. Two such ranking strategies that have received major attention are the PageRank [4] and HITS [22]. Achieving high PageRank has become a sort of obsession for many companies' IT departments, and the *raison d'être* of spamming companies. Some estimates indicate that at least 8% of all pages indexed is spam [10] while experts consider web spamming the single most difficult challenge web searching is facing today[18]. Search engines typically see web spam as an interference to their operations and would like to restrict it, but there can be no algorithm that can recognize spamming sites based solely on graph isomorphism [3].

First, however, we need to understand *why* spamming works beyond the technical details, because spamming is a social problem first, then a technical one. In this paper we show its extensive relationship to social propaganda, and evidence of its influence on the evolution of search engines. Our approach can explain the reasons why web spamming has been so successful and suggest new ways of dealing with it. Finally, we present a framework for the long-term approach to web spam.

Background. Web spamming has received a lot of attention lately [2, 3, 10, 11, 15, 17, 18, 20, 23, 26, 27, 30]. The first papers to raise the issue were [27, 18]. The spammers' success was noted in [2, 9, 10, 11, 19].

Characteristics of spamming sites based on diversion from power laws are presented in [10]. Current tricks employed by spammers are detailed in [14]. An analysis of the popular PageRank method employed by many search engines today and ways to maximize it in a spamming network is described in [3]. TrustRank, a modification to the PageRank to take into account the evaluations of a few seed pages by human editors, employees of a search engine, is presented in [15]. Techniques for identifying automatically link farms of spam pages were presented in [38, 1].

A comprehensive treatment on social networks is presented in [36]. The connection between the Web and social networks was explicitly noted in [24, 31] and implicitly used in [4, 22]. In fact, Kleinberg’s work explores many of these connections (e.g., [21]). Identification of web communities was explored in [23, 12]. The effect that search engines have on page popularity was discussed in [7].

The rest of this paper is organized as follows. The next section gives an overview of the problem of information reliability and web spamming. Section 3 has a short introduction to the theory of propaganda detection and the next section 4 discusses the relationship between the webgraph and the trust social network. The following section 5 analyzes the evolution of search engines as their response to spam. Finally, the last section has the conclusions and a framework for the long-term approach to web spam.

1.2 Web Spam

The web has changed the way we inform and get informed. Every organization has a web site and people are increasingly comfortable accessing it for information on any question they may have. The exploding size of the web necessitated the development of search engines and web directories. Most people with online access use a search engine to get informed and make decisions that may have medical, financial, cultural, political, security or other important implications in their lives [9, 35, 19, 26]. Moreover, 85% of the time, people do not look past the first ten results returned by the search engine [33]. Given this, it is not surprising that anyone with a web presence struggles for a place in the top ten positions of relevant web search results. The importance of the top-10 placement has given birth to a new “Search Engine Optimization” industry, which claims to sell know-how for prominent placement in search results and includes companies, publications, and even conferences. Some of them are willing to bend the truth in order to fool the search engines and their customers, by creating web pages containing web spam [10].

Spammers attack search engines through text and link manipulations:

- **Text spam:** This includes repeating text excessively and/or adding irrelevant text on the page that will cause incorrect calculation of page relevance; adding misleading meta-keywords or irrelevant “anchor text” that will cause incorrect application of rank heuristics.

- **Link spam:** This technique aims to change the perceived structure of the webgraph in order to cause incorrect calculation of page reputation. Such examples are the so-called “link-farms,” domain flooding (plethora of domains that re-direct to a target site), page “awards,” (the spammer pretends to run an organization that distributes awards for web site design or information; the awarded site gets to display the “award”, an image linking back to awarding organization, effectively increasing the visibility of the spammer’ site), etc.

Both kinds of spam aim to boost the ranking of spammed web pages. So as not to get caught, spammers conceal their actions through cloaking, content hiding and redirection. Cloaking, for example, aims to serve different pages to search engine robots and to web browsers (users). For a comprehensive treatment of the spamming techniques, see [14].

Since anyone can be an author on the web, these practices have naturally created a question of *information reliability*. An audience used to trusting the written word of newspapers and books is unable, unprepared or unwilling to think critically about the information obtained from the web. A recent study [13] found that while college students regard the web as a primary source of information, many do not check more than a single source, and have trouble recognizing trustworthy sources online. In particular, two out of three students are consistently unable to differentiate between facts and advertising claims, even “infomercials.” Very few of them would double-check for validity. At the same time, they have considerable confidence in their abilities to distinguish trustworthy sites from non-trustworthy ones, especially when they feel technically competent. We have no reason to believe that the general public will perform any better than well-educated students. In fact, a recent analysis of internet related fraud by a major Wall Street law firm [9] puts the blame squarely on the questionable critical thinking skills of the investors for the success of stock fraud cases.

1.3 On Propaganda Theory

On the outset, it may seem surprising that a technical article discusses social propaganda. This is a subject that has been studied extensively by social scientists and might seem out of the realm of computing. However, the web is a social network, influenced daily by the actions (intentional or otherwise) of millions of people. In that respect, web researchers should be aware of social theories and practices since they may have applicability in their work. We believe that a basic understanding of social propaganda can be valuable to technical people designing and using systems that affect our social interactions. In particular, it can be useful to researchers that study Web Spam. We offer here a brief introduction to the theory of propaganda detection.

There are many definitions of propaganda, reflecting its multiple uses over time. One working definition we will use here is

Propaganda is the attempt to modify human behavior, and thus influence people's actions in ways beneficial to propagandists.

Propaganda has a long history in modern society and is often associated with negative connotation. This was not always the case, however. The term was first used in 1622, in the establishment by the Catholic Church of a permanent Sacred Congregation *de Propaganda Fide* (for the propagation of faith), a department which was trying to spread Catholicism in non-Catholic Countries [37]. Its current meaning comes from the successful Enemy Propaganda Department in the British Ministry of Information during WWI. However, it was not until 1938, in the beginning of WWII, that a theory was developed to detect propagandistic techniques. For the purposes of this paper we are interested in ways of detecting propaganda, especially by automatic means.

First developed by the Institute for Propaganda Analysis [25], classic Propaganda Theory identifies several techniques that propagandists often employ in order to manipulate perception.

- **Name Calling** is the practice of giving an idea a bad label. It is used to make people reject and condemn the idea without examining the evidence. For example, using the term “miserable failure” to refer to political leaders such as US President George Bush can be thought of as an application of name calling.
- **Glittering Generalities** is the mirror image² of name calling: Associating an idea with a “virtue word”, in an effort to make us accept and approve the idea without examining the evidence. For example, using the term “patriotic” to refer to illegal actions is a common application of this technique.
- **Transfer** is the technique by which the propagandist carries over the authority, sanction, and prestige of something respected and revered to something he would have us accept. For example, delivering a political speech in a mosque or a church, or ending a political gathering with a prayer have the effect of transfer.
- **Testimonial** is the technique of having some respected person comment on the quality of an issue on which they have no qualifications to comment. For example, a famous actor who plays a medical doctor on a popular TV show tells the viewers that she only uses a particular pain relief medicine. The implicit message is that if a famous personality trusts the medicine, we should too.
- **Plain Folks** is a technique by which speakers attempt to convince their audience that they, and their ideas, are “of the people,” the “plain folks”. For example, politicians sometimes are seen flipping burgers at a neighborhood diner.
- **Card Stacking** involves the selection of facts (or falsehoods), illustrations (or distractions), and logical (or illogical) statements in order to give an

² Name calling and glittering generalities are sometimes referred to as “word games.”

incorrect impression. For example, some activists refer to the Evolution Theory as a theory teaching that humans came from apes (and not that both apes and humans have evolved from a common ancestor who was neither human nor ape).

- **Bandwagon** is the technique with which the propagandist attempts to convince us that all members of a group we belong to accept his ideas and so we should “jump on the band wagon”. Often, fear is used to reinforce the message. For example, commercials might show shoppers running to line up in front of a store before it is open.

The reader should not have much trouble identifying additional examples of such techniques used in politics or advertising. The next section discusses the relationship of propaganda to web spam, by first describing the similarity of social networks to the web graph.

1.4 The Webgraph as a Trust Network

The web is typically represented by a directed graph [6]. The nodes in the webgraph are the pages (or sites) that reside on servers on the internet. Arcs correspond to hyperlinks that appear on web pages (or sites). In this context, web spammers’ actions can be seen as altering the contents of the web nodes (mainly through text spam), and the hyperlinks between nodes (mainly through link spam).

The theory of social networks [36] also uses directed graphs to represent relationships between social entities. The nodes correspond to social entities (people, institutions, ideas). Arcs correspond to recommendations between the entities they connect. In this context, propagandistic techniques can be seen as altering the trust social network by altering one or more of its components (i.e., nodes, arcs, weights, topology).

To see the correspondence more clearly, we will examine some of the propagandistic techniques that have been used successfully by spammers: The technique of testimonials effectively adds a link between previously unrelated nodes. Glittering generalities change the contents of a node, effectively changing its perceived relevance. Mislabeled anchor text is an example of card stacking. And the technique of bandwagon creates many links between a group of nodes, a “link farm”. So, we define web spam based on the spammers actions:

Web Spam is the attempt to modify the web (its structure and contents), and thus influence search engine results in ways beneficial to web spammers.

Table 1.1 has the correspondence, in graph theoretic terms, between the web graph according to a search engine and the trust social network of a particular person. Web pages or sites correspond to social entities and hyperlinks correspond to trust opinions. The rank that a search engine assigns to a page or a site corresponds to the reputation a social entity has for the person. This rank is based on some ranking formula that a search engine is computing,

<i>Graph Theory</i>	<i>Web Graph</i>	<i>Trust Social Network</i>
Node	web page or site	social entity
weight	rank (accord. to a search engine)	reputation (accord. to a person)
weight computation	ranking formula (e.g., pagerank) computed continuously	idiosyncratic (e.g., 2 recommenders) computed on demand
Arc	hyperlink	trust opinion
semantics	“vote of confidence”	“recommendation”
weight	degree of confidence	degree of entrustment
weight range	[0 . . . 1]	[<i>distrust</i> . . . <i>trust</i>]

Table 1.1. Graph theoretic correspondence between the Webgraph and the Trust Social Network. There is a one-to-one correspondence between each component of the two graphs. A major difference, however, is that, even though a person may feel negative trust (distrust) for some entity, there is no negative weight for hyperlinks.

while the reputation is based on idiosyncratic components associated with the person’s past experiences and selective application of critical thinking skills; both are secret and changing.

This correspondence is more than a coincidence. The web itself is a social creation, and both PageRank and HITS are socially inspired ranking formulas. [4, 22, 31]. Socially inspired systems are subject to socially inspired attacks. Not surprisingly then, the theory of propaganda detection can provide intuition into the dynamics of the web graph.

PageRank is based on the assumption that the reputation of an entity (a web page in this case) can be measured as a function of both the number and reputation of other entities linking to it. A link to a web page is counted as a “vote of confidence” to this web site, and in turn, the reputation of a page is divided among those it is recommending³. The implicit assumption is that hyperlink “voting” is taking place independently, without prior agreement or central control. Spammers, like social propagandists, form structures that are able to gather a large number of such “votes of confidence” by design, thus breaking the crucial assumption of independence in a hyperlink. But while the weights in the web graph are assigned by each search engine, the weights in the trust social network are assigned by each person. Since there are many more persons than search engines, the task of a web spammer is far easier than the task of a propagandist.

³ Since HTML does not provide for “positive” and “negative” links, all links are taken as positive. This is not always true, but is considered a reasonable assumption. Recently, Google introduced the “nofollow” attribute for hyperlinks, as a tool for blog site owners to mark visitor opinions. It is very unlikely that spamming blog owners will use it, however.

1.5 Search Engine Evolution

In the early 90's, when the web numbered just a few million servers, the **first generation** search engines were ranking search results using the vector model of classic information retrieval techniques: the more rare words two documents share, the more similar they are considered to be. [32, 17]

According to the *vector model* in Information Retrieval [32], documents contained in a document collection D are viewed as vectors in term space T . Each document vector is composed of term weights w_{ik} of term T_k appearing in document D_i . These weights are computed as the normalized dot product of $tf_{ik} \cdot idf_k$, where tf_{ik} is the frequency of term T_k in document D_i , and idf_k is the inverse document frequency of term T_k in document collection D . Typically, idf_k is computed by a logarithmic formula so that this term will not grow significantly as the number of occurrences of T_k increase. Under this formulation, rare words have greater weight than common words, because they are viewed as better representing the document contents. The term weights are then normalized to fall on a unit sphere so that longer documents will not have an advantage over shorter documents:

$$w_{ik} = \frac{tf_{ik} \cdot idf_k}{\sqrt{\sum_{1 \leq k \leq t} (tf_{ik})^2 (idf_k)^2}}$$

In the vector model, document similarity $sim(D_1, D_2)$ between document vectors D_1 and D_2 is represented by the angle between them, and is computed as $\sum_{1 \leq i \leq t} w_{1i} \cdot w_{2i}$ cosine normalized:

$$sim(D_1, D_2) = \frac{\sum_{1 \leq i \leq t} w_{1i} \cdot w_{2i}}{\sqrt{\sum_{1 \leq i \leq t} (w_{1i})^2 \cdot \sum_{1 \leq i \leq t} (w_{2i})^2}}$$

A search query Q is considered simply a short document and the results of a search for Q are ranked according to their (normalized) similarity to the query. While the exact details of the computation of term weights were kept secret, we can say that the ranking formula R^{G1} in the first generation search engines was based in the following principle: the more rare keywords a document shares with a query, the higher similarity it has with it, resulting in a higher ranking score for this document:

$$R^{G1} = f(sim(p, Q)) \tag{1.1}$$

The first attack to this ranking came from within the search engines. In 1996, search engines started openly selling search keywords to advertisers [8] as a way of generating revenue: If a search query contained a “sold” keyword, the results would include targeted advertisement and a higher ranking for the link to the sponsor’s web site.

Mixing search results with paid advertisement raised serious ethical questions, but also showed the way to financial profits to spammers who started

their own attacks using **keyword stuffing**, i.e., by creating pages containing many rare keywords to obtain a higher ranking score. In terms of propaganda theory, the spammers employed a variation of the technique of *glittering generalities* to confuse the first generation search engines [25, pg. 47]:

The propagandist associates one or more suggestive words without evidence to alter the conceived value of a person or idea.

In an effort to nullify the effects of glittering generalities, **second generation** search engines started employing additionally more sophisticated ranking techniques. One of the more successful techniques was based on the “link voting principle”: Each web site s has value equal to its “popularity” $|B_s|$ which is influenced by the set B_s of sites pointing to s .

Therefore, the more sites were linking to a site s , the higher the popularity of s 's pages. Lycos became the champion of this ranking technique [28] and had its own popularity skyrocket in late 1996. Doing so, it was also distancing itself from the ethical questions introduced by blurring advertising with ranking [8].

The ranking formula R^{G2} in the second generation search engines was a combination of a page's similarity, $sim(p, Q)$, and its site's popularity $|B_s|$:

$$R^{G2} = f(sim(p, Q), |B_s|) \quad (1.2)$$

To avoid spammers search engines would keep secret their exact ranking algorithm. Secrecy is no defense, however, since secret rules were figured out by experimentation and reverse engineering. (e.g., [30, 27]).

Unfortunately, this ranking formula did not succeed in stopping spammers either. Spammers started creating clusters of interconnected web sites that had identical or similar contents with the site they were promoting, a technique that subsequently became known as **link farms**. The link voting principle was socially inspired, so spammers used the well known propagandistic method of *bandwagon* to circumvent it [25, pg. 105]:

With it, the propagandist attempts to convince us that all members of a group to which we belong are accepting his program and that we must therefore follow our crowd and “jump on the band wagon”.

Similarly, the spammer is promoting the impression of a high degree of popularity by inter-linking many internally controlled sites that will eventually all share high ranking.

PageRank and HITS marked the development of the **third generation** search engines. The introduction of PageRank in 1998 [4] was a major event for search engines, because it seemed to provide a more sophisticated anti-spamming solution. Under PageRank, not every link contributes equally to the “reputation” of a page $PR(p)$. Instead, links from highly reputable pages contribute much higher value than links from other sites. A page p has reputation $PR(p)$ which is calculated as the sum of fractions of the reputations of the set B_p of pages pointing to p . Let F_v be the set of links out of page v , $v \in B_p$. The reputation of a page is

$$PR(p) = \frac{1-t}{N} + t \sum_{v \in B_p} \frac{PR(v)}{|F_v|}$$

where t is the so-called “transport” factor and N is the total number of pages in the collection. That way, the link farms developed by spammers would not influence much their PageRank, and Google became the search engine of choice. HITS is another socially-inspired ranking which has also received a lot of attention [22] and is reportedly used by the AskJeeves search engine. The HITS algorithm divides the sites related to a query between “hubs” and “authorities”. Hubs are sites that contain many links to authorities, while authorities are sites pointed to by the hubs and they both gain reputation.

Unfortunately, spammers again found ways of circumventing these rankings. In PageRank, a page enjoys absolute reputation: its reputation is not restricted on some particular issue. Spammers deploy sites with expertise on irrelevant subjects, and they acquire (justifiably) high ranking on their expert sites. Then they bandwagon the irrelevant expert sites, creating what we call a **mutual admiration society**. In propagandistic terms, this is the technique of *testimonials* [25, pg. 74] often used by advertisers:

Well known people (entertainers, public figures, etc.) offer their opinion on issues about which they are not experts.

Spammers were so aggressive in pursuing this technique that they openly promoted “reciprocal links”: Web masters controlling sites that had some minimum PageRank, were invited to join a mutual admiration society by exchanging links, so that at the end everyone’s PageRank would increase. HITS has also shown to be highly spammable by this technique due to the fact that its effectiveness depends on the accuracy of the initial neighborhood calculation.

Another heuristic that third generation search engines used was that of exploiting “anchor text”. It had been observed that users creating links to web pages would come to use, in general, meaningful descriptions of the contents of a page. (Initially, the anchor text was non-descriptive, such as “click here”, but this changed in the late 1990’s.) Google was the first engine to exploit this fact noting that, even though IBM’s web page made no mention that IBM is a computer company, many users linked to it with anchor text such as “computer manufacturer”.

Spammers were quick to exploit this feature too. In early 2001, a group of activists started using the anchor text “miserable failure” to link to the official Whitehouse page of American President George W. Bush. Using what became known as “Googlebomb” or, more accurately, **link-bomb** since it does not pertain to Google only, other activists linked the same anchor text to President Carter, filmmaker Michael Moore and Senator Hilary Clinton.

Using the anchor text is socially inspired, so spammers used the propagandistic method of *card stacking* to circumvent it [25, pg. 95]:

Card stacking involves the selection and use of facts or falsehoods, illustrations or distractions, and logical or illogical statements in order to give the best or the worst possible case for an idea, program, person or product.

S.E.'s	Ranking	Spamming	Propaganda
1st Gen	Doc Similarity	keyword stuffing	glittering generalities
2nd Gen	+ Site popularity	+ link farms	+ bandwagon
3rd Gen	+ Page reputation + anchor text	+ mutual admiration societies + link bombs	+ testimonials + card stacking

Table 1.2. Changes in ranking by generations of search engines, the response of the web spammers and the corresponding propagandistic techniques.

The ranking formula R^{G3} in the third generation search engines is, therefore, some secret combination of a number of features, primarily the page's similarity, $sim(p, Q)$, its site's popularity $|B_s|$ and its the page's reputation $PR(p)$:

$$R^{G3} = f(sim(p, Q), |B_s|, PR(p)) \quad (1.3)$$

Search engines these days claim to have developed hundreds of little heuristics for improving their web search results [16] but no big idea that would move their rankings beyond the grasp of spammers. As Table 1.2 summarizes, for every idea that search engines have used to improve their ranking, spammers have managed quickly to balance it with techniques that resemble propagandistic techniques from society. Web search corporations are reportedly busy developing the engines of the next generation [5]. The new techniques aim to be able to recognize “the need behind the query” of the user. Given the success the spammers have enjoyed so far, one wonders how will they spam the fourth generation engines. Is it possible to create a ranking that is not spamable? Put another way, can the web as a social space be free of propaganda?

This may not be possible. Our analysis shows that we are trying to create in cyberspace what societies have not succeeded in creating in their real space. However, we can learn to live in a web with spam as we live in society with propaganda, given appropriate education and technology. We touch upon it in our concluding section.

1.6 Conclusions

In this paper we have argued that web spam is to cyberworld what propaganda is to society. As evidence of the importance of this analogy, we have shown that the evolution of search engines can be largely understood as the search engines' responses in defending against spam. We do not suggest here that web spam is the *sole* force behind the evolution of search engines, but that it

is a dominant one. New search engines are developed when researchers believe they have a good answer to spam because it directly affects the quality of the search results.

Further, our findings suggests that anti-spamming techniques can be developed by mimicking anti-propagandistic methods. In a followup paper [29] we present automatic ways of recognizing trust graphs on the web based on anti-propagandistic techniques. Our idea is to *propagate distrust* to a spamming network whenever one of them is recognized. In the next couple paragraphs we give a short description of our results.

We are considering trustworthiness to be a personal decision, not an absolute quality of a site. One person's gospel is another's political propaganda, and our goal is to design methods that help individuals make more informed decisions about the quality of the information they find on the web. Here is one way that people defend against propaganda in every day life:

In society, distrust is propagated backwards: When an untrustworthy recommendation is detected, it gives us a reason to reconsider the trustworthiness of the recommender. Recommenders who strongly support an untrustworthy recommendation become untrustworthy themselves.

This process is selectively repeated a few times, propagating the distrust backwards to those who strongly support the recommendation. The results of this process become part of our belief system and are used to filter future information. (Note that distrust is not propagated forward: An untrustworthy person's recommendations could be towards *any* entity, either trustworthy or untrustworthy.) Experimental results [29] from a number of such instances show our algorithm's ability of recognizing parts of a spamming network. Therefore, our work is complementary to the recent developments that recognize web spam based on link analysis [38, 1].

But what one should do once one recognizes a spamming network. This is a question that has not attracted the necessary attention in the past. The default approach is that a search engine would delete such networks from its indices or might downgrade them by some prespecified amount. Search engines are reportedly doing a fair amount of this [34, 10, 15]. A more effective way is *personalizing the web graph* a user sees, effectively increasing the task difficulty of a spammer to the level of a propagandist: As we mentioned, a spammer has an easier job than a propagandist because he/she has to influence the web graphs of a few search engines instead of the trust graphs of millions of individuals.

There are clearly cases where these approaches are appropriate and effective. But in general, both of these approaches require a universal agreement of what constitutes spam. Such an agreement cannot exist; one person's spam may be another person's treasure. Should the search engines determine what is trustworthy and what is not? Willing or not, they are the *de facto* arbiters of what information users see [34]. As in a popular cartoon by Ohman & Willis, a kid responds to the old man who has been searching his entire life for the meaning of life: "[...]if it's not on Google, you probably won't find

it.” We believe that it is the users’ right and responsibility to decide what is acceptable for them. Their browser, their window to cyberworld, should enhance their ability to make this decision. User education is fundamental: without it, people will largely trust what they see, regardless its credibility. People should know how search engines work and why, and how information appears on the web. But they should also have a trained browser that can help them determine the validity and trustworthiness of information.

References

1. Benczúr, A., Csalogány, K., Sarlós, T., and Uher, M. (2005). Spam Rank – Fully automatic link spam detection. In *Proceedings of the AIRWeb Workshop*.
2. Bharat, K., Chang, B.-W., Henzinger, M. R., and Ruhl, M. (2001). Who links to whom: Mining linkage between web sites. In *Proceedings of the 2001 IEEE International Conference on Data Mining*, pages 51–58. IEEE Computer Society.
3. Bianchini, M., Gori, M., and Scarselli, F. (2003). PageRank and web communities. In *Web Intelligence Conference 2003*.
4. Brin, S. and Page, L. (1998). The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1–7):107–117.
5. Broder, A. (2002). A taxonomy of web search. *SIGIR Forum*, 36(2):3–10.
6. Broder, A., Kumar, R., Maghoul, F., Raghavan, P., Rajagopalan, S., Stata, R., Tomkins, A., and Wiener, J. (2000). Graph structure in the web. *Comput. Networks*, 33(1-6):309–320.
7. Cho, J. and Roy, S. (2004). Impact of search engines on page popularity. In *WWW 2004*.
8. CNETNews (1996). Engine sells results, draws fire. <http://news.cnet.com/2100-1023-215491.html>.
9. Corey, T. S. (2001). Catching on-line traders in a web of lies: The perils of internet stock fraud. Ford Marrin Esposito, Witmeyer & Glessner, LLP. <http://www.fnw.com/archive/lies/>.
10. Fetterly, D., Manasse, M., and Najork, M. (2004). Spam, damn spam, and statistics. In *WebDB2004*.
11. Fetterly, D., Manasse, M., Najork, M., and Wiener, J. (2003). A large-scale study of the evolution of web pages. In *Proceedings of the twelfth international conference on World Wide Web*, pages 669–678. ACM Press.
12. Flake, G. W., Lawrence, S., Giles, C. L., and Coetzee, F. (2002). Self-organization of the web and identification of communities. *IEEE Computer*, 35(3):66–71.
13. Graham, L. and Metaxas, P. T. (2003). “Of course it’s true; i saw it on the internet!”: Critical thinking in the internet era. *Commun. ACM*, 46(5):70–75.
14. Gyöngyi, Z. and Garcia-Molina, H. (2005). Web spam taxonomy. In *Proceedings of the AIRWeb Workshop*.
15. Gyöngyi, Z., Garcia-Molina, H., and Pedersen, J. (2004). Combating web spam with TrustRank. In *VLDB 2004*.
16. Hansell, S. (2007). Google keeps tweaking its search engine. New York Times.
17. Henzinger, M. R. (2001). Hyperlink analysis for the web. *IEEE Internet Computing*, 5(1):45–50.

18. Henzinger, M. R., Motwani, R., and Silverstein, C. (2002). Challenges in web search engines. *SIGIR Forum*, 36(2):11–22.
19. Hindman, M., Tsioutsoulouklis, K., and Johnson, J. (2003). Googlearchy: How a few heavily-linked sites dominate politics on the web. In *Annual Meeting of the Midwest Political Science Association*.
20. Introna, L. and Nissenbaum, H. (2000). Defining the web: The politics of search engines. *Computer*, 33(1):54–62.
21. Kleinberg, J. (2000). The small-world phenomenon: an algorithm perspective. In *STOC '00: Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 163–170. ACM Press.
22. Kleinberg, J. M. (1999). Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5):604–632.
23. Kumar, R., Raghavan, P., Rajagopalan, S., and Tomkins, A. (1999). Trawling the Web for emerging cyber-communities. *Computer Networks (Amsterdam, Netherlands: 1999)*, 31(11–16):1481–1493.
24. Kumar, R., Raghavan, P., Rajagopalan, S., and Tomkins, A. (2002). The web and social networks. *IEEE Computer*, 35(11):32–36.
25. Lee, A. M. and Lee(eds.), E. B. (1939). *The Fine Art of Propaganda*. The Institute for Propaganda Analysis. Harcourt, Brace and Co.
26. Lynch, C. A. (2001). When documents deceive: trust and provenance as new factors for information retrieval in a tangled web. *J. Am. Soc. Inf. Sci. Technol.*, 52(1):12–17.
27. Marchiori, M. (1997). The quest for correct information on the web: hyper search engines. *Comput. Netw. ISDN Syst.*, 29(8-13):1225–1235.
28. Maulding, M. L. (1997). Lycos: Design choices in an internet search service. *IEEE Expert*, January-February(12):8–11.
29. Metaxas, P. T. (2009). Using Propagation of Distrust to find Untrustworthy Web Neighborhoods. In *Proceedings of the fourth international conference on internet and Web Applications and Services*, Venice, Italy, 24–28 May, 2009.
30. Pringle, G., Allison, L., and Dowe, D. L. (1998). What is a tall poppy among web pages? In *Proceedings of the seventh international conference on World Wide Web 7*, pages 369–377. Elsevier Science Publishers B. V.
31. Raghavan, P. (2002). Social networks: From the web to the enterprise. *IEEE Internet Computing*, 6(1):91–94.
32. Salton, G. (1972). Dynamic document processing. *Commun. ACM*, 15(7):658–668.
33. Silverstein, C., Marais, H., Henzinger, M., and Moricz, M. (1999). Analysis of a very large web search engine query log. *SIGIR Forum*, 33(1):6–12.
34. Totty, M. and Mangalindan, M. (2003). As google becomes web’s gatekeeper, sites fight to get in. In *Wall Street Journal CCXLI(39)*.
35. Vedder, A. (2000). Medical data, new information technologies and the need for normative principles other than privacy rules. In *Law and Medicine. M. Freeman and A. Lewis (Eds.), (Series Current Legal Issues)*, pages 441–459. Oxford University Press.
36. Wasserman, S. and Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge University Press.
37. Welch, D. (1999). Power of persuasion - propaganda. *History Today*, 49(8):24–26.
38. Wu, B. and Davison, B. (2005). Identifying link farm spam pages. In *Proceedings of the fourteenth international conference on World Wide Web*. ACM Press.