



gumgum[®]

ENTERPRISE



APPLIED COMPUTER VISION
IN THE ENTERPRISE



**Harvard
Business
Review**
ANALYTIC SERVICES

Highlights:

60%

OF RESPONDENTS SAY THE REVIEW AND ANALYSES OF IMAGES, VIDEOS, OR OTHER VISUAL DATA WILL BE VERY IMPORTANT TO THEIR ORGANIZATION'S PERFORMANCE IN TWO YEARS.

The research and summary content of this report was produced by Harvard Business Review Analytic Services; and the design was produced by GumGum.

*Design by Andrew Colin Beck
Cover by Daniel Hertzberg*



82%

OF RESPONDENTS SAY THAT THEY WILL BE ASSESSING, PILOTING, IMPLEMENTING, OR HAVE ADOPTED COMPUTER VISION CAPABILITIES WITHIN TWO YEARS.


41%

OF RESPONDENTS LACK A VISUAL DATA STRATEGY.



“IF SOFTWARE IS EATING THE WORLD... AI IS EATING SOFTWARE.”


Ophir Tanz
CEO/FOUNDER



gumgum[®]

Marc Andreessen famously once said that software is eating the world. If he were to revisit the topic today, he might declare that artificial intelligence is eating software. It's not too big a stretch to say that every company that hopes to thrive in the 21st century will need to become an applied AI company, one way or another. Among the most exciting and practicable areas of applied AI is computer vision. It's what my company does—and it's a good business to be in. The computer vision market is growing nearly 50 percent annually and projected to reach \$25 billion by 2023, according to Research and Markets.

“The technology can be found in everything from self-driving cars and drones to medical imaging devices and personal robots.”



It can be applied in the enterprise to reduce errors, increase productivity, combat fraud, improve marketing efficiency, and enhance the customer experience. Yet the vast majority of visual data available in the enterprise goes unused.

We partnered with Harvard Business Review to find out why. This report reveals a gap between computer vision ambitions and reality. Business leaders are awake to AI's potential for radically transforming their industries, but they don't know how to implement it. They don't understand how to collect, store, and manipulate data effectively. Or they're stymied by organizational silos or legacy systems. These might be acceptable excuses if there weren't such a variety of ways to bring computer vision to the enterprise.

Businesses can rely on AI-as-a-Service from major cloud vendors. For example, they can use Amazon Rekognition or Google Cloud Vision to identify and label images. Cloud services

are cost prohibitive at scale and offer few options for customization, so they're best suited to startups with limited computer vision needs or enterprises that have highly specific or controlled internal tasks.

Another option is to contract an AI company specializing in tasks appropriate to a given need. My company, GumGum, has various divisions delivering computer vision solutions to distinct industries. Dental practices, service organizations, or labs can, for example, hire us to automate tasks like x-ray analysis or margin marking for implants.

We're among an ever-growing number of vertical AI companies – like Avvir in construction, Innoviz in autonomous vehicles, and Standard Cognition in retail.


The third option is to build an in-house AI

stack. If a manufacturing company wants the robotic arms in its factories to identify faulty products and remove them from a conveyor belt, the company should develop an in-house computer vision solution. That's because the company already has the necessary deep subject matter expertise and large pools of data needed to train up machine learning model.


“Whichever AI model is selected, it must be implemented with vigor. Half-measures are the same as none at all where AI is concerned.”

It's a major commitment, but getting started is actually not as daunting as it appears. So don't be intimidated. The organizations that implement computer vision solutions first are the ones more likely to last.

Applied Computer Vision in the Enterprise



Applied computer vision is poised to make a major impact on businesses across industries.



However, as artificial intelligence (AI) capabilities have advanced, computer vision has matured to the point where it can enable machines to gauge things mere mortals can't, such as temperature or air quality. By incorporating deep learning, today's computer vision tools could get better at detecting patterns in images or other data over time. What's more, AI makes it possible for machines to process, categorize, and understand images and video at a scale and speed that would be impossible to accomplish otherwise.

Thus far, computer vision remains a largely untapped capability within most enterprises, according to a new survey

of 251 business leaders conducted by Harvard Business Review Analytic Services in November 2018.

However, the results also indicate that respondents expect computer vision to play a significant role in the near future and could lead organizations to make use of the growing volumes of images, videos, and visual data to reduce errors, elevate customer experiences, perform more predictive analytics, increase productivity, and reduce risk and fraud.

Computer vision has been around for more than two decades and is already incorporated into products and services, from self-driving cars and medical imaging devices to in-image advertising and social media platforms to mobile phones and factory automation. While many early applications required significant time, effort, and money to develop, the introduction of increasingly


sophisticated AI-driven algorithms could inspire a vast array of potential use cases across sectors.

“If you have enough data to train a system, you don’t need to devote thousands of man hours to develop algorithms, and you can move much faster,” says Michael J. Palma, research director of enabling technologies for IDC Research.



That could make developing computer vision applications more efficient by orders of magnitude. And the expectation is that that will open the door to many more use cases because it will shorten time-to-market and reduce development costs.”





round a third of respondents (32%) say that the review and analyses of images, videos, or other visual data is extremely important to their organization's performance today, while 29% say it was not very important and 39% rated its importance somewhere in the middle. Looking ahead just two years, the number of respondents that said the review and analyses of images, videos, or other visual data will be extremely important to their organization's performance nearly doubles to 60% with only a handful of respondents saying that it will not be very important.

(See Figure 1.)


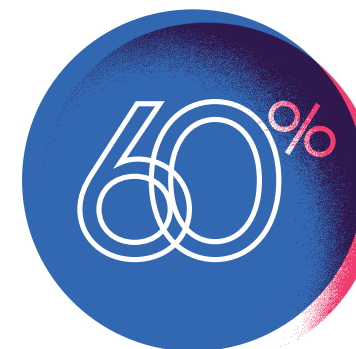


Figure 1:
Percentage of respondents who say the review and analyses of images, videos, or other visual data is very important to their organization's performance (8-10, on a scale of 1-10).

Growing Importance of Computer Vision:




TODAY



IN TWO YEARS


Today, the vast majority of images and video available to businesses—a potential treasure trove of visual data and intelligence—go unseen and unused. That's a missed opportunity. Three out of 10 respondents say they use very little of none of it, just under half (47%) say they utilize or analyze some of it, and just 13% say they are making use of the majority of it. That may be, in part, why the majority of respondents are beginning to at least evaluate the value of computer vision in the enterprise, with 72% of respondents indicating that they are either assessing, piloting, or implementing. Within two years, a quarter of respondents think they will have fully implemented computer vision technologies.

(See Figure 2.) As the number of embedded cameras that make up the world's visual sensor network continues to grow exponentially, computer vision could soon shift from a nice-to-have functionality to a competitive necessity, enabling organizations to mine photos, videos, and other visual data to streamline operations, improve product and service quality, elevate customer experiences, refine their marketing approaches, and more.



Computer Vision Adoption

Figure 2: Percentage of respondents describing their enterprise computer vision adoption today and in two years:



STATUS	TODAY	IN 2 YEARS
Assessing	25%	17%
Piloting	19%	20%
Begun implementation	20%	20%
Have not looked into implementing / Will not implement	18%	4%
Have adopted at scale	8%	25%
Don't know	10%	13%

Computer vision capabilities can also enable companies to harness visual data they've been collecting for years but have been unable to analyze quickly and effectively.

76%


In fact, more than three-quarters (76%) of those respondents who said they are either piloting, beginning to implement or have fully adopted computer vision technologies say they are able to utilize or analyze the majority of their visual data.

Experiences of Early Adopters

Certain industries are out front in terms of computer vision adoption thus far. The greatest percentages of early explorers and adopters, according to the survey, are in the

technology, consulting, health care and life sciences, and manufacturing sectors.


Figure 3: Computer Vision Leaders by Industry. Percentage of respondents in each industry who said they are piloting, beginning to implement, or have adopted computer vision.




67% TECHNOLOGY



53% CONSULTING SERVICES




52% HEALTH CARE / PHARM / LIFE SCIENCES



52% MANUFACTURING



47% BUSINESS AND PROFESSIONAL SERVICES



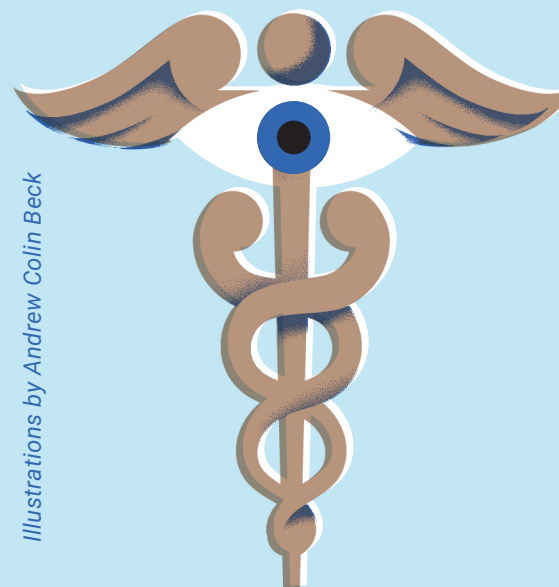
40% GOVERNMENT / NON-PROFIT

Healthcare⁰¹



One of the largest dental service organizations in the U.S. and the University of Pennsylvania Dental School both have plans to begin large-scale application of computer vision products, starting with one that will annotate dental x-rays.

“We can train it to do anything that an expert can,”

“In dentistry, we have a bunch of unstructured data that we’re not using to its full potential, particularly things like photos and digital scans,” says Kyle Stanley, a Beverly Hills, Calif.-based dentist and professor involved in the project. “Dental x-rays are one of the biggest and most important sources of unstructured data that we have.” The computer vision-driven application scans the x-ray for cavities, charts existing restorations, notes the patient’s anatomy, and highlights infection and gum disease. “We can train it to do anything that an expert can,” says Stanley. The goal is to one day correlate that visual data with other structured and unstructured data—perhaps even a patient’s full medical history—to create a more accurate and complete picture of his or her dental health. “That will help to build trust with patients,” says Stanley. “It also helps with medical liability because it means we’re less likely to miss something.”



Illustrations by Andrew Colin Beck



Marketing & Advertising⁰²

Computer vision applications are already helping corporate marketers and agencies see more clearly.

“Both social networks and the web have become a lot more visual,”

says Charlotte Cochrane, a digital marketing consultant who has deployed computer vision applications in her current role as well as her previous one at a marketing agency. “So the ability for companies to read and understand visual sentiment has had to change. That’s what we call visual intelligence.” One beer maker, for example, used computer vision to help it better understand its placements in photos on social media. The application was able to comb through hundreds of thousands of pictures—something that would have taken an inordinate amount of time for human marketers to do—and confirmed a theme. An overwhelming number of the photos included a dog, inspiring the company’s marketers to launch a dog-focused social media campaign. “It resonated,” says Cochrane, “and the results were powerful.” Looking ahead, as companies seek to analyze even more visual marketing data, “computer vision is going to be critical,” she says.

Manufacturing⁰³



Computer vision can help elevate robots from simple, caged-off automatons into more intelligent machines capable of collaborating with their human counterparts. Veo Robotics is applying computer vision, 3D sensing, and AI to create high-performance industrial robots that can safely work side-by-side with humans.

“Our goal was to make interaction between machinery and humans safe, particularly with big robots,”

SAYS VEO ROBOTICS’ PRESIDENT AND CEO PATRICK SOBALVARRO.

“Computer vision was the best technology available for accomplishing this.”

THE TECHNOLOGY ENABLES GREATER PERCEPTION; THE ROBOTS CAN RECOGNIZE AND RESPOND APPROPRIATELY TO OBJECTS AND OBSTACLES IN THEIR PATH, FROM DEBRIS ON THE FLOOR TO A HUMAN IN ITS WAY.




Construction⁰⁴

Computer vision is also being used to solve a long-standing problem in building construction: keeping plans up to date. Rework accounts for as much as 10% of total construction cost, some of which could be


prevented with more frequent inspections and updates.

Start-up Avvir provides automated construction-verification-as-a-service, using laser scanners and computer vision to keep their building information models up to date and providing real-time as-built construction

models. “The primary problem is that construction never goes according to plan,” says Avvir founder and CEO, Raffi Holzer. But it takes an inordinate amount of human effort to continually update building information models. As a result,



owners and contractors often work from inaccurate and out-of-date plans when managing their multi-million dollar projects. Avvir's computer vision-enabled solution models. "Some of our customers were trying to solve this manually, scanning their construction sites and having someone visually compare that to building plans," says Holzer. That's what a large technology company had been doing, paying its general contractors millions of dollars to keep its data center building plans current. Thanks to computer vision, they've cut those costs by 90%. "And once the computer takes over the task," says Holzer, "it can also be more comprehensive, finding mistakes that a human might miss."



INTERNET OF EYES

METHODOLOGY:

A total of 251 respondents drawn from the HBR audience of readers (magazine/newsletter readers, customers, HBR.org users) completed the survey.

PARTICIPANT PROFILE

SIZE OF ORGANIZATION

- 53% 10,000 or more employees
- 39% 1,000- 9,999 employees
- 8% 500-999 employees

SENIORITY

- 14% executive management/board members
- 43% senior management
- 33% middle managers
- 9% other grades

KEY INDUSTRY SECTORS

- 18% Technology
- 11% Manufacturing
- 10% Financial Services
- 8% or less other sectors

JOB FUNCTION

- 13% general/executive management
- 12% R&D/Innovation/Product Development
- 11% IT
- 8% or less other functions

REGIONS



- 50% North America
- 26% Europe
- 13% Asia Pacific
- 6% South/Central America
- 4% Middle East/Africa
- 1% Other

Figures may not add up to 100% due to rounding.

BARRIERS

to

COMPUTER VISION



“**A**


says Steven Kuyan, managing director of New York University’s Tandon School of Engineering’s Future Labs. While those may sound like mundane problems to solve, AI-enabled computer vision can yield breakthroughs by doing work that would be impossible to do with human involvement only. Indeed, there are a number of functional areas of business—across all industries—that could be transformed by computer vision adoption. The functional areas respondents expect will see the most value from applied computer vision are operations and production, marketing, risk management, supply chain, IT, and last but not least, sales.

“**THERE ARE SOLID OPPORTUNITIES FOR R.O.I. IN A LOT OF ENTERPRISE APPLICATIONS,**”

says IDC Research’s Palma. “There’s the headline-grabbing stuff like autonomous vehicles or tracking shoppers in a retail store.

But there are also other applications ranging from improving operational processes to transforming how employees interact with systems to safety and security management.

nything that requires a tremendous amount of human labor is a problem that potentially can be solved with computer vision.”



However, only a fraction of these functions is actually using computer vision today.

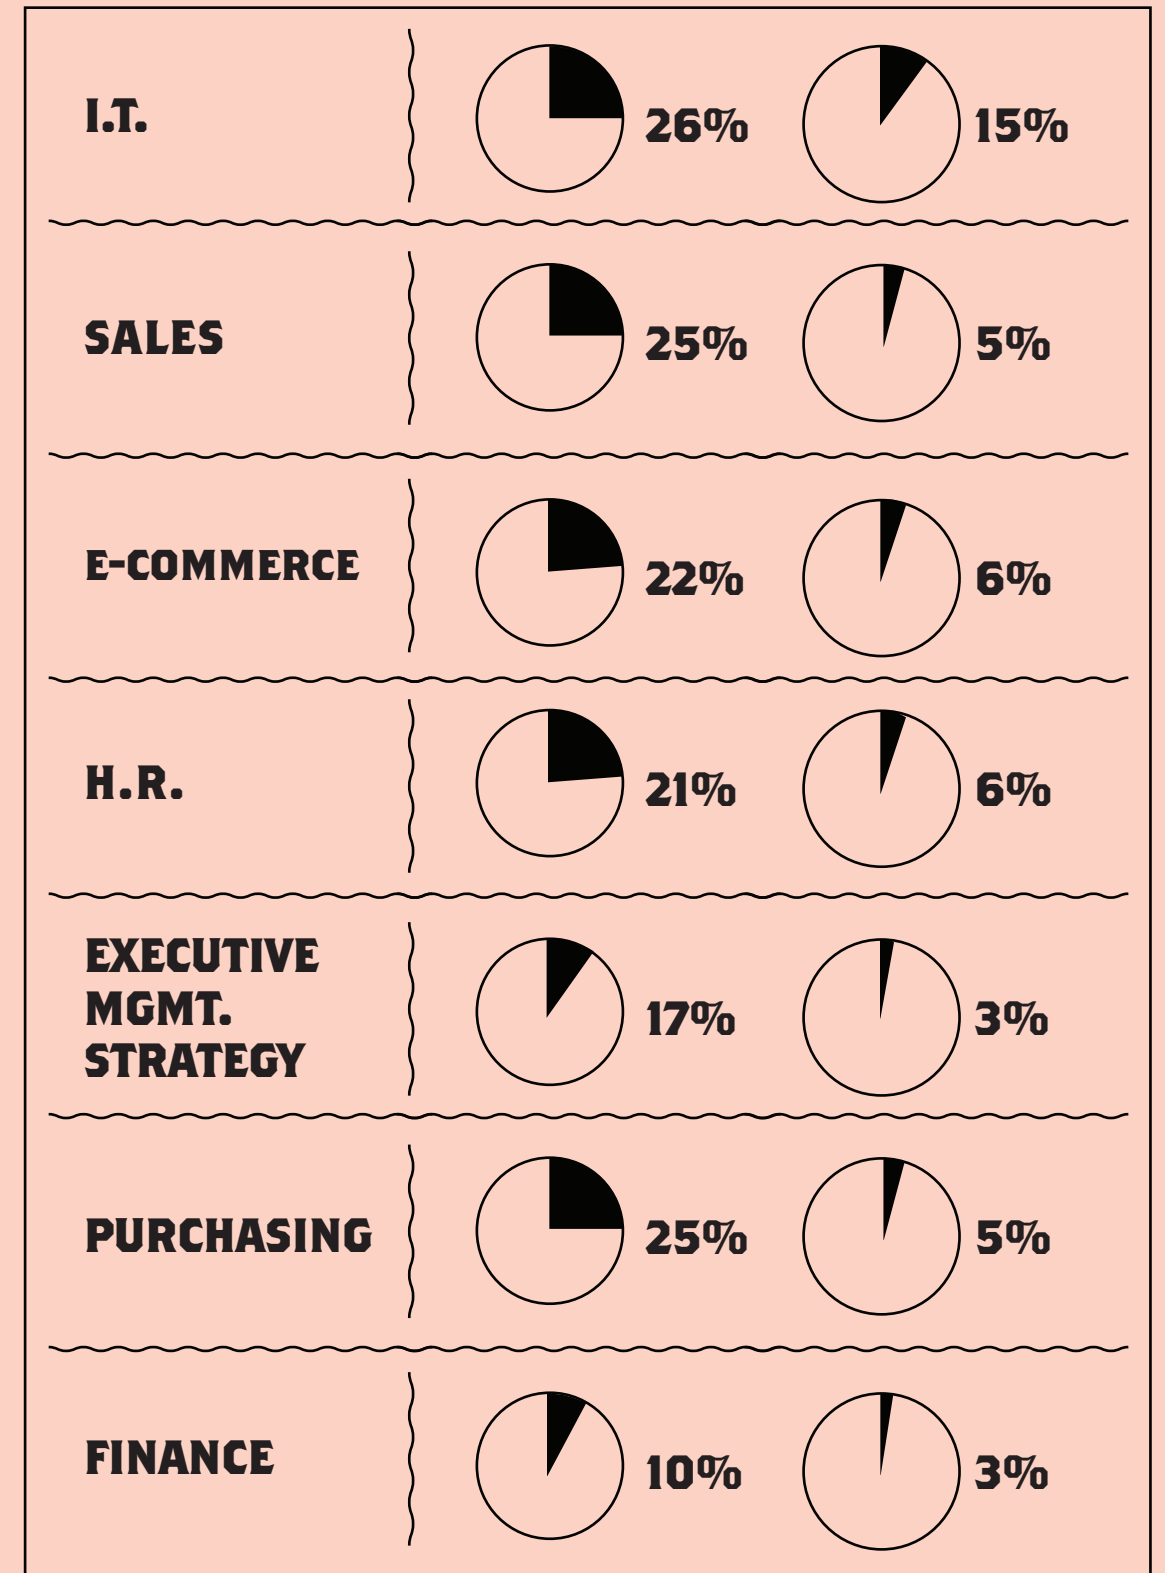
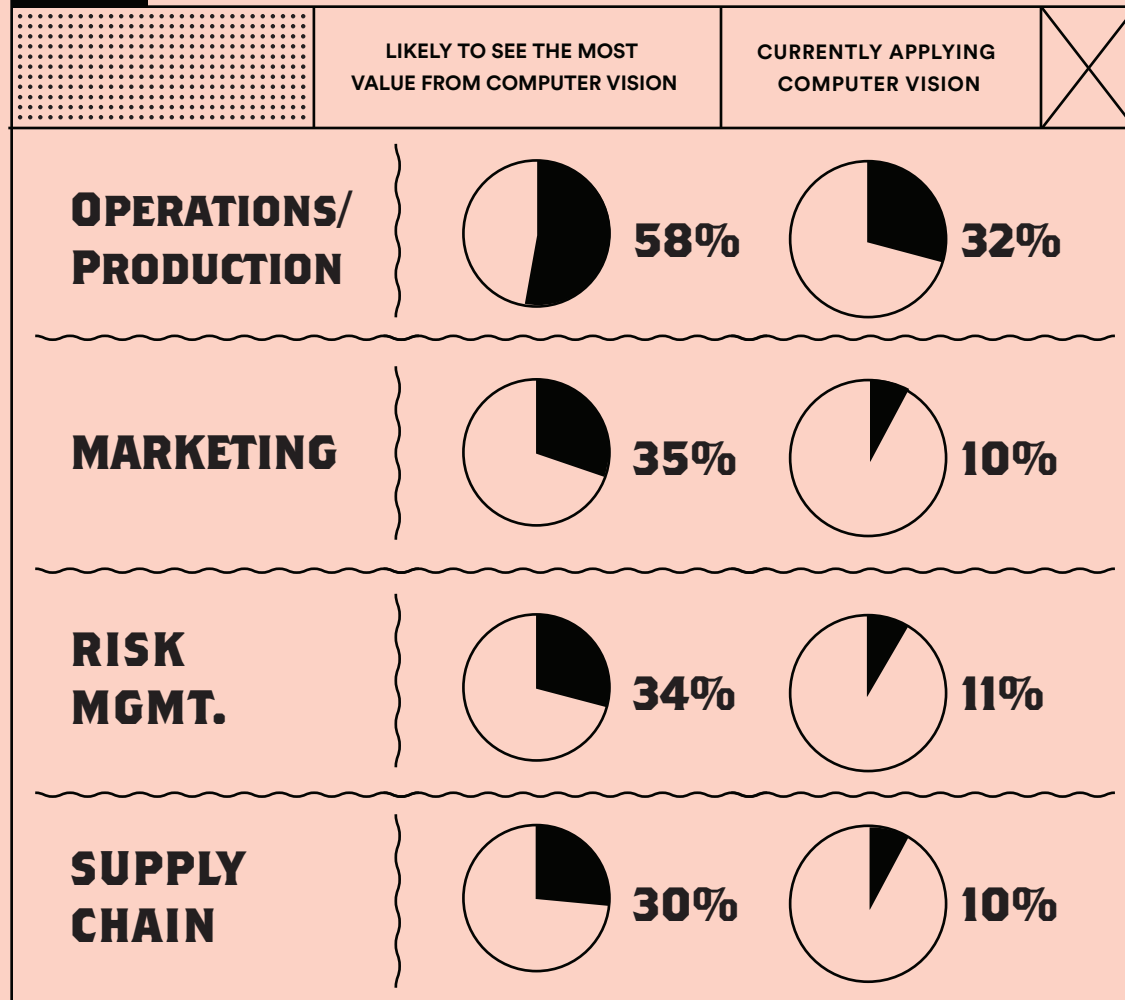
“

E

VERYONE WANTS TO DEPLOY IT,” SAYS NYU’S KUYAN, “BUT HOWEVER, ONLY A FRACTION OF THESE FUNCTIONS IS ACTUALLY USING COMPUTER VISION TODAY. NO ONE KNOWS HOW TO DO IT.”

FIG. 4

The Expectation-Application Gap
Percentage of respondents who said each function was likely to see the most value from computer vision versus the percentage who said that function was currently using computer vision.

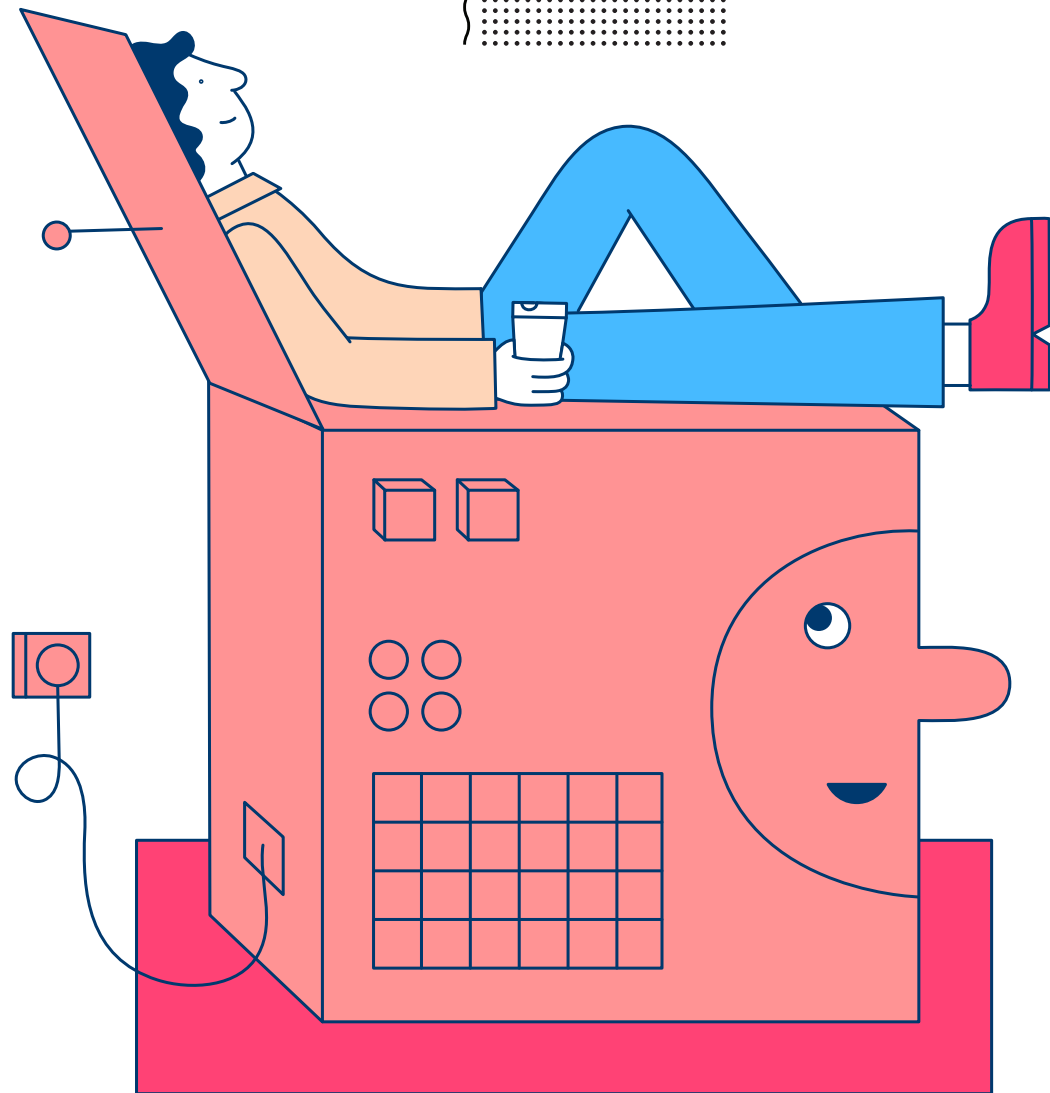


BIGGEST

HURDLES

to


ADOPTION



42%

OF RESPONDENTS SAY LACK OF UNDERSTANDING OF COMPUTER VISION CAPABILITIES AND TECHNOLOGIES WAS THEIR BIGGEST HURDLE.

Others included lack of strategy, organizational siloes, data issues, legacy systems and architecture, and a dearth of computer vision talent. (See Figure 5.) “It all starts with lack of education about the subject. In the digital era, there are so many things happening at once with so many buzzwords being thrown around,” says Cochrane, the digital marketing consultant. “Understanding what computer vision is and how it can be applied is step one.”



INTERNET of EYES

The biggest obstacle to greater computer vision adoption is a fundamental one: lack of knowledge—about how the technology works, how it might transform functions or businesses, and the technological and data foundation required for it to work.

FIG.5

BIGGEST HURDLES TO ADOPTION THE TOP SIX OBSTACLES TO COMPUTER VISION ADOPTIONS, ACCORDING TO RESPONDENTS *(next page)*

.01	Lack of understanding of computer vision capabilities and technologies	
.02	Lack of strategy	
.03	Organizational silos	
.04	Planning data collection, storage, security, and overall infrastructure	
.05	Legacy systems/architecture	
.06	LACK OF COMPUTER VISION SKILLS	

10100101011010100101
0110110100000101101
01010111110101010100
00010101101010101111
001010010001010010

28

AAA *Another significant hurdle is getting access to the massive volumes of data necessary to train a computer vision application—a task that currently has to be undertaken for each specific application of the technology.*

“A BIG CHALLENGE FOR EVERY COMPANY IS GETTING ENOUGH DATA TO CREATE A MODEL THAT WILL BE ACCURATE ENOUGH TO HAVE IMPACT,”

Indeed, “the whole notion of labeling data so it’s a much harder process than most companies can imagine,” NYU’s Kuyan says. “Getting to useful data is important.”

A Way Forward for Computer Vision

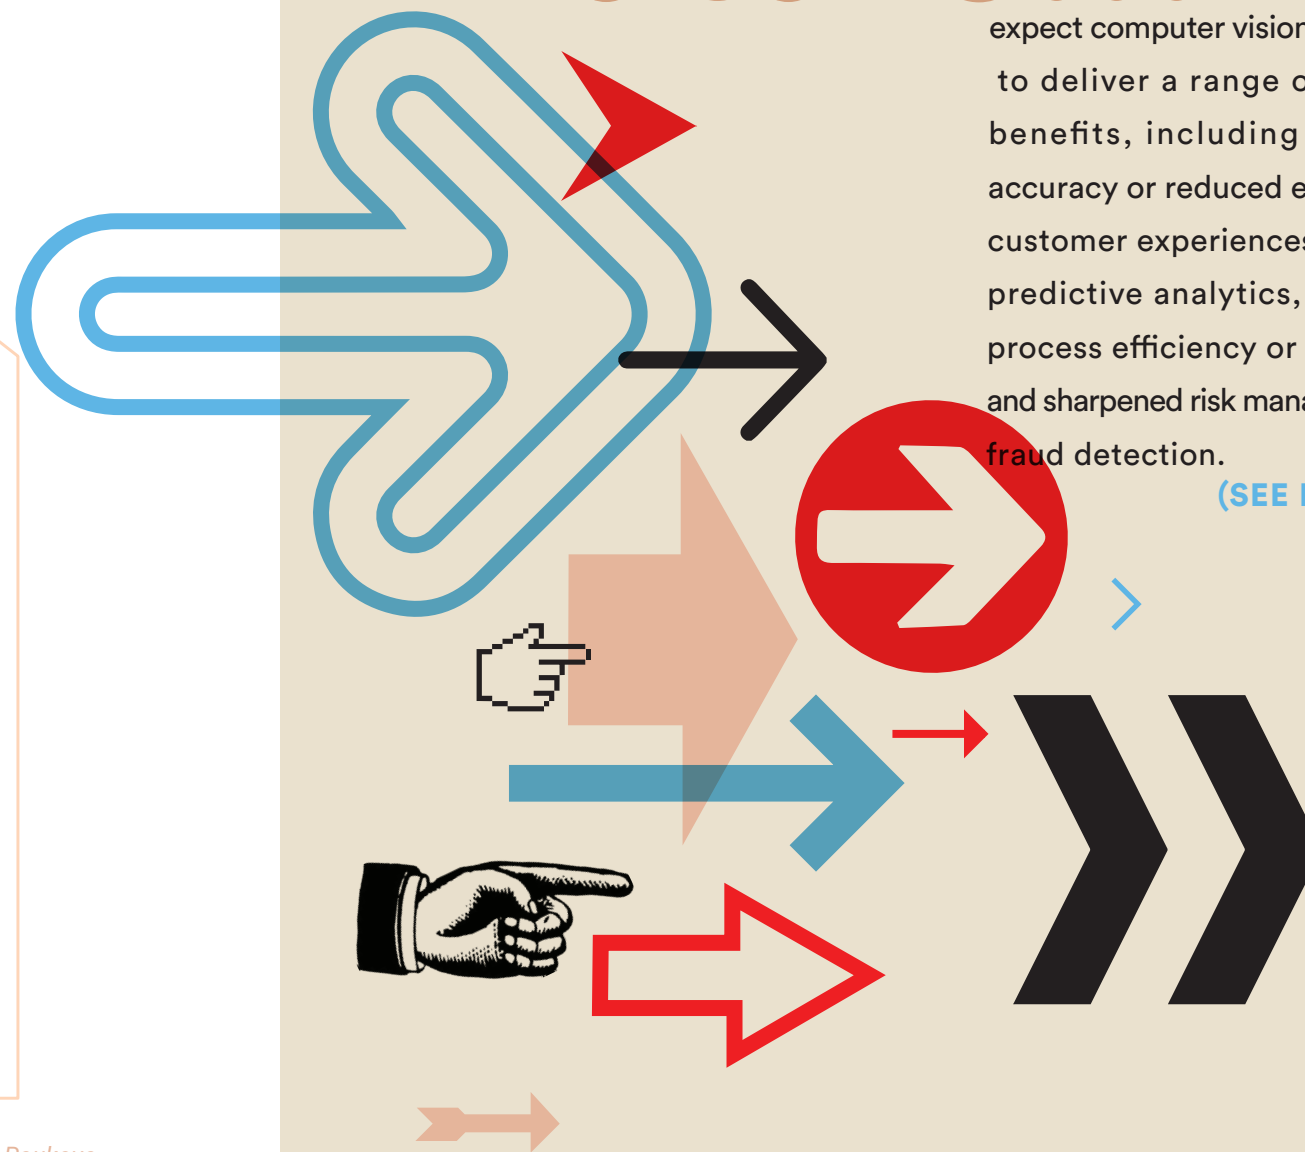


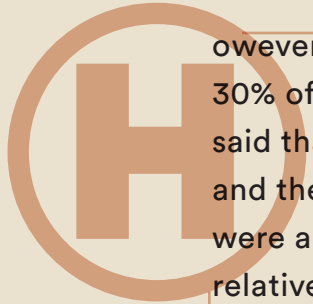
→ Illustrations by Martina Paukova

Looking into the Future...

There is significant impetus to address those issues standing in the way of greater adoption of computer vision. Respondents expect computer vision capabilities to deliver a range of business benefits, including improved accuracy or reduced errors, better customer experiences, enhanced predictive analytics, increased process efficiency or productivity, and sharpened risk management and fraud detection.

(SEE FIGURE 6)






However, today, just 30% of respondents said that visual data and their analyses were a high priority relative to other priorities within their organization's overall digital strategy, 21% described it as a very low priority, and around half rated it somewhere in the middle. What's more, four out of 10 respondents (41%) said they have no visual data strategy.


(fig6)



(fig7) THE STATE OF VISUAL DATA STRATEGY



Companies that want to take advantage of computer vision in the near future will need to develop a greater understanding of the capabilities. Looking at successful applications outside of one's industry or function to get a better sense of what has been accomplished can be a great place to start. "It's important to understand what it means and how it can be applied across industries, not just your own, and start connecting the dots," says Cochrane.



→ Any use case should begin with the business problem that needs to be solved, not the technology itself. “It can’t be computer vision for the sake of computer vision,” says Cochrane. While taking a strategic view of computer vision capabilities is important, having a standalone visual data or computer vision strategy or team is not advisable. “Don’t silo it,”

“IT’S A MISTAKE TO HAVE A VISUAL DATA STRATEGY ON ITS OWN OR HAVE A COMPUTER VISION TEAM ON THE SIDE. YOU NEED TO INTEGRATE IT. OTHERWISE IT WILL NOT SEE ITS FULL POTENTIAL.”

says Cochrane. Some companies may decide to invest in

customer-facing applications while others will see more benefits from back-office use cases or operational solutions. “The right problems are going to be different from organization to organization,” says Holzer.

→ Attracting talent that understands computer vision applications—or partnering with companies that can—is critical. “You have to have the right team in place that understands how to do this,” says Stanley. Companies must also be thoughtful about



their computer vision use cases, taking the time to determine where it will have the most value for the business. “Opportunities abound, and they’re very disperse,” says Palma. “You need to know what’s most important to you, where the greatest reward or returns will be, and focus the investment there.”


→ Once clear computer vision use cases have been identified, companies can then begin to assess their software options, data requirements, infrastructure needs, and—perhaps even most importantly—determine what business processes will need to change to apply the resulting visual intelligence. “It’s important to consider not just how you’re going to feed the data into computer vision,” says Palma, “but how you’re going to use it on the other side.” Then it’s time to test the application, learn from those that don’t yet deliver results, and scale those that do.

→ There is significant work ahead if companies want to implement computer vision effectively. “It will take

investment, knowledge, experience, and a lot of experimentation,” says Palma. Companies that develop a better comprehension of computer vision and its potential applications in the enterprise, create a roadmaps for use case identification, testing, and learning, and address other impediments to implementation stand to reap significant benefits.

“HARVESTING, PROCESSING, AND ANALYZING VISUAL DATA IS ONLY GOING TO BECOME IMPORTANT, AND NOW IS A GOOD TIME TO BEGIN TO EXPLORE IT,”

says Palma. Those that don’t are likely to miss some significant opportunities. “There are always early adopter, mainstream users, and laggards with any new technology,” he says. “But this is one of those areas where if you lag, there will be consequences.”





gumgum

www.GumGum.com/InternetofEyes