

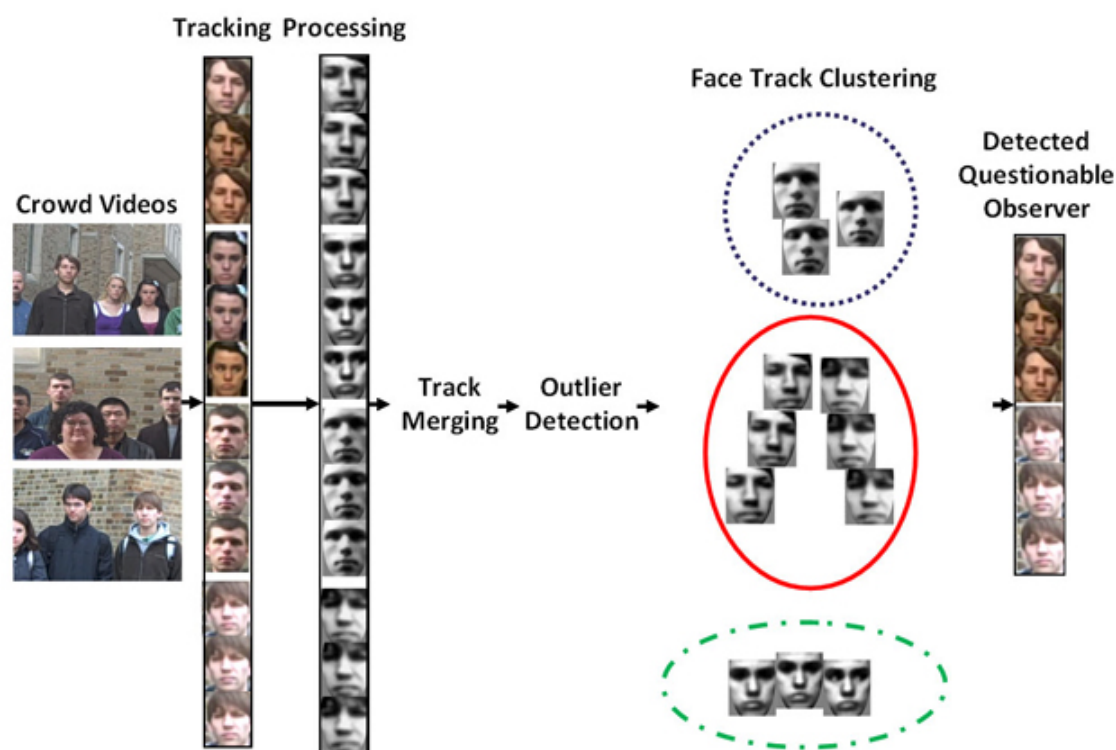


Where ideas and people meet

## Crime-Fighting Technology Spots Lingerin Arsonists in the Crowd

BY LAKSHMI SANDHANA Mon Jan 24, 2011

### The Questionable Observer Detection Algorithm

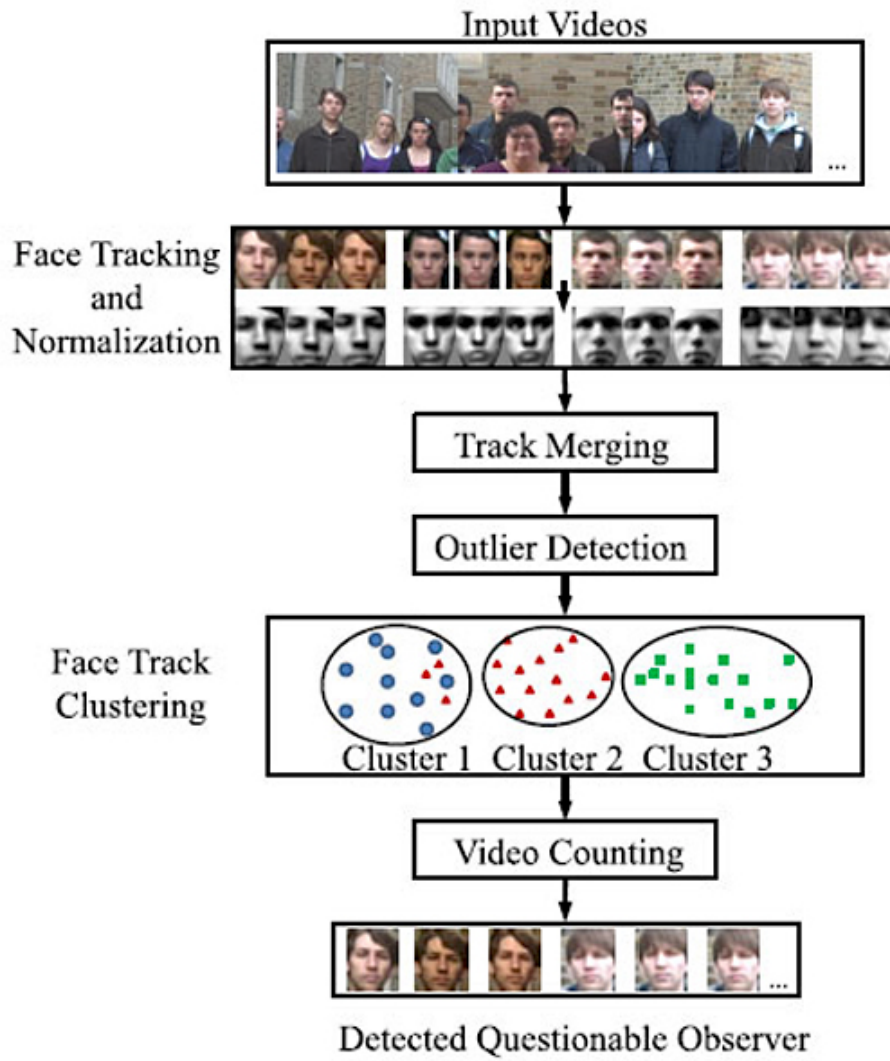


It's not easy to spot an arsonist in a crowd. Or, rather, it hasn't historically been easy for crime fighters. But if you have a "Questionable Observer Detector," or QuOD, on hand, it may not be that hard.

Currently under development by researchers at Notre Dame University, this advanced crime-fighting tool can pull out the faces of possible suspects when presented with multiple videos of lingering crowds at related crime scenes.

The idea here is that the criminal may be hanging out in the crowd that gathers after the event. In fact, arsonists are known to lurk within the crowd at the scene of a burning building; they are, in effect, watching their work in progress. Kevin W. Bowyer, a computer vision expert at Notre Dame, hit upon the idea of creating a practical tool to process all the available video clips of similar events, to see if it might spot someone turning up unusually often as part of the crowd, giving authorities a possible lead.

"If a suspicious person shows up in 10 videos, it should be enough to 'sound an alarm,' and tell the operators to look into who the person is and why they are showing up all the time," Bowyer tells *Fast Company*.



This is a hard problem to crack since it involves recognizing a face in various videos, which are shot most often by bystanders holding a standard handheld video camera. Besides poor resolution, snow, rain, and different lighting conditions, the perpetrators could have cut their hair, worn glasses, or be disguised in infinite ways. Bowyer’s most difficult challenge, however, was to create a viable 3-D facial recognition engine that didn’t rely on a backend database of existing images, but which could instead pull out individual faces from different videos and compare them.

“You have to group similar-looking people across video clips without knowing ahead of time whether anyone will actually show up in more than one clip,” says Bowyer.

The team’s approach is to create "face tracks" for every person in a video, repeating the process for all the video clips. The resulting face tracks are compared to see if any of the faces match each other or look similar. If the system spots a match, it creates a group for that person and any subsequent matching face tracks are added to the individual’s group.

A person is deemed "suspicious" if they appear in too many videos, the number being a variable figure that can be set by officials depending upon the number of input video clips.



Complicating factors in the experimental dataset. Top row: images of two questionable observers taken under varying illumination conditions. Middle row: images of another two questionable observers making distinct facial expressions in different videos. Bottom row: instances where subjects were occluded by other crowd members or their own body parts.



Face images from the Flip videos. Top row: blurry images recorded under full zoom. Bottom row: instances where the automatic exposure and white balancing adjustments changed facial appearance within the same face track.

While this sounds great in theory, in practice it’s a baffling problem that’s beyond the scope of current facial recognition technology, which needs proper lighting, good video resolution and face presentation to correctly identify people. “It will have to tackle two enormous challenges where previous companies have failed with facial recognition in the past--video quality and scalability across multiple cameras, camera resolutions, range and fields of view,” says Scott Schnell, President and CEO of VideoIQ, a company specializing in intelligent video surveillance.

The problem gets amplified further when you have video clips featuring huge crowds of people such as those likely to gather at a subway terminal or busy city street. “The challenge here is the horsepower and time required to process all of the data and arrive at a list of questionable observers,” says Timothy J. Whall, Chief Executive Officer of Protection 1, a home and business security monitoring company.

These aren’t insurmountable issues though. With processing speeds increasing rapidly, Bowyer’s getting closer to his dream of creating a reliable tool to present authorities with a neat list of possible suspects given the videos of, say, all the car bombings in Iraq or all the fires in Manhattan over a one-year period.

The team is presently at work on resolving false negatives, situations where the criminal might go undetected because they either looked very different in all the video clips or they didn’t make enough appearances in them to raise an alert. “There was a bad guy there, but the system never told the investigator to look at them, so the person doesn't get caught,” says Bowyer.

Currently the software takes about five hours to process 14 short videos, in order to detect possible suspects, and the team hopes to have a prototype at the end of a year. Once perfected, Bowyer hopes it will catch the interest of officials at the Federal Bureau of Investigation (FBI) and the Intelligence Advanced Research Projects Activity (IARPA).

So, what does this mean for all the rubber-neckers of the world? Some day will citizens have to worry about being arrested as a possible suspect just for stopping to stare like everyone else? Not unless you have a penchant for visiting and ogling multiple crime scenes. Criminals, however, will have to watch their backs--while crimes stoppers watch their faces.