accenture

ACCENTURE FEDERAL SERVICES

HOW COMPUTER VISION WILL HELP GOVERNMENT SEE FURTHER

WHEN IT COMES TO IMAGERY, volume is both a blessing and a curse. Thanks to satellites, drones, dash cams and body worn cameras, and IP-enabled monitoring and surveillance, government and commercial enterprises are enjoying a boom time in terms of both still images and video footage. New York City alone sought to roll out 18,000 body worn cams last year, while sales of IP video cameras are growing at 20 percent a year. For federal agencies, visual data delivers situational awareness, early warning, identity information, operational insight and a host of other valuable inputs.

At the same time, this sheer avalanche of information threatens to overwhelm traditional analytics approaches. As the intelligence community already knows, and others in government are fast learning, it's simply beyond the ability of humans to watch and evaluate this nonstop flood of visual data. The human eye cannot pick up every subtle cue; human attention cannot remain fully focused and on task for as many hours as the job requires; nor is the present analyst workforce adequate to keep pace with the escalating demand. Something better is needed to assist humans, to help them focus on the moments and images requiring their scrutiny.

A sub-specialty within the broader field of Artificial Intelligence (AI), "computer vision" promises to cut through the clutter. With machines trained to understand visual images in the same way as humans do, government agencies could better leverage their investment in visual data while streamlining their own processes and freeing human talent for higher-level work.

While computer vision has been evolving for some time, it has lately come to the fore thanks to low-cost computing, an increase in training data and higher processing power. Here we will give a high-level overview of the technology behind computer vision. We'll look at how computer vision is being deployed in the private sector, consider emerging federal use cases, and chart a path forward for government agencies looking to make the most of this fast emerging capability.

TECHNOLOGY REVIEW

Broadly speaking, computer vision is all about enabling machines to "see" images as humans do, to process and identify objects and even activities within a still frame or a video sequence. Ideally, the computer – empowered by AI – would understand what it is seeing; assess the situation and identify suspect or out of the ordinary sequences; and advise users of needed actions.

Broadly speaking, computer vision is all about enabling machines to "see" images as humans do.

The specifics around this include capabilities like object detection, which combines object classification ("it's a truck") or object identification ("it's a Volkswagen") with object location (where is the object in the image). Other techniques include motion analysis, to determine where and how objects are moving in the video; 3D scene reconstruction, to create 3D models based on one or more 2D images acquired from different viewpoints around the scene; and image segmentation, through which algorithms partition the image into semantic regions (e.g. cars, pedestrians, buildings, roads, background, etc.). A capability receiving increased attention is object tracking, in which the computer picks up an entity and follows it through the scene, across multiple cameras.

To support all this, you traditionally needed sophisticated hardware: Advanced cameras, lenses, sensors, and processing chips to glean the visual input. However, the cost and complexity of getting started today is dropping dramatically. Then the AI takes it deeper. Numerous complex algorithms come into play, processed through neural networks that attempt to mimic some of the functions of the brain's visual system.

Under the hood, a convolutional neural network parses data though multiple layers of kernels with increasing complexity. The key is that the actual numeric values of these kernels are learned from the data during the training process, as opposed to being manually engineered. Recurrent neural networks help machines to learn data causality given an ordered set of images by keeping track of the content of previously seen images. This enables computer vision systems to exploit temporal context and learn about object actions, behaviors, trajectories and other temporal patterns.

With so many moving parts, it isn't surprising that we see a broad range of industry players engaged in the development of computer vision.

Companies like Qualcomm are working the optics and camera hardware side, adding depth perception and other interpretive tools to make cameras themselves more intelligent and capable. On the processing side, Nvidia, Amazon, Microsoft, Facebook, and Google all are pushing the algorithmic envelope. Academia is in the game too, with the Rochester Institute of Technology among the leading players. Stanford University, the University of Toronto, MIT, Berkeley and the University of Montreal also boast strong programs.

The University of Michigan also is deeply engaged, with computer vision a natural outgrowth of the school's efforts to untangle the complexities around self-driving cars, which will depend heavily on computer vision for their safe operations. The electrical engineering lab there is developing algorithms that will enable machines to perform real-world visual tasks such as autonomous navigation, visual surveillance, and content-based image and video indexing.

Another capability of computer vision systems, and an advantage with respect to the human visual system, is that they can operate on different image modalities besides the visible spectrum. Algorithms can perform object recognition in the infra-red bands, or in other modalities like radar-based imaging, like synthetic aperture radar (SAR). This, for instance, allows them to process data to detect objects at night time or "see" through clouds.

These nuts and bolts help us to understand how computer vision might work: Smart processing enables computers to ingest and interpret complex data from sophisticated visual sensors. That's how it works. But what can it do? Here's where the fun starts.

SEEING THE ACTION

Computer vision is already hard at work. AI-enabled software is keeping us safe on the streets, policing the Internet, and simplifying commercial transactions.

In public safety, law enforcement dash cameras are supported by smart software that can read license plates on moving vehicles, looking for stolen cars or supporting Amber Alerts.

On the Internet, YouTube uses algorithms to scan millions of hours of video content in search of inappropriate language, nudity and copyright violations. Automation and AI are vital to this effort, which would certainly swamp even an army of human observers.

In the financial sector, Spanish bank BBVA allows new customers to open an account by uploading a photo of their ID and taking a selfie. Computer vision analyzes the photos to confirm identity. In healthcare, Gauss Surgical offers a real-time blood monitor that analyzes pictures of surgical sponges in order to track blood loss during surgery.

Accenture has also worked with a major insurer to use computer vision and other aspects of AI to automate vehicle damage claims processing. When receiving photos of a damaged car, the insurer can automatically detect the level of damage and use it to, for example, order spare parts and possibly detect fraud. As a result, this company is realizing 90 percent accuracy in automated claims analysis, reducing travel requirements of adjustors, and using an intelligent claims processing system that allows agents to spend their time on the subset of claims where their expertise is needed. Now the client is expanding this capability to other lines of insurance such as home and valuable personal property.

On the commercial side, Amazon has integrated a computer vision element into its prototype Amazon Go stores. By leveraging computer vision, sensors and other technologies, the retailer enables customers to exit the store without stopping at a cash register, with purchases charged automatically to their Amazon accounts.

Accenture recently worked with a large engineering services firm to use computer vision to optimize how they manage and maintain production equipment. By actively analyzing the operations of their maintenance facilities, they were able to streamline the process and double asset utilization across a $35 billion inventory.

Smart software can read license plates on moving vehicles, looking for stolen cars or supporting Amber Alerts.

It's not hard to imagine how these emerging use cases could impact federal agencies on both the military and civilian sides.

Take for instance the field worker documenting a construction site, conducting a restaurant inspection, or charting a family's domestic conditions. Computer vision can be trained to look for change states, documenting the evolution of a visual landscape over time and alerting a human analyst when things don't look right.

For example, Accenture is working with a European government agency to create an automated land registry. By comparing property records with aerial images, they can detect discrepancies and non-compliance.

This same capacity might be brought to bear in efforts to assess and remediate in the wake of natural disasters. A change-state assessment might enable computers to tell emergency responders at a glance where the water is deepest and how fast it is rising. Computer vision might likewise be able to determine which roads and bridges are still passable, or how many buildings have collapsed.

It's easy to see how the GSA might be interested in this, with over 9,600 owned or leased buildings under its watch. Together with the Federal Protective Service, they must ensure that security and well-being of more than one million tenants and visitors every day. It may be that computer vision, when layered on top of existing security-camera infrastructure, could improve both physical security and facility operations.

Backed by AI, computers can monitor a facility 24/7 and can be taught to recognize signs of trouble.

Perhaps the most talked-about role of computer vision is in support of self-driving cars, which will need to have eyes on the road – human or otherwise – if they are to navigate safely. Government has skin in the game, with responsibility over highway design and wayfinding. Transportation officials will need to incorporate an understanding of computer vision as they seek to make the roads safe for tomorrow's auto-piloted vehicles.

As computer vision becomes more sophisticated, it could serve an ever-expanding role in national security.

We're already seeing iterations of computer vision in use by the military as an important competent of C4ISR (Command, Control, Communications, Computer, Intelligence, Surveillance, and Reconnaissance). Computers can scan endless hours of drone and satellite footage, highlighting changes, discrepancies and suspicious activities. It's an invaluable service: Unlike human analysts, the computers don't get drowsy or distracted.

The military has shown early interest in taking such capabilities even further through the application of automation protocols. In late 2017, for instance, a team of researchers from the Air Force Research Laboratory won the Large-Scale Movie Description Challenge at the 2017 International Conference on Computer Vision in Venice, Italy. In that challenge, researchers deployed AI-driven systems to automatically create simple written descriptions of short clips taken from commercial film footage.

From this descriptive capability, the military would like to expand to a predictive capability. As a kid in school you may have seen a sequence of pictures: A child has a ball, child throws the ball, dog chases the ball. What happens next? (Dog brings it back.)

Kids can do this without having to think about it. Computers can't yet reason – but if they could, the military repercussions might be significant. Computer vision at that level could change the game dramatically, taking AI beyond a surveillance function to provide more enhanced situational awareness. What will it take to get there? Let's consider some important next steps for government.

ROAD TO COMPUTER VISION

As with any such evolution, it helps to start by thinking about infrastructure.

In computer vision, government IT leaders have a head start (at least for non-classified workloads), insofar as all the major cloud providers have tools in place to support such an implementation. For example, AWS, Google, IBM and Microsoft have APIs to make ingesting your visual data affordable and scalable, freeing IT leaders to focus on the business value of the effort.

Internally, there may still be some technical hurdles to cross. The cloud can eat all the video you want to feed it, but you'll need sufficient internal resources to take in what computer vision has to offer in return. These (possibly sizable) outputs will require scalable systems, as well as the ability to integrate the newly generated intel into your existing systems. What processes will you need to put in place in order to ingest, store and ultimately utilize the new information?

These are the technical issues. At a deeper level, IT leaders and business line owners need to go beyond the mechanical details, to think about the AI itself and consider what it is they are asking computer vision to do, and why.

First there's the question of return on investment: What will be the actual business benefits of automation and will they justify the cost? If an agency already has video in its toolkit, what will be the practical benefits of leveraging AI against that data? If you aren't collecting video already, will the outputs of computer vision justify the expense of a video deployment? Few can afford to do video for its own sake; you'll need to build the business case.

Then there are issues around personnel and processes. Who's to take on this new task? In all likelihood, responsibility for computer vision won't fall to the IT leaders alone. In addition to business line leaders, who will need to make the operational case, computer vision will require the active participation of subject matter experts.

If you're going to be reviewing images of a domestic situation, especially in jurisdictions that require advance permission prior to being videotaped, the proper person to train the AI is the social worker who visits that home, as this person will best understand the complexity and potential risk involved. This holds true with virtually any potential use of computer vision: The human expert has to be at the center of the situation. If we want the computers to see what we see, that information can be conveyed most effectively by those who are most experienced at looking.

Security considerations also factor in. Do the right analysts have access to the right intel? Are the outputs of computer vision being appropriately classified? If you're handling secret or top-secret imagery, does that put added constraints on your processing capacity? Similar restrictions exist around the use of personally-identifiable information (PII), especially for minors.

What this points to is the need for responsible AI policies within agencies employing computer vision. These policies should promote the ethical, transparent and accountable use of the technology. For example, the potential for bias in facial recognition technology is receiving fair scrutiny. Furthermore, there is limited consensus about the public use of computer vision to identify individuals without their consent.

Overall, experts advise patience in any initial foray. It isn't hard to teach a computer to know what a tree looks like, generally speaking. But there are many different kinds of trees and they will look wildly different in still photos and in videos, in sunlight and shadow, on the plains or on a mountainside. Even learning at the accelerated speeds made possible by AI, it will take the computers some time to go beyond the general, to be able to see with the kind of detailed specificity that will unlock the full potential of computer vision.

Regardless of use case, this will be a moving target. Learning is evolutionary. We, as humans, keep getting better the more we do a task, and we never stop learning. Machines should also have these capabilities. Agencies likely will need to periodically adjust the training models for computer vision as new targets of interest emerge or new business goals come into focus, without forgetting what has already been learned. Even once it's up and running, computer vision will require ongoing care and feeding.

LOOKING AHEAD

Market analysts say there's considerable economic energy behind the development of computer vision. Researchers at MarketsandMarkets say the computer vision market will top $17.3 billion by 2023, while Grand View Research puts the figure at $18.2 billion by 2025.

Federal agencies can begin to embrace the new capabilities by starting with what they have on hand. They can use computer vision to index, highlight and annotate existing images in order to make them searchable; in doing so, they'll be training the AI on how to read and interpret relevant future imagery. In practical terms, object detection and change detection will likely be the first applications of computer vision. They're achievable even with present-day technology.

For those who can successfully cross the threshold and implement video analytics, the next few years promise to bring significant advances.

In competitions, computer vision can recognize and classify objects with a 3 percent error rate, already better than the typical 5 percent error rate among human analysts. Especially for mission-critical uses, the technology will be getting even more precise. The engineers also are looking to further spoof-proof computer vision, training the machines to distinguish a fake image from the real thing.

All this fine tuning will help to enhance what is already a powerful capability. With its speed, accuracy and persistence, computer vision promises to lighten the load on human analysts while producing higher-quality intelligence.

COMPUTER VISION: A GLOBAL VIEW

Around the world, governments are turning to computer vision as a way to increase public safety and enhance citizen service.

The Singapore Government turned to video to better understand, prepare for and react to future events. Singapore is using a video analytics service platform to apply computer vision and predictive analytics to surveillance video feeds for six government agencies. It reports greater than 80 percent accuracy using just five video cameras in one of Singapore's busiest metro interchange stations.

In the United Kingdom, Accenture has worked with the London Borough of Newham to demonstrate uses of computer vision in public safety. The project delivered an interactive cartographic view of the borough to serve as the front-end to an advanced video analytics and automated alerts solution. Interconnected cameras helped to identify crowd build-up and supported suspicious-object detection in real time.

As governments look for innovative ways to ensure citizen safety, technologies such computer vision make play an integral role in the digital arsenal.

Contact the authors

BRYAN RICH
Managing Director
Applied Intelligence Lead
Accenture Federal Services

/in/bryan-rich-3632471b/

IAN MCCULLOH, PHD
Managing Director
Chief Data Scientist
Accenture Federal Services

/in/mcculloh/

Contributors

MARC BOSCH RUIZ, PhD.

JOHN CONLEY

DAVID NATION

NILANJAN SENGUPTA

MAGGIE SMITH

DAVID STEPP

About Accenture

Accenture is a leading global professional services company, providing a broad range of services and solutions in strategy, consulting, digital, technology and operations. Combining unmatched experience and specialized skills across more than 40 industries and all business functions—underpinned by the world's largest delivery network—Accenture works at the intersection of business and technology to help clients improve their performance and create sustainable value for their stakeholders. With 477,000 people serving clients in more than 120 countries, Accenture drives innovation to improve the way the world works and lives. Visit us at www.accenture.com.

About Accenture Federal Services

Accenture Federal Services, a wholly owned subsidiary of Accenture LLP, is a U.S. company with offices in Arlington, Virginia. Accenture's federal business has served every cabinet-level department and 30 of the largest federal organizations. Accenture Federal Services transforms bold ideas into breakthrough outcomes for clients at defense, intelligence, public safety, civilian and military health organizations. Learn more at www.accenturefederal.com.