



Transparency Note for Live Video Analytics (LVA)

Last Updated 6/1/2020

Contents

- What is a Transparency Note? 3**
- Introduction to LVA 3**
 - The Basics of LVA..... 3
 - Example Use Cases..... 4
 - Characteristics and Limitations..... 5
- Resources to learn more..... 6**
- Contact us 6**
- About this document 6**

What is a Transparency Note?

An AI system includes not only the technology, but also the people who will use it, the people who will be affected by it, and the environment in which it is deployed. Creating a system that is fit for its intended purpose requires an understanding of how the technology works, its capabilities and limitations, and how to achieve the best performance. Microsoft's Transparency Notes are intended to help you understand how our AI technology works, the choices system owners can make that influence system performance and behavior, and the importance of thinking about the whole system, including the technology, the people, and the environment. You can use Transparency Notes when developing or deploying your own system or share them with the people who will use or be affected by your system.

Microsoft's Transparency Notes are part of a broader effort at Microsoft to put our AI Principles into practice. To find out more, visit [Microsoft's Responsible AI Resources](#).

Introduction to LVA

Live Video Analytics (LVA) provides a platform for you to build intelligent video applications that span the edge and the cloud. The platform offers the capability to capture from CCTV and IP-based cameras, record, analyze live video, and publish the results (video and/or video analytics) to Azure services (in the cloud and/or the edge).

LVA is designed to be a pluggable platform, enabling you to focus on building solutions quickly, with the platform taking care of the complexity of building and running a live video pipeline.

Note on preview technologies: LVA is currently offered as a public preview service (as of June 1, 2020).

Preview status does not permit production use, so customers should not build on top of LVA while offered in public preview. Preview status does, however, provide an ideal opportunity to think deeply about a system's considerations for responsible AI and how to design for them.

The Basics of LVA

Live Video Analytics allows you to ingest video from CCTV and surveillance cameras, optionally run AI on the video, then send that video and resulting AI data on to other systems such as cloud storage or a video management system (VMS).

You decide what cameras to connect to, what AI you want to run on the video, and where you want the results to be sent or stored. You decide what AI you want to run based on the situation and can use any AI you choose: Microsoft Cognitive Services, open-source, licensed models, or even models you create yourself.

You use the combination of the video and the data from the AI models and integrate into existing solutions to alert the user when certain conditions are met. For example, if an unauthorized person is trying to use an entry door, it can send an alert to your video monitoring center. You can save a photo of the person, notify local security officers of the incident, and add that person to a list of known offenders to watch for in the future.

Example Use Cases

Live Video Analytics can be used in multiple scenarios across a variety of industries that use video monitoring systems. Some examples include:

- Monitoring empty shelves (Retail). A grocery store has cameras pointing at the cereal aisle, connected with LVA, and using a custom-built model to detect when rice puffs are running low. The LVA output is sent to their re-stocking system and notifies the stockroom to put more product on the shelves.
- Detecting trespassing (Retail, Manufacturing, Public Safety). A construction site places several cameras around the perimeter. After close of business hours, when a person is detected in those spaces, a system plays a recorded message in that area to leave the premises immediately.
- Confirming PPE use (Manufacturing, Health Care). Workers at a building site are required to wear hard hats at all times. A combination of PPE detection and face detection can identify that individual 1 did not wear their helmet. The supervisor reviews the notifications and talks to the employee.
- Counting vehicles at intersections (Public Safety, Transportation). A low-traffic intersection needs more repairs than expected, so the city sets up a camera and AI system to count the number of cars to get a more accurate estimate of actual usage, which reveals that the city's initial estimates of intersection traffic volume were too low.
- Identifying suspicious activities and objects (left bag, weapon detection) (Public Safety). A person in a bus terminal leaves a bag unattended for 5 minutes without returning. Local officers are dispatched, and the area is evacuated until backup can arrive.
- Searching for a lost child (Public Safety, Retail, Stadiums). A parent and child are separated at a baseball game. The parent finds security, and they scan a recent picture of the child. An AI system correlates features like the child's last location, clothing, physical description, and photograph to detect the child at a food stand, and security reunites the family.

Some common customer motivations for using LVA include:

- Automate or supplement some aspects of monitoring done by humans. Instead of a human watching 20 screens to look for bad activity, the system could highlight screens with possible issues. The human could then direct more attention on those screens.
- Improving human oversight to be more efficient and focus on critical events and activities. The operator has defined a detected weapon as more important than a vehicle failing to stop at a traffic light. The system therefore draws a yellow box around the traffic violation but a flashing red box around a weapon.
- Increase the number of cameras in use to cover critical areas. A dry cleaner has two cameras but would like to add five more to cover other areas and save the video for 30 days to the cloud. They connect the cameras to a deployment of LVA, configure it to send to their Azure Blob Storage account, and delete after 30 days.
- Gather more data and intelligence for additional analytics and trends. A city puts a camera at 100 intersections, counts cars for 30 days with an AI system, and correlates the data with weather, holidays, and sporting events. This leads to a more accurate understanding of how traffic density changes with other factors.

Characteristics and Limitations

Video monitoring can be a valuable tool for public safety, physical security, employee safety, and business efficiency. But video monitoring can also present challenges for personal privacy and ethical use that need to be considered with any solution.

When using LVA as part of your solution, below are some considerations to keep in mind:

- When outputting frames from the video to send to an AI model, it is suggested to send only the minimum number of frames that allow the AI model to process efficiently. Often, even 2-3 frames per second can be sufficient, and 10+ frames might overload and delay the AI result.
- Most AI models are tuned to perform best on certain processors such as CPU, GPU, FPGA, etc. Where possible, deploy LVA onto edge hardware that is optimized for the AI model.
- Typically, surveillance cameras are mounted in a high position, angled down, with a wide field of view. This might mean a person or face is composed of only a few pixels. Keep this in mind when selecting AI models so you understand what the model is capable of processing at an appropriate success rate.
- Surveillance cameras are also often mounted outside or with exterior views, so lighting, weather, and other atmospheric effects can often dramatically change the type and quality of video that is being captured.
- Review the Live Video Analytics product documentation for current quotas, limitations, and recommendations for load balancing number of cameras per edge device, etc.

Below are some general topics around video monitoring to keep in mind when using any solution.

You should:

- Respect an individual's right to privacy, and only ingest videos for lawful and justifiable purposes.
- Commit to respecting and promoting human rights in your design and deployment of your video monitoring solution. Microsoft publishes its commitments to human rights in the [Microsoft Global Human Rights Statement](#).
- Ensure your solution is secure and has adequate controls to preserve the integrity of your content and prevent any unauthorized access.
- Provide a feedback channel that allows users and individuals to report issues with the service.
- Obtain appropriate legal advice to review your solution, particularly if you will use it in sensitive or high-risk applications.
- Be aware of any applicable laws or regulations that exist in your area regarding notification and consent from those who will be subject to video monitoring.
- Keep a human in the loop. Do not use any solution as a replacement for human oversight and decision-making.
- Allow for real-time human intervention in the solution to prevent harm.
- Fully vet and review the capabilities of any AI model you are using to understand its capabilities and/or limitations.
- If you use AI in your broader system, consider how your AI solution aligns with [Microsoft's Responsible AI principles](#).

Resources to learn more

Responsible AI resources:

[Microsoft Responsible AI principles](#)

[Microsoft Responsible AI resources](#)

[Microsoft principles for developing and deploying facial recognition technology](#)

[Microsoft Azure Learning courses on Responsible AI](#)

Product documentation:

<https://docs.microsoft.com/en-us/azure/media-services/live-video-analytics-edge>

Contact us

Azure Media Services Help (external): amshelp@microsoft.com

About this document

© 2020 Microsoft Corporation. All rights reserved. This document is provided "as-is" and for informational purposes only. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it. Some examples are for illustration only and are fictitious. No real association is intended or inferred.

Published: June 1, 2020

Last updated: June 1, 2020