



Web-Based Child Pornography: Quantification and Qualification of Demand

Chad M.S. Steel, Virginia Polytechnic Institute and State University, USA

ABSTRACT

While the supply of child pornography through the World Wide Web has been frequently speculated upon, the demand has not adequately been explored. Quantification and qualification of the demand provides forensic examiners a behavioral basis for determining the sophistication of an individual seeking child pornography. Additionally, the research assists an examiner in searching for and presenting the evidence of child pornography browsing. The overall search engine demand for child pornography is bounded as being between .19 and .49%, depending on the inclusion of ambiguous phrases, with the top search for child pornography being "lolita bbs". Unlike peer-to-peer networks, however, the top child pornography related query ranks only as the 198th most popular query overall. The queries on search engines appear to be decreasing as well, and the techniques employed are becoming less reliant direct links to content.

Keywords: Child Pornography, Forensics, Search Engines

INTRODUCTION

One of the most frequent cases encountered by digital forensic examiners is the possession of child pornography. A major challenge for the examiner is characterizing the activities of the suspect – understanding how the suspect seeks out child pornography helps to meet that challenge. Because the web activities of a suspect are one of the most common features analyzed in an examination, understanding the demand for web-based child pornography is important.

A second reason for the study of web-based demand for child pornography is to provide an academic basis for statements related to the quantity of the activity. Those proposing particular legislation targeted at combating child

pornography may tend to exaggerate the activities, whereas defense counsels may downplay the relationships between child pornography and particular search terms. By performing similarity analysis, we can provide accurate information to both parties without prejudice.

Major Contributions

The primary contribution of this article is the overall quantification of the demand for child pornography. By analyzing both Google and Dogpile data, the demand was found to be between .19 and .49% of all searches. This cannot be extrapolated to a particular percentage of users, but does show that one out of every 200 to 500 searches is related to child pornography.

DOI: 10.4018/jdcf.2009062405

Secondarily, the top search terms used by child pornographers were identified and quantified. The terms showed a distinct dichotomy between sophisticated searchers who are fluent in the language and nature of the subculture and unsophisticated users. This dichotomy can be exploited by forensic examiners to focus their searches based on user classification.

In terms of trends, the overall demand for child pornography using web-based searches appears to be declining and not increasing. While this does not mean that the actual demand for all child pornography is on the decline, it does show a potential shift toward other methods of acquisition.

Other areas that are frequently linked to child pornography were also explored—specifically the linkage between the terms associated with naturism/nudism and child pornography and that between bbs-associated searches and child pornography. Many naturism/nudism-related searches were found to focus on visual depictions, but not necessarily of children. On the other hand, the now antiquated term “bbs” was found to be strongly associated with child pornography.

Prior Art

Originally, child pornography on the Internet was traded largely over newsgroups through Usenet. Users were able to upload files semi-anonymously and from around the globe, and others could similarly download the content semi-anonymously. Initial studies by Rimm (Rimm, 1995) showed vast amounts of Usenet posts were dedicated to the trafficking of child pornography, though later analysis revealed significant flaws in his approach (Sigel & Sauer, 1999) (Hoffman & Novak, 1999). A follow-on study by Mehta was comprehensive and showed a much more accurate picture of Usenet-based trafficking in child pornography—specifically citing a figure of 15% of traffic being child-pornography related (Mehta, 2001). While Usenet is still active and growing—recent traffic rates are up to 4.5TB a day (Giganews)—it has

been eclipsed by peer-to-peer and web-based traffic (Sandvine, 2008).

Early studies of peer-to-peer traffic were performed by the US Government Accountability Office (GAO) consisted of 42-44% child pornography (GAO, File-Sharing Programs: Peer-to-Peer Networks Provide Ready Access to Child Pornography, 2003) (GAO, File-Sharing Programs: Child Pornography is Readily Accessible over Peer-to-Peer Networks, 2003), though there are methodological issues with the study.

In 2006, Hughes et al conducted a study of the traffic being shared on the Gnutella network, a popular peer-to-peer network using the Gnutella protocol, focused on illicit content. The Hughes study used humans to classify search content, and used a small subset of queries (Hughes, Walkerdine, Coulson, & Gibson, 2006). A follow-on study by Hughes et al identified 1% of the queries on the Gnutella network as having child pornography. This study used an automated analysis on the same dataset and found a prevalence of child pornography amongst the queries of approximately 1% (Hughes, 2008). Another larger study of peer-to-peer networks confirmed the prevalence statistics of the Hughes study and found the single most prevalent query on the Gnutella network was child-pornography related (Steel, 2008).

No comprehensive studies of child-pornography demand on the web currently exist, but studies of the demand for pornography overall have been published by Beitzel et al. Beitzel et al analyzed a very large query log from a major search provider and categorized the content of the queries. The initial study found that approximately 10% of the queries were related to pornography (Beitzel, Jensen, Chowdhury, Grossman, & Frider, 2004). A follow-on study using automated analysis concluded that 7.9% of queries in 2007 were seeking pornography (Beitzel, Jensen, Lewis, Chowdhury, & Frider, 2007). As child pornography is a subset of all pornography, it stands to reason that the prevalence of child pornography would be substantially less than the 7.9% figure.

Notes on Research

Child pornography possession is against the law within the United States with no provisions for research present in the statute (United States Government). Because of this, no illicit material was downloaded as part of this research. The research focused exclusively on those searching for child pornography, but did not and could not verify what those individuals ended up finding. As such, no effort to quantify the supply of child pornography was made. Additionally, there was no way to identify intent within the search terms, so a conservative approach was taken toward the classification. Ambiguous terms where the content sought could not be accurately classified are called out in the text as appropriate.

For the keywords used by child pornographers, a small sampling is provided in this article. The samples provided are readily known to those seeking child pornography, and the sharing of a few terms is likely to benefit forensic examiners more than those seeking

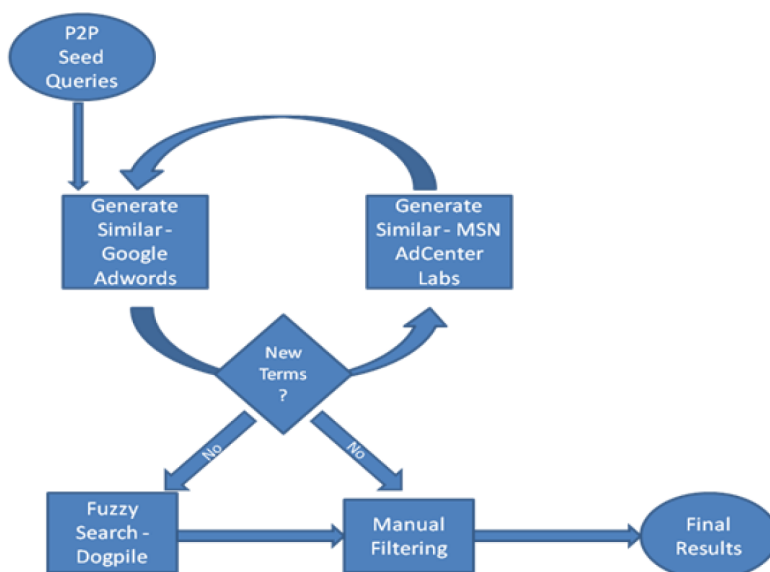
contraband. As with previous research, a more complete listing of child pornography related queries will be made available to law enforcement upon email request to the author.

As a final note, this research does not distinguish between child pornography and child erotica. In some jurisdictions, child pornography requires a sexual act as opposed to just nudity, which may be considered erotica. Except where noted in the section on nudism and naturism, child pornography-related queries as identified in this article are the superset of child pornography and child erotica queries.

METHODOLOGY

Data was collected from three different search engines -- Microsoft MSN, Google, and Dogpile were all used to generate the results. This data was mined for child-pornography related terms, and then manually cleansed to remove ambiguous terms. The overall process is shown in Figure 1.

Figure 1. The methodology employed for obtaining child pornography queries utilized similarity generation features of both Microsoft and Google, and applied the results to a custom corpus obtained from Dogpile.



Over 70% of the total searches conducted online are conducted on either the Google or MSN search sites (comScore.com, 2008), and Dogpile is a popular search engine aggregator, where the queries are distributed to and results received from all of the major search engines. Dogpile's aggregation makes it a good source for cross-site search engine validation, and allows it to make up for potential demographic differences present in the usage of the other search platforms.

Data on the most common child-pornography related queries was generated using the similarity tools present on both Google and MSN. Google Adwords is used to provide similar phrases to a seed phrase, along with the number of queries per month for each of those subphrases. The individual results are not filtered, and are generated from queries from the same users from which the seed query was associated (Google.com). MSN AdCenter Labs provides a similar interface into MSN search data with their group association tool (MSN AdCenter Labs).

An initial list of keywords associated with child pornography was generated from the top 10 keywords found in peer-to-peer networks (Steel, 2008). This list was used as a seed for both the MSN and Google similarity tools. The union of the results from both tools was then ranked according to their respective monthly volumes, and the top ten from these were used to seed another round of queries.

After each round, the ambiguous queries not primarily associated with child pornography were removed. Ambiguous queries included those strongly associated with child pornography that also had other meanings (e.g. "cp" is strongly related to "child pornography" but also "cerebral palsy" and the UNIX copy command) as well as those related to pornography but not conclusively to child pornography (e.g. "teen sex movies"). These were determined by their own similarity indexes where non-obvious ("underage teen sex" and "kiddy cat" are examples of obvious child pornography and non-child pornography queries, respectively). Any queries where more than 50% by count

of the similar queries were child-pornography related were considered child pornography queries for initial selection purposes. This process was repeated until no new queries were added to the corpus.

The final corpus was ranked according to search volume and manually segmented into likely and unlikely child pornography-related terms. Terms not exclusively associated with child pornography were pruned at this point, including those which were included as similar above.

Once the final Google query set was established, the same terms were utilized with the Dogpile corpus to provide independent verification of the results. The results of both were then analyzed both qualitatively and quantitatively.

RESULTS

Demand Prevalence

The overall prevalence of child pornography queries can be fairly accurately determined indirectly from the Google statistics and directly from the Dogpile statistics. The results of both are shown in Table 1. Both datasets had a significantly lower percentage of queries related to child pornography than found on peer-to-peer networks (Steel, 2008), with Dogpile having a slightly higher rate at .35% as opposed to Google's .19%. The reason for the discrepancy is likely due to limitations in Google's aggregate statistics for a particular term. With Dogpile, basic fuzzy searching was utilized to find minor misspellings. Additionally, the Dogpile corpus was processed with a context-aware punctuation parser, whereas the Google engine was unable to return terms such as R@ygold, which was condensed to rygold or r ygold for some of the tools. Google may also be less frequented by child pornographers due to the knowledge that search queries are tracked. Because of the access to the raw search statistics, the Dogpile results are likely more representative of the

Table 1. Percentage of overall queries related to child pornography

Search Engine	Percentage of Queries
Google	.19 - .27
Dogpile	.35 - .49

actual demand for child pornography on web search engines.

There are a few limitations in the prevalence results presented. Because of the manual filtering, there is little conflation expected in these numbers, though a small percentage may be reporters/researchers/law enforcement. Additionally, because the similarity tools from two distinct search engines were used, the likelihood of missing a disjoint cluster of terms is fairly low, though a small number of misspellings or very specific terms related to an individual site or series may not have been caught. Overall, the numbers reported are conservative – the inclusion of the ambiguous terms strongly associated with child pornography raises the prevalence numbers to .27% and .49% respectively for Google and Dogpile and represents an upper bound of the prevalence.

Top Searches

The top search queries associated with child pornography were taken from the Dogpile corpus, which provided an exact match of queries and presents a more accurate reflection of the top searches than Google. The queries are shown in Table 2 below.

The top search queries can be easily classified into a dichotomous typology based on the level of knowledge of the subculture as either unsophisticated or sophisticated. The searches for “underage nudity”, “preteen sex”, “child porn”, “underage porn”, and “preteen nude” represent unsophisticated efforts by those without extensive knowledge of child pornography. Based on the results likely to be returned by these searches, the individuals conducting them are not likely to be successful in obtaining access to child pornography.

The second typology of searches is represented by the queries “lolita bbs”, “pthc”, and “ls magazine”. These terms imply a greater knowledge of child pornography and should be considered “sophisticated” searches by individuals who have considerable knowledge of the child pornography subculture.

While “lolita” is a term used by both unsophisticated and sophisticated searchers, the modifier “bbs” implies a high degree of sophistication. In this context, “bbs” is an acronym for

Table 2. Top 10 child pornography phrases searched on Dogpile

Query	Count
lolita bbs	142
underage porn	135
preteen nude	51
Pthc	50
underage nudity	46
loli hc bbs	45
underage thai girls	38
preteen sex	35
child porn	33
ls magazine	33

“Bulletin Board System”, a term that evolved from the period when child pornography was traded over semi-private systems requiring a modem to dial directly in. In Internet terms, “bbs” denotes message boards or forums where information on various topics can be posted. As noted above, sophisticated users will search for forums that contain pointers to transient web-sites hosting child pornography, knowing that as soon as search engines index those sites they have a very limited lifespan (Jenkins, 2003). Further details on the links between “bbs” and child pornography are noted below.

The term “pthc” is slang for pre-teen hardcore. This is the single most common term searched for in peer-to-peer networks, and represents an advancement of knowledge over searching for the non-acronym terms. (Steel, 2008) Similarly, “ls magazine” refers to a child pornography magazine published in the 1970’s and indicates a historical knowledge of child pornography, implying either an individual old enough to remember the magazine or one that came across the terms through extended exposure to the subculture.

The remaining terms, “loli hc bbs” and “underage thai girls” are overly specific and are likely anomalous. Both terms appeared at temporally condensed intervals, whereas


the other terms appeared spread out over the period monitored. These are likely the result of an individual repeating a search multiple times and/or a program which automatically generated the queries.

Trends

Though the web-based prevalence of child pornography is undeniably a problem, there is a tendency in the press to conflate (and sometimes fabricate) growth trends in this arena. Articles which tout the growth of web-based pornography frequently cite tangential statistics, such as the number of arrests or the number of complaints (Scherer, 2005). Many of these statistics have been repudiated (Libertus.net, 2008), but no actual analysis of the growth of web-based child pornography has been completed. By using Google Trends and the top child pornography queries above, we can trend the demand for child pornography over the past 4 years, the results of which are shown in Figure 2. Trends for the other top child pornography-related terms did not have sufficient volume for Google to return data.

In general, the search trends show a decrease in child pornography searches conducted over the past four years. Specifically,

Figure 2. Child pornography search trends



the Google web-based searching appears to have declined – this may indicate a trending away from Google, a changing of search terms, a decline in overall demand, or a trending away from web-based demand. Overall, a 59.3% average decline in searches occurred for the four year period, with only the term “pthc” showing an increase. Similar to the 68.4% growth in search requests for the term “xxx”, the demand for “pthc” increased 63.5%. This particular growth may reflect the emergence of a “new” term – prior to February 2004, this particular term did not register on Google at a high enough volume to track. This particular term peaked in December 2005, and the latest results show a decline of 52.5% from that peak, consistent with the overall decline.

A trending away from Google toward other search engines is not likely given the growth of general pornography searches as shown by the comparison term “xxx,” and the general dominance of Google in the search marketplace. It is possible that demands by the United States Department of Justice for Google search data dampened the demand on that engine, but the timelines do not coincide as this only came to light in 2006, when the downward trend had already begun (Gonzales, 2006).

The likelihood of search terms having changed is again slight. There is a general downward trend amongst all terms, both sophisticated and unsophisticated. Additionally, several of the popular terms and words in the terms date to the 1970’s and 1980’s, such as “ls magazine” and “bbs”. This would appear to indicate little temporal change.

In terms of overall demand for child pornography, it is possible that an increase in the number of prosecutions and greater public awareness have decreased the overall demand. Alternatively, it is possible that Internet-based child pornography seekers have switched to peer-to-peer and other mediums and away from web-based searching. Determining which possibility is more likely would require a temporal analysis of peer-to-peer demand and is beyond the scope of this article.

Naturism and Child Pornography

Naturism (and in this context nudism) is a lifestyle in which nudity is encouraged while participating in non-sexual recreational and mundane activities. Naturist families may encourage their children (of all ages) to participate in the same activities while unclothed. One of the controversies surrounding naturism is that it exploits children, either directly by participation or indirectly by later through photographs of the activities. The determination of exploitation of naturist pictures posted to the web is not disputed – naturist photos have been found in many child pornography forensic examinations – but whether those seeking naturist and/or nudist picture are likely looking for child pornography has not been examined.

The top 20 search terms containing naturism and its permutations are shown in Figure 3. The majority of the terms related to nudism and naturism are ambiguous and cannot be directly associated with child pornography. The only unambiguous query was “underage nudist”, which shows a clear link to obtaining illicit content. Ambiguous terms included “nudist pictures”, “nudist pics”, “nudist videos”, “naturism free gallery”, “nudist gallery”, and “nudist photos”. These imply a search for images of those practicing naturism as opposed to seeking information on the lifestyle. Given the stigma on picture taking within many naturist communities, this implies non-naturists seeking the content, though it cannot be directly associated with pictures of children. Similarly, the queries “young nudist”, “nudist teens”, “teen nudist”, and “teenage nudist camp” imply the searcher is looking for younger (though not necessarily underage) nudist information. It cannot be ascertained whether or not those particular searches are seeking images or just information on the lifestyle.

Based solely on the search results, naturism queries cannot be said to imply a search for child pornography. There is evidence in the entire list of terms that those seeking child pornography look for images of nudists, but the majority of the content searches related to

Table 3. Most frequent nudism and naturism queries

Query	Count
nudist	77
young nudist	38
family nudism	30
naturism free gallery	23
nudism	20
nudist teens	20
nudist gallery	18
teen nudist	18
naturist	17
nudist girls	14
nudist pictures	13
nudist pics	12
nudist videos	11
naturism	11
nudist photos	10
teenage nudist camp	10
underage nudist	10
naturism young couples	9
russian nudist	9
nudist camps	9

naturism cannot be definitively associated with child pornography.

Child Pornography Bulletin Boards

To avoid detection, web-based child pornography has been moved to a bulletin-board based model of delivery. Those in the class of sophisticated searchers are more likely to search for boards related to child pornography than directly for the child pornography itself. While there are many forums present on a multitude of topics on the web, the term “bbs” has a primary association with child pornography.

Table 4 shows the top 10 queries associated with the term “bbs”. While the acronym is also associated with the terms “babies” and “bare-butt spanking”, “bulletin board system”

Table 4. Top terms related to “BBS”

Query	Count
lolita bbs	142
loli hc bbs	45
pedoland bbs pthc	28
pthc bbs gallery	27
tiny angels bbs	20
preteen bbs	18
bbs	17
dark bbs	17
sun bbs	16
lolitas bbs	16

appears to be the most common acronym in this context based on the returned results.

Of the top 10 results, the only term that is ambiguous is the singular use “bbs”. Given the context of the other terms, it is reasonable to assume that at least a portion of those searching for this term are seeking child pornography. All of the other terms are child-pornography related. Both “dark bbs” and “sun bbs” refer to bulletin boards that were associated with child pornography. The original boards are no longer around, but their names are still used as terms of art by sophisticated searchers.

While “bbs” is strongly associated with child pornography, other terms such as “message board” and “forum” showed no significant association. None of the top terms associated with either imply child pornography searches as was present with “bbs”, showing there is a large degree of linguistic specificity associated with the subculture.

Supply Prevalence

It is impossible to determine the overall prevalence of child pornography supplied over the web. Most of the web-based child pornography is believed to be served using transient sites and subpages on large, free hosting providers. Because of their transient nature, the limited links to these pages, and the disguising of those

links as detailed by Jenkins, the likelihood of them being indexed by a search engine is small (Jenkins, 2003) and quantifying them would be a Sisyphean task. As a result, many of the advanced searches appear to seek out forums and bulletin boards associated with child pornography that may contain pointers to the transient sites.

DISCUSSION

While child pornography is available on the web and there is a demand present, it is to a much lesser degree than shown on peer-to-peer networks and Usenet. (Steel, 2008) (Mehta, 2001) This is likely due to the perception of anonymity present using those technologies as opposed to the use of web, as well as a greater degree of prevalence. Because the lifespan of web-based child pornography sites that are publicly available would be extremely short, at any given time the available volume is small. Because of this, a forensic examiner is likely to find websites extracted from browsing history are no longer live, requiring cache reconstruction or similar mechanisms to show their original contents. As such, a more valuable technique than trying to reconstruct full browsing histories may be to show the individuals were seeking child pornography by way of search terms utilized.

Another key factor in the use of the peer-to-peer vis-à-vis the web to seek child pornography is best evidenced by the rank of the most popular child pornography-centered search. On the most common peer-to-peer network, the top ranked term was child pornography related. On the web, the highest ranking term was 198th amongst all search terms. While web-based child pornography is certainly available, the relative difference in frequencies of search may mean that a detailed analysis of peer-to-peer history might possibly yield greater fruit and be of a higher priority than a detailed browsing history.

The old practice of hosting on a temporary location and posting it to a bulletin board has largely been replaced by BitTorrent. By placing

a torrent link and finding an “open” tracker, the same result can be achieved without needing either the bandwidth or hosting space associated with the old method. Additionally, by using BitTorrent the original source can go away once saturation occurs. The same technique of posting the password days later to an encrypted file can be used to further hide the identity of the original distributor with less risk. This can be especially useful when doing a temporal forensic analysis. Showing a visit to forum-based website, followed by the download of a torrent file, the appearance of a new file, and a secondary visit to the website can be indicative of this behavior. The high co-occurrence of the term “bbs” with child pornography-related terms underscores the importance of forums to child pornographers and as a result forensic examiners.

The nature of peer-to-peer networks v. the World Wide Web similarly changes the search strategy. There is no concept of forums, so child pornography is sought directly instead of indirectly. Searches on peer-to-peer networks were more likely to search for terms likely to be present in filenames, whereas web based searches were more topical. The topical nature provides insights into the suspect’s behavior for exploitation during interviews as well as putting together potential victim profiles.

Another behavior evidenced in web-based child pornography searchers is the attempt to find seed pages. Like forums, seed pages do not generally contain child pornography themselves but may contain links to other sites which do. A common current modus operandi for child pornographers is to search for a child-pornography related term, find an aggregator site with thumbnails that link to movies, pictures, or other aggregator sites, then follow links depicting younger individuals until their desired content is found. Often, the initial aggregation pages contain only adult pornography, with one or two images of clothed children mixed in, clicking on which leads the seeker closer to the illicit content. Because the aggregators are dynamic in nature, a forensic analysis looking

for the original search term is often the most valuable action.

As a cautionary note, the presence of searches for terms sometimes popularly associated with child pornography (e.g. naturism and nudism) do not necessarily indicate the seeker is a child pornographer. If an individual is known to possess child pornography, they may be an indication of searches for more content (or worse, possible victims), but for the purposes of obtaining probable cause to search a machine the presence of those searches alone may not be sufficient.

One estimate that cannot be accurately extrapolated from searches is the number of users searching for child pornography on the web. Because individual users may search multiple times (as evidenced by two of the top search terms in the Dogpile results), there is no way to know the total number of users searching, and there is no way to know if child pornographers search more or less than average users, there is no accurate way to estimate the total number of individuals searching for child pornography.

FUTURE WORK

The decline in demand for child pornography through the use of search engines cannot be adequately explained without comparison based on the demand changes in peer-to-peer, newsgroup-based, and other distribution mechanisms. Revisiting the earlier Usenet studies and in several years repeating the peer-to-peer studies would provide a common ground for comparison of relative demand. Additionally, a comparison of the relative demands of different countries might provide insights into the global (or local) nature of the decline.

While the supply of child pornography would be difficult to quantify, it could be better qualified. An examination of the lifecycle of each method for child pornography distribution would provide insight into the likelihood of availability over time. Additionally, qualifying current trends in distribution would allow for

future studies to examine the subculture shifts as technology changes.

Finally, the construction of filtering/monitoring tools as well as forensic discovery tools based on the results of this research should be constructed and tested. The filtering/monitoring would be a direct application of the query lists generated, and existing techniques for automated detection of child pornography could incorporate linguistic parsing as well as image analysis.

CONCLUSION

Many of the queries became child pornography-related only through modifiers. While “cp” is ambiguous, “xxx cp videos” is strongly related to child pornography. Another difficulty in the research is the non-transitive nature of similarity. Because of this, “pre-teen” may be strongly associated with “pics”, but “pics” is not necessarily strongly associated with “pre-teen”. This shows a need for entire query analysis for monitoring systems, as opposed to singleton keyword matching to reduce false positives.

Additionally, because some of the terms are very short, forensics work is better off searching for phrases where possible. As an example, assuming random ASCII data on a 1TB hard drive, “bbs” would appear on **average** 65,536 times and “pthc” 256 times. As such, strings that have a minimum of five characters will produce a reasonable number of false positives on current drives.

Finally, the behavior of those searching for child pornography can be used to differentiate automated behavior (caused by malware) to that of a human by using a forensic reverse Turing test. It would highly unlikely that automated code entered search terms into a search engine, read forums related to child pornography, pieced together links that are written in a way to make non-human identification difficult, downloaded files from a website or through a torrent, then revisited the site later to obtain a password.

REFERENCES

- Beitzel, S. M., Jensen, E. C., Chowdhury, A., Grossman, D., & Frider, O. (2004). Hourly Analysis of a Very Large Topically Categorized Web. *SIGIR '04: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval* (pp. 321-328). Sheffield, United Kingdom: ACM.
- Beitzel, S. M., Jensen, E. C., Lewis, D. D., Chowdhury, A., & Frieder, O. (2007). Automatic classification of Web queries using very large unlabeled query logs. *ACM Transactions on Information Systems*, 25(2). comScore.com. (2008, June). *comScore Releases June 2008 U.S. Search Engine Rankings*. Retrieved July 19, 2008, from <http://www.comscore.com/press/release.asp?press=2337>
- GAO(2003). File-Sharing Programs: Child Pornography is Readily Accessible over Peer-to-Peer Networks.
- GAO. (2003). *File-Sharing Programs: Peer-to-Peer Networks Provide Ready Access to Child Pornography*. Government Accountability Office. Giganews. (n.d.). *Usenet Experiences Unprecedented Growth*. Retrieved February 23, 2009, from Giganews.com: <http://www.giganews.com/news/article/newsfeed-growth.html>
- Google.com. (n.d.). *Google Keywords Tool*. Retrieved July 17, 2008, from <https://adwords.google.com/select/KeywordToolExternal>
- Hoffman, D., & Novak, T. (1999). *A Detailed Critique of the TIME Article: 'On a Screen Near You: Cyberporn (DeWitt, 7/3/95)*. Retrieved December 2, 2007, from <http://sloan.ucr.edu/cyberporn/time.dewitt.htm>
- Hughers, D. R.-C. (2008). Supporting Law Enforcement in Digital Communities through Natural Language Analysis. *Proceedings of IWCF'08*. Washington, D.C.
- Hughes, D., Walkerdine, J., Coulson, G., & Gibson, S. (2006). Peer-to-peer: is deviant behavior the norm on P2P file-sharing networks? *Distributed Systems Online*, 7.
- Jenkins, P. (2003). *Beyond Tolerance: Child Pornography on the Internet*. New York: NYU Press.
- Libertus.net. (2008). *Statistics Laundering: false and fantastic figures*. Retrieved January 20, 2009, from <http://libertus.net/censor/resources/statistics-laundering.html>
- Mehta, M. (2001). Pornography Usage in Usenet: A Study of 9,800 Randomly Selected Images. *Cyberpsychology & Behavior*, 4, 695-703. doi:10.1089/109493101753376641
- Motion to Compel. Google, 98-CV-5591 (Northern District of California 2006).
- MSN AdCenter Labs. (n.d.). *Keyword Group Detection Tool*. Retrieved July 16, 2008
- Rimm, M. (1995). Marketing Pornography on the Information Superhighway: A Survey of 917,410 Images, Description, Short Stories and Animations Downloaded 8.5 Million Times by Consumers in Over 2000 Cities in Forty Countries, Provinces and Territories. *The Georgetown Law Journal*, 83, 1849-1915.
- Sandvine. (2008). *2008 Analysis of Traffic Demographics in North-American Broadband Networks*. Sandvine Incorporated.
- Scherer, R. (2005, August 17). Child porn rising on Web, prompting U.S. action. *Christian Science Monitor*.
- Sigel, L., & Sauer, G. (1999). *Critique of Rimm Article on Online Pornography*. Retrieved November 9, 2007, from <http://sloan.ucr.edu/cyberporn/sigel.sauer.critique.htm>
- Steel, C. M. (2008). Child Pornography Supply and Demand in Peer-to-Peer Networks. *International Family Violence and Child Victimization Research Conference*. Portsmouth, N.H.
- United States Government. (n.d.). Certain activities relating to material involving the sexual exploitation of minors. *U.S. Code, 18 USC §2252*.

Chad M.S. Steel (M'98) Steel holds bachelor's and master's degrees in computer engineering from Villanova University and is currently pursuing a PhD in computer science at Virginia Polytechnic Institute. He has served as the chief security officer and managing director of a Fortune 100 corporation, been the head of IT Investigations at a Global 100 corporation, and taught computer forensics in Penn State Great Valley's graduate program. He has written multiple articles and a popular book in the field of digital forensics. He currently does consulting work in computer forensics in the Washington, D.C. area.