Narayanan A, Zevenbergen B. No Encore for Encore? Ethical questions for web-based censorship measurement. *Technology Science*. 2015121501. December 15, 2015. http://techscience.org/a/2015121501

No Encore for Encore? Ethical questions for web-based censorship measurement

Arvind Narayanan and Bendert Zevenbergen

Highlights

- A computer science research project in January 2015 executed code on the web browsers of unsuspecting users to detect censorship worldwide including in China and Iran
- This raises the ethical issue of should researchers be permitted to surreptitiously alter the behavior of Internet-connected devices in order to gain scientific data?
- We analyze this issue from the ethical, benefit-harm, consent, transparency, and legal perspective
- We find that Encore (1) does not yet fulfill the standard definition of having human subjects, (2) generates significant positive benefits with some potential harms that can be mitigated, (3) could have increased its transparency in affirming consent, and (4) does not violate any US laws though that may not be true for all the countries tested in the study

Issue	Question
Ethics	• Who are the stakeholders? • Are there human-subjects?
Benefit vs. Harm	• What are the benefits of the research? • What are the harms? • Can the harms be mitigated?
Transparency & Consent	• Do the researchers need to get informed consent from the users? • Are the researchers transparent enough?
Laws	• Did the researchers violate any laws in the US and internationally?

Key issues and questions about the Encore study analyzed in this paper

Abstract

A pair of computer scientists recently developed a clever way to measure Internet filtering and censorship worldwide, including countries such as China and Iran. Their system, named Encore, does this by executing a snippet of code on the web browsers of people who visit certain web pages— without the consent of those individuals. It caused a minor furor over research ethics in the computer networking and Internet measurement research communities.

Analysis summary: We analyze this conundrum through the lens of established ethical principles, while keeping in mind the peculiarities of Internet and big data research: its global reach, large scale, and automated nature. We also comment on the unusual model that computer scientists use for ethical oversight. We hope that the questions we raise will be useful for researchers facing similar dilemmas in their own work, as well as for students of research ethics, both in technical disciplines and in fields such as law and philosophy.

Introduction

```
<iframe src="//encore.noise.gatech.edu/task.html"
      width="0" height="0"
      style="display: none">
</iframe>
```

Anyone who administers a web page can copy-paste the above snippet into the source code of the page. It comes from the Encore project at the Network Operations and Internet Security Lab, now at Princeton and formerly at Georgia Tech. Its effect is to inject an invisible element into the page, which will then instruct the visitor's browser to download and execute a piece of code [1]. The code in question performs *censorship measurement*: it further instructs the visitor's browser to access content from one of various potentially filtered websites—again invisibly—and report back to the research team's server whether the access attempt was successful. By aggregating data from visitors to websites that deploy this measurement code snippet and inferring these visitors' locations based on their IP addresses, researchers can obtain an accurate and up-to-date view into web filtering worldwide.

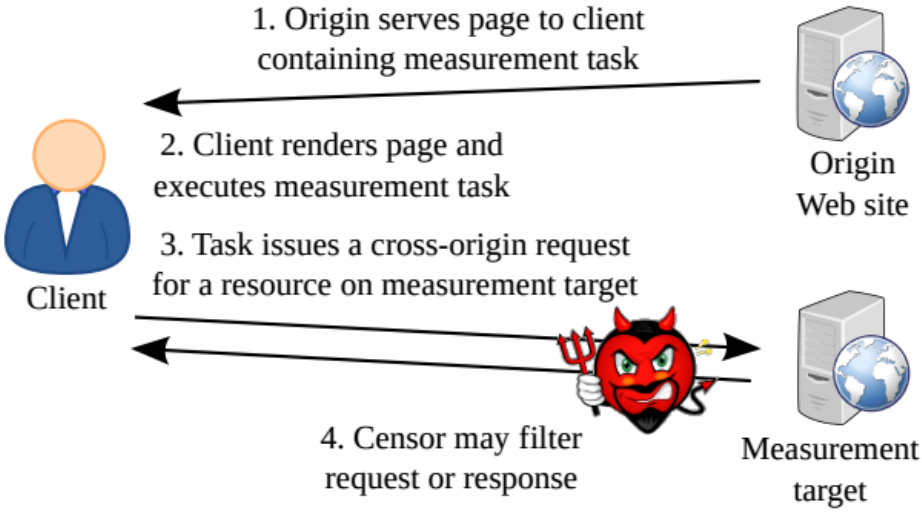

Figure 1. How Encore works. Reproduced from Burnett and Feamster's paper [2].

The researchers, Sam Burnett and Nick Feamster, used this technique to conduct measurements for a period of seven months, as of January 2015, via installations by at least 17 volunteers. They recorded measurements from 88,260 distinct IP addresses in 170 countries, with China, India, the United Kingdom, and Brazil each reporting at least 1,000 measurements, and more than 100 measurements from Egypt, South Korea, Iran, Pakistan, Turkey, and Saudi Arabia.

Encore revealed valuable information about the censorship activities of these governments, but it did so by altering the behavior of computers in ways that users were probably not anticipating and had not consented to. Encore is thus one example of a growing ethical conundrum for computer science and computer security researchers: should researchers be permitted to surreptitiously alter the behavior of Internet-connected devices in order to gain scientific data about the behavior of users and networks? If they should be allowed to in at least some cases, what are the criteria for determining proper and improper uses of these techniques, and who should enforce such standards? Encore also throws into sharp relief the conflict between two objectives: building automated, large-scale, globally applicable measurement tools and carefully analyzing ethical issues with consideration to all relevant stakeholders, laws, norms, and social, cultural, and political contexts.

Background

Technical Background

The architecture of the Internet, and the web in particular, affords a variety of ways of observing the behavior of networked devices on a large scale without the cooperation of users. Indeed, the multi-billion-dollar online ad targeting industry is built on this idea. Invisible "third parties" track our devices as we browse the web to build profiles of our

interests and behavior; the average top-50 website contains 64 tracking mechanisms [3]. Meanwhile, analytics firms track people in physical spaces such as shopping malls based on the WiFi, cellular, and other emanations from their smartphones [4].

Computer science researchers also make creative use of methods that allow observing devices without affirmative user consent. These studies have led to insights on the state of computer security, the economics of online advertising and of spam campaigns, Internet censorship and filtering around the world, and more.

The most intrusive of these studies, technologically speaking, are those that exploit unpatched security flaws to turn users' devices into observation points. A well-known one is *Spamalytics,* a study where researchers took over control of a botnet—a network of machines infected with malware and controlled by a single operator—to modify and study the spam campaigns that originated from the infected machines [5]. In another instance, an anonymous researcher or researchers *created* a botnet named Carna by infiltrating more than 400,000 routers and other devices whose default passwords hadn't been changed, and used the botnet to study essentially the entirety of Internet-connected devices [6, 7].

Other studies are non-intrusive: they simply eavesdrop on network traffic without interfering with devices. In computer networking, analyzing traffic data for improving performance and testing new protocols is standard practice and arguably essential. Such studies typically make use of data provided by Internet Service Providers (ISPs) that can be staggeringly large in size. For example, a 2014 study of IPv6 adoption utilized (among others) a dataset of traffic statistics that covered an estimated 33–50% of all Internet traffic for 2013 [8]. This type of research generally looks at network traffic in the aggregate rather than the behavior of individual users or devices, but there are exceptions. One study used the traffic metadata of millions of users, including a campus network, to study the economics of online advertising [9]. The study did this by inferring the information that advertisers collected about individuals, as well as how expensive the ads shown were, and then analyzing the relationship between the two.

Peer-to-peer networks are particularly amenable to non-intrusive study. Since such networks route information among peers rather than to and from designated servers, a researcher can simply set up one or more peers and hop on board, without needing any special privileges such as cooperation from an ISP. Researchers have used this method to study the BitTorrent file-sharing system, the Tor anonymity network, and the Bitcoin cryptocurrency network [10, 11, 12].

A burgeoning category of research lies in between these two in terms of intrusiveness: methods that use active probing of devices in some way but not exploitation of any security holes. These techniques are both technically and ethically fascinating, and include the Encore study.

An archetypal example of active probing is *network scanning* for security assessment of networks [13]. The ZMap research tool allows performing fast scans on an Internet-wide scale [14]. Network scanning has a long history, but a variety of new techniques are stealthier. *Idle port scanning* uses "side-channel attacks" to bounce traffic off an Internet-connected device in order to make measurements of other devices [15]. These side-channel attacks are different from exploitation of security bugs. The researcher doesn't take control over devices in any way; the bugs that allow side channels are common to many different implementations, and may be inherent to the protocol specification.

Encore similarly makes use of unintended effects inherent in the architecture of the web rather than a bug in any specific browser. The "same-origin policy" used in programming web browsers seeks to quarantine content from different domains even when loaded side-by-side on the same page, but there are limits to the effectiveness of this protection. Recent research at Princeton created an interesting twist to Encore's research methodology, showing how to deploy measurements through online advertisements [16]. The researcher simply purchases ad impressions—available cheaply by the thousands—and delivers the measurement code as part of the ad. This technique allows targeting by geography and demographics and can reach any user without relying on deployment by a website reachable by the user.

Encore is part of the small but growing research area of censorship measurement, an area that sees a handful of significant publications each year. In the United States, funding for censorship measurement comes from the National Science Foundation, indirectly from the State Department through its funding of censorship circumvention research, and from a few companies and philanthropic organizations.

The most basic objective of censorship measurement is compiling data on what is censored or filtered, when, and for which users. A prominent example is the Harvard Berkman Center's Herdict project, which aims to crowd-source and aggregate data about web filtering [17]. The Tor project's Open Observatory of Network Interference (OONI) has a similar aim; it provides a downloadable script that users can run [18]. Such data collection is sometimes straightforward but can require technical innovation and research. Encore, of course, is one example; another is ConceptDopplr, a tool that incorporates a way to efficiently probe a keyword-based blocking system to discover the set of all blacklisted keywords [19].

Another objective of censorship measurement is understanding the technical mechanisms by which censorship operates. Here are some questions on which researchers have been able to shed light: Do governments operate filters in a centralized way at Internet routers at the nation's borders, or in a decentralized way closer to the users [20]? How quickly are censors able to remove content from microblogging sites [21]? Does censorship operate purely by blocking or removal of content, or are performance degradation and modification of content also part of the picture [22]? Do censors have the technical infrastructure to examine the entire contents of Internet traffic ("deep packet inspection") without slowing it down, or do they only look at the metadata [23]? What types of collateral damage does censorship cause

[24]? So far, the bulk of the computer science research on censorship measurement addresses this class of questions.

Moving from the realm of computers and networks to the realm of people and governments, political scientists are interested in Internet censorship in terms of the motives of censors, the impact of censorship on freedom of speech, and so on [25]. Measurement directly or indirectly helps answer these questions. In 2013 Harvard researchers analyzed millions of social media posts to show that censorship in China allows government criticism but silences collective expression [26].

For obvious reasons, it's hard to measure censorship from a vantage point outside the country or countries of interest. There are some interesting and important exceptions. When censorship of online posts happens after the publication of posts *and* if researchers can obtain the content before the censors do, measurement can happen from the outside [26]. In another instance, a hacktivist group leaked 600 gigabytes of log files of Internet filtering devices used in Syria, allowing researchers to gain insights into censorship in that country [27]. Similarly, an anonymous ISP in Pakistan provided researchers access to a trove of data that enabled analysis of Pakistani censorship [28].

Outside such exceptions, collecting data about censorship requires the participation of volunteers "on the ground"—volunteers who might expose themselves to some risk and whose numbers limit the scale of measurements. Encore straddles these two categories: it avoids the need for researchers to recruit volunteers and is easily scalable, but the individuals whose devices are used face potential risks. Encore also provides a geographically fine-grained view of measurement, which researchers in the field value [29]. Further, since Encore turns regular Internet users into measurement vantage points, it avoids the problem of censors being able to detect and disable measurement units.

Ethical Oversight by Program Committees

Which of the research projects we've looked at should be considered human-subjects research? Questions of this sort have long been contentious in computer science. Human-subjects research at institutions that receive federal funding in the United States is subject to Institutional Review Board (IRB) oversight. IRBs approve research proposals based on investigators' efforts to account for and mitigate risk to participants. Typically research that poses little risk to individual human subjects is categorized as "exempt" from extensive oversight. In practice, however, much of such computer science research operates without IRB involvement. Historically computer science and engineering considers itself to be researching human-less systems, and university IRBs are typically geared toward regulating biomedical and social science research. When IRBs encounter computer science research, there is often mutual confusion.

Acknowledging that there are ethical and legal questions regardless of whether their activities involve human subjects, researchers have sought an alternative way to ensure that

published research is justifiable on scientific ethics grounds. Several sub-communities, including computer security, networking, and Internet measurement—which collectively encompass all of the research described above—appear to be converging on conference program committees as the oversight mechanism. What's a program committee? The most prestigious research in computer science is published in the proceedings of conferences rather than journals; each iteration of each conference selects a program committee to carry out peer review.

The appropriateness of ethical gatekeeping by program committees is a topic of continual debate in the community [30]. Supporters of this model argue that technical domain expertise is critical for ethical review and that each subcommunity must evolve its own norms by adapting ethical principles to the specific domain. The program committee process might help subcommunities evolve those norms because of the flux of members among committees, as opposed to IRBs that operate in relative independence from each other.

On the other hand, the system has numerous shortcomings, some inherent and others potentially fixable. First, the review happens after the research is complete. Unlike IRBs, there is no process for advance or continual review. The uncertainty induced by the potential rejection of research by program committees might lead researchers to abandon some research ideas or entire areas of research—especially research that pursues methodological innovation—even if, on balance, the research would have been found to be ethically acceptable had it proceeded. In cases where the putative harm arises from *conducting* the research rather than its publication, the retrospective ethical review fails to prevent that harm.

Second, program committee members are domain experts and rarely include any members with scholarly expertise in research ethics or ethics in general. Third, since they are formed and disbanded for each conference, they lack institutional memory—whether about specific research projects or about decision-making critieria and procedures. So far, they have operated without consistency in ethical standards and with ad hoc decision-making processes. Indeed, to our knowledge, among the conferences where the papers referenced above appeared, not a single one has published rules or guidelines for what qualifies as ethical research as part of the call for papers!

The Encore paper was submitted to ACM SIGCOMM 2015, a prestigious conference on computer networking. After heated debate, the committee accepted the paper for publication, but with a "signing statement" at the top of the paper, an unprecedented move [31].

The committee's ethical objections stemmed from several arguments, outlined in the public review of the paper [32]. First, third-party requests used for ad tracking, the committee held, at least notionally reflect the user's intent, whereas Encore's requests do not. Second, users downloading censored URLs might face repercussions if they live in a regime without due

process. Third, the committee believed that most users for whom censorship is an issue would be unlikely to consent to Encore's measurements.

Methods

Several analytical frameworks for scientific ethics and regulation have been created by U.S. government commmissions. For example, the Belmont Report [33] concerns scientific and medical research involving human subjects. This report, issued by the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research in 1979, established "respect for persons," "beneficence" and "justice" as the guiding principles of research ethics. The Common Rule [34] is the federal regulation that tasks IRBs with reviewing research to ensure it meets those principles. The 2012 Menlo Report [35] builds on the Belmont Report and translates the scientific ethics research principles into the computer science and network engineering domain. There are many other frameworks and guidelines for this type of research [36]. In our analysis here we'll roughly follow the Menlo report in terms of structure and the set of principles used. Here are the ways we examine the Encore project.

First, we provide an ethical inspection. Ethics guidelines typically recommend something akin to the following from the Menlo report: "it is first necessary to perform a systematic and comprehensive stakeholder analysis." In this study, this inspection raises questions of who the stakeholders are and whether Encore is human-subjects research.

Second, we provide a beneficence analysis. The principle of *beneficence* concerns the goal of the welfare of research participants and the balancing of probable harms. Ideally, the principle requires a systematic identification of the probability and magnitude of risk as well as benefits for the stakeholders, a subsequent iterative analysis of minimizing risk and maximizing benefits through the research design, and finally plans to mitigate identified risks and any unforeseen harms that materialize. In this study, this analysis raises questions about identifying Encore's potential benefits and harms, identifying the benefits of the research, minimizing risk of harm, and mitigating harm.

An important aspect of the Belmont report is respect for persons, law, and the public interest. We examine these principles in terms of informed consent, transparency, and accountability, as described below.

Third, we review informed consent in the context of Encore. Research ethics requires treatment of individuals involved in a study as autonomous persons. In traditional human subjects research, the investigator (ideally) approaches participants before the data collection begins, explains the research, and seeks their consent. Seeking informed consent is not always feasible, as is frequently the case in network measurement research, and certain proxies for consent are sometimes deemed appropriate by the research community. Researchers can seek consent from a representative authority while debriefing research

subjects after completion of data collection, or an IRB can, prior to the research, waive informed consent requirements completely. And according to the Common Rule, in cases where the researcher does not intervene in the life of an individual person to gather data, and there is no reasonable expectation of privacy, no form of consent is required. For example, anonymous observations of public activities, textual research, and examination of public records and other publicly accessible databases (even if access requires payment) are forms of research considered exempt from consent, even though they may reveal sensitive data about persons.

Fourth, we review transparency and accountability. Research ethics guidelines typically stress the importance of transparency of projects to serve the principles of accountability and meaningful informed consent. Additionally, guidelines recommend: "Debriefing is typically required when deception is used in order to mitigate harm resulting from loss of trust in researchers by those subjects who were deceived."

Finally, we assess legal compliance. Laws and policies regarding censorship and accessing unlawful or undesirable content on the Internet vary widely across jurisdictions; sometimes they may not be codified into law or may be subject to interpretation by political officials.

Analysis

Who are the stakeholders?

Any Internet user worldwide can stumble upon the invisible Encore script and carry out a censorship measurement. When an unsuspecting Internet user's browser sends a request to a potentially censored website, as instructed by the Encore code, the user's IP address may be recorded by the server hosting that website, as well as by many intermediaries and potentially unknown third parties. Most significantly, government-mandated censorship systems may also record and try to identify persons who access a censored website, although this is an assumption and may vary significantly by country. The Encore research team also records such measurements, which include the user's IP address.

Trying to identify the stakeholders immediately reveals a conflict. As mentioned above, ethics guidelines typically recommend something akin to the following from the Menlo report: "it is first necessary to perform a systematic and comprehensive stakeholder analysis." Yet the worldwide scale of Encore means that analyzing all potential stakeholders individually is infeasible. Worse, the principle is inherently at odds with the goal of *scalability* in computer science and engineering. Scalability is a goal that the Encore authors emphasize; in this context, it means that the team can expand the set of measurement targets by simply adding more machine resources and without multiplying the researcher effort required [37].

The dogma of computer science (and the technology industry) enshrines scalability as a virtue. Web companies regularly boast about their ratio of users to engineers, which can be over a million to one [38]. Similarly, all else being equal, a research project that scales better

is considered superior. In contrast, in fields where research involves experimenting on people, researchers aim to *minimize* the number of subjects necessary to measure a given effect with statistical rigor. When research that uses automated methods affects people, even if indirectly, we see a clash between these two paradigms.

The Menlo report briefly acknowledges the issue, stating: "Even a simple link traffic characterization study could involve millions of computers used by humans who are not themselves the direct subjects of research." This tension is a theme to which we will repeatedly return.

Is Encore human-subjects research?

Unsuspecting Internet users across the globe generate research data for the Encore project. Does the reliance on these humans mean that the Encore project constitutes human-subjects research in the traditional sense, analogous to fields such as medical research or psychology? Although networking researchers typically see themselves as conducting research on technical systems, the Internet is more properly understood as a sociotechnical system in which humans and technology interact. Experiments on the Internet will likely also include data collection about the behavior of humans, or affect their environment.

Neither the Princeton nor the Georgia Tech IRB considered Encore to be human-subjects research. Under the operational definition from the Common Rule that IRBs use, a human subject is a living individual about whom an investigator obtains "(1) Data through intervention or interaction with the individual, or (2) Identifiable private information." Should Encore's collection of IP addresses classify it as human-subjects research?

The question of whether or not IP addresses constitute personally identifiable information (PII) is a well-worn debate [39]. Buchanan et al. note:

"The Office for Human Research Protections has not issued a formal statement on whether IP addresses are considered to be personally identifiable information for purposes of the HHS protection of human subjects regulations at 45 CFR Part 46. However, for purposes of the HIPAA Privacy Rule, the HHS Office for Civil Rights has opined that an IP address is considered to be a direct identifier of an individual. Other European data regulations consider IPs as identifiers, and as such fall under the realm of the EU Data Directives (1995, 2006). This presents a challenge for international research and should be considered carefully by researchers and boards [40]."

Can researchers design Encore's data collection to generalize collected IP addresses so that they are no longer personally identifying, yet still carry out their measurement and analysis objectives? This is an open question.

Under a narrow interpretation, the data that Encore collects is not about the individual but rather about the behavior of censorship systems. On the other hand, the definition of PII

stems from medical and behavioral research, and probably did not anticipate the investigator's actions causing *other* parties—in Encore's case, the censor—to collect data about the individual. The Menlo Report advises the investigator to "respect individuals who are not targets of research yet are impacted" and says that "human subject research should now be considered as 'human-harming research'—so the internet users may not be subjects per se, but they can still experience harm due to the research being conducted."

Garfinkel is a proponent of the view that much of computer security research should be viewed as human-subjects research [41]. He proposes what he calls the human test: "would the experiment be useful if the data were generated by a random process and not by a human?" It is not obvious how to apply this test to Encore. If it were practically possible—which, unfortunately, it is not—to replace the humans whose devices Encore uses for measurement by robots that visit websites in a random fashion, Encore would work very well, and in some ways better than it currently does since biases in measurement times and so on would be minimized.

Identifying Potential Benefits and Harms

Given Encore's global scale, it will be tough for a small research group to adhere fully to the requirements of the beneficence principle. For example, before risks and harms can be identified, they must first be defined. However, due to the complex, dynamic, and innovative nature of the Internet, it is difficult to concretely define the harms for each Internet user, or even for regional groups of Internet users. The norms and attitudes of identified stakeholders with regard to accessing censored content differ greatly around the world, along with the type of censored content or possible enforcement actions. These are influenced by political, religious, historical and other social factors and are difficult—if not impossible—to quantify into a solid assessment of risks for each individual user.

Benefits of the research

Internet censorship measurement researchers argue that "whilst filtering and censorship can, to an extent, be open and transparent, their nature tends towards secrecy [42]." Measurement helps illuminate censorship—both its motivations and the technologies behind it. Understanding the motivations behind censorship yields valuable insights in political science, such as the fact that censorship in China allows government criticism but silences collective expression that may spur collective action [43]. Illuminating censorship techniques enhances the ability to create effective censorship circumvention tools [44].

A view of Internet censorship as harmful to citizens subjected to it is implicit in much of censorship measurement research. Many see censorship as violating human rights—the freedom of speech and more specifically the freedom to seek, receive, and impart information [45].

Computer scientists and engineers also have technical concerns. Architecting networks to allow censorship and filtering by governments and intermediaries violates the "end-to-end" principle, a key design philosophy of the Internet. As far back as the year 2000, Clark and Blumenthal warned that as the end-to-end design erodes, the "Internet might lose some of its key features, in particular its ability to support new and unanticipated applications [46]." Internet engineers also raised this concern, among others, in response to the Stop Online Piracy Act (SOPA) and PROTECT IP Act (PIPA) bills in the United States, which proposed to allow the government to block copyright-infringing websites: "Censorship of Internet infrastructure will inevitably cause network errors and security problems. This is true in China, Iran and other countries that censor the network today; it will be just as true of American censorship [47]."

An unequivocally negative view of censorship is not universal. Bambauer argues that "widespread censorship on-line is not necessarily bad" and that the legitimacy of censorship should be assessed not by what is blocked but rather the transparency and accountability of decision-making regarding censorship [48]. Chu and Cheng view the "Western" lens of individual autonomy and equality as inappropriate for Chinese society and evaluate Chinese online censorship from the perspective of "raising the moral level of both the state and society [49]."

At any rate, scholars have raised critical questions about data science, especially "big data," [50] that apply to censorship measurement as well. For example, since some types of censorship are much easier to detect than others, measurement results may produce a biased picture of the state of censorship and the direction of its movement. Moreover, stripped of cultural and political context, data is hard to interpret. For instance, it may be easy to find correlations between Internet filtering and news events, but causal attribution is far trickier. Similar concerns have been raised in the field of international development [51]. Censorship measurement researchers should be aware of this debate.

To conclude the discussion of benefits, let us recall the principle of justice, which entails that burdens as well as benefits be fairly and equitably distributed. For example, participants in a study who run the risk of harm should also benefit in the longer run from the research findings. From the Menlo report: "Each person deserves equal consideration in how to be treated, and the benefits of research should be fairly distributed according to individual need, effort, societal contribution, and merit."

Harm: does Encore present more than minimal risk?

The Encore paper itself considers harm primarily in terms of a comparison between Encore usage and regular web browsing. The Common Rule captures this type of comparison in the notion of *minimal risk* [52]: "minimal risk means that the probability and magnitude of harm or discomfort anticipated in the research is not greater in and of itself than those

encountered during daily life or during the performance of routine physical and psychological examinations or tests [53]."

The authors argue that normal web browsing exposes users to the same risks that Encore does, saying "the prevalence of malware and third-party trackers itself lends credibility to the argument that a user cannot reasonably control the traffic that their devices send" and "laws against accessing filtered content vary from country to country, and may be effectively unenforceable given the ease with which sites (like Encore) can request cross-origin resources without consent."

It is true that the average web user today is not in a position to effectively control third-party requests that their browser makes. Tracking technologies often go to great lengths to be stealthy, and publishers are often oblivious to the tracking technologies deployed on their properties [54]. Furthermore, online trackers make requests to yet other third parties, just as Encore does. In fact, these "chains" of trackers can be half-a-dozen (or more) deep [55]. While trackers may not necessarily make requests to censored domains, they beget other risks such as exposing users to surveillance agencies that monitor Internet traffic [56, 57]. Finally, advertisements themselves—and not just advertising networks—can make cross-origin requests to arbitrary domains. The bar to serving ads is much lower than the bar to becoming an advertising network.

However, there are several caveats to this argument and nuances that we should note. First, current online tracking practices are deeply at odds with users' expectations [58, 59, 60]. According to Nissenbaum's theory of contextual integrity, "what people care most about is not simply restricting the flow of information but ensuring that it flows appropriately [61]." According to the Association of Internet Researchers (AoIR)'s ethical guidelines, researchers must ask "What are the ethical expectations users attach to the venue in which they are interacting, particularly around issues of privacy [62]?" Arguably, both Encore and much third-party tracking today equally flout these expectations. In such an environment, there is a risk of an "ethical race to the bottom." Credentialed researchers and respected academic organizations arguably should not participate in and facilitate a race to the bottom even if advertisers feel obliged to do so—their tools may be similar, but their ethical obligations need not be.

Second, the probability and magnitude of harm may depend on the type of censored website. For social media sites such as Facebook, Twitter, and YouTube, which are frequently censored, an Encore measurement might not stand out in any way, since widgets from these websites (such as Facebook's Like button) are encountered extremely frequently in regular web browsing. The status is less clear with other websites such as news sites, also frequent targets of censorship. The nature and magnitude of harm may also depend on the reason the website was censored. A pattern of repeated access to specific religious websites deemed sensitive and censored will likely be viewed differently from accesses to Facebook or Twitter. It is difficult to generalize about the feasibility of enforcing laws across different regimes with

different technological capabilities, real-world enforcement resources, and, more fundamentally, different levels of respect for the rule of law. There is little information available on the likelihood and severity of persecution for simply accessing (or attempting to access) blocked domains, although of course citizens of many countries face such risks for online *writing* [63]. Tor Project leader Roger Dingledine notes that there is "little reprisal against passive consumers of information [64]." On the other hand, we know that the NSA monitored visits to pornography websites as part of a plan to "discredit radicalizers." [65]

Third, the focus on harm to individuals doesn't account for other types of harms that might result. For example, the Encore authors argue that "more widespread measurements like Encore become, the less risky they are for users" by making cross-origin requests to censored domains a commonplace occurrence. On the other hand, the censors might conceivably respond by shutting down Internet connectivity altogether.

Mitigating Harm

The Encore researchers limited the set of URLs that the script induced users to measure. All such URLs came from the list that Herdict asks its users to test. The current version of Encore tests only Twitter, Facebook, and YouTube, the rationale being that these domains are accessed regularly and automatically by most users' web browsers in the course of normal web browsing. In the section on informed consent, transparency and accountability below, we discuss other (actually used as well as potential) methods to mitigate harm.

The need for harm mitigation should inform the research design process, especially in research areas where established norms don't exist. From the AoIR guidelines:

"Ethical decision-making is a deliberative process, and researchers should consult as many people and resources as possible in this process, including fellow researchers, people participating in or familiar with contexts/sites being studied, research review boards, ethics guidelines, published scholarship (within one's discipline but also in other disciplines), and, where applicable, legal precedent."

The document further provides several questions for investigators to consider:

"How are the concepts of "vulnerability" and "harm" being defined and operationalized in the study? How are risks to the community/author/participant being assessed? How is vulnerability determined in contexts where this categorization may not be apparent? Would a mismatch between researcher and community/participant/author definitions of "harm" or "vulnerability" create an ethical dilemma? If so, how would this be addressed?"

Informed consent, transparency and accountability

The SIGCOMM program committee reviewing the Encore paper stated that the main ethical concern with the research would be mitigated if those who deployed Encore obtained

informed consent from users. The authors argue against both the feasibility and desirability of obtaining informed consent and provide three related arguments. First, they say that it would be impractical since it would require teaching users "nuanced technical concepts... across language barriers" that would "dramatically reduce the scale and scope of measurements." The challenges of communicating the necessary technical information to a global set of participants again highlights the tension between the scalability imperative and established ethical norms. Second, they argue that Encore with informed consent would be essentially equivalent to existing alternatives (such as, presumably, Herdict), forfeiting the benefits of the novel measurement architecture. Third, they say that informed consent may even increase risk to users by removing plausible deniability. However, we must consider that if there is no rule of law or guarantee of a fair trial with an independent judiciary in the censoring country, plausible deniability may not be enough to protect a user.

In terms of transparency, the Encore website contains a statement at the bottom "Visitors of this page have performed XXX measurements of Web filtering" and provides links to a page with additional information and the ability to opt out of future participation. Encore is meant to be deployed by other website operators; accordingly, the FAQ contains the question "Do I need to inform my site's visitors about Encore?" whose response begins:

"Although we cannot provide legal advice, we believe that you are not *required* to inform your site's visitors about Encore or obtain their consent before collecting measurements. That said, Encore's installation instructions explain how your [sic] can inform your visitors of Encore's presence and allow them to disable Encore entirely." (emphasis in original)

The focus appears to be on legal compliance over ethical obligation.

There are several possibilities for strengthening notice. The notice and opt-out link provided to users could be more prominent, perhaps in the style of the EU "cookie law" notices. Encore could require website operators who deploy it to give the same type of notice. Encore's FAQ could be expanded to include an explanation of the risks and benefits of the research as well as the technical concepts necessary to fully understand these.

Legal compliance

Although the Encore team is based in the U.S., the measurement actions occur in browsers of Internet users worldwide, which makes the issue of jurisdiction unclear.

In terms of compliance with United States computer law, Encore appears to be in the clear. U.S. cybersecurity law expert Jonathan Mayer notes that "While the scope of computer abuse law remains deeply unsettled, courts have converged on two baseline principles. First, circumventing a security protection on a remote system is illegal [66]. Second, when a system's owner explicitly revokes permission, remote access must cease [67]." He argues that both Encore and the ZMap tool discussed in Section 2 "unambiguously abide by these guidelines" since "both measurement approaches take advantage of known, intentional

software functionality." He continues, "As for respecting system owner preferences, the main deployment of both platforms is accompanied by a straightforward opt-out mechanism. If a system's owner revokes permission, research data collection immediately terminates [68]."

A global study of Internet censorship law and policies—as well as other applicable bodies of law such as privacy and data protection law—would be a near-impossible task for a legal researcher, let alone a team of computer scientists. Enumerating all possible (albeit remote) legal risks to Encore users is similarly infeasible. For example, the Falun Gong organization is banned in China; perhaps visits to a Falun Gong website may be interpreted as support for their cause [69]. Or perhaps Encore measurements are interpreted to constitute an act of espionage by helping a foreign power to map the national filter.

Since a thorough worldwide legal study is infeasible, Encore researchers cannot be certain that the measurements they induce do not constitute a violation of any local law. Ethicists would advise not putting people in a position where they could be perceived to have broken a law. In exceptional cases, however, researchers could develop an ethical justification that a law (or a type of law) is not in the public interest. The researchers must then demonstrate that they accept responsibility for their actions and the consequences, and have the necessary mitigation strategies in place [35].

Discussion

In conclusion, Encore makes for a fascinating case study that presents a thick web of considerations and no easy answers. While the scale of today's Internet and datasets is giddying to researchers and companies alike, the ethical responsibility that comes with it is rather sobering. Our analysis reveals a complex interplay between the technical design of the experiment and its potential risks and benefits.

As of this writing, Encore is very recent work, and there is an ongoing debate about its ethics, the broader question of norms for ethical research in network measurement, computer security, data science, and other disciplines, as well as the meta-question of how these disciplines should exercise ethical gatekeeping. We invite you to join the conversation.

References

1. The code is located at http://encore.noise.gatech.edu/task.html and can be viewed using your web browser's source code viewing feature.

2. Burnett S, Feamster N. Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests. SIGCOMM '15. August 17, 2015. http://conferences.sigcomm.org/sigcomm/2015/pdf/papers/p653.pdf

3. Angwin J. The Web's New Gold Mine: Your Secrets. The Wall Street Journal, July 30, 2010.http://www.wsj.com/articles/SB10001424052748703940904575395073512989404

4. Euclid Analytics. Accessed August 11, 2015. http://euclidanalytics.com/

5. Kanich C, Kreibich C, Levchenko K, Enright B, Voelker G, Paxson V, Savage S. Spamalytics: An Empirical Analysis of Spam Marketing Conversion. CCS '08. October 27, 2008. http://www.icsi.berkeley.edu/pubs/networking/2008-ccs-spamalytics.pdf

6. Internet Census 2012: Port scanning /0 using insecure embedded devices. Carna Botnet. Accessed August 11, 2015. http://internetcensus2012.bitbucket.org/paper.html

7. Krenc T, Hohlfeld O, Feldmann A. An Internet Census Taken by an Illegal Botnet – A Qualitative Assessment of Published Measurements. ACM SIGCOMM Computer Communication Review. July 2014. http://www.net.t-labs.tu-berlin.de/papers/KHF-ICIBQ-14.pdf

8. Czyz J, Allman M, Zhang J, Iekel-Johnson S, Osterweil E, Bailey M. Measuring IPv6 Adoption. SIGCOMM '14. August 17, 2014. http://nsrg.eecs.umich.edu/publications/sigcomm14_ipv6.pdf

9. Gill P, Erramilli V, Chaintreau A, Krishnamurthy B, Papagiannaki D, Rodriguez P. Follow the Money: Understanding Economics of Online Aggregation and Advertising. IMC '13. October 23, 2013. http://conferences.sigcomm.org/imc/2013/papers/imc184s-gillAemb.pdf

10. Pouwelse J, Garbacki P, Epema D, Sips H. The Bittorrent P2P File-Sharing System: Measurements and Analysis. Department of Computer Science, Delft University of Technology, the Netherlands. Accessed August 11, 2015. http://www.cs.unibo.it/babaoglu/courses/cas04-05/papers/bittorrent.pdf

11. Winter P, Köwer R, Mulazzani M, Huber M, Schrittwieser S, Lindskog S, Weippl E. Spoiled Onions: Exposing Malicious Tor Exit Relays. Karlstad University, Sweden, SBA Research, Austria, and FH Campus Wien, Austria. Accessed August 11, 2015. http://www.cs.kau.se/philwint/spoiled_onions/pets2014.pdf.

12. Miller A, Litton J, Pachulski A, Gupta N, Levin D, Spring N, Bhattacharjee B. Discovering Bitcoin's Public Topology and Influential Nodes. University of Maryland, College Park, Accessed August 11, 2015. https://cs.umd.edu/projects/coinscope/coinscope.pdf

13. Nmap ("Network Mapper"). Accessed August 11, 2015. https://nmap.org/

14. Durumeric Z, Wustrow E, Halderman J. ZMap: Fast Internet-Wide Scanning and its Security Applications. Proceedings of the 22nd USENIX Security Symposium. August 2013. https://zmap.io/paper.pdf

15. Ensafi R, Park J, Kapur D, Crandall J. Idle Port Scanning and Non-interference Analysis of Network Protocol Stacks Using Model Checking. Department of Computer Science, University of Mexico, Accessed August 11, 2015. https://www.usenix.org/legacy/event/sec10/tech/full_papers/Ensafi.pdf

16. Zimmerman P. Measuring Privacy, Security, and Censorship Through the Utilization of Online Advertising Exchanges. A Thesis Presented to the Faculty of Princeton University in Candidacy for the Degree of Master of Science in Engineering. June 2015. ftp://ftp.cs.princeton.edu/techreports/2015/988.pdf

17. Herdict. Berkman Center for Internet & Society at Harvard University. Accessed August 11, 2015. https://www.herdict.org/

18. Filasto A, Appelbaum J. OONI : Open Observatory of Network Interference. The Tor Project and the University of Washington. Accessed August 11, 2015, https://www.usenix.org/system/files/conference/foci12/foci12-final12.pdf

19. Crandall J, Zinn D, Byrd M, Barr E, East R. ConceptDoppler: A Weather Tracker for Internet Censorship. 14th ACM Conference on Computer and Communications Security. October 29, 2007. https://www.cs.unm.edu/~crandall/concept_doppler_ccs07.pdf

20. Xu X, Mao Z, Halderman J. Internet Censorship in China: Where Does the Filtering Occur? Department of Computer Science and Engineering, University of Michigan. Accessed August 11, 2015. https://web.eecs.umich.edu/~zmao/Papers/china-censorship-pam11.pdf

21. Zhu T, Phipps P, Pridgen A, Crandall J, Wallach D. The Velocity of Censorship: High-Fidelity Detection of Microblog Post Deletions. 22nd USENIX Security Symposium. August 2013. http://arxiv.org/ftp/arxiv/papers/1303/1303.0597.pdf

22. Burnett S, Feamster N. Making Sense of Internet Censorship: A New Frontier for Internet Measurement. ACM SIGCOMM Computer Communication Review. July 2013. http://www.sigcomm.org/sites/default/files/ccr/papers/2013/July/2500098-2500111.pdf

23. Wilde T. Knock Knock Knockin' on Bridges' Doors. The Tor Project. January 7, 2012. https://blog.torproject.org/blog/knock-knock-knockin-bridges-doors

24. Dainotti A, Squarcella C, Aben E, Claffy K, Chiesa M, Russo M, Pescape A. Analysis of Country-wide Internet Outages Caused by Censorship. 2013 IEEE. Accessed August 11, 2015. http://www.caida.org/publications/papers/2014/outages_censorship/outages_censorship.pdf

25. Morozov E. The Net Delusion: The Dark Side of Internet Freedom. PublicAffairs. 2012. http://www.amazon.com/The-Net-Delusion-Internet-Freedom/dp/1610391063

26. King G, Pan J, Roberts M. How Censorship in China Allows Government Criticism but Silences Collective Expression. American Political Science Review. May 2013. http://gking.harvard.edu/files/censored.pdf

27. Chaabane A, Chen T, Cunche M, Cristofaro E, Friedman A, Kaafar M. Censorship in the Wild: Analyzing Internet Filtering in Syria. IMC '14. November 5, 2014. http://conferences.sigcomm.org/imc/2014/papers/p285.pdf

28. Khattak S, Javed M, Khayam S, Uzmi Z, Paxson V. A Look at the Consequences of Internet Censorship Through an ISP Lens. IMC '14. November 5, 2014. http://conferences.sigcomm.org/imc/2014/papers/p271.pdf

29. Wright J, de Souza T, Brown I. Fine-Grained Censorship Mapping Information Sources, Legality and Ethics. Oxford Internet Institute and Oxford University Computing Laboratory. Accessed August 11, 2015. http://static.usenix.org/event/foci11/tech/final_files/Wright.pdf

30. Kenneally E, Bailey M. Cyber-security Research Ethics Dialogue & Strategy Workshop. ACM SIGCOMM Computer Communication Review. April 2014. http://mdbailey.ece.illinois.edu/publications/ccr-2014.pdf

31. SIGCOMM's call for papers had a requirement (for the first time in 2015) that authors "attest that their work complies with all applicable ethical standards of their home institution(s), including ... policies on experiments involving humans". The actual ethical review went well beyond checking for such compliance.

32. Byers J. Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests – Public Review. Department of Computer Science, Boston University, Accessed August 11, 2015. http://conferences.sigcomm.org/sigcomm/2015/pdf/reviews/226pr.pdf

33. The Belmont Report. U.S. Department of Health and Human Services. April 18, 1979. http://www.hhs.gov/ohrp/humansubjects/guidance/belmont.html

34. Federal Policy for the Protection of Human Subjects ('Common Rule'). U.S. Department of Health and Human Services. Accessed August 11, 2015. http://www.hhs.gov/ohrp/humansubjects/commonrule/

35. The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. U.S. Department of Homeland Security Science and Technology Directorate, Cyber Security Division. August 2012. https://www.caida.org/publications/papers/2012/menlo_report_actual_formatted/menlo_report_actual_formatted.pdf

36. See, for example: Markham A, Buchanan E. Ethical Decision-Making and Internet Research: Recommendations from the AoIR Ethics Committee. approved by the AOIR general membership. December 2012. http://aoir.org/reports/ethics2.pdf ; and Zevenbergen B, Brown I, Wright J, Erdos D. Ethical Privacy Guidelines for Mobile Connectivity Measurements. Oxford Internet Institute, University of Oxford. November 2013. http://www.oii.ox.ac.uk/research/Ethical_Privacy_Guidelines_for_Mobile_Connectivity_Measurements.pdf

37. There are other, related, meanings of the word scalability. In particular, it can refer to the ability of a computer system to handle greater and greater loads by adding proportionally many computing units in parallel. This sense of the term is simply a sound engineering principle.

38. Bosworth A. Facebook Engineering Bootcamp. Facebook Engineering. November 19, 2009. https://www.facebook.com/notes/facebook-engineering/facebook-engineering-bootcamp/177577963919

39. McIntyre J. Balancing Expectations of Online Privacy: Why Internet Protocol (IP) Addresses Should Be Protected as Personally Identifiable Information. DePaul Law Review. Spring 2011. http://via.library.depaul.edu/cgi/viewcontent.cgi?article=1151&context=law-review

40. Buchanan E, Aycock J, Dexter S, Dittrich D, Hvizdak E. Computer Science Security Research and Human Subjects: Emerging Considerations for Research Ethics Boards. Journal of Empirical Research on Human Research Ethics: An International Journal. June 2011 http://www.orau.gov/communityirb/Resources/EmergingConsiderationsForResearchEthicsBoards.pdf

41. Garfinkel S. IRBs and Security Research: Myths, Facts and Mission Creep. Naval Postgraduate School & Harvard University. April 7, 2008. https://calhoun.nps.edu/bitstream/handle/10945/40330/garfinkel_IRBs_and_Security_Research.pdf?sequence=1

42. Wright J, de Souza T, Brown I. Fine-Grained Censorship Mapping Information Sources, Legality and Ethics. Oxford Internet Institute and Oxford University Computing Laboratory. Accessed August 11, 2015. http://static.usenix.org/event/foci11/tech/final_files/Wright.pdf

43. King G, Pan J, Roberts M. How Censorship in China Allows Government Criticism but Silences Collective Expression. American Political Science Review. 2013. http://core.ac.uk/download/pdf/28941335.pdf

44. The Encore authors say, “Researchers, activists, and citizens aim to understand what, where, when, and how governments and organizations implement Internet censorship. This knowledge can [...

45. La Rue F. Report of the Special Rapporteur on the promotion and protection of the right to freedom of opinion and expression. United Nations General Assembly, Human Rights Council, seventeenth session. May 16, 2011. http://www2.ohchr.org/english/bodies/hrcouncil/docs/17session/A.HRC.17.27_en.pdf

46. Clark D, Blumenthal M. Rethinking the design of the Internet: The end to end arguments vs. the brave new world. Version for TPRC submission. August 10, 2000. http://dspace.mit.edu/bitstream/handle/1721.1/1519/TPRC_Clark_Blumenthal.pdf?sequence=1&origin=publication_detail

47. Cerf V, et. al. “An Open Letter From Internet Engineers to the United States Congress,” December 15, 2011. https://www.eff.org/files/internet-engineers-letter.pdf

48. Bambauer D. Censorship V3.1. 18 IEEE Internet Computing 26, Arizona Legal Studies Discussion Paper No. 12-28. September 9, 2012. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2144004&download=yes

49. See Chapter 1, Cultural convulsions: examining the Chineseness of cyber China. of Herold D, Marolt P (eds.) Online Society in China: Creating, celebrating, and instrumentalizing the online carnival. Routledge, 2011. http://www.tandf.net/books/details/9780415838221/

50. boyd d, Crawford K. Critical Questions for Big Data. Information, Communication & Society. 2012. http://www.tandfonline.com/doi/pdf/10.1080/1369118x.2012.678878

51. Taylor L, Broeders D. In the name of Development: Power, profit and the datafication of the global South. Geoforum. August 2015 http://www.sciencedirect.com/science/article/pii/S0016718515001761

52. Proposed Revisions to the Common Rule for the Protection of Human Subjects in the Behavioral and Social Sciences. National Academy of Sciences. 2014 http://www.ncbi.nlm.nih.gov/books/NBK217976/

53. Code of Federal Regulations. U.S. Department of Health and Human Services. January 15, 2009. http://www.hhs.gov/ohrp/humansubjects/guidance/45cfr46.html

54. Angwin J. Meet the Online Tracking Device That is Virtually Impossible to Block. ProPublica. July 21, 2014. https://www.propublica.org/article/meet-the-online-tracking-device-that-is-virtually-impossible-to-block

55. Li Z, Zhang K, Xie Y, Yu F, Wang X. Knowing Your Enemy: Understanding and Detecting Malicious Web Advertising. CCS '12. October 16, 2012. http://dl.acm.org/citation.cfm?id=2382267

56. Kamdar A, Reitman R, Schoen S. NSA Turns Cookies (And More) Into Surveillance Beacons. Deeplinks. Electronic Frontier Foundation. December 11, 2013. Accessed August 11, 2015. https://www.eff.org/deeplinks/2013/12/nsa-turns-cookies-and-more-surveillance-beacons

57. Englehardt S, Reisman D, Eubank C, Zimmerman P, Mayer J, Narayanan A, Felten E. Cookies That Give You Away: the Surveillance Implications of Web Tracking. WWW 2015. May 18, 2015. http://www.www2015.it/documents/proceedings/proceedings/p289.pdf

58. Ur B, Leon P, Cranor L, Shay R, Wang Y. Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising. CyLab, Carnegie Mellon University. April 2, 2012. http://www.futureofprivacy.org/wp-content/uploads/Smart-Useful-Scary-Creepy.-Perceptions-of-Online-Behavioral-Advertising-.pdf

59. McDonald A, Cranor L. Americans' Attitudes About Internet Behavioral Advertising Practices. WPES '10. October 4, 2010. http://www.aleecia.com/authors-drafts/wpes-behav-AV.pdf

60. Turow J, King J, Hoofnagle C, Bleakley A, Hennessy M. Americans Reject Tailored Advertising and Three Activities That Enable It. Technology | Academics | Policy (TAP). September 2009. https://www.techpolicy.com/TechnologyAcademicsPolicy/media/document-library/Americans-Reject-Tailored-Advertising---And-Three-Activities-That-Enable-It---September-2009-Survey-for-the-Rose-Foundation.pdf

61. Nissenbaum H. Privacy in Context Technology, Policy, and the Integrity of Social Life. Stanford University Press. 2009. http://www.sup.org/books/title/?id=8862

62. Markham A, Buchanan E. Ethical Decision-Making and Internet Research: Recommendations from the AoIR Ethics Committee. approved by the AOIR general membership. December 2012. http://aoir.org/reports/ethics2.pdf

63. Freedom on the Net 2014. Freedom House. Accessed August 11, 2015. https://freedomhouse.org/report/freedom-net/freedom-net-2014#.VfWZz2TBzGc

64. Dingledine R. Tor and circumvention: lessons learned. The Tor Project. Accessed August 11, 2015. https://www.iacr.org/conferences/crypto2011/slides/Dingledine.pdf

65. Greenwald G, Grim R, Gallagher R. Top –Secret Document Reveals NSA Spied on Porn Habits As Part of Plan To Discredit 'Radicalizers'. Huffington Post. November 26, 2013. http://www.huffingtonpost.com/2013/11/26/nsa-porn-muslims_n_4346128.html

66. See, e.g., Facebook, Inc. v. Power Ventures, Inc., No. C 08-05780 JW, 2010 WL 3291750, at *5-12 (N.D. Cal. July 20, 2010).

67. See, e.g., Craigslist Inc. v. 3Taps Inc., No. CV 12-03816 CRB, 2013 WL 4447520, at *5 (N.D. Cal. Aug. 16, 2013).

68. Jonathan Mayer, personal communication.

69. The full list of URLs ever measured by Encore is listed at http://encore.noise.gatech.edu/urls.html It does not include any websites related to the Falun Gong organization

---

Authors

Arvind Narayanan is a computer science professor at Princeton. He advised a Master's thesis, described in Section 2, that utilized a similar methodology to the Encore project.

Bendert Zevenbergen is a Ph.D candidate and researcher at the Oxford Internet Institute, where he studies the intersection of law, ethics, social science, and the Internet. Along with a colleague at OII, he first brought certain ethical concerns to the Encore authors' attention, resulting in a significant change to the design. This case study is the result of a dialogue between us.

This case study was first written for the Council on Big Data, Ethics, and Society. Funding for this Council was provided by the National Science Foundation (#IIS-1413864). For more information on the Council, see: http://bdes.datasociety.net/

We are grateful for useful feedback from Nick Feamster, Jacob Metcalf, Matt Salganik, Stuart Schechter, Joss Wright, members of the BDES council, and anonymous reviewers.

Referring Editor: danah boyd

Citation

Narayanan A, Zevenbergen B. No Encore for Encore? Ethical questions for web-based censorship measurement. *Technology Science*. 2015121501. December 15, 2015. http://techscience.org/a/2015121501

Data

Under review for data sharing classification.