

Illustra Video Intelligence Analytics

Introduction

This guide is intended for use by those administering Illustra cameras for configuring Video Intelligence Analytics. It provides guidance and advice on configuration as well as in-depth explanations of the Video Intelligence software features. Video Intelligence is a video-based object detection and tracking engine using computer vision algorithms to analyze video.

Video Intelligence Analytics are pre-installed on Illustra Pro Gen2 and Pro Gen3 cameras.

Object Detection Rule

The simplest type of event that can be detected is the appearance of an object in an area, or region of interest (ROI). Whenever an object, such as a person or a vehicle enters the ROI or appears in the ROI, it is detected as a new object in that ROI. If the object leaves the region and comes back, it will be detected again.

This event type can be used to define trip wires or protection zones. Draw a thin threshold ROI or cover a whole zone with the ROI, and the camera will detect whenever a person or object crosses the threshold or enters the zone.




Fig. 1— Detect when people enter the area around product displays.

The Object Detection event type has just one parameter besides the ROI, and that is **Overlap (%)**. The overlap parameter refers to the way detected objects related to the ROI. It is the percentage of the object that must be inside, or overlap the ROI in order for it to be counted. A high value, such as 100%, means the object has to be completely inside the ROI. But a low overlap, such as 10%, means the object merely needs to touch the ROI.




Fig. 2— In these two pictures the yellow region is the ROI. Higher Overlap: the object must be mostly inside the ROI to trigger an alarm (left). Lower Overlap: the object just needs to touch the ROI (right). The Overlap percentage refers to the amount of the red box that is inside the yellow ROI.

The overlap parameter can be used to fine-tune results. With the Object Detection rule set low, the setting will product the most results. As you increase the setting, the number of results will decrease. So it acts like a sensitivity setting. Higher overlap settings will ignore objects that are only partially within the region, including shadows cast by objects and objects passing by or occluding the ROI but not completely within it.

The Object Detection rule also has the **Color** parameter. This parameter can be used to detect only objects of certain colors. See Color Signatures, in a later section, for further details.

Direction Rule

The Direction event type is used when you want to detect people, vehicles or other objects moving in a certain direction. Because the Video Intelligence engine detects objects on every frame of video, over time it builds up a “track” of how an object moves in the scene. The main parameter of this rule is **Direction**, either North, South, East or West. This is the direction of motion that an object has to be moving to be detected. These compass directions are relative to the camera field of view (North is up, and West is left).

When drawing an ROI for direction events, draw a path on the ground. For best results, draw a clear, simple path, such as a straight line or an arc. Avoid paths with sharp angles. Make sure your path covers only the trajectory that the person or object needs to follow in order to be considered a hit. If you draw your path so that it covers more than the trajectory you are interested in, then people who only partially follow the trajectory will not be considered hits.



Fig. 3— If a person started at the bottom of the path and then turned right into the first office door, the person would not be detected because the person did not cover the majority of the yellow ROI drawn on the floor.

The Direction rule also has an **Overlap (%)** parameter. Usually you want to keep this low, because you are looking for objects that follow the path defined by the ROI, and are not necessarily inside it. But in some cases the overlap parameter might be useful to fine-tune results.

The other parameter of the Direction rule is **Traversal Time (secs)**. This parameter lets you set a maximum amount of time an object can take to cross the ROI. If you leave the duration '0' secs, any object moving through the ROI in the defined direction will be detected. If you define a duration of '1' second or more, then only objects that move through most of the ROI in that time will be detected. Objects that take longer than the duration to move through the ROI will not be detected. The lower the setting, the faster the objects have to move to be detected. This gives you a way to avoid detecting objects moving too slowly to be of interest.

The direction rule also has the **Color** parameter. This parameter can be used to detect only objects of certain colors. See Color Signatures in a later section for further details.

Linger Rule

The Linger event type is used to detect people or objects that linger in a region of interest for a certain length of time. An object is lingering if it simply stays in the ROI. The object can move around or stay still, as long as it stays in the ROI. The Linger rule has the **Overlap (%)** parameter, which applies to the entire time the object is lingering. The **Linger Time (secs)** parameter is used to define how long an object must linger before it is detected. Although you can set the Linger Time quite high, in practice the rule works best with times less than a minute or two.

The Linger rule also has the **Color** parameter. This parameter can be used to detect only objects of certain colors. See Color Signatures, in a later section, for further details.

Dwell Rule

The Dwell event type is very similar to the Linger type. The only difference is that while the object is in the ROI it must remain mostly still and not move around. A small amount of motion is acceptable but no more. For example, a person standing in line or a person looking at products on a shelf. If you are thinking about people dwelling, then they are not taking a step. If you are thinking about a car dwelling, then its wheels are not moving. Every dwelling object is also a lingering object, but not the other way around.

The Dwell rule has the **Overlap (%)** parameter, which applies to the entire time the object is dwelling. It also has the **Dwell Time (secs)** parameter. Like Linger you can set this quite high, but in practice it works best when it is 30 seconds or less. An object that dwells for a very long time might be considered an abandoned object.

The Dwell rule also has the **Color** parameter. This parameter can be used to detect only objects of certain colors. See Color Signatures, in a later section, for further details.

Perimeter Protection Rule

The Perimeter Protection rule is used to detect objects that cross a perimeter to enter a protected area, or that linger along the perimeter for a period of time. This rule uses a special ROI with two areas: the protected area and the perimeter area. Generally the perimeter area is along one or more sides of the protected area. For example, a camera with a view of a fence where the protected area might be inside the fence and the perimeter area might include the fence and some area in front of it.

The Perimeter Protection rule has a **Linger Time (secs)** parameter. When an object lingers in the perimeter area for this amount of time it will trigger the rule. If an object goes through the perimeter area and into the protected area it will trigger the rule. In this case it does not matter how long it was in the perimeter area.

The Perimeter Protection rule also has minimum and maximum object size parameters. You can use these parameters to filter out false alarms from objects that are too large or too small to be interesting, such as a small animal or a large shadow.




Fig. 4— A perimeter fence view. The yellow region is the protected area, green is the perimeter area, red is the max object size and magenta is the min object size.

Enter Rule

The Enter event type is used to detect objects that enter the camera field of view through a doorway or threshold. Draw an ROI that completely contains the doorway or threshold. The Video Intelligence engine looks for objects that appear inside the ROI and then later move outside it. Make sure your ROI contains the whole doorway, including when it slides or swings open. The door itself is a moving object but if you draw the ROI so that the door stays inside it, then the door will not trigger a false enter hit.




Fig. 5— The yellow ROI complete contains the door.

If the wall around the door contains any glass, then include that in the ROI as well. You want to make sure that any object coming through the doorway appears inside the ROI. If the ROI did not contain the glass, then objects might be first seen through the glass outside the ROI and would not trigger a hit.

The Enter rule has the Overlap (%) parameter. This applies to the object when it appears in the ROI. Higher overlap rules tend to work better. If the field of view is such that people coming through the door may be grouped together, then consider adding a bit of "runway" in the ROI leading up to the door. The idea is to make the ROI big enough so that the people separate before moving out of the ROI. That way each one will trigger a hit.

The Enter rule can also detect objects entering the scene from the edge of the field of view. Draw the ROI up against the edge so that objects appear inside the ROI and then later move out of it.




Fig. 6— The Enter rule can be used on the edge of the field of view.

The Enter rule also has the **Color** parameter. This parameter can be used to detect only objects of certain colors. See Color Signatures, in a later section, for further details.

Exit Rule

The Exit event type is the opposite of the Enter rule. It is used to detect objects that exit the camera field of view through a doorway or threshold. The Video Intelligence engine looks for objects that are outside of the ROI, then move into the ROI and disappear from view.

The same ROI drawing tips for Enter events also apply to Exit.

The Exit rule also has the **Color** parameter. This parameter can be used to detect only objects of certain colors. See Color Signatures, in a later section, for further details.

Crowd Formation Rule

The Crowd Formation event type is used to detect when more than a certain number of objects are detected in an ROI. This rule works best in scenes with a wide field of view. Objects must be seen individually before they enter the group in order to be counted. Groups of objects that are always together and move together might not be counted individually.

This event type has the **Overlap (%)** parameter; only objects satisfying this setting are counted in the group. Use the **Minimum Crowd Size** parameter to set how large the group has to be to trigger a hit. You will get an alert when a crowd grows above the given size, and another alert when it drops back down again.

Queue Length Rule

The Queue Length event type is used to detect when a queue or line of people or cars is short, medium, long, empty or not empty. This rule uses a special ROI that has three different areas: short, medium and long. You draw each area separately, but together they should indicate the area where the queue forms.



Fig. 7— Queue region with short (green), medium (yellow) and long (magenta) areas.

Once you draw the three areas of the queue, the length of the queue is determined by using Object Detection in each area. When no object is detected in any of the areas, the queue is empty. If an object is detected in the short area, but not in the medium area, then the queue is short. If objects are detected in the short and medium areas but not in the long area, then the queue is medium. If objects are detected in all the areas, then the queue is long. The queue is not empty if it is short, medium or long.

Queue region drawing is very flexible by design. You can draw regions for queues that curve and bend, that go behind occlusions or zig-zag through an area. However, you generally want the three regions that make up the queue to be adjacent and non-overlapping for best results, as depicted in Fig. 7.

Use the **Overlap (%)** parameter to adjust how objects are detected in each of the three areas. Choose the length of queue you want to trigger: **Empty, Not Empty, Short, Medium or Long**. The alert rule will generate an alarm when the queue transitions into and out of the chosen length.

Try to avoid excessive reflection and shadows in the queue area. If the flooring is reflective, try to angle the camera in such a way as to avoid reflections in the field of view. If strong shadows are present in the queue area, try to adjust the lighting to reduce their impact if possible.

Abandoned / Removed Rule

The Abandoned/Removed event type is used to find objects that are abandoned in a region of interest, or removed from a region of interest. The Video Intelligence engine detects this by looking for changes in the background of the camera scene. When an object is left behind, or taken away, a change is made to the background of the scene. It is this change that is detected. Use the **Overlap (%)** parameter to define whether the object has to be completely inside the ROI or not when it is left behind or removed.

Sometimes an object stops moving but only temporarily. When the Video Intelligence engine detects an object being abandoned or removed it will continue to monitor the area for some time to make sure it is not temporary. It may take as many as 20 seconds to make a decision.

The Abandoned/Removed event type can also be used to detect scene changes, such as tampering. For example, installing a skimmer on an ATM machine can be detected as an abandoned object (the skimmer). Writing graffiti on a wall can be detected as well, as can damaging or removing signage. Any time a semi-permanent change is made to the background of the scene it can be detected with the Abandoned/Removed event.

The **Amount Changed (%)** parameter lets you define what percentage of the region must change before it is detected. For example, at 50% a change must be detected in half the ROI. Smaller changes will be ignored. When a change takes place, Video Intelligence measures the area of the change and compares it to the area of the ROI. If it is more than the Amount Changed an alert will trigger.

If you may also set the **Within (secs)** parameter, this is the amount of time over which the change can take place. Each time a change is detected Video Intelligence measures the area. Each changed area is added until the total area exceeds the Amount Changed threshold. But if Within (secs) elapses before exceeding the threshold, then the total amount is reduced by the aged-out amount. For example, with 50% Amount Changed and 300 Within (secs), then when all the changed over a 5 -minute period amount to 50% or more of the ROI, an alarm will trigger.



Fig. 8— When using Abandoned/Removed to monitor assets it is best to draw an ROI closely fitting the display fixture. When you do this, the Amount Changed matches the assets taken. Use a higher Overlap so that changes near the fixture are ignored.

Color Signatures

Each object that is detected by the Video Intelligence engine is given a color signature in the metadata. During the time that the object is visible on camera it is analyzed to determine which of the following colors it appears close to: red, blue, green, yellow, black, white, and gray. On every video frame the analysis is repeated so that over time if the object changes position and shows different colors those are recorded too. The color signature of the object is the set of all observed colors for it.

Each of the Video Intelligence event rules that apply to individual objects has a Color parameter that can be used to filter the results to one or more colors. When no color is selected then color filtering is disabled and all objects satisfying the rule will trigger regardless of color. Choose one or more color values to only trigger results that have at least one of those colors in its signature.

When using color signature it is important to remember that the camera is using colors as they appear in the video stream, and not necessarily how they look in real life. If the camera color calibration is bad then that is what is actually yellow in the real world may appear green to the camera. These objects will have a green signature.

Some colors perceived by the naked eye might not correspond to colors registered by the camera or displayed by the monitor. There might be several technical reasons for this. The best way to alleviate this problem is to perform a color calibration (or white balance) of the camera. It is very beneficial to do white balance calibration of any camera to be used with color filters.

Shiny objects (or highly reflective surfaces) seem to increase yellow content in the color signatures. One cause of this is the artificial indoor light sources that contain a yellow hue. This is true for any off-white light source in the environment. The effect might be indiscernible to the naked eye but can have a biasing effect on the color filter. Fluorescent lighting introduces greenish hues, incandescent lighting a reddish-orange hue, and sunlight introduces a bluish hue.

Darker colors tend to trigger black or gray filters, which is expected. As object get gradually darker they lose their color hue. Consequently, they get assigned gray and/or black labels. Using artificial light sources to illuminate darker areas of the scene will counteract this effect. Shades of gray may trigger white, gray, and/or black, and often two at a time. This is expected since gray is a highly subjective color.