

Law and Technology

The Rise of the Robo Notice

Examining the conflicting claims involving the use of automated tools in copyright-related notice-and-takedown procedures.

MOST INTERNET PROFESSIONALS have some familiarity with the “notice and takedown” process created by the 1998 U.S. Digital Millennium Copyright Act (the DMCA). Notice and takedown was conceived to serve three purposes: it created a cheap and relatively fast process for resolving copyright claims against the users of online services (short of filing a lawsuit); it established steps online services could take to avoid liability as intermediaries in those disputes—the well-known DMCA “safe harbor”; and it provided some protection for free speech and fair use by users in the form of “counter notice” procedures.

The great virtue of the notice and takedown process for online services is its proceduralism. To take the most common example, if a service reliant on user-generated content follows the statutory procedures, acts on notices, and otherwise lacks specific knowledge of user infringement on its site (the complicated “red flag” knowledge standard), it can claim safe harbor protection in the event of a lawsuit. Services can make decisions about taking down material based on substantive review and their tolerance for risk. They may also adopt technologies or practices to supplement notice and takedown, though the law makes no such demands beyond a requirement for repeat infringer policies. The resulting balance has enabled a



relatively broad scope for innovation in search and user-generated-content services. As one entrepreneur put it in our recent study of these issues, notice and takedown was “written into the DNA” of the Internet sector.

This basic model held for about a decade. In the last five or six years,

however, the practice of notice and takedown has changed dramatically, driven by the adoption of automated notice-sending systems by rights holder groups responding to sophisticated infringing sites. As automated systems became common, the number of takedown requests increased dramatically.

For some online services, the numbers of complaints went from dozens or hundreds per year to hundreds of thousands or millions. In 2009, Google's search service received less than 100 takedown requests. In 2014, it received 345 million requests. Although Google is the extreme outlier, other services—especially those in the copyright 'hot zones' around search, storage, and social media—saw order-of-magnitude increases. Many others—through luck, obscurity, or low exposure to copyright conflicts—remained within the "DMCA Classic" world of low-volume notice and takedown.

This split in the application of the law undermined the rough industry consensus about what services did to keep their safe harbor protection. As automated notices overwhelmed small legal teams, targeted services lost the ability to fully vet the complaints they received. Because companies exposed themselves to high statutory penalties if they ignored valid complaints, the safest path afforded by the DMCA was to remove all targeted material. Some companies did so. Some responded by developing automated triage procedures that prioritized high-risk notices for human review (most commonly, those sent by individuals).

Others began to move beyond the statutory requirements in an effort to reach agreement with rights holder groups and, in some cases, to reassert some control over the copyright disputes on their services. These "DMCA+" measures included the use of content filtering systems, such as YouTube's ContentID and Audible Magic. They included special takedown privileges for "trusted" senders; expanded profiling of users; restrictive terms of service; modifications to search functionality; and other approaches that stretched beyond the DMCA's requirements.

The strength of these measures depended on a variety of factors, but most importantly the service's leverage with rights holders. Smaller services sometimes ceded control of takedown to rights holder groups, in some cases via direct back-end administrative access. Larger services with more leverage sought to retain control but, in some cases, adopted blocking or content filtering systems as part of

As automated systems became common, the number of takedown requests increased dramatically.

their risk assessments. Rights holder groups worked to turn incremental concessions into perceived industry norms that could filter back into law through litigation.

Today, for online services with significant copyright exposure (particularly search, cloud storage, music, and video services), the DMCA's procedural balance is, practically speaking, obsolete. The notice and takedown process has been widely supplanted by practices that shift the balance of protection from users to rights holders, and that hold out the statutory remedies as, at best, last resorts for user and services. For services outside the major copyright conflict zones, on the other hand, notices remain relatively infrequent and DMCA Classic process is still generally the rule, built around the substantive, human review of notices.

One major concern in our work is to understand the balance of power in this divided ecosystem. Does the adoption of DMCA+ measures by large players drive out DMCA Classic over time—forcing an ever-wider range of sites to manage floods of automated notices or adopt content filtering and other strategies? Some of the smaller services we spoke with (in interviews and surveys conducted in 2014) clearly expressed such concerns. From their perspective, traditional DMCA compliance appears fragile—dependent on rights holder forbearance or inattention that could change at any time.

A second important question is whether automated notices do, in fact, reliably target infringing material. As rights holders are quick to point out, automated systems overwhelmingly

target either file-sharing sites that can be plausibly described as "dedicated to infringement" (in the current enforcement lingo), or search engine results that lead to those sites. For rights holders, such targeting generally justifies the minor sins of misidentification associated with these systems, such as the targeting of dead links or pages. The likelihood of damage to important protected expression, in these cases, is relatively low. But the same practices are also brought to bear against a wider array of sites, including those with predominantly non-infringing uses and active notice and takedown policies. In these cases, the risks of collateral damage are higher.

The accuracy of automated notices has become a flashpoint in copyright enforcement debates, marked on one side by rights holder claims to reliability and procedural safeguards (such as human spot checks) and, on the other side, by a steady stream of anecdotal examples of bad takedowns. One goal of our work is to examine these conflicting claims. To this end, we undertook a close examination of a representative sample of the 108 million takedown requests sent to Google Search during a six-month period in 2013. Google, it is worth noting, is both a pioneer of automated triage and a major investor in human review. On average, it takes down 97.5% of requests. This percentage proved, in our surveys, to be on the low end for companies that receive large numbers of automated notices.

Our work suggests the situation is, predictably, complex. Nearly all of the 1,800 requests in our sample targeted file-sharing sites—in principle, the easy targets with low potential for non-infringing content. In practice, approximately 1 in 25 (4.2%) of these requests were fundamentally flawed—simply targeting the wrong content. Nearly one-third of requests (28.9%) had characteristics that raised concerns about the validity of the claim. This number included 7.3% targeting content that could make a reasonable claim to one or more fair use criteria and 10.6% that led to dynamic search results or aggregator pages that made identifying the targeted content difficult or impossible. (This form of identification has

COMMUNICATIONS APPS

Access the latest issue, past issues, **BLOG@CACM**, News, and more.



Available for iPad, iPhone, and Android



Available for iOS, Android, and Windows



Association for Computing Machinery

The accuracy of automated notices has become a flashpoint in copyright enforcement debates.

not fared well in recent court cases). Despite the resources Google invests in this process, almost none of these notices, statistically speaking, receive substantive examination. Although the practices are not directly comparable, several DMCA Classic services (which often do look closely at the targeted content) reported rejecting much larger proportions of notices, in some cases exceeding 50%.

Automation has transformed notice and takedown from a targeted remedy for specific acts of infringement into a large-scale strategy for limiting access to infringing content. At one level, this looks like a fruitless enterprise: an endless game of whack-a-mole against a resilient and multi-jurisdictional file-sharing community. Yet, as the big services grow bigger and are pushed beyond the DMCA, important parts of the user-generated-content Internet ecosystem shift toward rights holder control.

And so the process continues to escalate. This year may see well over one billion DMCA takedown requests, further marginalizing the procedural protections afforded by the DMCA. The adoption of content filtering by video and music services such as YouTube, SoundCloud, and Vimeo represents a profound shift in the direction of enforcement—no longer basing takedown on the rights holder's identification of unauthorized uses after the fact, but on the a priori filtering of user activity in collaboration with rights holders. For many DMCA Classic services, such filtering represents new and potentially unsustainable costs and a threat to the wide latitude for expression on which

their user communities are built. For the DMCA+ services, filtering is a response to a complex array of pressures, from staying on the right side of perceived enforcement norms, to exercising some control in keeping content up, to creating—in the case of ContentID—a new payment model for artists.

There is no easy way around these dilemmas. Stronger liability for reckless or malicious notice use might be a good step in curbing the worst notice practices, which can include deceptive or predatory behavior. But such changes are currently a dead letter in U.S. copyright politics. Nor is there appetite on the part of the Internet companies to revisit the core DMCA compromises, for fear of undermining the general principle of safe harbor protection. And so despite Congressional interest in copyright reform more generally, there would seem to be little chance of significant policy change in this area. The rights holder companies are slowly winning on enforcement and the largest Internet companies have become powerful enough to fend off changes in law that could threaten their core business models. The ability of large companies to bear the costs of DMCA+ systems, moreover, has become a source of competitive advantage, creating barriers to entry on their respective terrains. This will not be news to those watching the business consolidation of Web 2.0. It is bad news, however, for those who think both copyright and freedom of expression are best served by clear statutory protection and human judgment regarding their contexts and purposes. In important parts of the Internet sector, robots have taken those jobs and they are not very good at them. **C**

Joe Karaganis (joe.karaganis@gmail.com) is vice president at The American Assembly, a public policy institute based at Columbia University.

Jennifer Urban (jurban@law.berkeley.edu) is assistant clinical professor of law at the University of California, Berkeley, and director of Berkeley Law's Samuelson Law, Technology & Public Policy Clinic.

The authors, with Brianna Schofield, have produced a study of notice and takedown available at <http://takedownproject.org/>.

Copyright held by authors.