

Minutes from TRECVID 2008 Event Detection Planning Telecon

February 27, 2008

Jon Fiscus, NIST
Martial Michel, NIST
Paul Over, NIST
Travis Rose, NIST
Sadiye Guler, intuVision
David Eichmann, U. Iowa
Mert Dikmen, UIUC
Jingen Liu, UCF
Ciarán Ó Conaire, DCU
Heather Simpson, LDC
Stephanie Strassel, LDC
Murtaza Taj, QMUL

TRECVID Event Detection 2/27 telecon agenda

- Role call
- Old business
 - Discussion of previous calls (if any)
- New business
 - Discuss the new evaluation plan
 - Discuss the "reverse direction" event.

The "reverse direction" event is of interest, however we would like your input into defining it. A discussion about this event in particular will be helpful.

- Schedule next telecon

Review old business.

Paul Over will send information (username and passwords) for each dataset.
David Eichmann -- will set up a mirror with the same username/password.

New business.

Group discussion of the new evaluation plan.

Ciarán Ó Conaire:

Regarding the derivation: there is a cost for missed detection (MD) and false alarms (FA) - these cost values will determine how systems perform.

Jon Fiscus:

True - these parameters define the target application.

Previously, we've either said that MD is worse than FA, or vice versa.

We want to choose a ratio, for example a 10 MD : 1 FA ratio, but we also want to decide this as a community (understanding that the target application is not the same as the real application).

Dave Eichmann:

Given this is an airport security scenario, if looking for weapons, etc., isn't the cost of a miss is high?

This would suggest a target application with high recall vs. high precision (perhaps we should weight MD higher).

This also has an impact depending on number of candidates we generate.

Although FAs are a nuisance, it may be the case that it is "better safe than sorry" (i.e., penalizing missed detections more).

Sadiye:

Yes, number one concern is FA rate.

Systems often do not rely on just one modality (such as video).

There is the problem with having a high FA rate and "crying wolf", which becomes annoying.

Paul:

Are the types of events such that they could be catastrophic if missed?

For example, not worrying about large crowds, rather, is the environment more constrained?

Sadiye:

The systems we have dealt with include perimeter security.

Airports have multiple security check-points, for example, where these are deployed.

Paul:

Has NIST tried to find out the target rates from DHS, etc.?

Airport settings may be different from the target application.

Need the full operating curve -- if picking one set of constants seems wrong, we may be tempted to pick two or three points.

So then what do systems to optimize to?

Jon:

Yes, NIST is investigating what are realistic FA and MD rates.

We may choose one primary characteristic point as the primary condition, and allow others to be explored as contrasts.

Paul:

Regarding the events, we expect these to have similar numbers of occurrences. How similar will these frequencies be?

Jon:

Expect to have within 2 orders of magnitude.

In STD (Spoken Term Detection), for example, there was the following:

-- high frequency vs. low frequency terms

-- bimodal distribution, and each distribution had different characteristics

Paul:

If the event distributions are different, we can analyze them separately.

Can we compare systems if we do this?

Jon:

Yes, we've developed a technique for statistical analysis of DET curves. This will appear in subsequent documentation.

Will NIST give the Beta value expressed in the metric?

Yes - plan to post this in the next couple of weeks.

Ciarán:

Regarding an earlier version of the evaluation slides (20080110-trecvid-event-plan.pdf): it states that systems would emit a confidence score (for decision-error tradeoff analysis). Will average precision still be used?

Paul/Jon:

-- we'll only define a single metric so that all systems optimize to the same error tradeoff.

Jingen Liu:

About the video:

-- can we get video with a clean background?

Jon:

-- this is not available because we are using operational data, which includes only limited calibration info, etc.

Group discussion of reverse direction event.

two problems:

- how to tell when someone reverses vs. milling around?

- how much of an angle does a person need to turn to be considered a "reverse"?

[[See below for the full discussion]]

Outcome:

The "reverse direction" event will be dropped due to the difficulty of defining it adequately. Attention will be focused instead on the event that deals with a person moving through a door opposite to the normal flow of traffic, where "normal flow" is defined.

Action items (NIST):

- Add list of things to resolve to the evaluation plan.

- Investigate the relative importance of false alarms vs. missed detections in airport security operation.

- Propose a value for Beta that weights false alarms vs. missed detections.

- Next part of evaluation plan will include details of the alignment of system output to ground truth.

- Post Corpora for download.

Action items (Participants)

- Complete evaluation data agreements to receive data.

Telecon adjourn

Discussion of the reverse direction event

David Eichmann:

given a vector of direction, we need to detect a change:

-- how about 90 degrees?

Jon:

This may not be enough of an angle, for example, when turning to approach the elevator (as in camera 3).

Paul:

We were thinking of something similar, for example, at least 135 degree turns.

Stephanie, LDC:

Keep in mind the annotation burden here; we want to define it so there is good consistency.

Heather, LDC:

What if it's only considered a cross their path?

Dave:

Someone could simply be doing obstacle avoidance.

Sadiye:

Event is a little ambiguous; for example, diverting to go visit a shop is not necessarily going to be a "reverse".

Paul:

How long the person is moving in the opposite direction may be a factor.

Can suggest a change direction event, for example, turning away from a checkpoint.

Dave:

Not more ambiguous than "chatting", as in the event "a group forms".

Heather:

In fact, we've defined the "group forms" event to include communication.

If we define "reverse direction" as going back the way you came it would be easier.

Paul:

Also, moving away from something like a checkpoint is not supported by these camera views.

What about entering and exiting from the same entry point?

Stephanie:

Annotation burden may be difficult given the windowing needed for annotating this dataset:

-- may interact with consistency, so we need to be careful

Sadiye:

Reverse direction depends on the context:

-- there are a lot of benign examples in the middle of a terminal

-- may be more important around secure areas