Ambient.ai

Entering the Era of Computer Vision Intelligence:

How Modern, Privacy-Designed AI is Transforming Safety & Security Operations

Late evening, an empty campus building. Short, windy storms scattered throughout the day. An alarm is triggered in a security operations center (SOC). Something suspicious is detected in the north-west part of the building. While a security officer is dispatched to the location of the alarm, the actual tragedy happens: someone collapses from a seizure on the opposite side of the building. Nobody was there to help. And the alarm? It was a false alarm from a damaged tree, falling across a building entry door due to heavy winds.

How many false alarms have you responded to throughout your career? Probably more than you can count.

According to recent surveys, physical security incidents have increased during the pandemic. Active shooter incidents in the United States, tracked annually by the FBI, have also increased. In 2019, there were 28 active shootings, whereas in 2020 the number jumped up to 40. This is the highest number in the last 20 years. The risk of confusion between real danger and false alarms has never been higher.

The Impact of Covid-19

Physical Security Incidents, 2021

Do you believe there will be more, fewer or the same number of physical security incidents in the year ahead?

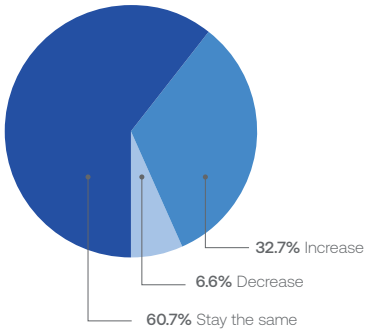

Source: *Provigil Survey, 2020*

Active Shooter Incidents, 2016-2020

Aggregate data from the United States compiled by the FBI on an annual basis.

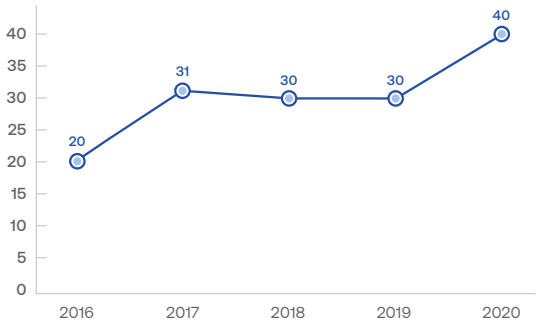

Source: *U.S. Department of Justice, FBI, Active Shooter Incidents in the United States, 2020.*

If there are thousands of alarms that go off every day, it's just impossible to check all the surveillance feeds, even with video verification. Wouldn't it be better to have an "AI brain" for the physical security operations center? Software that can constantly monitor every single camera and alarm endpoints, and send notifications *only when there is a situation that genuinely calls for it?*

We invite you to explore the evolution of AI technology and what modern computer vision intelligence platforms are capable of. Each level of technology builds on those that came before it. Each generation of researchers and engineers stand on the shoulders of the giants. Thanks to their hard work, we are now able to significantly minimize false alarms while preventing real harm.

The Boy Who Cried Wolf

The First Level of AI in Physical Security

Artificial Intelligence (AI) technology has experienced several serious transformations over the years. As AI started to take shape as a field in the 1960s, researchers started to pursue using computers to mimic human vision systems. At first researchers tackled still images. Computers, just like animals, "see" the world differently from humans. They count the number of pixels in an image and try to discern borders between objects by measuring shades of color and estimating the relative distance between objects.

As the first video analytics systems were introduced, AI for physical security mostly meant motion based detection. Algorithms were trained to compare pixel changes scene by scene in moving video in order to assess motion.

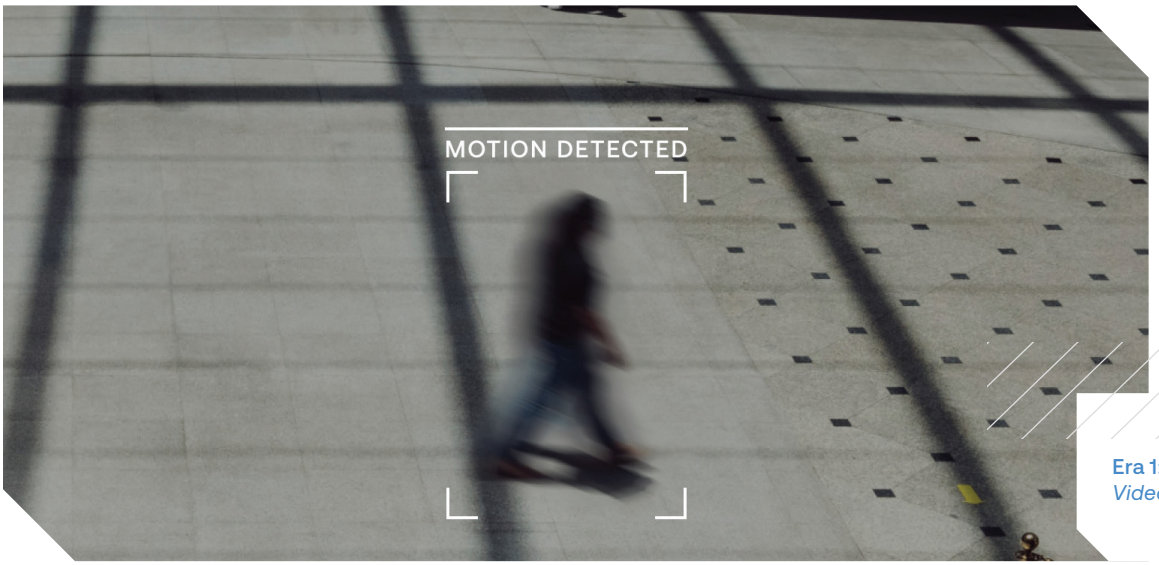

Era 1: *Motion Video Analytics*

While you can imagine that early motion detection systems were challenged with speed and accuracy, the technology evolved to include more advanced techniques and algorithms to improve commercial viability of these products. One widely adopted technique was **background subtraction**.

This technique identifies the background of a scene as the parts which stay relatively static over the series of frames in a video representation of the scene. This background can then be subtracted from the original frames to identify dynamic objects in the scene.

Very often, these algorithms make incorrect predictions, especially under poor environmental conditions, like nighttime, wind and rain. Even moving tree leaves can disrupt algorithmic predictions and trigger false detections. Shadows of moving objects can affect lighting and can also disrupt algorithms. Or, if the color background and foreground are similar, the algorithm would classify the foreground as background.

In order for the technology to work, it requires a substantial amount of manual effort with data to extract useful information from images.

This stage of research represents the first level of AI in physical security, what we will refer to here as the era of Motion Detection. Video analytics based on Motion Detection made many promises that didn't result in real business value.

Smart but Not Intelligent

The Second Level of AI in Physical Security

In 2012, computer vision took a huge leap forward, thanks to deep neural networks. These neural networks offered more than the classic machine learning (ML) algorithms. **Deep neural network algorithms** enabled what the industry commonly refers to as deep learning. It made object detection and even the detection of different attributes within those objects with a high level of predictive accuracy possible. Solutions built on object detection algorithms permeated many business areas, including physical security. **Convolutional neural network (CNN)** algorithms have become widely used to analyze images and video streams.

The design of CNNs was inspired by the human visual cortex. The artificial neuron receives input data like an image. Then it applies filters to an image according to rules that it has learned from training data. Eventually, by applying the same filter repeatedly, it builds out a feature map of an image and 'understands' what is represented on that image.

One of the first CNN architectures, **AlexNet, showed 84.7% accuracy** during a machine learning challenge, ImageNET. This was a revolutionary result of that time which started the revitalization of the AI field and the neural network approach, particularly known as the 'AI summer'. By comparison, the first level of AI algorithms **showed only 58% accuracy**.

After strides in classifying entire images with deep neural nets, the next extension was to be able to detect objects localized in regions of an image. **RCNN (Region-Based Convolutional Neural Network)** models identify regions in an image that contain objects of interest and group together semantically contiguous regions based on scale, color, enclosures, and textures. They use this information to extract patterns, objects, and features from an image and are commonly used for object detection in surveillance camera images. They can be trained to identify different types of objects, like cars and people.

While these AI algorithms have significantly improved the applicability of this technology to physical security and advanced video analytics, problems still persist. RCNNs are slow. On average, it takes 47 seconds to analyze a test image.

Era 2: *Deep Learning Video Analytics*

Moreover, these models do not understand the context of a situation. Detecting only objects tells us very little about what is happening in a scene and whether or not the event is suspicious from the standpoint

of security and safety. Context is key in building a truly intelligent system. In the realm of physical security, the context of time, location, environment, activity and movements along with the risk history of a specific space must blend with computer vision detection algorithms to realize actual perception.

The Era of True Intelligence

The Third Level of AI in Physical Security

True computer vision intelligence comes with the ability to understand cues that humans recognize immediately. Today, this kind of drastic improvement is possible with AI that has a near-human-level perception. The ability to comprehend contextual aspects of an image or video plays an integral role in preventing threats without relying on data that compromises individual privacy.

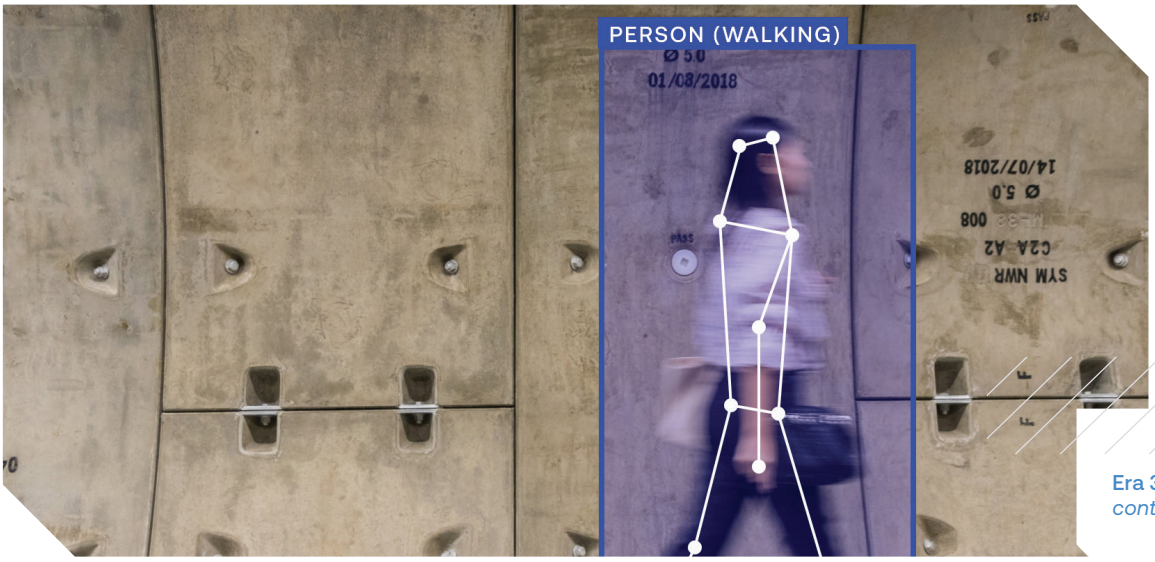

Era 3: *AI understands context in real-time*

FOR EXAMPLE:

1. **Human activity.** Understand human activities like jumping, drawing a weapon, opening a door, or leaving a building with a laptop. For example, a person breaking into a secure room through a window will be recognized as a security threat.
2. **Human-object interaction.** Identify when a human is interacting with another human or an object. For example, a person taking a box or a laptop out of a secure room will be recognized as a security threat.
3. **Contextualized events.** Consider the context of events as part of each assessment. For example, a person entering a secure room at 2am on December 25th would be more likely to set off an alarm than a person entering the same room during working hours on a work day.

TO ILLUSTRATE:

If you install an era one motion detection video doorbell in your house and fail to program regions of interest, it sends 20 alerts an hour, every time the scene changes with moving foliage, rain etc. and triggers motion detection. With an era two device, if you specify an explicit alarm, like “alert me only if there is presence of a person or a car”, the number of alarms would be reduced to 5 per day, but it would still be an unnecessary amount. This is characteristic of the previous generations of video analytics technology. What we really want to know is when a group of people show up at the front door in the middle of the night versus someone just passing by the front street. Computer vision intelligence platforms can understand such contextualized activity and distinguish them from the normal pattern of life. They use the context of

the situation to automatically understand real threats and warn you about the threats. A computer vision intelligence platform would only send an alarm for individuals approaching your door in the middle of the night and trying to break through your door, for example.

How Ambient.ai Computer Vision Intelligence Works

Ambient.ai analyzes and contextualizes information, historical patterns, and real-time activities. It provides alerts with a very high signal-to-noise ratio.

To give a sense for how Ambient.ai works in action, here are two images.

Whiteboard Session *2021 Physical Security Incident Predictions*

On the left, a person is drawing on a whiteboard, and on the right another man is drawing on a garage door. They both appear to be drawing, but context completely changes how we **perceive** this kind of situation. Ambient.ai's algorithms were trained to recognize hundreds of threat signatures. A threat signature is a pattern of events that are indicators of compromise or contextual patterns of suspicious behavior. Understanding a threat signature allows Ambient.ai to distinguish between these two examples and raise warnings for the situation on the right.

The intelligence of the Ambient.ai platform goes beyond real-time threat identification. It also produces descriptive data or metadata from video footage to speed any reporting & analysis needs that you have in the organization. This makes it possible to perform mission-critical assignments, like investigations, 20X faster. For example, you can conduct a smart search of more than a 100 threat signatures like 'a person loitering for an extended period of time outside an auxiliary entrance' or 'a person brandishing a weapon in a crowded room' and within seconds find all footage of that behavior that was captured on camera, no matter where on campus or when that behavior was displayed. The AI will understand what you are looking for and queue up videos immediately if they exist.

Our engineering team has leveraged cutting-edge computer vision technology to build an advanced platform that works well in real-world scenarios. Our AI is trained on a wide variety of threat signatures across varied environments and maintained by our threat detection analysts. They provide human-in-the-loop feedback to ensure accuracy in threat detection and automated response. Their feedback also enables constant learning and improvement of our system in the field faster than ever was possible. We provide security operations centers with applications for real-time risk alerts, forensic search, mobile dispatch, and operational dashboards.

Since our AI is trained to immediately deliver value, these applications start providing relevant recommendations to teams on day one, acting as the privacy-aware 'AI brain' for your existing security systems.

Why Ambient.ai is a Privacy-by-Design AI product

Our team has always believed that AI & computer vision will transform how we analyze video content. It's just a matter of how & when. Having worked at some of the largest technology companies in the world, we know that other companies are developing solutions. But many of them have and will take the easy approach and deploy facial recognition or other privacy-invasive technologies without regard to the impact on society.

From the beginning, Ambient.ai knew there was another way. We were inspired by the work of former Information Privacy Commissioner of Ontario, Ann Cavoukian, who introduced privacy by design as an approach to systems engineering. The work that she started was formalized in 1995 in a report on privacy-enhancing technologies by a joint team of the Information and Privacy Commissioner of Ontario (Canada), the Dutch Data Protection Authority, and the Netherlands Organisation for Applied Scientific Research in 1995. The privacy by design framework was adopted by the International Assembly of Privacy Commissioners and Data Protection Authorities in 2010. It has become standard guidance for privacy-aware software development and has been incorporated into privacy regulations like the European Union's General Data Protection Regulation (GDPR).

We've followed privacy by design principles since the first line of code. Our product architecture, product design and company policies all follow the basic principles of privacy by design. Those principals are:

Proactive not Reactive; Preventative not Remedial. We anticipate in advance privacy vulnerabilities and address them in our system design.

Privacy as the Default Setting. Our product automatically protects an individual's privacy.

Privacy Embedded Into Design. Our software architecture is designed to be privacy compliant at its core.

Full Functionality. Positive-sum, not zero-sum. By using context-based and a pattern recognition approach, Ambient.ai detects threats without compromising anyone's privacy. We deliver both privacy and safety & security without trade-offs.

End-to-end Security. Full lifecycle protection. From the data collection process to deployment of machine learning models into production, we follow the strict and consistent guidelines of data management.

Visibility and Transparency. Keep it open. All our stakeholders are aware of our privacy measures and how exactly we apply them. We've gained the trust of Fortune 500 companies.

Respect for User Privacy. Keep it user-centric. Our product is built for people to ensure their safety and security. We put our users first before anything else and always inform them on how our technology works and how exactly it protects their privacy.

Learn more about Ambient.ai

Ambient.ai proactively monitors every single camera, correlates threats across cameras, and delivers site-level intelligence. It cross-correlates activities on surveillance streams with the data from other building endpoints like PACS and other sensors to automatically perform video verification and detect new previously undetected threats.

By adding contextual knowledge, Ambient.ai takes computer vision technology to the level of an autonomous system with near-human perception behind every single camera in the enterprise. The system that never gets tired and proactively alerts you of real threats to your security team 24x7x365.

To learn more about ambient.ai visit our website at www.ambient.ai.