Innovative Solutions Canada Program

Challenge EN578-170003/32: Automated redaction of video recordings for the purposes of Access to Information requests Challenge

Attachment 3
Questions and Answers #4 to #11

This document contains questions and answers related to this challenge.

Question #4:

The statement reads “In addition, before the video footage can be released, it needs to be redacted so that other individuals **and objects** in the scenes cannot be identified, in order to protect the privacy of others.”
Can you please provide an exhaustive list of potential objects that would require redaction?

Response #4:

Personal identifiable information of others who are not part of the initial request must be redacted. This could include Faces, License Plates, Name tags, etc, and could also include where the personal identifiable information is part of a reflection (mirrors, windows, shiny chrome objects, etc). In the end, it is impossible to identify all the potential objects, as this is dependent of the scene being redacted. However, it is important for a system to allow an operator to select objects of interest and then for the system to redact these objects from the scene.

Question #5:

Essential Outcome #4 reads “be able to blur the overall scene in order to remove any background private information while maintaining the basic context of the scene.”
We do not understand what this criteria achieves in addition to Essential Outcome #1. Essential Outcome #1 requires the redaction of all instances of all personal identifiable content (whether in the foreground or background). Once this is achieved, why would blurring the overall scene be required?

Response #5:

This is one of the options available to an operator of the system for redacting complicated scenes that may have a large number of objects that contain personal identifiable information. It is not a matter of one (Essential Outcome #4) or the other (Essential Outcome #1), but rather both methods need to be available to an operator of the system, and an operator/systems can choose which method to apply for a video depending on the content.

Question #6:

Essential Outcome #5 reads “minimise manual intervention after identifying initially a particular individual/object in a video feed”. Is this an individual/object that we are redacting or keeping?

Response #6:

Yes to both, depending on the approach taken. If you tell the system what to keep and blur the rest, then the identified object/individual must be kept throughout the video scene with minimal intervention from the operator. However if the object is tagged as being redacted then it must remain redacted throughout the video scene with minimal intervention from the operator.

Question #7:

Could it be possible to elaborate on point 6 Essential Outcome results? Is it a matter of creating a single video sequence from various sequences that is, putting them end to end to create a single video?

Response #7:

In many cases the request involves more than one video file. For example, a requester may be asking for all the video of them while they were being processed through the border clearance process at an airport on a particular day. This would involve a set of multiple videos where each video would be from a separate surveillance camera stream. The solution must be able to manage this request as a set of videos and deliver the results as a set of videos. It is a matter of being able to manage multiple videos in a single request and deliver the final output as a set of these multiple redacted videos.

Question #8:

Do we want the solution to store the information at a protected level B or does the information, during the achievement of the challenge, have to be stored at a protected level B?

Response #8:

The resulting solution must be able to handle protected B data. The initial system under development for the challenge does not need to host protected B information. However, eventually it will be important for the system under development to host protected B information in order to test the system in an operational context.

Question #9:

How many video sequences will you be able to provide for the training of the various algorithms for object recognition, image classification and tracking?

Response #9:

We have about a dozen unclassified videos that can be used as examples for testing possible solutions. Some of these could be used for training, however, the ones used for training could not be used for evaluating the system performance. It may be more beneficial to leverage open source video content for training, and then use the CBSA unclassified videos for evaluation.

Question #10:

Is it possible to know the exact software currently used to perform these tasks?

Response #10:

Manual redacting using Adobe, Input Ace, and /or AmpV

Question #11:

What are the restrictions of using the cloud in connection with this challenge and the protected level B of the data?

Response #11:

The cloud environment must be able to host protected B level information, and must have in place all the needed security controls to ensure that the information is protected. This includes dat at rest in the cloud environment as well as data in motion through the cloud environment.