In Search of the Perfect Filter:
Indexing Theory Implications for Internet
Blocking and Rating Software Products

Alvin M. Schrader
School of Library and Information Studies
University of Alberta
alvin.schrader@ualberta.ca

Growing public concern about controversial content on the Internet has inspired a political and technological challenge to open access. The political challenge is represented in a myriad of coercive legislative initiatives by politicians and government actors across the globe to control the technology of access to the Internet. The technological challenge takes the form of an increasing array of commercially-available software products that claim to filter or to rate expressive content on the Internet.

Almost everyone invokes technology-based arguments to buttress their views on Internet filtering and rating products, whether they are for or against them. Advocates and product owners present them as benign, neutral, and highly reliable means of enabling control; they respond to criticisms of software imperfections by countering that current limitations will be overcome as the technology improves, and are optimistic that technological solutions are just around the corner.

Even critics who are opposed to these products on intellectual freedom grounds tend to take refuge in arguments about their technological limitations (some go so far as to say "scientific" limitations). To date, public discourse is fixated on the technology of access to the Internet.

This paper approaches the issue from neither a technological nor a philosophical perspective, but rather is grounded in the foundational framework of indexing and retrieval theory. It explores the application of indexing principles to the tasks of identifying, describing, regulating, and prohibiting Internet content. Even though the purpose of Internet filtering and rating products is to control and prevent access to information and images rather than to facilitate their access, the intellectual operations of identifying and describing expressive content are similar to those employed in conventional indexing systems for subject representation.

The result of the critique explicated here is a view of the limitations inherent in both language and its indexing that render perfect control over expressive content a theoretical impossibility in any communications medium. Internet filtering and rating products are no exception. The limitations issue from the unsolvable problems of ambiguity in language, reading, and indexing. Internet filtering and rating products will not work at the performance levels claimed by their owners and marketers because of the permanent variability in human cognitive and communicative processes, not on account of "the technology".

Introduction

Converging communication technologies offer mesmerizing potentialities for global access to local culture (more precisely, they create electronic access to digital versions of ideas, information, stories, images, and sound). At its deepest level, the Internet is praised as one of the most democratic media yet to be invented, "a babble of millions of ungoverned and ungovernable voices" (Book and Periodical Council 1997, 14). Or, as Karen Schneider described it recently, the Internet brings together "the good, the bad, the ugly, the inaccurate, and the outdated" (Schneider 1997b, xiii).

What troubles many people is precisely this sort of juxtaposition. The Internet makes accessible not only the generally positive aspects of recorded cultures but also all of the elements that may be considered repugnant and offensive by someone, somewhere, at some time. Such offence is typically prompted by representations of sex, sexuality, nudity, violence, hatred, profanity, religious belief, political ideology, gender, class, race, and power relations that vary from those finding favour within one's own system of customs and values. And for parents, exposure of children and adolescents to these objectionable representations magnifies the clash of cultures. As a consequence, what follows close on the heels of euphoria over the liberating possibilities of the Internet are fear and panic.

But this is not a late-20th century phenomenon. Throughout recent history, every new technology for public communication—from the printing press in the 15th century to the major 19th and 20th century inventions of photography, motion pictures, sound

recordings, radio, television, video, and the Internet—has been subjected to ingenious forms of control and censorship. Hard-won battles for unfettered access are fought anew with each new development in communications technology. Be it church, state, lobby group, or individual citizen, the response has been an effort to control access to ideas, images, and sounds that threaten those with power and influence. Efforts to censor occur at the social, political, and economic levels, and take one of two classic approaches: attempts to exert control over content, and attempts to exert control over audiences.

It would appear that the impulse to control is as true of the Internet today as it was of the printing press more than 500 years ago. While politicians around the world are engaged in legislating varying degrees of suppression of Internet content, the marketplace has also responded to Internet fears, particularly in North America. Over the past two or three years, a bewildering array of software products has appeared on the U.S. and Canadian markets that offer a variety of filtering and rating options that claim to be able to control and regulate expressive content on the Internet.

These filtering and rating products are being presented, both in the marketplace and the political arena, as benign, objective, and highly reliable means of regulating, restricting, and prohibiting undesirable materials and images. Typical product claims are couched in the rhetoric of child protection and parental guidance. But while much of the initial advertising rhetoric has focused on ensuring safety and appropriateness for young people in the family home, entrepreneurs are not unaware of the immense profit potential that these products have in business and institutional applications, and some of them have targeted library audiences in particular.

These technologies, however, are being adopted in the home, by libraries and schools, in government, and by business without adequate investigation and evaluation. For the most part, observations to date about their performance have been anecdotal and atheoretical.

The purpose of this paper is to describe and critique these emerging technologies for regulating, restricting, and prohibiting access to ex-

pressive content on the Internet. Several bodies of theory and principles that form part of the foundational knowledge of library and information studies offer a powerful framework for thinking about the theoretical feasibility of controlling Internet content through technology. These bodies of thought are intellectual freedom, reader response theory, and indexing and retrieval theory. This paper focuses on the third of these theoretical elements, indexing and retrieval, to inform a critical analysis of claims made by and on behalf of Internet filtering and rating products.

Software Products for Filtering and Rating Expressive Content on the Internet

At last count, more than 40 software products were available commercially, configured, variously, for individual workstations, local area networks, remote vendor servers, ISPs, and other arrangements. They offered an array of software options for controlling and suppressing expressive content on the Internet:

> • stoplists of "bad words", "bad phrases", and even "bad syllables", that is, keyword blocking that requires the identification in painstaking detail by either the product owner, the customer, or a third party of every conceivable synonym and euphemism;

> • lists of "bad sites", identified through stoplists, which block at various levels—the domain or host level, the IP address, the directory level, or the file level;

> • lists of "bad topics", which organize objectionable sites according to broad subject categories devised by the product owner;

> • content rating systems, which use the technology of PICS (Platform for Internet Content Selection) to block access to all unrated sites and/or to those carrying particular ratings assigned either voluntarily by site owners on the basis of a self-administered ratings questionnaire or by third party reviewers, human or robot; and

- lists of "bad resources", which block access to services such as telnet, FTP, Internet chat, or newsgroups.

To achieve these goals, product owners rely almost exclusively on automation rather than on human eyes and brains—marketing claims to the contrary. Software robots, either existing or customized mechanical search engines referred to as web crawlers or web spiders, are used to search for and identify unacceptable Internet content. Some products employ a small staff in addition who review a similarly small proportion of sites that have been identified to some threshold level by the software robots (Censorware Project 1997). A glimpse into the general functioning of these robots is afforded by this description of CyberPatrol's customized search software, Cyber Spyder:

> Cyber Spyder visits the sites and creates a report including 25 characters before and 25 characters after each occurrence of the keywords used in a particular search. The researchers start by reviewing this report. If necessary, the sites are visited and viewed by a human being before being added to the CyberNOT list. If not necessary, the sites are not viewed or added. For example, if the context of the word "breast" was the proper way to prepare chicken, that is a good indication that the site doesn't meet the CyberNOT criteria (Censorware Project 1997).

The current language of Internet searching is essentially metaphorical, whether that searching is for information and images deemed objectionable or for more socially conforming expression: searching, browsing, surfing, navigating, identifying, analyzing, reviewing—these are all human cognitive processes. A more precise representation of Internet searching, however, would describe it as an automated process, a computerized scoring algorithm based on pattern or character recognition. Internet filtering is nothing more than exact-match character recognition in a free text environment. Nonetheless, one product describes its software as a "state of the art, context sensitive phrase filtering 'engine' to identify objectionable web sites," claiming that it is 90 percent effective "without even knowing where the objectionable material is" (CyberSitter 1997).

The magnitude of the task that the owners of filtering and rating products have undertaken is formidable. With 1,000 to 3,000 new sites appearing each day in the U.S. alone, the Internet is a dynamic phenomenon that leaves product owners shooting at moving—and indeed, speeding—targets. A recent estimate of the size of the Internet suggests that there are over 300 million web pages currently accessible to casual browsers, a number expected to grow by 1,000 percent in the next few years. Of this total, the most comprehensive index of scientific information alone covers only 34 percent of all indexable pages (HotBot), while other search engines perform even more poorly, AltaVista at 28 percent coverage, Northern Light 20 percent, Excite 14 percent, and Lycos at 3 percent (Lawrence 1998).

An alternative approach that one or two products offer is a programming "shell" that allows customers to create their own stoplist of prohibited words, phrases, or sites. Various products also offer other capabilities: downloading updated stoplists manually or automatically, sometimes as often as weekly; restricting access to everything except a list of sites containing content deemed suitable for children; blocking all unrated sites; developing, editing, or adding a customized list of offensive words or sites; obtaining a report of all sites visited or site violation attempts; issuing warnings; restricting access to the computer if a certain number of forbidden sites is accessed or even shutting down the Internet connection; restricting access based on time of day or on total Internet time used; blocking email; and, blocking out-bound transmission of credit card numbers, family name, home address, and telephone number.

Although age-appropriateness is the official concern of product owners and marketers, there is no standardization in software capabilities or policies. While most have a pre-programmed stoplist of terms, sites, and topics, few permit customers to view their stoplist and, indeed, the majority regard them as highly valuable commercial trade secrets; curiously, some of the most secretive manufacturers permit customers to view their lists of periodic updates to the master pre-programmed stoplist. Only a few permit customers to disable the stoplist.

Most of the products organize their targeted sites into subject groupings or categories, but these categories are not standardized across

the marketplace, so there is no uniform classification system or authority control, no MARC-like record for classifying and describing a site. Some products have only three or four categories, while others have up to 30; the typical product has 10 to 15 categories. Some products frequently assign each undesirable site to multiple categories.

Typically, product categories reach far beyond sex and violence. As Schneider has observed, the categories "read like a laundry list of human concerns, with some venal sins thrown in" (Schneider 1998, 37). One product has, for example, in addition to four categories for sex and violence (violence/profanity, partial nudity, full nudity, and sexual acts), eight other categories: gross depictions, intolerance, satanic/cult, drugs/drug culture, militant/extremist, sex education, questionable/illegal and gambling, and, alcohol and tobacco. And another product assigns sites according to one or more of the following 13 subject categories: adult and mature subject matter of a sexual nature, pornography or adult oriented graphics, drugs or alcohol, illegal activities, gross depictions or mayhem, violence or anarchy, hate groups, racist groups, anti-Semitic groups, advocating of intolerance, computer hacking, advocating violation of copyright laws, and, any site publishing information interfering with the legal rights and obligations of a parent or the product's customers.

In contrast, another categorizes sites under the following five headings only: adult/sexually oriented, adult/violence, gay/lesbian activities, advocating illegal/radical activities, and, advocating hate/intolerance. And yet another uses only four categories: sexually explicit, violence/hate speech, drugs/alcohol, and gambling.

This cursory overview shows that product manufacturers are targeting divergent materials based on divergent criteria, a pattern reflected in the wide range of Internet sites currently blocked, with claims from as low as 15,000 to as high as 138,000 (Oder 1997, 41). One Internet rating product also reports that it has rated 1.5 million URLs in a mere six months (NetShepherd 1997).

As an aside, none of the products notifies the owners of Internet sites that theirs has been blocked or rated, and it is only by accident that anyone discovers that their site has been targeted; complaints

about inappropriate categorization are frequently ignored or dealt with very slowly. Sites critical of CyberSitter, for example, have recently sued the product owner Solid Oak Software for blocking them (Schneider 1997b, 114), and others have threatened to take legal action against CyberPatrol for assigning their sites to what they allege are unjustifiable and defamatory categories (Meeks and McCullagh 1996).

Indexing Theory Implications for Filtering and Rating Internet Content

What is the relevance of indexing and retrieval theory to these software products? Even though their purpose is to control and prohibit access to information and images rather than to facilitate their access, the intellectual operations involved in identifying and describing expressive content for targeting are similar to those employed in conventional retrieval systems—as are the problems and challenges.

The goal of indexing in conventional retrieval systems is to represent information, that is, to provide a systematic guide to the contents of information records. More generally, the goal is to name information, to gather together ideas into categories so that everything on a subject can be identified. In order to accomplish this end, the indexer must decide first what concept or characteristic is to be represented, then what name to give the concept or characteristic, and finally how to organize the designated names (descriptors) into a searchable database.

There is also a very special feature of indexing operations and retrieval systems: materials both for and against a subject are regarded as being "about" the same topic and are therefore normally gathered together under the same classification number and under the same index term. Additional linguistic devices connect and control related terminology to maintain consistency and avoid redundancy, hence the concept of authority control. So, for example, in order to provide access to the literature on abortion as an ethical issue, both pro-choice and pro-life materials are classified under the same number in the *Dewey Decimal Classification* (179.76 Abortion under

179.7 Respect and disrespect for human life), and in the *Library of Congress Subject Headings* under the generic "Abortion—Moral and ethical aspects". Even an index term such as "Pro-life movement" will encompass oppositional critiques. In short, indexing and retrieval systems are a kind of "fuzzy system," designed to accommodate, however imperfectly, vagueness in language just as easily as exactitude, and partial truth in human understanding just as easily as received wisdom.

Like language itself, indexing operations for retrieval are ambiguous. Since indexing operations are bounded by language, the representation of subjects is similarly ambiguous, susceptible to nuance, imprecision, inconsistency, cultural variation, and unpredictable change over time. These operations of representation, then, pose immediate problems for effective identification, control, and retrieval of materials. The problems concerned are the concepts of aboutness, specificity, consistency, exhaustivity, relevance, and universality.

Aboutness is the central problem of indexing: how does one decide what a text is about? How does one determine context? Coextensiveness is the match of aboutness with its representation in the descriptor system: "the extent to which the index term reflects the precise content of the item of information as compared with assigning the item to a preformed class that may or may not be reasonably coextensive with its actual topic" (Milstead 1984, 143).

Specificity is the opposite problem: how does one decide which aspects of a text will be represented, and with which terms? Specificity is a characteristic of the language, and if one cannot express the specific, then one can only resort to a broader level of linguistic representation.

Consistency is also involved in the problem of indexing specificity, because synonymous terminology must be identified and controlled, and hence one element in the need for authority control. Specificity is also one aspect of coextensiveness.

Exhaustivity prompts the question of how many concepts from a document will be represented. Relevance is related to exhaustivity

:n posing the challenge of identifying and indexing only those aspects that people want information about. So the representation of a document must exhaust all topics of likely interest to users but at the same time be parsimonious in omitting topics likely to be irrelevant to them.

Finally, the assumption of universality challenges all indexing operations. Universality is the pervasive—and mistaken—belief that there is a one-to-one correspondence, an absolute link, between concept and descriptor that transcends not only culture, ideology, and time but age and reading differences as well. Universality assumes the existence of one and only one cultural perspective—and generally it is the indexer's.

With Internet stoplists of words, phrases, sites, and topics, the problems of identifying and describing concepts are similar to those encountered in free text searching in conventional retrieval systems. Free text searching rests on the assumption that what a text is about can be succinctly captured in individual words and phrases, discrete strings of characters. Free text searching is therefore context- and concept-free, permitting no human intermediary to impose an intellectual structure for effective retrieval. It is word focussed, not concept focussed. Thus, even when the goal is to block or rate expressive content on the Internet rather than to represent it, the problems approximate those involved in free text searching in conventional retrieval systems.

Those problems are ubiquitous. In one type of filtering product option, for example, "x's" or blank spaces are substituted for the offending word or phrase. The result of this literalness is to make gibberish of the text. Schneider has documented an instance of this in a search she undertook to verify the OCLC record for *Our Tribe* by Nancy Wilson (Schneider 1997a). With the filtering product in operation, the title was shown as:

Our tribe : $b   folks, God, Jesus, and the Bible /
        $c Nancy Wilson.

And there were odd blanks in the 650 fields:

15 650 0 $x Religious aspects $x Christianity.
15 650 0 $x Biblical teaching.

Similarly, a title on human sexuality by the well-known Bishop John Shelby Spong was represented as:

Living in sin? : $b a bishop rethinks human   /
      $c John Shelby Spong.

And yet the 650 field was:

13 650 0   Sex $x Religious aspects $x Christianity.

But the most egregious free text blocking makes the offending words, sites, and topics disappear, utterly invisible to searchers so that they are completely unaware that suppressed information even exists. The only way one would discover this would be to search for already known items, sites, or topics. For example, in the "Society and Culture" category in Yahoo! is a heading for sexuality, but searching through one filtering product makes this topic simply vanish from the subject listing (Carroll and Broadhead 1996, 568-569). In another case, the product temporarily prevented access to the entire library web site of the Archie R. Dykes Medical Library because corporate policy blocks wholesale the topic of homosexuality, and therefore of course the term "dyke" (Chelton 1997). The spokesperson for another product that also attempts to suppress information about homosexuality has been quoted as saying: "We filter anything that has to do with sex. Sexual orientation [is about sex] by virtue of the fact that it has sex in the name" (quoted by Meeks and McCullagh 1996).

Targeting "sex", however, also blocks the newsgroup dedicated to Star Trek's Captain Jean-Luc Picard, alt.sexy.bald.captain, the NASA site marsexplorer.com, the works of poet Ann Sexton, sexual harassment sites, sexual abuse sites, sites about gender (sex) discrimination in the workplace, and sites providing information about sexually transmitted diseases.

Consistency in the assignment of index terms is a well-known problem in retrieval research. Inter-indexer consistency studies show

over and over again that there is a great deal of variation in levels of agreement among indexers on their assignment of terms representing the subject content of a text. Consistency among indexers ranges from a minimal 4 percent to 82 percent (Markey 1984, 155-160). Research also shows that greater levels of consistency are achieved when indexers choose terms from a controlled vocabulary, with consistency scores then ranging from 34 percent to approximately the same ceiling (Markey 1984, 161).

In the light of these long-known patterns, why would Internet technologies achieve the much higher rates of indexing consistency, exhaustivity, specificity, and certainty that their advertising rhetoric claims? Their scripts and those of their advocates are full of unqualified assurances of prevention, protection, safeguarding, 24-hour monitoring, child safety, empowering parents, peace of mind, keeping an eye on everything, complete control, truly complete solutions, very sophisticated methods of controlling access, categories rated to a fine level of granularity in filtering.

Even when the promotional claims are more qualified, however—when, for example, the promise is to provide for the "relative safety" of children exploring the Net, or there is an explicit disclaimer of responsibility either for accuracy or completeness of information or for errors or omissions (CyberPatrol 1997a)—the vast majority of parents and other consumers will not long remember the qualifiers or the "fine print", if indeed they notice any of it in the first place.

In reality, the new technologies do not live up to their promises at all. In a recent small-scale study conducted by *Consumer Reports* of 22 easy-to-find websites that had been judged by investigators to be inappropriate for young children, not one of the four most common software blockers—CyberPatrol, CyberSitter, NetNanny, and SurfWatch—blocked all of the sites. NetNanny failed to block any of the 22 sites, while 14 were blocked by CyberSitter, 16 by CyberPatrol, and 18 by SurfWatch; and only 3 sites were blocked by Internet Explorer (Is your kid 1997, 30). These rates are far below the levels that parents and other consumers have been led to expect. Another small-scale study by *PC World* found marginally better performance: two of the five

products tested were effective in blocking all ten of the adult-oriented sites in the evaluation (Internet filters 1997).

Anecdotal testimony abounds on listservs, websites, and in print. Some products block feminist sites such as NOW, the National Organization for Women, alt.feminism, soc.support.pregnancy.loss, and soc.support.fat-acceptance, and feminist newsgroups. Also blocked are the sites of Planned Parenthood, the reference book publisher Sinauer Associates, information about sexual dysfunction, and Kierkegaard, the sites of the Quakers, the American Association of University Women, and the U.S. Central Intelligence Agency. The entire HotWired domain is blocked by one product, as is a site dedicated to the safe use of fireworks. One product temporarily blocked the important Holocaust archive and anti-revisionist resource site Nizkor, because it contained "hate speech" (Wallace 1997c).

Additional targets have been newsgroups of pagans, naturists, and other alternative culture activists such as misc.activism.progressive, misc.headlines, misc.health.aids, misc.health.alternative, talk.abortion, talk.euthanasia, talk.politics.drugs, alt.atheism, alt.teens, and 220 newsgroups prefixed alt.support (Thompson 1998, 7). One product has blocked the entire Echo ISP, New York City's oldest online community, another articles about AIDS and HIV from clarinet's AP and Reuters, another the University of Newcastle's computer science department site, and another all mailing lists run out of cs.colorado.edu.

Lesbian and gay sites are also regularly targeted by several products, sites such as the Queer Resources Directory, the Critical Path AIDS Project, and the HIV Info Center; gay political and journalism newsgroups such as clari.news.gays, alt.journalism.gay-press, alt.politics.homosexual, soc.support.youth.gay-lesbian-bi; and chatrooms for gays and lesbians.

In the most extensive analysis yet undertaken of a particular product, "Blacklisted by CyberPatrol: From Ada to Yoyo", a group of writers and Internet activists calling themselves the Censorware Project have documented a large number of sites that CyberPatrol has blocked inappropriately, and concludes that correcting errors is

a low priority for the product owners (Censorware Project 1997b). The report's title comes from two of the blocked sites discovered to be listed in the product's "full nude" and "sexual acts" categories: one, the site of a political committee in Ada, Michigan concerned with responsible township government; the other, a student run server at Monash University in Melbourne, Australia that has 5,000 student accounts and over 6,000 web pages.

The Censorware Project identified hundreds of unjustifiably blocked sites. Among them were an award-winning pet care site; the site for a Nike poster advertising shoes named after Penny Hardaway; the MIT Project on Mathematics and Computation site; a site about bodybuilding products and protein bars; the "Sunset Strip" neighbourhood of GeoCities that contains 90,000 individual sites about rock, grunge, punk, techno, and the alternative music scene; the National Academic of Clinical biochemistry site; a server at the Chiba Institute of Technology in Japan; the site of the U.S. Army Corps of Engineers Construction Engineering Research Laboratories; *Mother Jones* magazine online; the Envirolink site, a clearinghouse of environmental information; the HIV/AIDS Information Center site of the *Journal of the American Medical Association*; and numerous gay sites including one that advertises gourmet coffees, teas, food, and gifts, and another, one of the most contentious blocks, "West Hollywood", which is a GeoCities neighbourhood that contains over 23,000 gay and lesbian sites with over 50,000 webpages.

Over 50 Internet server machines were also blocked outright, including one, members.tripod.com, that has almost 800,000 members and contains 1.4 million webpages blocked on all twelve of CyberPatrol's categories. Well over 300 newsgroups were also blocked, including 100 groups in the entire rec.games.* hierarchy, 220 groups in the alt.sup* hierarchy, the soc. support* hierarchy, many groups in the soc.* hierarchy, and the alt.cyber* tree (Censorware Project 1997b).

The most comprehensive study to date, The Internet Filter Assessment Project (TIFAP), was coordinated by Schneider and involved 40 volunteers who tested 13 products on 100 questions over a six-month period April-September 1997 (Schneider 1997b). All of the

products were found to hinder information retrieval, blocking innocuous occurrences of words and many sites with information similar to what would be found in libraries—but at the same time permitting access to offensive sites. One product blocked "good sites" 5 to 10 percent of the time while pornographic sites slipped through about 10 percent (Schneider 1998, 37). The poem "pussycat, pussycat" was blocked consistently, but "Roger" was never blocked although it is Australian slang for penis. Sites for "Alternative Journals" and "Activist Groups" were blocked by two products. Sites for hate groups, press releases on sex offenders, an interview with Leslea Newman, a list of jockeys, safe sex information—all were blocked, among many others (Schneider 1997b, 93).

But why such low rates of indexing consistency, exhaustivity, specificity, and certainty by Internet filtering and rating products? The answers are found in the essential ambiguities that language presents for both free text searching and subject identification. Every language has a multitude of synonyms and antonyms and euphemisms, puns and double entendres, and in English there are generally slippery terms like "objectionable", "patently offensive", "degrading", "harmful", "morally dangerous", and "pornographic". There are a number of notoriously slippery legal concepts in North America—"obscenity", "undueness" (as in undue exploitation of sex), and "community standards"—as well as the newer concepts being lobbied for legislative sanction, among which the most prominent in the U.S. Congress are "indecency" and "harmful to minors".

And there are homographs, words with multiple meanings, of which only one meaning is sexual or otherwise controversial, such as gay, pansy, fairy, queer, cock, cherry, shag, crab, wiener, woody, bitch, and curse, and there are literally hundreds if not thousands of other terms, many slang, that refer to sex, sexuality, and genitalia. In 1996, for example, the word "couple" was added to one product stoplist, resulting in part of the White House site being temporarily blocked because "couples" appeared in a reference to the Clintons and the Gores. The same product also blocked the site for Super Bowl XXXI and a hockey site because of news that a player had been sidelined due to a groin injury (*Newsletter on Intellectual Freedom* 1997, 29). And one of the earliest and most widely publicized

examples was the blocking by America Online of the word "breast" from some areas of its service. Other homographic variations consist of terms used as proper names such as Butts, Dykes, Gay, and those in which the string 'sex' appears.

We also witness new terms invented, such as "Bobbitt", "cyberporn", and former Prime Minister Pierre Trudeau's "fuddle-duddle", while older terms are twisted into new meanings, such as "rock and roll", which was originally African American slang for sexual congress, "snow" for cocaine, "political correctness", and "family values". Pejorative epithets are sometimes appropriated by marginalized minorities and turned into affirmations of group pride, terms such as "queer" and "nigger". Terms also go out of fashion, such as "French letter" for "condom", "hooch" for alcohol, and "cats" for men.

Then there is the problem of English usage outside the U.S., particularly the problem of slang—"randy", and "the full Monty" are British examples with which American stoplist producers are unlikely to be familiar. And there are foreign terms imported into English in some regions, as in the example of "merde" in a book by Geneviève Edis that is entitled *The Complete Merde! The REAL French You Were Never Taught in School*, or the Hebrew term "schmuck", which is extremely pejorative slang for penis.

There are also inherent category problems in all indexing systems. As Trinh T. Minh-Ha has written, categories always leak (Trinh 1989, 119). Complex concepts do not fit into simple compartments. Is all violence of the same kind? Is a punch the same as an execution? Is sex the same as nudity? Is erotica the same as the sexually explicit? A recent example is the boast of CleanNet that it distinguishes easily between hard-core pornography and sites supporting "art and literature"—yet blocks access to the swimsuit edition of *Sports Illustrated* (Jones 1998).

Internet filtering and rating products also impose, inevitably, cultural and ideological agendas. Some of them acknowledge the existence of a value system or some sort of socio-political framework that informs their filtering and rating judgments, while others are ambivalent; one cautiously states that its objective is to filter sites that "may not be suitable for all audiences" (NetNanny, undated).

Others appear confused in their public utterances. CyberSitter, for example, says that it makes no apology for its choices:

> We don't simply block pornography. That's not the intention of the product. The majority of our customers are strong family-oriented people with traditional family values. Our product is sold by Focus on the Family because we allow the parents to select fairly strict guidelines" (quoted by Meeks and McCullagh 1996).

But at the same, it denies that it has an agenda, claiming to block only those sites that "meet a pre-defined criteria...without exception" and that it "has no agenda of any kind, unless you consider protection of children a hidden agenda" (CyberSitter 1997).

CyberPatrol acknowledges that an explicit value system informs its operations: "In evaluating a site for inclusion in the list, we consider the effect of the site on a typical twelve year old searching the Internet unaccompanied by a parent or educator". And its blocking criteria pertain to "advocacy information: how to obtain inappropriate materials and or how to build, grow or use said materials. The categories do not pertain to sites containing opinion or educational material, such a the historical use of marijuana or the political situation in Germany during the 1930s and subsequent World War II" (CyberPatrol 1997a, 1997b).

Whether the agenda is hidden or explicit, however, indexing decisions are nonetheless an interplay of indexer judgment and the text. The naming of information through indexing is of necessity personalized: there must be a perspective, a viewpoint, a cultural framework. This is clearly evident, for example, in CyberPatrol's reference standard of the hypothetical twelve-year-old, in the nuanced phrasing of its blocking criteria based on a level of "possibly objectionable content" (CyberPatrol 1997a) and "potentially objectionable material" (CyberPatrol 1998), and in its identification of "sex education" as a censorable subject category.

CyberPatrol is not alone in having an ideological framework that far exceeds mainstream concerns about inappropriate sex and vio-

lence on the Internet. WebSense, for example, has 29 categories, only three of which are sex-related, and of the 12 assigned by CyberPatrol, only three are sex-related (Schneider 1998, 68). CyberPatrol's subject category of "gross depictions" also illustrates how wide ranging its value system is: gross depictions are "crudely vulgar or grossly deficient in civility or behavior or which show scatological impropriety" and include depictions of maiming, bloody figures, autopsy photos, or indecent bodily functions (CyberPatrol 1997b). The corporate mindset at CyberPatrol casts its umbrella wide enough with this operational definition to block the sites of animal rights groups because some show pictures of syphilis-infected monkeys and dogs tossed into garbage dumps. And the broad sweep of its "militant/extremist" category includes sites run by gun rights and gun advocates groups such as the Silicon Valley chapter of the National Rifle Association.

Moreover, since the naming of information for retrieval involves control, an exercise of power is involved (Olson 1996, 4-5). The ethical exercise of that power requires continuous self-awareness and self-reflection in the decisions that the indexer makes. It also requires that decisions and their rationale be made public and subjected to public debate.

The urgent need for disclosure and scrutiny introduces an even more disturbing aspect of the exercise of power and ideology by some software products. In retaliation for criticism on the Internet, CyberSitter's "bad word" list was revised last year to block access to sites containing the phrase "Don't Buy CyberSitter!" as well as to the site for Peacefire, a student organization opposing Internet censorship in any form. CyberSitter has also blocked the site for *The Ethical Spectacle* (www.spectacle.org), apparently in retaliation for the Webzine's criticism of the company that owns the program, Solid Oak Software; and in fact, even communicating with CyberSitter to ask questions about its blocking of the site prompted a message from the product owner, Solid Oak Software, Inc., accusing the inquirer of harassment or political motivation (Shallit 1997).

And CyberPatrol has similarly gone to extreme lengths to suppress criticism, blocking sites that oppose it or its approach to content

control. A recent example of the product's attitude towards criticism is the blocking of web pages pertaining to *Sex, Laws, and Cyberspace* by Jonathan Wallace and Mark Mangan, published by Henry Holt in 1996 (Wallace 1997). One product or another has also blocked the sites of the Electronic Frontier Foundation's censorship archive, the League for Programming Freedom at MIT, the MIT Student Association for Freedom of Expression, which is a group that opposes software patents, and a site critical of America Online, alt.aol-sucks.moderated.

Summary

What the principles of indexing and retrieval tell us is that Internet filtering and rating technologies are theoretically unworkable, that the essential ambiguities of language, reading, and subject representation ensure the failure of automated searching for objectionable content. The problems of identifying and describing such content for purposes of control and prohibition are intractable: new sites, new terms, new issues, the world cacophony of languages, variable interpretations of meaning, variable perceptions of offensiveness, variable perceptions of age appropriateness, infinitely variable descriptor terms, texts in languages other than English, foreign language words adopted into English, culture-specific values and priorities (for example, marked differences in American and European attitudes to violence, sex, and nudity), and even regional spelling variations.

These ambiguities and dynamics prevent blocking and rating strategies from ever being successful in controlling the world of ideas at levels of consistency, exhaustivity, specificity, and certainty that would be sophisticated enough to satisfy critics, reassure parents, and relieve librarians and teachers of unpleasant encounters with complainants. Human language is just too unstable, words and meanings are just too indeterminate, too elastic, too mutable, too imperfect. As one critic has put it, "safe-only access can not happen because individual perceptions of safe are as varied as the number of sites on the Internet" (Crosslin 1998, 52).

And yet the very names of the software products—nanny, patrol, shepherd, sitter, watch—conjure up images of unqualified protec-

tion, safety, guidance, and comfort. NetNanny advertising, for example, says: "NetNanny is watching when parents aren't."

Instead of fulfilling these explicit advertising promises, however, what the new products offer is the illusion of success—an illusion that comes with a high price tag. One price is a false sense of security. And its twin is a false sense of confidence that all appropriate information will still be retrieved when one searches the Internet.

Another price is intellectual freedom. Since indexing for any retrieval system is about the control of ideas as much as it is about access, the dangers to intellectual freedom are always imminent. The crude, paternalistic strategies adopted by filtering and rating products should serve to remind us that authority control keeps some voices out just as easily as it lets others in. When blocking and rating decisions are made by unknown third parties with unknown qualifications and unknown ideological agendas, the danger to public debate is palpable. With a broad sweep, these products indict all representations of violence, sex, hatred, and other targets as equally bad, and as especially bad for young people.

They attempt to impose an "objective", highly simplified standard of description and measurement on what is essentially a complex and highly variable matter of personal tastes, individualized family values, highly variable perceptions of age appropriateness, and widely varying thresholds of social tolerance. In spite of the denials of some product owners and marketers, any operation that identifies words, phrases, topics, and sites for blocking is of necessity imposing an ideological agenda or value system.

The great irony of Internet filtering and rating products is that sites arguing against the perceived enemy will also be restricted or suppressed: anti-gay sites, anti-feminist sites, anti-drug sites, anti-smoking sites, sites about recovery from sexual abuse, anti-sex education sites, anti-abortion sites, and even sexual abstinence sites.

To sum up the lessons of indexing theory, Internet filtering and rating software appears to represent a worst case scenario in the representation of information. These commercial operations dis-

sect text and images on the basis of pre-determined corporate value systems into selected parts, which are then highlighted for prohibition or access restrictions. Wholeness and context are sacrificed; ideas are pigeonholed; integrity of text and respect for reader are ignored in favour of a single, uniform standard of "safe" words and "safe" ideas irrespective of age or maturity. One four-letter word becomes more important than 400 pages of story. Margaret Laurence is said to have called this "snippet censorship," the practice of basing one's judgment of a work on excerpts, offending words or phrases, and scenes lifted out of context (Carver 1997). The result is content targeting that is vague, overbroad, and incapable of adequate discrimination by age and topic.

The new products take judgment, control, and accountability away from indexing and cataloguing professionals in favour of third-party commercial interests that rely on faceless and unaccountable software robots for exact character matching. Automatic searching for objectionable content has no more chance of attaining an acceptable level of success than has computerized translation—or even computerized spellchecking. As Schneider succinctly describes them, Internet filters are "mechanical tools wrapped around subjective judgment" (Schneider 1997b, xiv). And Meeks and McCullagh (1996) have written: "Technology is no substitute for conscience."

What society needs far more than commercial exploitation of naive fearmongering is funding support for world-wide classification of digital resources in the great tradition of the theory and practice of library and information studies.

References

Book and Periodical Council. 1997. Freedom to Read Week Kit. Toronto, Ontario: Book and Periodical Council, 13-14.

Carroll, Jim and Rick Broadhead. 1996. *1997 Canadian Internet Handbook*. Scarborough, Ontario: Prentice Hall Canada.

Carver, Peter. 1997. "Battle over novel" (letter to the editor). *Globe and Mail*. February 7.

Censorware Project. 1997a. "Blacklisted by CyberPatrol: from Ada to Yoyo." December. Internet WWW page, at URL: < http://www.spectacle.org/cwp >

Censorware Project. 1997b. "Blacklisted by CyberPatrol: from Ada to Yoyo. Problems." December. Internet WWW page, at URL: < http://www.spectacle.org/cwp/problems.html >

Chelton, Mary K. 1997. "Internet names and filtering software." March 4. Email message from < cheltonm@esumail.emporia.edu >

Crosslin, Donna. 1998. "Unsafe at any modem speed" (letter to the editor). *American Libraries.* 29 January: 52.

CyberPatrol. 1997a. "Overview: the CyberNOT block list." May 11. Internet WWW page, at URL: < http://www.cyberpatrol.com/cp_block.htm >

CyberPatrol. 1997b. "CyberPatrol CyberNOT list criteria." May 11. Internet WWW page, at URL: < http://www.microsys.com/cyber/cp_list.htm >

CyberPatrol. 1998. "Internet acceptable use guide." Internet WWW page, at URL: < http://www.cyberpatrol.com/cpc.aug.htm >

CyberSitter. 1997. "Frequently asked questions about CyberSITTER." Internet WWW page, at URL: < http://www.solidoak.com/cyberfaq.htm >

Internet filters. 1997. "Internet filters: The smut stops here. Or does it? Screening five top web filters." *PC World.* October. Also available at URL: < http://www.pcworld.com/softwar.../articles/oct97/1510p078a.html >

Is your kid. 1997. "Is your kid caught up in the Web? How to find the best parts—and avoid the others." *Consumer Reports.* May: 27-31.

Jones, David. 1998. "CleanNet continues to push for Burlington Public Library censorship." March 3. Email message from < djones@insight.dcss.McMaster.ca > to < efc-talk@insight.dcss.McMaster.ca >

Lawrence, Steve. 1998. "Millions of web pages overwhelm search engines." Email message from < thomaqui@vpl.vancouver.bc.ca > to < libs_all@vpl.vancouver.bc.ca > See also Steve Lawrence and C. Lee Giles, "Searching the World Wide Web." *Science.* 280 April 3, 1998: 98-100.

Markey, Karen. 1984. "Interindexer consistency tests: A literature review and report of a test of consistency in indexing visual materials." *Library and Information Science Research.* 6: 155-177.

Meeks, Brock N. and Declan B. McCullagh. 1996. "Keys to the kingdom." CyberWire Dispatch. July. Internet WWW page, at URL: < http://www.eff.org/pub/Publications/Declan_McCullagh/cwd.keys.to.the.kingdom.0796.article >

Milstead, Jessica. L. 1984. *Subject Access Systems: Alternatives in Design.* Orlando, Florida: Academic Press.

NetShepherd. 1997. "NetShepherd responds to the EPIC report 'Faulty Filters'." December 2. Internet WWW page, at URL: < http://www.netshepherd.com/fsEpicResponse.htm >

Newsletter on Intellectual Freedom. 1997. March: 29.

Oder, Norman. 1997. "Filtering and its contradictions." *Library Journal.* 122 May 1: 41-42.

Olson, Hope A. 1996. "The Power to Name: Marginalizations and Exclusions of Subject Representation in Library Catalogues." Ph.D. diss., University of Wisconsin-Madison.

Schneider, Karen G. 1997a. "Filters, homosexuality, responsibility, and so forth." Email message from < schneider.karen@epamail.epa.gov > to < web4lib@library.berkeley.edu > {The WWW for Libraries mailing list}.

Schneider, Karen G.. 1997b. *A Practical Guide to Internet Filters.* New York: Neal-Schuman. See also "The Internet Filter Assessment Project" at URL: < http://www.bluehighways.com/tifap >

Schneider, Karen G.. 1998. "Figuring out filters: A quick guide to help demystify them." *School Library Journal.* 44 February: 36-38.

Shallit, Jeffrey. 1997. "Query." January 23. Email message from < terminator@solidoak.com > to < Shallit@graceland.uwaterloo.ca >

SurfWatch. 1998. "SurfWatch filtering criteria." < http://2221.surfwatch.com/ filteringcriteria/index.html >

Thompson, Ken. 1998. "When in doubt...Filter! Filter! Filter!" *Social Responsibilities Round Table Newsletter.* March: 6-7.

Trinh, T. Minh-Ha. 1989. *Woman, Native, Other: Writing Postcoloniality and Feminism.* Bloomington: Indiana University.

Wallace, Jonathan D. 1997a. "Subject 1: Solid Oak blocking software and Ethical Spectacle." January 19. Email message from < jw@bway.net >

Wallace, Jonathan D.. 1997c. "Purchase of blocking software by public libraries is unconstitutional."

The Ethical Spectacle. November 9. Internet WWW page, at URL: < http:// www.spectacle.org/cs/library.html >