



Time for a check-up:

The brand safety CRISIS one year later

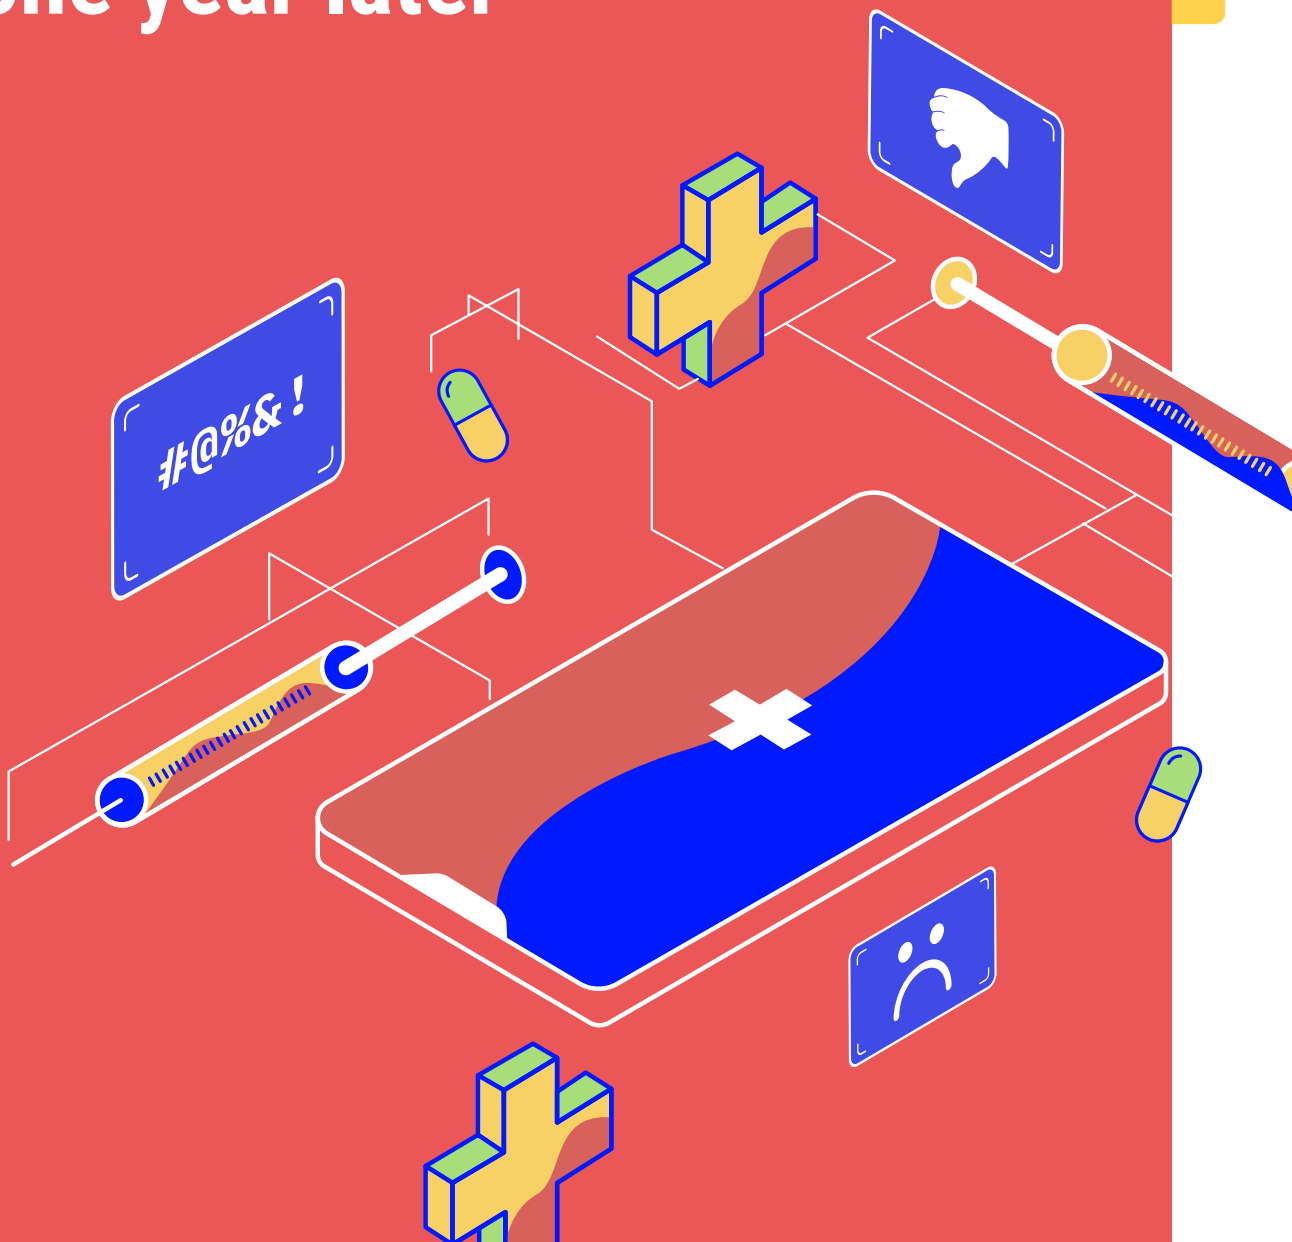




Table Of Contents

Brands are taking matters into their own hands

p. 03

“A lot of media clients we work with have someone who in their job description is to be the de facto brand safety officer.”

The platforms are fighting back – and it’s working


p. 04

“Social platforms have certainly taken brand safety seriously. They’ve taken notice that there was work to do and they’re addressing it.”

Marketers are inoculating themselves against favorable audiences

p. 09

“I think the market has over-corrected. But I don’t blame the marketers for it—I blame the tools, systems and ecosystems that got us to the point where there’s too much risk for brands in not over-correcting.”



THE CRISIS

ONE YEAR LATER

In late 2017, the sky was falling. Brands panicked, pulling ads from prominent platforms, berating social media executives in private calls and begging their agency partners to find a solution. Daily headlines heightened the hysteria: Controversial influencers attacking loyal social followers, consumer data breaches and a phantasmagoria of brands finding themselves adjacent to violence, nudity and fake news.

According to our 2017 survey, 90 percent of digital media pros saw brand safety as a serious problem. And the ecosystem is still hazardous.

But too much of the media coverage surrounding brand safety has missed a huge part of the story: Over the past year, a slew of successful treatments have been used to treat the problem, including promising technologies such as image recognition and natural language context detection.

By and large, marketers say these solutions are working. But at the same time, they're worried that certain treatments — especially preventative measures like blacklisting and whitelisting — may be overly restrictive, serving to inoculate their brands against favorable audiences even as they're protected from deadly adjacencies.



FOR Media Buyers

DUE TO Platform & partners risks

PRIMARY Brand RX

ONE YEAR LATER, ARE BRANDS SAFER?

Marketers and social platforms have been quick to recognize the dangers of inadvertently exposing themselves to violence, nudity, extremism and other safety hazards. Since the brand safety epidemic hit more than a year ago, they've fought deadly adjacencies with urgent treatments such as blacklisting, whitelisting, third-party measurement, brand safety specialists and AI-driven technologies.


While treatments have proven effective overall, some have given rise to entirely new health risks, including reduced audience pools and weaker targeting capabilities.

Patient reports improved symptoms and reduced exposure to brand safety hazards. Cites treatments such as AI-driven technology, better hiring practices and a healthier social media environment. But many symptoms persist, and some treatments have given rise to debilitating side effects such as smaller audiences and weaker targeting.

SIGNED

In concert with Digiday, we surveyed 274 industry pros from brands, agencies and publishers in November 2018. From our findings, it's clear that most marketers are at least somewhat pleased with the steps that prominent social media platforms have taken to address brand safety since 2017. Many have also introduced brand safety measures of their own and made key hires.

But vulnerabilities remain, and many of those treatments have had grave side effects. So what preventative measures are marketers taking to stay protected? And are they working? Let's find out.



INDUSTRY PROFESSIONALS WHO SAY BRAND SAFETY IS A SERIOUS PROBLEM FOR THEIR MARKETING EFFORTS:



TYPES OF BRAND-UNSAFE CONTENT THAT INDUSTRY PROS SAY THEY'VE MOST COMMONLY BEEN EXPOSED TO:



TOOLS THAT INDUSTRY PROS DEEM TO BE THE MOST EFFECTIVE AT PREVENTING BRAND SAFETY PROBLEMS:



INDUSTRY PROS WHO SAY THEIR ABILITY TO REACH THEIR DESIRED AUDIENCES HAS DECLINED ALONG WITH THEIR USE OF TOOLS SUCH AS WHITELISTS AND BLACKLISTS:



INDUSTRY PROS WHO USE IMAGE RECOGNITION TO PREVENT BRAND SAFETY ISSUES:



BRANDS ARE TAKING MATTERS INTO THEIR OWN HANDS

THE OUTBREAK

“Brand safety” isn’t exactly a new risk factor — but it wasn’t much-discussed before March 2017. That’s when the initial diagnosis took place: High-profile brands like Verizon and L’Oréal appearing alongside YouTube videos from the likes of ISIS and Neo-Nazis.

Hundreds of brands, terrified of such deadly adjacencies, threatened to pull their ads — and many made good on those threats.

But most remained — or returned.

It was the same story on other platforms — Facebook, Google Search, Twitter and more. The ubiquity of the pandemic made staying out of the game a non-starter, despite the risk. In the past year, many marketers have done what they could to take matters into their own hands.

Of course, this doesn’t necessarily re-

THE FIRST ROUND OF TREATMENT

quire new resources, but the refocusing of existing ones.

“A lot of media clients we work with have someone who in their job description is to be the de facto brand safety officer,” explained Joe Barone, Managing Partner of Brand Safety for GroupM. These specialists tend to emerge from companies’ existing marketing teams, according to Barone.


But with brands and agencies such as Bank of America, IPG Mediabrands and GroupM itself hiring new blood as designated brand safety officers, it’s safe to say that companies are taking a multitude of approaches to treat the problem. Either way, dedicating personnel resources indicates a clear seriousness in terms of solving the problem.

So just how much of an impact are these roles having?

Notably, 63 percent of marketers believe these specialists have only helped a little, while 31 percent say they’ve delivered a moderate amount of help.

Only 5 percent think these specialists have helped a lot.

- Yes, a little
- Yes, a lot
- Yes, moderately



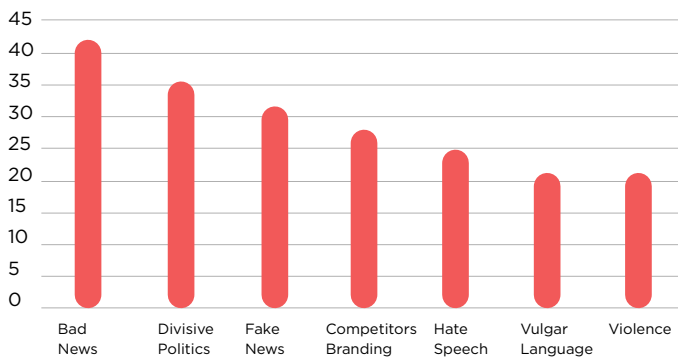
METHODOLOGY

The GumGum Brand Safety Crisis survey interviewed 274 industry professionals between October 27 and November 25, 2018. Respondents predominantly worked for brands, agencies, online publishers or technology providers.

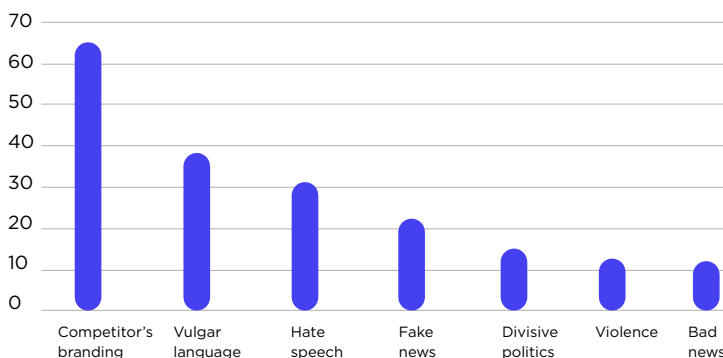
THE PLATFORMS ARE FIGHTING BACK — AND IT'S WORKING



2017



2018



THREAT VECTORS

When it comes to which types of content marketers fear most, the ground is shifting dramatically. Two years removed from the 2016 election, worries about news-related adjacencies have in fact waned.

Meanwhile, by an overwhelming margin — 63 percent — marketers' say that adjacency to a competitor's branding is now the most common brand-unsafe exposure they've experienced. That's a dramatic increase from last year, when the number was 28 percent.

"When we're trafficking our ads we're not going to put a Ford ad next to a GM ad and we're not going to put a McDonald's ad next to a Burger King ad," explained Blake Sabatinelli, CEO of the video-heavy news publisher Newsy. "We want to make sure that that brands can stand out on their own."

"It's the same thing that exists in traditional television," he added. In other words, the fear of being grouped with a competitor is not exactly a new one for brands — in fact, it's one of the oldest. But it's never been a more pressing concern than it is now.

In a digital marketing environment increasingly dominated by visual media, brands are increasingly concerned about making sure a logo or other piece of signage isn't sharing the screen with rival branding, thereby neutralizing its value or even helping a competitor. As advertisers drop huge sums of money to proliferate still images and viral video clips across the digital and social ecosystem, letting such a mistake slip through the cracks might cause it to be seen thousands or even millions of times.

So as marketers devote a growing portion of their budgets to digital visual media content, it's no surprise that they're becoming more concerned about competitive adjacencies.

Nonetheless, that leap from 28 percent to 63 percent, all in the space of a single year, is still shocking. In part, that shift is due to the fact that concerns about other types of brand-unsafe exposures have decreased. The question is, why?

The shifting landscape signals a massive effort made by the social media platforms themselves. Indeed, if you ask

THE RUSH TO TREATMENT

marketers, social platforms are doing a pretty good job of cleaning house.

It's an easy story to miss as the media breathlessly reports Facebook's or Google's latest missteps. But the industry's position — on brand safety at least — is clear: The epidemic is being treated effectively, and the Facebooks and Googles of the world are playing a significant role.

Across the board, respondents rated the brand safety efforts of prominent social platforms highly.



“Social platforms have certainly taken brand safety seriously,” said Barone. “They’ve taken notice that there was work to do and they’re addressing it.”


Facebook, for instance, has given marketers the ability to see where their ads might appear on Instant Articles or in-stream videos so that they can block certain content categories or publishers from those ad placements. And YouTube has made it far easier for marketers to create whitelists that limit their ads to a pre-selected group of trusted content owners.

But perhaps most importantly, the likes of Facebook, YouTube and Twitter have increasingly opened themselves up to far more third-party measurement, making

it easier for agencies and vendors to observe where ads appear on the platforms.

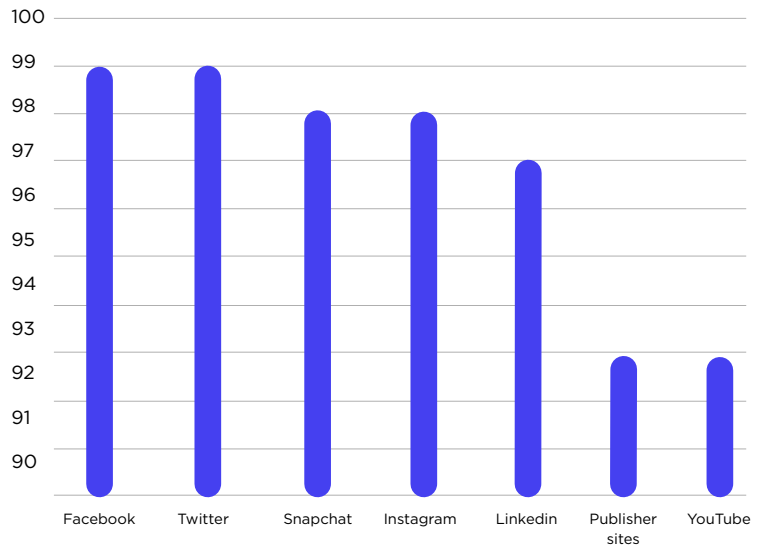
“We’re not all the way there yet,” continued Barone. “But most of our clients absolutely see movement, and they see a commitment to restoring confidence in digital media.”

Publishers feel the same way. “What you’re seeing out of these platforms now, and will continue to see over the next 18 to 24 months,” said Sabatinelli, “is greater accountability and greater reporting responsibility.”



WHEN IT COMES TO IMPROVING ON BRAND SAFETY OVER THE PAST 12 MONTHS, WHICH OF THE FOLLOWING PLATFORMS’ EFFORTS WOULD YOU RATE AS FAIR OR BETTER?

2018



TWITTER: A CLEAN BILL OF HEALTH



Twitter, in particular, has overhauled its approach to brand safety and turned its reputation around in the process.

In 2017, only one percent of respondents cited Twitter as the most brand-safe social platform platform. In 2018 percent, the number was 45 percent — a significant plurality of respondents.

Last year, LinkedIn came out on top. Of course, we're likely not experiencing a fall from grace on the part of the career platform; the shift likely speaks well of Twitter, not poorly of LinkedIn — only four percent of respondents called the latter a brand-unsafe platform this year.

For Twitter, its progress is a symptom of a larger trend: Social platforms have spent much of the past year booting controversial accounts (e.g. Infowars) from their platforms, and scrubbing out brand-unsafe content like violence and nudity. Tumblr, for example, made waves when it finally decided to ban porn. Some users may not have approved — but marketers did.

WHICH OF THE FOLLOWING PLATFORMS IS CURRENTLY THE MOST BRAND SAFE?



“These platforms were allowing everything,” said Matt Rivitz, founder of the media advocacy organization Sleeping Giants. “Richard Spencer was a verified Twitter user and it was no holds barred. The racists had sort of taken over taken over Twitter, Facebook and YouTube, and [the platforms] weren’t doing a thing about it. But then advertisers said, “we’re going to pause for a little bit until you can get this figured out.”

Twitter, in addition to introducing more third party measurement, has devoted significant attention to eliminating bots and fraudulent accounts. “They’ve done a lot to try to avoid the content adjacency issues,” said Barone.

“They’re also a smaller platform,” he

pointed out. “Sometimes Google and Facebook can create a lot of shade for those platforms platforms because they take so much of the heat.”

None of these treatments occurred spontaneously — they were a response to pressure from marketers who rated Twitter as a brand-unsafe platform in 2017.

“Good platforms are ad-driven,” said Rivitz. “They’re there to make money. The only way they change is with some public outcry and some internal debates over what’s right for their business.”

Twitter listened to the market — and it’s probably no coincidence that its stock value soared in 2018 after a series of rosy earnings reports.



MARKETERS ARE INOCULATING THEMSELVES AGAINST FAVORABLE AUDIENCES

PUBLISHER PARANOIA

As casual observers focus their attention on social media, marketers are focusing elsewhere. According to our survey, they think that independent websites threaten them with toxic adjacencies more than any other platforms. A significant plurality, 45 percent, told us that they consider publisher sites to be the most brand unsafe ad platforms — and 65 percent said that display ads on publisher sites are the most unsafe kinds of placements.

Facebook, the most feared social media platform, came in at a distant second with 26 percent. And YouTube, the platform that kickstarted the brand safety crisis in the first place was in third with 11 percent.

Clearly, for marketers, the threat vectors are changing dramatically. Just one year ago, 34 percent said Facebook was the riskiest platform for brands. At that time, publisher sites placed second with 27 percent, while YouTube sat at third with 15 percent.





THE RISKS OF OVERTREATMENT

Marketers have turned to various tools to address their ills — and lately, those tools reflect an increased fear of certain publishers. Over the past year, 55 percent of marketers say they're turning to direct relationships with trustworthy publishers to prevent brand safety issues before they happen — significantly more than any other technique. Last year, the number was only 39 percent.

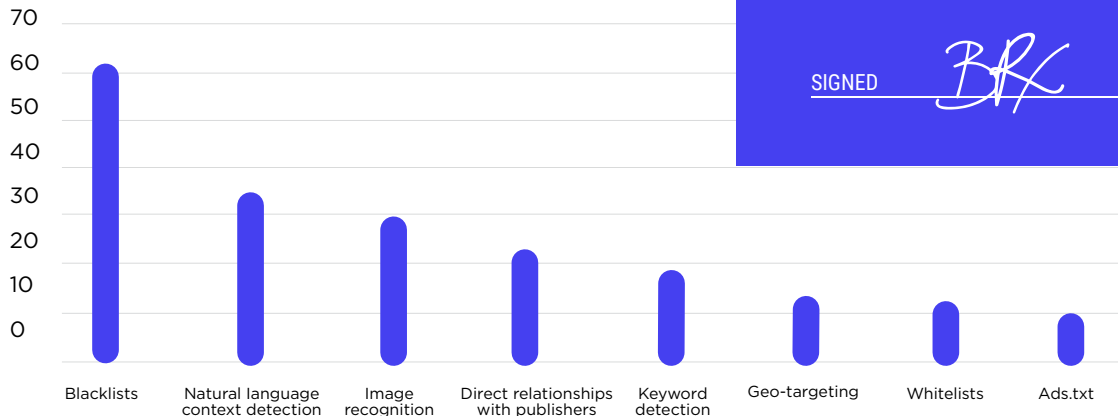
Yet marketers say that other strategies have proven more helpful. 62 percent of respondents said that blacklists — pre-prepared lists of sites where marketers' ads can't run — have been their most effective brand safety treatment over the past year, up from 50 percent a year ago.

But there are some huge shortcomings. For one thing, marketers often have to develop blacklists reactively. Sometimes they'll only know to blacklist certain sites once they've already been exposed to brand-unsafe content — and by that point much of the damage has already been inflicted. Preventative treatment is always a better option.

Even when used preemptively, blacklists can accidentally inoculate brands against favorable audiences. Sometimes parts of a site might be ideal for a brand, and what the brand really needs is simply to avoid the other parts.

"Blacklists are a very effective tool," said Barone. "But the problem with a blacklist is you don't know what you don't know."

IN THE PAST YEAR, WHICH BRAND SAFETY SOLUTION HAS BEEN MOST EFFECTIVE (PICK TWO)?



FOR *Media Buyers*

DUE TO *Platform & partners*

PRIMARY *Brand RX*

ARE YOU OVE-MEDICATING?

Blacklists and whitelists are undeniably effective treatments. But when overused, they can inoculate brands against favorable audiences along with brand-unsafe content.

The treatment shouldn't be as bad as the illness

While most brands have found whitelists and blacklists to be effective at treating their primary symptoms, the side effects have often been severe. "The problem with a blacklist is you don't know what you don't know," said Joe Barone, managing partner of brand safety for GroupM. "And we're very big users of whitelists. But sometimes a whitelist can carve away some meat with the fat."

Patient reports fewer brand-unsafe adjacencies, but also suffers from diminished return on ad spend. Recommend image screening and natural language context detection to supplement targeting capabilities, protect against brand-unsafe adjacencies and diminish reliance on restrictive whitelists and blacklists.

SIGNED *BRX*



FOR *Media Buyers*

DUE TO *Platform & partners*

PRIMARY *Brand RX*

NEWS-RELATED THREAT VECTORS

“Fake news” is still a deadly adjacency, but brands don’t worry about it like they used to. Social media platforms and publishers have taken numerous steps to treat the problem in the past year. Toxic accounts like Infowars have been purged. Agencies, third-party measurement specialists and brands have been allowed to look under the hood and see where ads are appearing. And brands and advertisers have taken to blacklisting controversial sites like Breitbart.

Know your partners

Many advertisers have formed closer relationships with news publishers they deem particularly accurate and reliable. “There are no second chances anymore,” said Blake Sabatinelli, CEO of Newsy. “If you make a factual mistake you’re done. Our job is to make sure that everything is vetted over and over again.”

Avoidance isn’t a cure-all

There’s a whole lot of news out there, and not all of it is warm and fuzzy. Divisive politics, violence, corporate malfeasance — it’s all bound to make brands uncomfortable. The problem is, such news items often draw the most engaged users. “Brands need to make their own decisions around what sorts of topics or conversations they want to be around,” said Ryan Pauley, chief revenue officer at Vox Media. “But it’s getting harder to avoid some of these conversations. If you start to narrow it down too much you’re not going to have any awareness.”

SIGNED

Even whitelisting — a technique, used by 11 percent of respondents, in which brands reserve their ads for a pre-approved list of safe sites — runs the risk of keeping brands out of favorable environments by restricting them to a handful of sites. “Sometimes a whitelist can carve away some meat with the fat,” said Barone.

Even as marketers have become more confident that their brands are protected, they’ve become less confident that they’re reaching their desired audiences. In fact, 69 percent said that using brand safety technologies and techniques, including blacklisting and whitelisting, has created an inability to reach specific audiences. That’s a shocking increase from last year, when the number was 30 percent.

“I think the market has over-corrected,” said Ryan Pauley, chief revenue officer at Vox Media. “But I don’t blame the marketers for it — I blame the tools, systems and ecosystems that got us to the point where there’s too much risk for brands in not over-correcting.”

Pauley thinks emerging technologies might help solve the problem. (Vox Media, for example, has a tool called “Conversational Intelligence,” which uses natural language processing and machine learning to help brands assess tone and context.)

“So far,” said Pauley, “the conversation has been about opting out of the bottom five percent of publishers, instead of about what it means to opt into brand safety. I think as we head into 2019, we should ask, what is the context? There’s content on the page, and there’s someone reading it — what’s their mindset? And what are some more sophisticated tools we can use to find out, and achieve [brand safety]?”



PREVENTION AND TREATMENT

DIRECT RELATIONSHIPS WITH PUBLISHERS

Buying ad space directly from a reputable publisher allows marketers to head off brand safety concerns at the very start of a relationship. Moreover, a direct relationship makes it far easier for brands to make or request adjustments when brand safety concerns arise. Fifty-five percent of respondents say they're turning to direct relationships as a preventative measure, while 54 percent said they've turned to such relationships after having already been exposed to a brand safety threat. Last year the numbers were 39 percent and 34 percent, respectively.

PRESCRIBE FOR:

Direct relationships help maintain overall health and quality of ad buys and allow for quick regimen-adjustments.

SIDE EFFECTS MAY INCLUDE:

Occasional exposure to hard news; premium pricing.

BLACKLISTS

Created by agencies and implemented by demand side platforms, blacklists are designed to exclude specific publisher sites based on the content they distribute. Twenty-nine percent of respondents said they'd implemented blacklists as a preventative measure, and 28 percent said they did so as a corrective measure. But a whopping 62 percent said they found blacklists to be more effective at mitigating brand safety issues than any other practice.

PRESCRIBE FOR:

Preventing brand exposure to hazardous environments.

SIDE EFFECTS MAY INCLUDE:

Diminished audience pool; less accurate targeting.

IMAGE RECOGNITION, AKA COMPUTER VISION

Image recognition is an AI-driven technology that uses neural networks to identify and sort imagery. Among other uses, it can detect brand-unsafe content and automatically keep marketers out of unsafe environments. Yet despite the omnipresence of images and videos across the publisher and social media ecosystem, only 21 percent of respondents are using image recognition technology as a preventative measure, while 16 percent said they'd used it to address brand safety exposures after they'd occurred. That's still an increase over last year, when the numbers were 12 percent and ten percent, respectively.

PRESCRIBE FOR:

Monitoring images and videos for brand-unsafe content; targeting favorable audiences in safe environments.

SIDE EFFECTS MAY INCLUDE:

None known.

PREVENTION AND TREATMENT

NATURAL LANGUAGE CONTEXT DETECTION

Often used in conjunction with image recognition, natural language context detection is the use of natural language processing to detect and avoid brand-unsafe text. Employed as a preventative measure by 21 percent of brands, and to address existing brand safety crises by 17 percent. Last year the numbers were 13 percent and eight percent, respectively.

PRESCRIBE FOR:

Prevention of unwanted adjacencies to hazardous content; supplement to an image recognition regimen.

SIDE EFFECTS MAY INCLUDE:

None known.

THIRD PARTY MEASUREMENT SOLUTIONS


Some respondents pointed to third-party measurement partners like Integral Ad Science and DoubleVerify to keep them protected, with 20 percent saying they turn to such companies to prevent brand safety exposures and 14 percent doing so after an exposure has already occurred. This is a significant drop from last year, when the numbers were 39 percent and 28 percent, respectively.

PRESCRIBE FOR:

Monitoring vital statistics.

SIDE EFFECTS MAY INCLUDE:

None known.





FOR *Media Buyers*

DUE TO *Platform & partners*

PRIMARY *Brand RX*

INNOVATIVE THERAPIES

While restrictive blacklists and whitelists can insulate brands from reaching favorable audiences, it's undeniable that they're addressing an urgent brand safety need. "I think the market has resulted in the need for these blunt force instruments," said Vox Media's Ryan Pauley. "I think marketers are doing what they need to to take back control there." The question is, how can marketers thread the eye of the needle, hitting their targets and protecting their brands at the same time? The answer likely lies in technology.

INNOVATIVE THERAPIES

Computer vision technology and natural language context detection can fuel contextual targeting, and protect brands from hazardous adjacencies, without blacklisting a site altogether. "Computer vision can tell you if an image or video has nudity inside of it, if there's a Swastika in the corner, anything," said Cambron Carter, head of computer vision at the tech provider GumGum. Used in conjunction with natural language context detection — which can perform similar functions with brand-unsafe text — such technologies can steer brands toward or away from specific parts of a site without totally avoiding it.

SIGNED

Image recognition, perhaps the most promising technology on this list, received a boost from last year; it's now used by 21 percent of marketers to prevent brand-safety episodes, up from 12 percent. Moreover, 16 percent of marketers said they'd turned to image recognition solutions after experiencing a brand safety episode, up from 10 percent a year ago.

Those are still very small numbers. Marketers are missing out on a hugely effective tool in the fight against negative adjacencies — one that could lessen their dependence on restrictive blacklists and whitelists.

Today's most effective image recognition technology can analyze and identify objects with an extreme degree of accuracy. That can include brand-unsafe images such as nudity, violence, or even a competitor's logo.

Such tools are especially effective when used in combination with natural language processing, which can detect brand-unsafe text. "It's not all image recognition," explained Cambron Carter, Head of Computer Vision for the tech provider GumGum. "It's multi-modal. If you have another stream of information, there are situations where one fails and the other succeeds."

For instance, an airline wouldn't want its ad to appear next to a news story that contains the phrase "crash" or "terrorism." Nor would it want its ad to appear next to a news story that features a photo of a plane wreckage. Yet it's entirely possible that a news story about a plane crash will contain only brand-unsafe imagery or brand-unsafe text — not both. By using natural language context detection and image recognition in conjunction, brands can cover all their bases.

Natural language context detection has seen an increase in usage very similar to that of image recognition. 21 percent of marketers now use it to prevent brand safety issues, while 16 percent use it to address brand safety issues after they've already occurred. One year ago, the numbers were 13 percent and eight percent, respectively

Image recognition technology and natural language processing tools can also recognize brand-favorable content. For instance, the tools might point a luxury car brand in the direction of a site geared toward auto-enthusiasts.

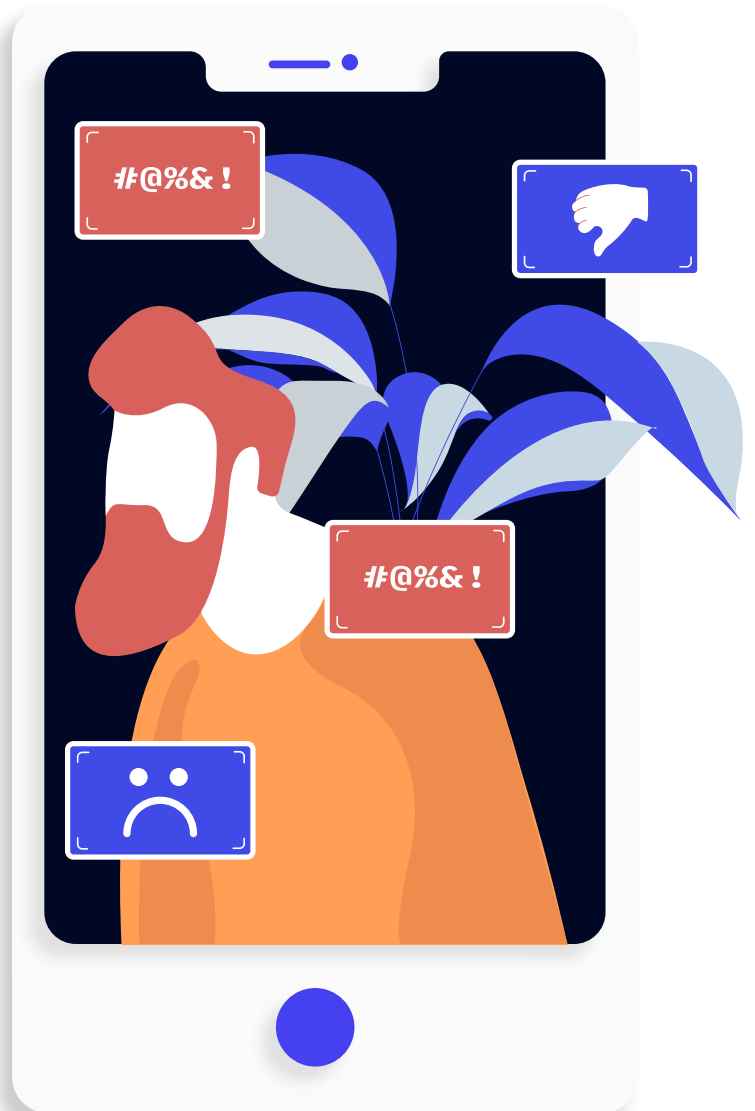
Crucially, today's image recognition tools are integrated into major DSPs and have the ability to comb through thousands of publisher sites in real-time. In many ways, image recognition is a cure while blacklisting and whitelisting are merely palliative measures; it enables marketers to venture into seemingly unsafe publisher territory without fear.

While the number of marketers using image recognition and natural language processing tools has increased only modestly, many more respondents now cite those tools as effective brand safety treatments. Thirty-five percent and 31 percent of respondents, respectively, said that natural language context detection and image recognition are amongst their most effective tools when it comes to mitigating brand safety concerns. A year ago, 18 percent cited image recognition and 13 percent cited natural language context detection.

The results seem strange at first — how can the number of marketers who find these tools effective have increased more dramatically than the number of marketers who've adopted them? There's only one explanation: Many of the respondents who were already using the tools a year ago are now more convinced of their value.

So how do you explain the significant majority of marketers that aren't using such technology? It may merely be the result of an awareness gap or perhaps budget concerns, but either way, it's a missed opportunity.

"It's absolutely true that we need to move beyond a semantic analysis of brand safe content or even what's starting to be called sentiment analysis," said Barone. The brands that have made use of image recognition tools have often seen impressive results.



LINGERING SYMPTOMS AND EMERGING HEALTH HAZARDS

The industry is hardly experiencing full remission. A slew of marketers still say they've found brand-unsafe content — ranging from hate speech to violence and pornography — adjacent to their marketing content within the past year. If that weren't the case, they wouldn't be rushing to use restrictive tools that limit their audience reach and hamper their targeting capabilities.

Moreover, new brand safety threats loom on the horizon. One particularly troubling example is the so-called “deepfake,” an emerging form of visual content that uses AI technology to create genuine-looking phony videos. The technique has already been used, for instance, to superimpose the face of a celebrity onto the body of a pornographic actress.

Some university computer science departments have already developed algorithms that can automatically detect deepfake videos. But according to Siwei Lyu, director of the Computer Vision and Machine Learning Lab at SUNY Albany's College of Engineering and Applied Sciences, those algorithms are mostly effective at identifying shoddier efforts — though such efforts comprise around 99 percent of deepfake videos, Lyu said.


It's the other one percent that digital media pros need to worry about. “I have seen cases where people go into those videos and fix little problems frame by frame,” Lyu said. He estimates that a sophisticated programmer with the right type of computer can already develop such a video in one to two weeks. “That

becomes very difficult for the algorithm to detect,” he said.

So far, deepfake videos have mostly been a rare curiosity. Major companies like Google and Facebook have mostly been effective at detecting and eliminating them. (They've had more staying power in the more outflung reaches of the internet, like pornographic websites.) But they're growing in sophistication, and it remains to be seen whether they ultimately start slipping through the cracks on more mainstream platforms.

DARPA, the technological research division of the U.S. military, is already spending \$68 million to sponsor a research program designed to create more sophisticated deepfake-detection tools. The program is slated to conclude in 2020; at that point, Lyu says, it wouldn't be surprising to see companies of all sorts begin to use those tools commercially.

“I think there is definitely a potential market for detecting fake videos,” said Lyu. Given the grave risk that fake videos could pose to brands, it won't be particularly surprising if ad tech becomes a significant part of that market.





THE ROAD TO FULL RECOVERY

—

Last year, brand safety was an epidemic; now, thanks to the efforts of everyone involved, it's on the road to containment. Social media platforms have taken great pains to enable better measurement and boot controversial accounts while brands have placed more emphasis on tools designed to shield them from dangerous adjacencies — before they happen.

Marketers are cautiously optimistic. By and large, they like what they've seen. And many organizations have begun elevating brand-safety specialists to shore up their immune systems.

But many marketers are struggling to find the balance between reasonable preventive measures and overtreatment. As a result, they're often insulating themselves from potentially receptive audiences. That's an unacceptable side-effect that has endured since last year and

will persist as long as marketers overuse these tactics.

The key lies in the right technology. Gaining an increased understanding of tools like image recognition and natural language context detection will be an integral part of treating the whole problem instead of just a part of it. Marketers need a way to both protect themselves and avoid dangerous contexts and reach their full potential audiences.

They're not there yet. But the cure may be in sight.





gumgum[□]

—
TO LEARN MORE ABOUT GUMGUM

@GumGum; gumgum.com

DIGIDAY | CUSTOM

©DIGIDAY MEDIA 2018