

Over the course of this report, we'll take you on a visual journey from the dark, unseen world inhabited by those without computer vision to a vibrant, open world just waiting to be unlocked.

Unseen

Seen



PART ONE: The visual World

unseen

For marketers, visual content is the primary canvas for telling stories and influencing consumers. Images, videos, and everything between (and beyond) dominate websites, social platforms and apps. Marketers and publishers are loath to post messages without attaching some sort of visual asset.

That's both a blessing and a curse. Visual formats provide marketers with innumerable ways to bring their messages to life. But with images and videos populating every corner of the digital landscape, it's tough for them to ensure their efforts stand out or reach the right people. This visual clutter has left marketers poking around in the dark, struggling to reach their targets or to derive meaningful insights from their campaigns.

But in recent years, a number of new tools have emerged to help marketers take control of their visual strategy. One of those is computer vision, a technology that processes and understands visual information.

Digiday surveyed more than 300 industry pros from agencies, brands and publishers and conducted a series of interviews with marketing and technology experts to reveal the strategies and struggles of media in a visual world and how computer vision is helping marketers see clearly. GumGum and Digiday Media surveyed 379 digital media professionals from July 19 to 22, 2018. Respondents worked for agencies, brands, publishers, social media platforms and technology providers.



Let's start here: 80 percent of media pros said visual content is either "very" or "somewhat" important to their marketing strategy. Absolutely no one said it wasn't important at all.

"Maybe back in the day you would just do a text post on Twitter or Facebook," said Erin Rech, head of digital at the IPG-owned agency Initiative. "But at this point there is nothing we're doing that isn't visual."

Aside from time-tested tools like GIFs and vertical videos, marketers have inevitably turned to more cuttingedge technologies to get an edge. For instance, there's...

- 360-degree video: Videos that let viewers pan through every direction of the recording by moving, touching or clicking on a screen
- Virtual reality: A computergenerated simulation of a threedimensional environment
- Augmented reality: A technology that can overlay digital images onto real-world environments
- Wearables: Devices with digital screens that users wear on their bodies (e.g., smart watches)
- Computer vision: An umbrella term for any technology that allows a computer to process and analyze imagery with a high degree of understanding. Also referred to as image recognition

The first four have clawed their way into the popular consciousness. With prominent products and services like the Oculus Rift, Apple Watch and Pokémon Go asserting themselves in the marketplace (and garnering mountains of media coverage), it's easy to see why. But computer vision, used in products ranging from self-driving cars to medical imaging devices, boasts the widest array of applications. Nonetheless, it's had perhaps the least amount of buzz, especially when it comes to media and marketing. Meanwhile, those other visual media technologies actually rely on computer vision. For instance, Snapchat's augmented reality filters use computer vision to size up the shape and dimensions of faces and bodies.

If you want to dive into the technical details, today's computer vision tools utilize "deep learning," a type of machine learning that can teach computers to detect patterns in images or other data. "Since 2012, deep learning has been working its way through every part of computer vision, especially the branch called 'recognition,'" explained Serge Belongie, a computer vision expert in Cornell University's computer science department. "Recognition deals with identifying faces, vehicles, pedestrians, products on grocery shelves, plants, animals — and it can do a remarkably good job on these tasks."

That ability has also proven remarkably effective at boosting the scope and precision of digital marketing campaigns. At least, it has for the marketers who know what to do with it. But many don't.



operated

platforms

WHY MARKETER ARE EMBRACING CONTENTION TO

WHY MARKETERS CONTEXTUAL TARGETING

unseen

Nobody would question that a robust visual content strategy is a prerequisite to sales and engagement. "Clearly the digital landscape has become very focused around imagery in the last four or five years," said Tom Dunmore, a former technology journalist and cofounder of the U.K. agency Mediablaze. Dunmore estimates that more than three billion images are posted online every day. "That's how people communicate," he said.

As a result, marketers have learned to communicate in the same way.

How are they using visual content? According to our survey:

- 72 percent of marketers post images and videos on their owned-andoperated platforms
- 29 percent say they use videos and images in posts on social media platforms
- 36 percent say they use augmented social media graphics, like lenses and geofilters
- Yet only 19 percent use images or videos in display ads

So marketers overwhelmingly stick to their own platforms and social media accounts instead of relying on display ads. This allows them to maintain control of their own content instead of unleashing it into the wild and facing off against countless competing efforts. It also allows them to control all the technological infrastructure underpinning their efforts, and to customize visual content to the platforms they control.

How digital marketers use visual content



(e.g., lenses or

geofilters)

Other major concerns include ensuring viewability and brand safety when placing programmatic ads. When running ads on outside sites and apps, it's much harder to prevent unfavorable adjacencies, or to reach engaged users who are likely to dwell on an image for more than a split second.

But there's a big problem here: By ignoring many of the platforms where consumers are active, marketers miss out on the ability to expand brand awareness beyond that inner circle of core devotees. Just because consumers aren't active on a brand's sites or social platforms doesn't mean they're inherently unreceptive to the brand's messages.

Here's where computer vision tools can be particularly effective. "It's an ever-evolving space," said Rech. "What [marketers] want to think about is, how is the technology valuable to the problems we're trying to solve?"

Good question.

Just because consumers aren't active on a brand's sites or social platforms doesn't mean they're inherently unreceptive to the brand's messages.



WORKING IN THE DARK WHY MARKETERS ARE

WHY MARKETERS ARE EMBRACING CONTEXTUAL TARGETING

unseen

 \bigcirc

-

When it comes to visual media, marketers deal with a slew of problems on a day-to-day basis. For instance, according to our survey:

- 42 percent of marketers said that insufficient viewability challenged their ability to raise brand awareness
- 48 percent of marketers said that insufficient viewability was a challenge when it came to driving e-commerce
- 41 percent of marketers said that contextual irrelevance was their biggest concern in terms of getting users to engage with content
- 46 percent said contextual irrelevance was their biggest overall concern when it came to generating revenue

For all four issues (viewability, raising brand awareness, getting users to engage with content and generating revenue) about 30 percent of marketers said that ineffective demographic targeting was a major problem. And it's no coincidence that demographic targeting is losing cachet across the digital ad industry, for myriad reasons. The basic census data that marketers rely on has limited value, conveying superficial information like age or gender. Demographics are not destiny; they don't tell marketers what users are interested in, what motivates them or what their biggest concerns are.

The European Union's recently enacted General Data Protection Regulation (GDPR), along with California's new Consumer Privacy Act, are also major impediments to effective demographic targeting. Now facing a dramatically higher threshold for what qualifies as opt-in user consent, marketers have been hindered in their efforts to collect demographic information from a wide range of users.

"We know everything that people can do online," said Sean McInerney, group vice president of technology at ad agency Huge. "It's all digital, and we've got the cookies and the tracking and everything else you do. But GDPR is specifically aimed at limiting that information."

That brings us to contextual targeting. It's not exactly a new technique, but it's made a significant resurgence. Essentially, it's all about placing marketing materials adjacent to content or in environments that are in some way related or complementary. It's a way to capture the attention of users who are already likely to be interested in the products or services being advertised.

Contextual targeting requires little to no demographic information and allows marketers to venture off their own platforms and social accounts with the confidence that they'll actually reach interested consumers.

Hugely effective technological solutions designed to help marketers hone their contextual targeting abilities already exist. So the question is, why did more than 40 percent of marketers tell us they were concerned about contextual irrelevance?

SHINING A LIGHT



unseen

Contextual targeting is nearly impossible to achieve manually. "The challenge is the sheer scale of the images that are out there," said Mediablaze's Dunmore. "You're talking about [scanning through] millions of photos a day, and there are billions out there. It's an amount of processing that makes that passive monitoring very difficult." To understand why computer vision tools are so valuable to marketers, you need to understand how they can address that issue.

Image recognition technology can help comb through millions of existing images and videos, then place contextually relevant marketing adjacent to the right ones in a fraction of the time that humans can do so. For instance, an advertiser marketing a luxury car might HOW COMPUTER VISION PUTS MARKETERS IN CONTEXT

use the technology to run an ad on a website geared toward auto enthusiasts. And it's a safe bet that consumers who are viewing such a site will be more likely to be receptive.

"Computer vision can help you build a visual index of key images that you would want to run ads next to based on what's appearing on screen," explained Rech.

And while Google may have a little ways to go before it fully perfects the self-driving car, the technology has pretty much arrived when it comes to marketing applications. "The performance curves are definitely above threshold for what industries need for that kind of automation to work," said Belongie.

You're talking about [scanning through] millions of photos a day, and there are billions out there. It's an amount of processing that makes that passive monitoring very difficult.

- Tom Dunmore, co-founder, Mediablaze.

Then there's the question of brand safety. It's an issue that's been at the top of most marketers' minds since early last year, when news broke that a slew of YouTube ads had been running against videos from extremist groups. Ask any digital advertiser and they'll tell you that unwanted adjacencies can be catastrophic.

- Only 12 percent of respondents said they already use image recognition in their marketing efforts
- A small but greater number of marketers already use augmented reality (15 percent) and virtual reality (14 percent)
- Only 16 percent said they planned to start using image recognition tools later this year

Automated computer vision tools can go a long way toward addressing the problem, ensuring that images and videos aren't placed alongside the wrong content. "If you say, for instance, 'I don't want to be next to any trucks,' [image recognition] would have a large degree of success with that," explained Huge's McInerney.

Yet technologies like AR and VR still get more notice from marketers.



Which visual media technologies are marketers using?





One issue is that some of the world's biggest tech giants are behind the most heavily publicized AR and VR tools. "Folks will look at Snapchat, Instagram and Facebook," said Jeremy Sigel, global svp of content and innovation at the agency Essence Global. "Those large media owners have made those tools accessible to consumers and for brand experiences."

Far too few marketers understand the sheer breadth of computer vision's applications; perhaps even fewer realize that computer vision powers many of the tools they're already using. By overlooking image recognition technology, marketers miss out on one of the key tools at their disposal to help them see the visual media landscape more clearly — and to make sure their efforts are noticed by the right consumers.

"There's an awareness gap," said Sigel. "The reality is that advertising and brands always lag behind consumers. There are a lot of brands today that are finally getting their bearings on social or on mobile, never mind being able to use image recognition technology to make sales. I'm sure there are a lot of CMOs in corner offices who don't even realize that's possible." While less than a majority of agencies, publishers and brands currently use image recognition in their marketing efforts (only 12 percent), it's worth taking a closer look at how exactly that growing group of adopters is taking advantage of the technology.

As it turns out, they're applying it to a fairly narrow range of uses. Most notably:

- Of the survey respondents who currently use image recognition technology, 59 percent said they use it to detect brand-unsafe content
- 44 percent said they use it to to determine the reach or value of visual branding

Yet only 10 percent said they use it for ad targeting. That's 10 percent of 12 percent. If marketers don't expand the scope of their ambitions, they'll be hard pressed to master the visual world.

If your company uses image recognition technology, how does it do so?





GETTING PAST TUNNEL VISION

HOW MARKETERS ARE USING COMPUTER VISION

unseen

##\$#%!

%

It's easy to see why most marketers who use computer vision are applying it to brand safety. Ever since last year's so-called "ad-pocalypse," when brands from L'Oréal to Verizon were shocked to learn that their ads had been running on racist and extremist YouTube videos, companies have been scrambling to protect their reputations.

Those concerns haven't abated. (In a November 2017 survey, Digiday found that 70 percent of brands, agencies and publishers were were taking brand safety seriously or very seriously.) Yet brands are only increasing the scope of their ambitions, targeting an everincreasing number of platforms, sites and apps. And as advertisers reach into every corner of the digital landscape, it only gets harder for them to keep track of where their marketing efforts wind up.

That's where image recognition comes in. "As there are more and more niche

opportunities to advertise, and as the advertising networks are finding ways into digital properties that either could or could not subscribe to the point of view that advertisers want to have related to their product, there becomes more need of the technology," said Huge's McInerney.

Applications like these can be tailored to different definitions of safety. "Two advertisers may both want to avoid content which contains nudity or violence," explained Cambron Carter, head of computer vision for the tech provider GumGum, which uses image recognition technology for contextual targeting and brand safety. "But one advertiser may consider a particular person or object to be unsafe for their brand whereas another may not." For instance, a vegan food brand likely wouldn't want its ads appearing next to a news story about a barbeque cookoff, while a burger brand might.



Then there's the 44 percent who use image recognition to determine the reach of visual branding. But what does that mean, exactly? Marketers want a sense of just how effective their marketing is, both on and off their own platforms. From broadcasts to social media, computer vision can help determine how viewable a company's branding is.

For instance, brands like GEICO and Endeavor have used the technology to determine how prominently their logos were featured on-screen during sporting events and whether or not those logos were sharing the screen with another brand's logo to manage competitive brand safety. And earlier this year, Goodyear placed its logo on the uniforms of Cleveland Cavaliers players and then used computer vision tools to determine how long-and how prominently-it appeared during game broadcasts. GumGum's Sports unit considered factors such as what percentage of the jersey patch was viewable and how prominently it shared the screen with with other brands, then determined that the logo accounted for \$20.6 million in media value across the Cavaliers' social accounts.

Goodyear also used the technology to determine how frequently its visual branding was shared across the overall social media landscape. Other brands, including Miller Lite, have done the same. With a more precise understanding of reach and viewability, marketers can more accurately assess just how valuable their branding efforts have been. Moreover, brands including Nike and Pabst Blue Ribbon have gone so far as to license widely shared photos that include their brand logos, then feature the photos in usergenerated content campaigns.

"If you're doing a partnership or sponsorship, you usually get earned media off of that through social and media coverage," explained Initiative's Rech. "[Image recognition] helps us evaluate the extra value that we're gaining when a brand is aligned with a platform, a sports team, or whatever... It's a way of quantifying something that used to just be thought of as general added value."

Clearly, these are highly valuable uses for the technology. The real question is, why are so many marketers stopping here?

When you take a deeper look into the numbers, the answer emerges: When deciding when to use computer vision tools, marketers are fixated on the future. What they don't realize is that the "future" is already here. [Image recognition] helps us evaluate the extra value that we're gaining when a brand is aligned with a platform, a sports team, or whatever... It's a way of quantifying something that used to just be thought of as general added value.

- Erin Rech, head of digital, Initiative.



ILLUMINATING MISCONCEPTIONS

COMPUTER VISION'S PRESENT AND FUTURE

seen

Marketers know full well that computer vision is a promising tool. According to our survey:

- Only 6 percent of respondents said it would never be possible to prevent contextually irrelevant ad placements with image recognition
- Only 3 percent said it would never be possible to prevent poorly targeted app and site placements with image recognition
- Only 7 percent said it would never be possible to prevent ads from appearing next to brand-unsafe content using image recognition

audience

Only 9 percent said it would never be possible to use image recognition to avoid proliferating content that fails to resonate with a brand's intended audience

Given how sparsely marketers are actually using the technology, that result seems baffling — until you dive deeper into the numbers:

- Only a small portion of respondents

 between 10 and 14 percent said
 it was already possible to address
 these pain points with computer
 vision.
- A similar number of respondents between 14 and 17 percent – said it would be possible "later this year."
- All other responses ranged from "one year from now" to "more than four years from now."



What problems do marketers think





Marketers are living in a world of misconceptions, waiting for the day when they can use a technology that's already useful – and totally within their reach.

"It seems futuristic," acknowledged Vikram Bhaskaran, Pinterest's head of market development. But Pinterest already uses image recognition to serve contextually relevant ads to users who have searched for specific images. "Visual discovery is already here," said Bhaskaran. "I think it is the quiet revolution that's happening in the background."

Pinterest isn't alone in incorporating image recognition tools into visual search. Last year, for instance, e-commerce platform eBay launched a tool that scans user-generated photos, then suggests products similar to those pictured. And H&M Europe's Image Search feature applies the same basic tools to fashion products, surfacing product suggestions based on users' Instagram photos.

"The visual signals themselves are used as a parameter for ad targeting," said Bhaskaran. "It's a leaned-in audience taking an action that puts them in a bucket that can be targeted."

Some major advertisers have also used image recognition to scan social media posts for targeting. For instance, Coca-Cola's Gold Peak brand recently identified people who posted social media images that featured glasses or jugs of iced tea, then served them ads on 40 apps and mobile sites after they left Facebook, Twitter or Instagram.



Then again, searches and usergenerated photos aren't actually necessary for effective contextual targeting. Consider GumGum's partners alone:

- L'Oréal and Maybelline both used image recognition to serve up contextually relevant ads adjacent to online images related to beauty and personal care
- In 2017, Canon used computer vision to place contextually relevant ads alongside online imagery related to cameras and videography. (The resulting placements more than doubled the average industry benchmark for user engagement.)
- In August, dairy-substitute brand Ripple Foods used computer vision to identify health, fitness, and organic living-related photos across numerous publisher sites, then placed ads on those pages

Many other brands, including Dick's Sporting Goods, Bounty, New Balance, Nike, Jaguar and SeaWorld have all used image recognition to fuel contextual targeting and prevent ads from running on sites and apps where audiences are unlikely to be interested. Ten percent of 12 percent, it turns out, is still a sizable number.

The technology can be directly integrated into major DSPs, giving it massive reach across publisher sites, apps and social sites like YouTube and Facebook. Companies and platforms like Google, AppNexus and MediaMath have all worked with image recognition tools, allowing the software to canvas the entire digital landscape in real-time with an eye-opening degree of control.





There's no denying that computer vision technology is still young. And indeed, there are areas where the technology is far from perfect, requiring a heavy human element to ensure that things don't go off the rails. Serge Belongie, the computer vision expert from Cornell, brought up the example of medical imaging. "If you wanted [computer vision] to recognize melanoma, looking at different grades and stages and different levels of skin pigmentation, you can't do that with [computer vision] alone," he said.

But herein lies the misconception that plagues the advertising industry. Marketers understand that the holy grails of medical imaging and self-

33

driving cars have yet to be perfected. But that understanding prevents them from recognizing that other uses have fully arrived.

"Image recognition is but one of many components in a very complicated symphony of sensors and processing which allow vehicles to navigate autonomously," said GumGum's Carter. "In medical imaging, data can be sparse and there is often disagreement about what constitutes a particular pathology or anatomy." (Though even here, computer vision specialists are making progress; in the next few months, GumGum itself will launch a tool that can recognize diseases and abnormalities in dental X-rays.) "Today's most effective examples of computer vision and machine learning are hyper-focused applications with very specific data and deliverables," Carter added. And while identifying subtle gradations in skin pigmentation may still be a tough task for computer vision tools, the technology is well up to the task of identifying, say, cars or racist signage. So long as marketers provide their tech partners with specific and detailed preferences, it's fairly easy for computer vision tools to target contextually relevant environments or avoid unfortunate adjacencies.

"Today's image recognition technology is powerful where data exists," said Carter. "Marketers can define and refine context based on the summation of simple concepts. That's why this method of targeting is so effective."

Yet according to our research, most marketers aren't using it. And the ones who are tend to ignore some of the technology's most effective functions.

"We have learned how to use AI and computer vision to harness smarter media," said Initiative's Rech. The next step, she said, is "using that to increase our ability to target precisely and efficiently to drive business results."

Today's most effective examples of computer vision and machine learning are hyper-focused applications with very specific data and deliverables,

- Cambron Carter, head of computer vision, GumGum

gumgum

TO LEARN MORE ABOUT GUMGUM

@GumGum; gumgum.com





