

ZDNet Undercover: The YouTube File

How good is Google at separating out unlawful uses of copyrighted uploads to its popular video-sharing site? Answer: Systematically better. But not perfect.

By Tom Steinert-Threlkeld

November 2008



YouTube video identification: The basics

YouTube

Business: Web-based hosting service which allows its users to upload, view and share video clips.

Problem: Separating out unlawful uses of copyrighted music and images contained in uploads.

Volume: Thirteen hours of video created by users uploaded every minute. Roughly 5 billion videos viewed by visitors to the site each month in United States.

Market share: Approximately 45 percent. Next biggest player has about one-tenth of YouTube's share.

Founded: February 2005

Headquarters: San Bruno, Calif.

Owned by: Google

Key players:

Eric Schmidt, CEO, Google

Engineered acquisition of YouTube when Google Video failed to get traction. Sought to stem illegal uploads to YouTube; Faces a \$1 billion lawsuit from Viacom. Aimed for system that could catch 80 percent or more of infringing uploads.

Chad Hurley, CEO & CoFounder, YouTube

Wants to make uploading a video "as simple as placing a phone call." Expects to narrow gap in quality between professionally produced and consumer produced video. Allowed Google engineers figure out how to identify videos that infringe on professional producers' rights.

David King, product manager, YouTube's video identification

Managed the launch of the Video ID system in Oct. 2007. Incorporated feedback from multiple content holders that tested the product, including Turner Broadcasting System and CBS Interactive (owner of ZDNet). Aims to automate all aspects of identifying infringing videos, removing them and letting original owners 'monetize' them, splitting ad revenue with YouTube.

Glenn Brown, product counsel, YouTube

Once said that identifying copyrighted content would be harder than identifying porn, according to a Valleywag report. Advocates that Internet companies couldn't grow if they didn't have a "safe harbor" from users' copyright infringements.

Key technologies:

YouTube's content identification and management system combines the following tools:

Audio fingerprinting. Source: Audible Magic, Los Gatos, Calif.

Video fingerprinting. Source: Google engineers, Mountain View, Calif.

Content policies include:

- Creating unique "hash" – or digital signature – of every video removed for copyright infringement;
- Blocking the repeat upload of removed videos, using the "hash";
- Terminating accounts of repeat infringers, based on notices filed under Digital Millennium Copyright Act;
- "Click of a mouse" notification of violations by content owners. ❖

YouTube's video ID system: Is 75 percent good enough?

By Tom Steinert-Threlkeld

Comedian Tina Fey returned shortly after Labor Day to her old haunt at Saturday Night Live and did a [sendup of Alaska Governor Sarah Palin](#) ("I can see Russia from my house.") as the new vice presidential candidate of the Republican Party. The skit was watched on Sept. 13, 2008 by more than 10 million TV viewers and by midnight had become the talk of households nationwide for what was seen as a dead-on portrayal of Palin. Of course, the clip was almost immediately uploaded to YouTube by its users.

And there, some versions of the clip stayed, eluding NBC, which produced and owned the skit, and YouTube, for nine more days. In its various postings, the skit gathered hundreds of thousands of views on YouTube that rightfully should have gone to NBC. Those views aren't much by broadcast standards. The same skit was viewed 7 million times on NBC.com.

Still, those views are notable, says Rick Cotton, general counsel for NBC Universal. NBC.com or Hulu.com, an NBC joint venture with News Corp., could have received those views. Hulu aims to monetize clips with ads from the likes of Splenda.

That something as visible as Fey's skit could slip through YouTube's monitoring system raises questions. How effective is YouTube's 13-month-old video identification process? The question looms as a \$1 billion suit filed by entertainment giant Viacom is expected to reach federal court.

NBC says that YouTube's system has "come on quite strongly" and snares about 75 percent to 80 percent of clips that are illegally uploaded. That tally is "improving month by month," according to Cotton. Indeed, two weeks after the first Fey appearance as Palin, a faux interview of her character with Katie Couric was blocked.

How YouTube improved its system over those last two weeks is unclear, but the process isn't foolproof by any stretch and could take years to perfect. YouTube declines to talk about the mechanics of its system or make any claim as to how effective its system is. That silence gives detractors such as Miles McNamee, vice president of business development at the [Copyright Clearance Center](#), an opening.



Sarah Palin could be easily found on YouTube and played back, two days after it first appeared on the NBC broadcast network's "Saturday Night Live" show on Sept. 13.



Two weeks later, attempts to upload complete versions of a follow-up rendition of Tina Fey as Sarah Palin that appeared on SNL on Sept. 27 were effectively blocked, over the next two days.

Takedown On Takedowns

Anvato, a Mountain View, Calif., supplier of ‘machine vision’ technology, says YouTube is only taking down a fraction of infringing clips uploaded to its site.

“I don’t think they’re doing anything. I think they’re full of it. I don’t think they’re trying at all,” says McNamee.

Safe copyright harbors

To some degree, YouTube doesn’t have to try. [The Digital Millennium Copyright Act](#) of 1997, creates a “safe harbor” for a video sharing service like YouTube. According to the DMCA, [YouTube is not liable for users’ infringing activities](#). Instead, YouTube’s responsibility is to “respond expeditiously to take the material down or block access to it,” once notified of an infringement.

Whether it was NBC’s inability to find copyright infringements or YouTube’s unwillingness to proactively look for them, one fact remains: Fey’s Palin clip was still available on YouTube on Sept. 15, when ZDNet brought it to the attention of NBC; and on Sept. 22. The clip was not available in unedited form on Sept. 24. Spotting such infringement is not that hard – at least in the eyes of film and TV programmer Viacom. In March 2007, Viacom filed a federal



suit against YouTube and its parent, Google Inc. [Viacom sought \\$1 billion in damages](#), saying, “YouTube has deliberately chosen not to take reasonable precautions to deter the rampant infringement on its site.”

How does Viacom know? It uses human eyes. Lots of them. Viacom employs between one and two dozen people at any given time just to watch videos uploaded to YouTube for infringement. Each of these individuals, known as “user-generated video reviewers,” can screen between 40 and 50 videos in an hour. Using this manual process, Viacom on Feb. 2, 2007 ordered the take down of more than 100,000 inappropriate uses of video content on its networks such as Comedy Central, Nickelodeon and MTV. By the time it filed its suit against Google in March 2007, the number had grown to 150,000. The suit led Google to develop its own technology for identifying videos (and music) that appear on YouTube and violate copyrighted material.

At this point, YouTube uses “multiple tools” to help content creators identify their content. But the key is that “content creators” must be proactive about the identification. Only with content creators’ help can YouTube match images on uploads to its site to videos that it is supposed to protect.

Here’s how it works: If a YouTube employee or executive sees an infringing clip on its site, there is no imperative to take it down. The content owner has to make a formal takedown request to remove the video. This process is not likely to satisfy Viacom’s lawyers. By the end of August 2008, Viacom’s screeners had identified a total of 375,000 postings of infringing video clips on YouTube. The monthly average is about 12,500, by its reckoning.

Takedown On Takedowns (continued)

NBC

Total infringing views found: 38 million

Total infringing videos found: 1,765

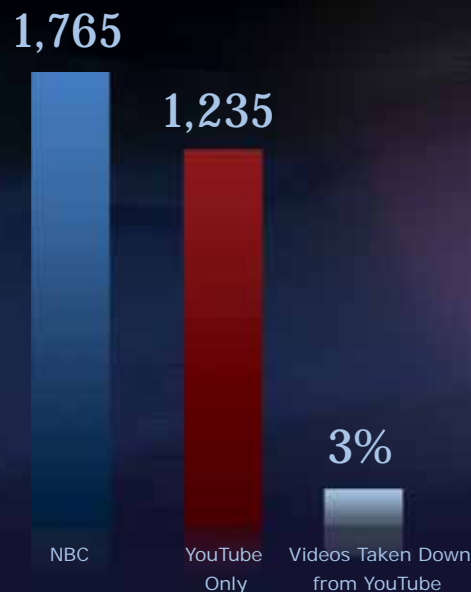
YouTube

Total infringing videos found: 1,235

Total videos taken down: 37 or 3%

Period: July 29 to Sept. 24

Sites surveyed: YouTube, Google, MetaCafe, Dailymotion



Video identification tools and how they work

For its part, YouTube says it has a bevy of tools at its disposal, but is reticent to quantify the success of its systems. YouTube says that it uses audio-filtering technology from Audible Magic, which can identify infringement of copyrighted songs or live music performances on soundtracks, as well TV show audio tracks. In addition, YouTube says it is using home-grown technology that watches patterns in images, to spot matches of copyrighted video as well. There are “many other content policies we enforce on the site,” the company says. YouTube declined further comment on how its systems work.

Early on, the YouTube system needed at least 30 seconds of a clip to identify it, one video identification expert contends. Rival technologies take as few as five seconds to make a match. YouTube says the system has improved and “is very effective.” YouTube wouldn’t comment on whether NBC’s effectiveness rate of 75 percent to 80 percent is correct.

Here’s what is known about YouTube’s approach. YouTube uses different techniques for fingerprinting audio and video files. It can also create “subfingerprints” for parts of a clip, according to the systems expert. Broadly speaking (and vastly simplifying), a fingerprint can detail whether a particular file is 10 percent yellow, 50 percent green, and 30 percent red, according to Vance Ikezoye, chief executive of Audible Magic. This approach is similar to what Audible Magic uses to distinguish audio tracks, say a song in an episode of ER versus a live performance. The most advanced approaches identify how human ears or eyes “perceive” colors, movement, and sounds, in context.

continued on page 7

Betting on Nexicon and help of its “natural-born content pirate”

Google, even amid Wall Street’s crisis of confidence, is worth \$113.1 billion, making it one of the 25 most valuable public companies in America. It pockets \$1.2 billion in profit on revenue of \$5.4 billion – every three months.

It wants a similar return on YouTube, for which it paid \$1.65 billion two years ago. That’s why perfecting a system that allows it and content creators to make money is critical.

Yet go-to company Google has allowed an outfit named Nexicon to make use of its video identification and monitoring system on behalf of rights holders Nexicon is run by a self-declared “natural born content pirate.”

Meanwhile, Nexicon’s financial statements are not current with the Securities and Exchange Commission and its shares are traded on “pink sheets,” where risky stocks that can’t meet the requirements of major exchanges are found.

Nexicon, however, is confident it can deliver:

“The volume and amount of digital assets (that) exist online, not just the originals but all the different variations that might exist online,” says Sam Glines, the company’s strategy and planning vice president. YouTube is “looking for service providers, someone unique such as Nexicon, that can provide one-stop shopping and handle this service for them,” he adds.

continued

YouTube bets on Nexicon, its “natural-born content pirate” (continued)

Nexicon, based in Malibu, Calif., now largely bases its business on mass data collection services as well as anti-piracy technology created by a Norwegian native, [Tommy Stiansen](#).

His company, called Pluto Communication when it was located in Bergen, Norway, merged with a forerunner of Nexicon, in order to go public. In the process, Nexicon asserts it has excised itself of its past – when it was a controversial [online seller of cigarettes known as Cyco.net](#).

“The only thing we share is a tax ID number,” says Glines. “The model has completely changed.”

The executive ranks have not completely changed, however. The company’s chief executive, Richard Urrea, also ran Cyco.net. Daniel Urrea is listed as chief financial officer [then](#) and [now](#).

The 32-year-old Stiansen is the company’s chief technology officer. He says he got his start in the business as a “[natural-born content pirate](#)” in Norway. His reference point? PirateBay, a supplier of illegally shared entertainment files based in Sweden, where he learned how to grab content that was not posted there by rights holders.

“The model has completely changed.”
–Sam Glines, Nexicon

Now, Stiansen is trying to “monetize” his knowledge about the way peer-to-peer and content-pirating techniques work.

Nexicon runs a high-volume data network that fights piracy by linking servers in northern Norway, Spain and the United States, by fiber connections. Its software can monitor “every single part of a peer-to-peer network,” according to Stiansen, and crawl the Web for other violations of copyrighted content. For the last four years, it has collected data, as well, on how much piracy is occurring in film and video industries.

Now, Nexicon can send out 95 million notices every day related to the digital copyright act. Three years ago, Nexicon entered into an agreement to supply “next-generation intelligence and data mining services” to the Motion Picture Association of America. This fall, it plans to introduce a service that allows content holders to send notices of infringement as soon as a download is initiated and seek compensation. In effect, peer-to-peer networking could be turned into an anti-piracy version of an online music store.

How well the fight is going is not clear, however. The company, while public, has not released any financial statements in years and [was losing money in 2005](#). Glines and Stiansen say the company is trying to bring its reports current with the U.S. Securities and Exchange Commission.

YouTube did not comment on why it struck the deal with Nexicon, which remains thinly traded, over the counter.

Under the YouTube agreement, Nexicon is acting as an intermediary or “outsourcer” for those content owners who want help monitoring and managing large amounts of music, video and other digital assets. Nexicon hopes to serve several film and content producers in the process.

Glines says YouTube – which, he says, approached Nexicon in July – provides automated takedown and monetization services to content owners, but does not have enough people in-house to provide assistance to content holders with a large volume of assets. Nexicon employs 10 people, seven of them full-time.

Nexicon says it has one “major studio” ready to use its services in dealing with YouTube, but has not disclosed its identity. In May, Nexicon similarly announced that a “major motion picture studio” had signed up for its [GetAmnesty anti-piracy program](#). But the identity of the studio was nowhere to be found in the [official announcement](#). ❖

Takedown On Takedowns (continued)

CBS

Total infringing views found: 9 million

Total infringing videos found: 2,163


YouTube

Total infringing videos found: 1,818

Total videos taken down: 473 or 26%

Period: July 29 to Sept. 24

Sites surveyed: YouTube, Google, MetaCafe, Dailymotion



continued from page 5

Audible Magic's approach is to sort through the complete aural spectrum of a song, extract key characteristics, and create a "fingerprint" that is a small, easy-to-transmit identification that can be matched against other copies of the same set of sounds (see sidebar pg. 13).

YouTube would not describe how its video identification system works even though it's based on fingerprinting. The system has to work well enough to automatically match even short segments of a video against the digital assets supplied by content owners such as NBC or CBS. Then the clips must be reviewed to gauge infringement.

YouTube's video ID system (official links)

- <http://googleblog.blogspot.com/2007/10/latest-content-id-tool-for-youtube.html>
- <http://googleblog.blogspot.com/2008/08/making-money-on-youtube-with-content-id.html>
- <http://www.youtube.com/t/contentid>
- http://www.youtube.com/t/video_id_about

YouTube may have developed its own fingerprinting system to help content owners identify uploads that infringe on their property. But Ikezoye says his company will soon begin fingerprinting video as well as audio files. One reason YouTube has been reticent to talk about its fingerprinting system is because it is still in its infancy, says

Mark Ishikawa, the CEO of BayTSP, Los Gatos, Calif. BayTSP helps Viacom and other content companies monitor top video-sharing sites for copyright infringement. Sturdy systems developed by Japan's NTT telecommunications giant, Europe's consumer electronics company Philips, and others take a dozen or more years to become functional. "If you put nine women in a room, you don't get a baby in a month," Ishikawa says. "The ultimate system cannot be created in this period of time."

In fact, the "ultimate system" has to include not just audio and video fingerprinting, but human eyes. At some point, unrecognized segments or pieces of a mashup have to be reviewed, Ishikawa says. But video identification technology is advancing rapidly. Anvato of Mountain View, Calif., says its approach, which tries to mimic how humans perceive moving images, needs only 20 seconds from a given TV show to identify any part of the episode. In its own scan of YouTube uploads from July 29 to Sept. 24, it detected 1,235 different infringing clips involving NBC prime-time shows. Of those, 97 percent were not blocked or taken down.

NBC's Cotton says he was unaware of the Anvato study and could not respond. YouTube says "no third party can know what policies a content owner has set regarding content claimed on YouTube." The content owner may elect to leave the clips up for promotional purposes or to bring in money from ads.

Following the money

How much money is at stake is not known. YouTube's ad deals are not disclosed, but the bulk of revenue probably

Takedown On Takedowns (continued)

ABC

Total infringing views found: 16 million

Total infringing videos found: 1,371


YouTube

Total infringing videos found: 1,093

Total videos taken down: 173 or 16%

Period: July 29 to Sept. 24

Sites surveyed: YouTube, Google, MetaCafe, Dailymotion



goes to content owners like Google's AdSense program, which takes 30 percent of the revenue for text ads. Similarly, in the iTunes store, the content owner gets 70 percent of revenue from a sale of a song. The store manager (Apple) gets 30 percent. Meanwhile, a broad attempt to monetize uploaded videos through ads that appear next to them, or even inside the clips, could change YouTube's and Google's status under the Safe Harbor provisions of the [Digital Millennium Copyright Act](#). Service providers like YouTube, which allows users to share content with each other, are not liable if they don't know about an infringement that is occurring or they act promptly to take down an infringing piece of content, once notified.

But the Safe Harbor protection is placed at risk if the service provider is gaining financially from apparent infringement. This is because of a provision in section 512c which says: "If the provider has the right and ability to control the infringing activity, it must not receive a financial benefit directly attributable to the infringing activity."

That provision means that YouTube or any other similar service provider must be proactive enough to ensure that it is not placing ads next to or inside any infringing content. However, even if the Safe Harbor protection is lost, that does not necessarily mean YouTube or another provider is guilty of infringement, says Evan Brown, an attorney with the intellectual property and technology group of [Hinshaw & Culbertson](#) in Chicago. A plaintiff such as Viacom "still would have to prove it," he says.

To prove that YouTube directly benefited from the infringing activity, the content holder would "have to show that

people were watching the video because it was infringing," adds Brown. More critically, the holder would also have to "prove whether the ad produced any revenue at all." That argument would be difficult to prove if Google only received a piece of the ad action when users actually clicked.

Nonetheless, the potential for making money, along with YouTube's passive approach to copyright enforcement, is what has content holders up in arms. Even though the Digital Millennium Copyright Act requires YouTube to act merely as a "data conduit" to avoid getting caught in the crossfire between content creators, copiers, uploaders, and consumers, the reactive approach to copyright enforcement does not sit well.

"The technologies exist to identify these things but [YouTube doesn't] seem to use the technologies unless someone complains," says Thomas H. Hamilton, chief information officer at the Copyright Clearance Center. "The technology is there, the compute power is there, and the dollars are there for them to be more proactive on this."

Google, for instance, recorded net income of \$4.8 billion on revenue of \$19.6 billion last year. That's a profit of about 25 cents on every dollar of revenue, after all taxes, interest, and expenses. Yet Google didn't deploy any serious video identification technology at the start of 2007, just before Viacom filed its suit. Google bought the profitless YouTube for \$1.65 billion in October 2006.

According to one video identification technology company executive tracking the case, the first YouTube video identification system was designed to protect Google's profit

continued on page 10

Sizing up the takedowns

NBC Universal says it believes YouTube's systems now catch 75 percent to 80 percent of all infringing video clips that get uploaded. YouTube won't comment for the record, but one insider confidently says that figure is low.

Then, comes Anvato, an independent supplier of technology that tries, using what is called machine vision, to replicate what the human eye sees. Anvato has zeroed in on NBC primetime content. And, its chief executive, Alper Turgut, says YouTube takes down only a small fraction – as little as 3 percent – of its infringing clips that get uploaded.

The Mountain View, Calif., company, which markets its own video identification technology to content owners, told ZDNet that it tracked 12 primetime NBC shows, including "30 Rock," "Heroes," "Scrubs" and "The Office," from July 29 to Sept. 24. In that period, clips that included infringing content from those shows were viewed a total of 38 million times on four video-sharing sites: YouTube, Google, MetaCafe and DailyMotion.

On YouTube, Anvato says it detected 1,235 different infringing clips, using its own technology. YouTube, Anvato CEO Turgut says they took down 37 of the clips, or 3 percent.

Anvato, founded by a group of Turkish entrepreneurs, uses "machine vision" technology to replicate how human eyes look at images and identify scenes, objects and movements inside a video. Turgut says the technology can sort through 1 million minutes of video from sites such as YouTube, Veoh or Dailymotion in under five seconds, using just six servers.

In the case of a half-hour NBC show, it needs only 20 seconds to "uniquely identify" content anywhere in that half-hour – even if it is not the same 20 seconds, Turgut says. There are enough continuing characteristics of characters, scenes and backgrounds to identify a potential violation, then review, refine and verify the violation.

But NBC general counsel Rick Cotton says he would have to review each instance to determine if a 3 percent or 38 percent takedown rate was out of whack – or in line with what NBC may have instructed YouTube to do.

Under YouTube's system, infringing content is not necessarily taken down, even if it is identified.

There are three choices, left up to the content owner:

- **Block:** Tell YouTube to take down the video.
- **Leave as is:** The content owner may consider replays of the clip to be promotion that benefits the original show. YouTube supplies statistics on usage.
- **Monetize:** YouTube agrees to place an ad next to the video on its site and split the revenue that comes in with the content creator.

So, in any of these cases, it's possible that the content owner elected to have the clips left up for promotional purposes or to make money. It is not necessarily a reflection on the YouTube system to not catch an infringing video in the first place.

Simply put, an outsider, such as Anvato, can't know what choice the content holder has made, YouTube notes. ❖

TAKEDOWN ON TAKEDOWNS

NBC shows tracked: 30 Rock, America's Got Talent, America's Toughest Jobs, Chuck, Heroes, Law & Order: Criminal Intent, Law & Order: Special Victims Unit, Lipstick Jungle, Medium, My Name Is Earl, Scrubs, The Office

CBS shows tracked: 48 Hours Mystery, Cold Case, Criminal Minds, CSI: Crime Scene Investigation, CSI: Miami, CSI: NY, Dexter, Ghost Whisperer, How I Met Your Mother, Jericho, NCIS, NUMB3RS, Rules of Engagement, Shark, Swingtown, The Big Bang Theory, The Mentalist, Two and a Half Men, Without a Trace, Worst Week

ABC shows tracked: According to Jim, Boston Legal, Dancing With the Stars, Desperate Housewives, Eli Stone, Lost, Men in Trees, Miss Guided, Opportunity Knocks, Samantha Who?, Ugly Betty

continued from page 8

margins. It used mathematical functions in the process of processing and compressing digital signals known as “wavelets” to identify duplicate videos that users were uploading, this technologist contends. The point: Save server capacity by only storing one copy.

This kind of saving would help preserve Google’s \$4.8 billion of net income. That annual profit is nearly three times the profit of Viacom, which still found enough money to pay dozens of workers to monitor YouTube uploads for possible violations. Viacom earned \$1.7 billion on sales of \$14.5 billion in 2007.

Here’s what rubs critics such as Hamilton the wrong way. Google is creating hundreds of server farms across the globe and is even considering data centers floating in the sea and alternative energy to power them. Yet its video identification systems aren’t as good as they could be. Hamilton says that Google has fostered an environment where copyright infringement is the norm. Google needs to be more proactive about copyright infringement.

A two-front war

Is it in Google’s best interest to be proactive? Independent suppliers of video identification technology, such as Anvato’s Turgut, say it doesn’t make sense for Google to catch more offending video than that which is forced upon them. Roughly 30 percent of viewership on sites like YouTube or Veoh is tied to copyright-infringing content, said Turgut. Police copyright better and that traffic “would vanish overnight,” Turgut adds.

Simply put, taking a hard line on copyright infringement could make it hard for Google to generate a return on YouTube. Former Bear Stearns analyst Robert Peck estimated that YouTube will take in about \$90 million of domestic ad revenue this year, or a 7 percent share of the online video advertising market. In July, Google served up about 5 billion online videos, almost half the U.S. tally, according to comScore.

To handle that, Google is probably spending hundreds of millions of dollars a year just on servers. “Their P&L is inverted,” one operator of data centers storing large multimedia files said in mid-September about YouTube’s profit-and-loss statement. YouTube declines to comment on its operational costs, but says it is motivated to make sure that content owners are treated properly.

One incentive for Google to play nice is a looming threat to YouTube’s model. Anvato and similar companies such as Contributor, Vobile, and Audible Magic have a dual role: They can provide video identification services to YouTube and content companies like CBS, NBC, ABC, Turner Broadcasting, and Viacom and also serve ads.

If Google’s rivals sign deals with major content holders for both video identification and ad placement services, the YouTube monetization model could be undercut. Why? YouTube has all the costs of serving up the infringing video, but little payoff, notes one chief technology officer of a company that works with rights holders.

To head off that looming competition, Google has been offering olive branches to content providers. When the search giant first tried to monetize offending videos it offered a 50/50 split of ad revenue, says one industry executive familiar with the proposal.

But content companies wouldn’t bite – remember they get 70 percent of revenue in the Apple iTunes precedent. According to one executive who has been watching negotiations with major rights holders, Google’s bid to garner ad revenue from high-quality video content has three looming issues: Who owns the advertiser, who gets what cut, and indemnification. According to this industry executive, Google wants the rights holders to agree not to take legal action for any infringing content on YouTube in the past three years. YouTube declined comment.

That legal action stipulation may explain why Google and YouTube haven’t been inking monetization deals with major content companies. YouTube’s main success story with a major broadcaster is CBS, which for two years has

continued on page 12

CXO CHECKLISTS

What You Should Ask About: Video Recognition Systems

If you're a CEO:

- ✓ How accurate is the technology?
Ask about:
 - a. Precision. If it gives 10 results, are all 10 results valid? Will you be wasting time on one out of 10?
 - b. Recall. If you get 10 results, is that all the results? Is the system capable of catching all the videos you want to spot?
- ✓ How are the results achieved? How does the system verify its results?
- ✓ Can you see the results yourself?
- ✓ Who has independently verified the results?
- ✓ What tests did you run, Mr. CIO?
- ✓ Does the system have a logical path that leads to producing revenue from the videos, Mr. CFO?
- ✓ How fast does the system work? Will it catch videos uploaded in the first 15 minutes after my TV show has ended?

If you're a CFO:

- ✓ How much will the system cost two years from now, as usage grows?
- ✓ How much human-checking will be required? How will those costs increase?
- ✓ How automated is the takedown process?
- ✓ How automated is the ad positioning and insertion process?
- ✓ Can I dictate the terms? How much control do I get?
- ✓ How does the system actually track revenue generated by ads associated with the videos?
- ✓ What metrics do I get?
- ✓ How comprehensive are the systems? That is, do they work not just against YouTube but all sites where you expect your company's content to be posted in volume? ❖

SOURCES: Anvato, Attributor, Turner Broadcasting

If you're a CIO:

- ✓ Did you run the Motion Picture Association of America's test videos against the system you're looking at?
- ✓ Did you run your own?
- ✓ How do the systems compare in side-by-side tests?
- ✓ Did you do final checking with your own eyes?
- ✓ How many "false positives" do you get out of 5,000 results? Out of 50,000? Out of 500,000? Out of 5 million?
- ✓ How much human-checking would that translate to on an annual basis?
- ✓ Is the system comparing apples to apples? The same video can be at the other end of multiple Internet addresses.
- ✓ How robust is the system in finding content through subterfuges, like changing the file format or rotating the angle of the picture slightly?
- ✓ How easily will it integrate with my other systems?

Video ID systems: The cost-benefit analysis

Video-identification systems aren't cheap, but content owners that want to prevent unauthorized use of their video the fees and charges may be worth it. This spreadsheet shows you how to crunch the numbers.

<http://whitepapers.zdnet.com/abstract.aspx?docid=393089>

And the revenue side of the equation

This spreadsheet details how a high-traffic distributor could implement video-id software to its advantage. The new software addresses risk, legal fees and ad revenue.

<http://whitepapers.zdnet.com/abstract.aspx?docid=393090>

continued from page 10

supplied content to a channel on the site. Of course, CBS Interactive (which owns ZDNet) retains advertising rights.

YouTube makes a deal

While YouTube hasn't made great inroads with traditional content companies, it has been dealing with smaller players. One of these smaller companies surfaced on September 11. Google said it will work with Nexicon, a Malibu, Calif. company that trades as an over-the-counter or "pink sheet" company, which is typically an illiquid form of stock to trade. Nexicon's mission is to help content owners generate revenue on YouTube or remove infringing content. The rights owner's digital assets will be stored on Nexicon servers, fingerprints generated and inserted into YouTube's content management system.

Nexicon, which last filed a Securities and Exchange Commission filing in December 2005, says it has one "major studio" ready to use its services in dealing with YouTube, but has not disclosed its identity. Also, the company will employ YouTube's content management system to match clips with ads and monitor unauthorized uploading.

It's unclear whether Nexicon will make much of a difference. Viacom declines to comment on how well YouTube

is performing in taking down violating videos. In the meantime, one individual close to the company says the pace of violations may be decreasing on YouTube "but it's not stopping."

And NBC?

"Our view is that it's getting better every month," Cotton says of YouTube's content recognition system. Takedown of videos is simply a matter of entering a code into a field on a screen and letting the system take over. There's been a "significant dropoff" in terms of full-length content getting broken up into 10-minute segments and posted on YouTube. Overall, posting of the full duration of SNL clips is down, he contends, even if the Fey-Palin skit slipped through and was available on America's most-used video-sharing site after more than a week, without authorization.

Cotton's proof point came in August, with the Olympics. NBC first filtered out visitors to NBCOlympics.com, to make sure only individuals using an Internet service provider in the United States could get to its content. Then it vigilantly watched for pirated material appearing elsewhere and ordered prompt takedowns. It sent copies of each stream of Olympics video to YouTube, Audible Magic, and Vobile. Fingerprinting caught about 80 percent of unauthorized uploads. Vobile also crawled the Web for violations.

The count? About 25,000 takedown notices. The result? Ninety-nine percent of the 3,600 hours of video programming (2,000 of it which only came online) accrued to NBC, Cotton says. Just one percent went to infringing sites.

Technology like that which YouTube has introduced means a broadband Internet ecosystem – with rules that can be acted upon instantly and largely through automation – is developing.

That ecosystem will consist of a marketplace of digital tools that includes competitive and complementary content detection services, monetization “facilitators,” and databanks storing information about copyright holders’ assets, says Scott Teissler, chief technology officer for [Turner Broadcasting System in Atlanta](#), which worked with YouTube on the development of its in-house video identification system.

“If I as copyright holder want detection methods y and z employed,” Teissler said, then “the companies who have those rival detection methods already (will be) used to working in harness at different aggregation sites,” such as YouTube or Dailymotion.

A marketplace of multiple suppliers also would eliminate what participants call the “fox in the hen house” problem: A content owner may not want the same company that identifies the infringing video to be the same one trying to monetize it. That content owner may want an independent company to figure out the violations going on at a video-sharing site – and it may want to handle the ad placement itself.

Even so, YouTube says that 90 percent of the content owners it works with now use its system to monetize clips.

Now, at least, the process is largely automated. But it’s unrealistic to expect that illicit uploads of copyrighted content will ever go away. Even if YouTube’s system gets refined over the next year, it’s not likely to catch or block every unauthorized posting of a Tina Fey comic take.

“The ability to say, gee, I found an infringing copy some place is always going to be true,” NBC’s Cotton said, just before the faux Couric-Palin skit aired. “It’s never going to be zero. But you can see the dial turning.” ❖

About the author: Tom Steinert-Threlkeld is a journalist who has looked at what media could become, rather than what it currently constitutes. Most recently, he served as editorial director of Broadcasting and Cable as well as Multichannel News magazines for Reed Business Information. Prior to Reed, he was vice president of the Enterprise Group of Ziff Davis Media, where he founded Baseline magazine and within four years made it a National Magazine Awards finalist for General Excellence. He also was the editor in chief of Inter@ctive Week.

Primer: Video identification 101

Until Viacom filed its lawsuit last year against YouTube, the Google video-sharing business relied on soundtracks to find clips uploaded by users that might infringe on TV programs and movies.

The primary tool is an audio recognition system from Audible Magic, of Los Gatos, Calif. The technology is still in use today by YouTube, in what is now an audio, video, and policy-driven expunging of illegal uploads of copyrighted content to its servers.

The system creates digital fingerprints for sounds, in a technological equivalent of how a musical score is written. A composer puts down notes that represent continuous sounds, but doesn’t put down the sounds themselves.

In the Audible Magic case, the notes are expressed mathematically as “coefficients.” These are numbers that indicate the amount or degree that the representation of the sound correlates to the characteristic of the original sound.

The Audible Magic approach, originally developed by a company called Muscle Fish, creates types called “mel-filtered cepstral coefficients.” These transform the original sounds into “spectral” representations that lie within the range of normal human hearing.



Snippets of sound are run through a mathematical function, known as a Fourier transform, that converts a signal (like a sound recording) into its frequency components. The results track how much energy is being expended in different frequencies, such as bass, midrange and treble. Those results are then run through a bank of filters, which exclude frequencies unrecognized by the human ear.

The output is a coarse representation of the shape of the sound, sometimes called a “spectral envelope.” The envelope is sent through one more Fourier transform, to produce the cepstral coefficients – “cepstral” being an intentional misspelling of the word “spectral.”

The advantage: Only a few coefficients are needed to create a fingerprint for each snippet of sound. This saves storage space, but also makes it easier to search soundtracks for any match against the snippet.

In response to pressure from content owners, YouTube debuted its home-grown version of video filtering a year ago, on Oct. 15, 2007.

The process was developed by Google and relies heavily on content owners to be proactive about protecting their assets. YouTube and Google don’t attempt to describe how they make the matches. But the task of a video identification system is largely to disregard irrelevant information in an image and concentrate on what is salient and unique.

The idea, pursued by such service suppliers as Attributor and Anvato, is to mathematically see what the human eye sees. The Anvato approach is to index objects that distinguish a particular snippet of video such as faces, cars, shapes and how they move.

All these items become part of a “perceptual signature” that can only be represented by the same movement in the same movie. How Paul Newman disables his knife-wielding enemy at the outset of *Butch Cassidy and the Sundance Kid* would be distinguished not just by the speed of his kick to the fellow’s groin, but the tilt of his hat, the way he ambles, maybe even the blueness of his eyes. (And the unusual features of his enemy’s face and the particular knife in his hand.)

The scene then gets reduced to a series of numbers – identifiers – that make sense of the video.

To get sufficient identification to keep track of any scene lifted from a TV show though, Anvato’s CEO Alper Turgut says only about 20 seconds of any episode is needed.

The process works in three stages. At high-speed, the system works in a blur, picking up potential violations. Then, batches of those potential violations are run across multiple processors at a slower speed with closer identification to establish infringement. Then, that final set of potential violations is run through the system one more time to verify that a clip is an exact match. If there’s a match the content owner must decide whether to leave it up on a video-sharing site or order it taken down.

YouTube’s original method of identifying the content of videos used mathematical functions when processing and compressing digital signals, according to one video identification systems expert. These functions, known as “wavelets,” identified duplicate videos that users were uploading, one technologist familiar with the progress of video identification systems contends. YouTube declines to describe the evolution of its process.

“To get sufficient identification to keep track of any scene lifted from a TV show, only about 20 seconds of any episode is needed.”

Tracking “wavelets” works this way: Colors are carried on different frequencies in each frame of a motion picture or TV show. If you start with an all-white frame and move across 30 frames to a completely black frame, you can stitch together the frames and create a continuous wave function, or wavelet, that plots the change of the color frequency.

The right way, video identification experts say, is a two-step approach that looks with more granularity inside the frames. The first step is to apply “frame-based” algorithms. These pull out colors, patterns and shapes in each frame, such as Newman’s foot (from our earlier example) and where it is positioned. The second is to apply “motion-based” algorithms. These look at the differences that occur from one frame to the next.

Both of these types of algorithms are used by different vendors of identification technology, with different flavors, much like search engines. None publish their secret sauce. But Google’s multi-level fingerprinting technology almost certainly applies both algorithms to identifying videos on YouTube, says one independent supplier of video tracking services.

Each vendor typically creates a series of “subfingerprints” that capture changes from frame to frame, in times ranging from one-hundredth of a second to a second. These “subfingerprints” encode shapes and picture element changes for an entire TV show or motion picture.

One company that does use both “frame-based” and “motion-based” algorithms is Attributor. The combination allows the company to sort through and see through a variety of feints that copy thieves use to elude detection.

Thieves make it hard to recognize the source of their video by changing its appearance or format. A high-definition stream may get broken down to web-quality video, with 80 or 90 percent of the original information tossed out. Thieves may change the format from MPEG4 to Flash, resulting in different bits being played back. They may rotate an image a degree or two, or change the contrast. They could also add or scroll subtitles across the screen, and even superimpose a logo.

The result degrades the quality and can change the format, at the same time. Depending on how it’s focused, the resulting frame may also include the edges of the TV or less of the picture than the original transmission.

These kinds of user-generated feints on professionally generated content are what video identification systems are designed to catch and see through. YouTube specifically demonstrates in its online demo that its identification system can determine what original content underlies a poor-quality upload that started with a camcorder recording of a copyrighted TV program.

YouTube’s system also creates a unique “hash,” or digital signature, that allows it to block attempts to repeat the upload of the same video file. Repeat infringers’ accounts are terminated. ❖