

SENTINEL

Intelligent Video Analytics using cameras for people tracking, counting and searching

Sentinel is a computer vision technology for precise human tracking using standard CCTV cameras.

It recognizes people in a video feed and is able to follow their movement through a venue over time. Applying object recognition techniques, Sentinel recognizes human shapes from images in a video stream based on their appearance, including features like **bulk, clothes color** and **height**. This enables Sentinel to discern people from the background and to enable implementing capabilities such as **people counting, tracking people's movements** and **searching for a likely match** among individuals previously "seen" by the video cameras.

Sentinel processes video streams from one or more cameras deployed around the monitored area. It is often possible for the system to use cameras already installed in a venue.

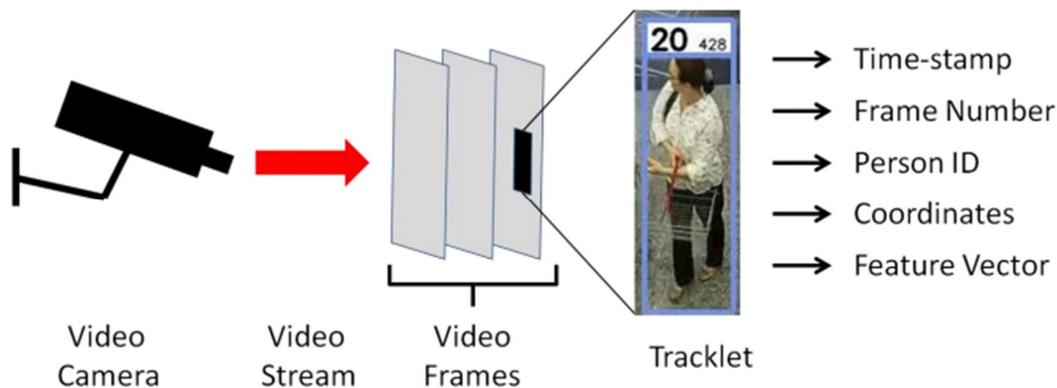
Sentinel solves a major problem: tracking people without using Tags and without the need of installing specific radio infrastructures throughout the venues. It precisely tracks people, allowing the development of intelligent analytic software and security applications.

It provides a reliable and accurate tracking. The videos from the cameras must be available for the Sentinel server, which can be on-site or a remote cloud, and the videos are processed in real time.

HOW IT WORKS

Recognizing people in video streams is done as follows:

1. Sentinel analyzes images from the video streams
2. It recognizes people and assigns each of them a **unique Person ID** and a **Feature Vector** that encodes the person's appearance
3. The actual image of each individual recognized by the system is saved as a **tracklet**. All tracklets are collected into a database, each of them identified by a **Tracklet ID**.



FEATURES

Sentinel detects, tracks and re-identifies people and makes the data available for external applications.

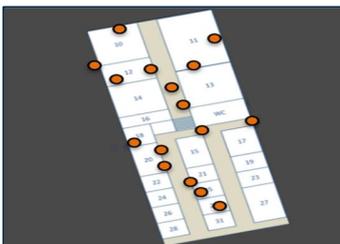
It is possible to integrate Sentinel into intelligent analytics software, and to use it for security and tracking applications. Sentinel is a flexible technology that can be adapted to multiple use cases, and can work in combination with other technologies, too, including facial recognition.



People Detection
Sentinel detects **people moving** inside an environment

```
3248 // resolution = 0
3249 // level = 4
3250 public int getState() {return mState();}
3251 /**
3252  * Return current latitude
3253  * @return current latitude
3254  */
3255 public Double getLatitude() {return mDataHolder.getLatitude();}
3256 /**
3257  * Return current longitude
3258  * @return current longitude
3259  */
3260 public Double getLongitude() {return mDataHolder.getLongitude();}
3261 /**
3262  * Return current altitude
3263  * @return current altitude
3264  */
```

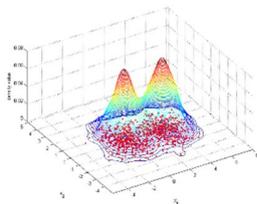
Integration
Sentinel provides **API** and **CSV** files for custom integration into external applications.



Location Tracking
Sentinel can provide the **location** of people as WSG-84 (latitude, longitude + level ID) and Cartesian Coordinates (x, y, z).



Crowd measuring and people counting
Sentinel enables people counting, crowd measuring and queue management.



Cross Camera Re-identification
Sentinel provides **Feature Vectors** to re-identify people across multiple cameras.



People Searching
Sentinel allows to search for specific individuals that have been previously detected.

RECOGNITION AND RE-IDENTIFICATION

Sentinel's functionality is **not face recognition** but a complementary technique that can recognize and follow individuals even when they face away from the camera.

Person ID	Each person is assigned a Unique ID , based on its current appearance. Each individual is also assigned a Feature Vector for re-identification.	The same person may be assigned different IDs at different times. If individuals leave the scene and re-enter it, new IDs may be assigned to them if their appearance is somewhat different because of varying lighting conditions. The Feature Vector should be used to re-identify people over time.
Crowds and People Counting	Sentinel can be used to monitor crowds and for people counting .	Recognizing people in a busy scene is challenging. Occlusion and groups where individuals cannot be picked apart are real possibilities, and applications must take into account that the counts provided are approximate
People Identification	Sentinel identifies human beings using AI and self learning algorithms.	Human's shapes' height should be at least 30px, in the video, to be recognized. Some anomalies may happen: Sentinel may recognize a human shape displayed on a poster as if it were an actual individual in a scene.

VIDEO STREAMS REQUIREMENTS

The videos analyzed by Sentinel must meet these requirements:

Digital / analog video	Sentinel can process video streams from digital and analog cameras. Analog camera streams must go through a DVR to provide remote access.
Color video	Visible color video streams are recommended. Black & White video Supported. Infrared video Not supported.
Resolution	There is no minimum, but as a general rule, the better the resolution, the better the detection and re-identification. HD is suggested.
Frame rate (frames per second – fps)	At least 7 fps; 10 fps or more is recommended
Scene lighting	The scene must have good visible lighting conditions, indoors and outdoors. Extreme lighting conditions like sun shining directly on the camera are not supported.
Video accessibility	Video stream must be accessible remotely through a URL.
Angles of view	Angles of view up to 90° are supported.

ON PREMISE SERVER REQUIREMENTS

Sentinel requires a powerful computing unit, for on-site installations, where Sentinel engine runs. These are the servers' minimum requirements for an installation with 6-7 cameras, running 24/7.

CPU	A Real Quad-core processor capable of producing at least 8 computing units. Xeon or i7 are suggested.
GPU	NVIDIA GPU Card (GTX 1080 or better)
RAM	16 GB
STORAGE	500 GB SSD (depending on the historical database)
O/S	At least Ubuntu 16.04

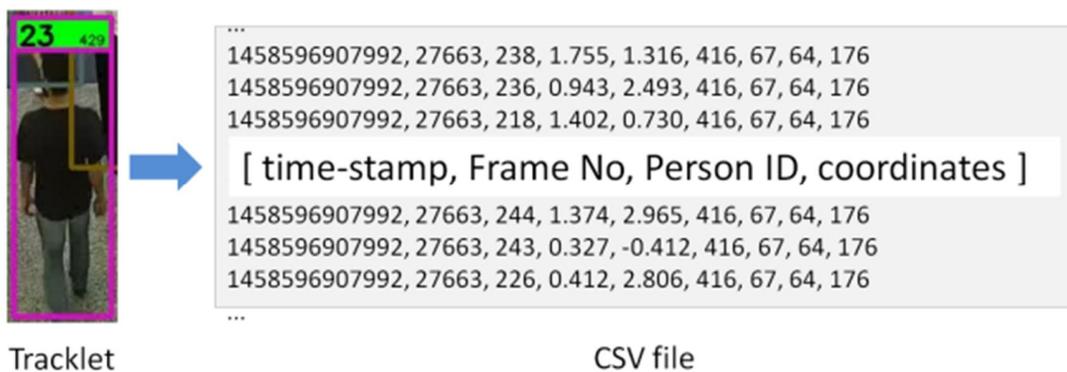
OUTPUT AND INTEGRATION (CSV, API)

Sentinel creates a record for each **tracklet**. Periodically, Sentinel outputs a **CSV file** containing data from records of people recognized by the system during the previous time window. The time window covered by a CSV file is configurable. By default, CSV files are output every 5 minutes. CSV records can be accessed via **API** by external applications.

The CSV record includes:

- **Timestamp,**
- **Frame Number,**
- **Person ID,**
- **Feature Vector,**
- **Coordinates** that identify the specific tracklet within a frame.

This data enables historical search of individuals whose appearance matches those in a given tracklet.



SYSTEM COMPONENTS

<p>Cameras One or more cameras (analog or digital) installed in the environment.</p>	
<p>Server Sentinel runs on a server that collects the video streams from the cameras. The server can be an on-site machine, or a remote server to which the cameras send the videos over an internet connection.</p>	
<p>API and CSV files API and CSV files allow the development of custom applications, dashboards and software on top of Sentinel.</p>	

SYSTEM ARCHITECTURE

- Sentinel runs on a local server or on a cloud server.
- The server receives video streams from one or more cameras.
- All streams contribute to Sentinel's tracklets database, each tracklet identified by a Tracklet ID.
- Each tracklet represents an instance of a person being detected, tracked and tagged with a Person ID.
- The sum of all active Person IDs at any one time enables people counting, and can also be used for building heat maps of human activity in the venue.

