



New AI Can Search For People By Height, Gender And Clothing

The invention of AI is comparable to the discovery of electricity: everyone in the world is using it for anything and everything. However, applications to control man are already more advanced than anything history has ever revealed. □ TN Editor

A team of AI researchers from India developed a tool to search for people in surveillance footage by height, clothing color, and gender. It's like a search engine that can find people in a video.

The scientists used deep learning (and Microsoft's COCO dataset) to train a convolutional neural network (CNN) how to recognize certain human features, called soft biometrics, using [computer vision](#).

Basically, you can tell this AI some details about the person you're looking for and it'll scour whatever video you give it. For example, a

request for “females wearing red shirts who are 153 cm tall” would, potentially, narrow down an entire video clip to just frames featuring people who meet that criteria.

According to researchers, the algorithm “correctly recovers 28 persons out of 41 in a very challenging dataset with soft biometric attributes.” Currently it only searches by height, torso (clothing) color, and gender.

At first glance the idea of identifying people in videos who meet relatively vague descriptions, and with accuracy that’s a little better than half, doesn’t sound like an important technological advance. But this early work shows plenty of potential. It’s worth asking what it would mean if accuracy could be improved beyond human capabilities.

There are scenarios where interested parties won’t know what they’re looking for in surveillance data in real time. This experimental CNN would be perfectly suited for use cases where we need to put together a timeline surrounding a specific individual, based on available historic surveillance footage.

Imagine a situation where a person is reported missing after two days. One might ask for footage from cameras in locations the individual was likely to have been in the vicinity of — like an oft-frequented gas station or campus the person attended. But after that, without any leads, it’s next to impossible to determine what footage to look at next. There could be millions of hours of video shot over two days within the confines of a particular set of city blocks.

However, if we could feed video to a [neural network](#) and let it narrow things down to a few hours of compiled footage, it would be possible to accurately track humans across multiple surveillance feeds.

This is exciting for numerous reasons. First, of course, the implications for finding missing persons or tracking suspected criminals are incredible. But perhaps just as important is the fact that this is a legitimate answer to the problem of [ubiquitous surveillance](#).

[Read full story here...](#)