FEATURE

A matter of resolution

Stephanie Fsadni learns about a University project aimed at facilitating the identification of criminal suspects from CCTV footage.

A team of University researchers is currently working on an algorithm that should facilitate and speed up the restoration of low-resolution facial images on CCTV videos.

Titled Deep-FIR (face image restoration), the project is adopting advanced artificial intelligence techniques to improve the performance of existing video forensic software packages adopted by forensic experts.

CCTV cameras cover large fields of view and typically capture very low-resolution facial images (around 10 to 40 pixels from side to side), making the identification of the subject of interest very difficult.

"Unlike what is featured in TV crime shows, restoring CCTV videos is a hard task and existing tools are quite limited," said Reuben Farrugia from the Department of Communications and Computer Engineering, who is leading the project together with Kenneth Camilleri from the Department of Systems and Control Engineering.

Methods used to restore frontal faces have had very limited success since the first resolution of a frontal facial image was increased in the year 2000, according to Farrugia.

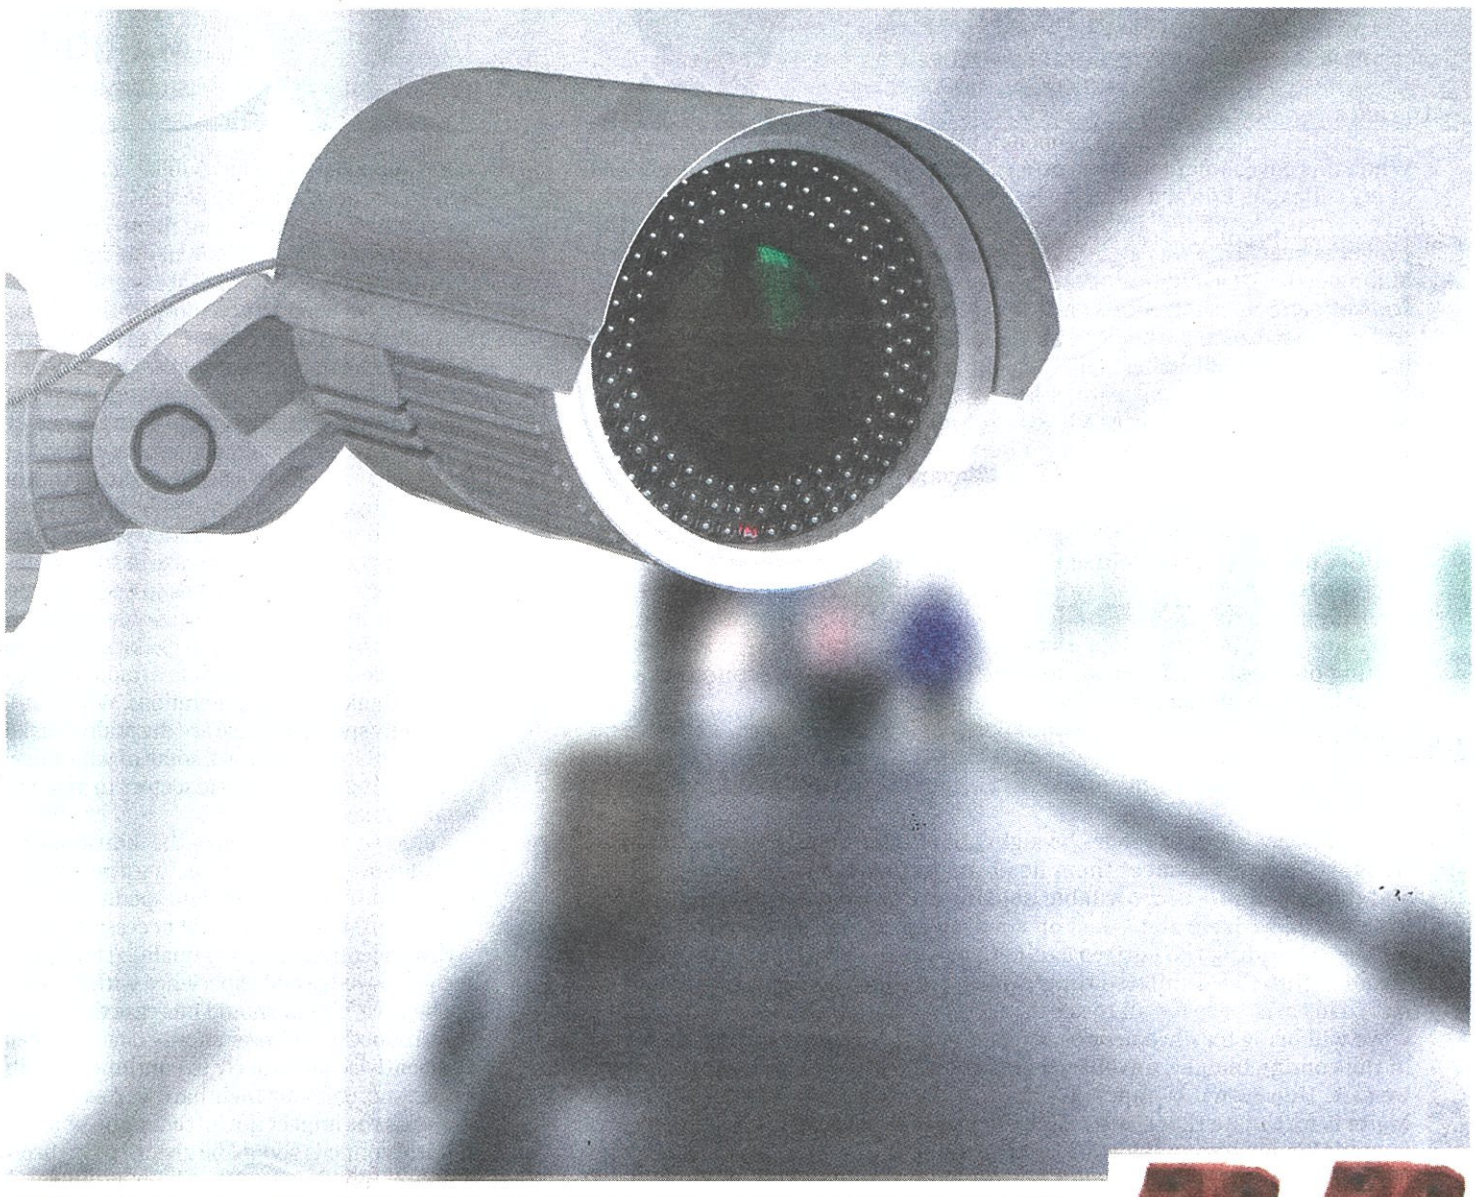

CCTV cameras cover large fields of view and typically capture very low-resolution facial images, making the identification of the subject of interest very difficult.

"Unlike what is featured in TV crime shows, restoring CCTV videos is a hard task and existing tools are quite limited"

"The culprit will not look straight into the camera. Moreover, all these methods completely ignore other distortions commonly present in CCTV video, that is artefacts caused by compression and also distortions caused by the camera lenses," he explains.

Farrugia started working on this "challenging problem" in 2015 after a visit to London's Metropolitan Police where he got a clear picture of the then state-of-the-art software technology used. He subsequently spent a year on sabbatical at INRIA – Bretagne Atlantique, a renowned French research institute, where, as part of the Sirocco team – which has a very good track record in image restoration – he tried to restore faces that were not looking straight into the camera and compressed images. This work was published in one of the leading journals in the field of image processing, *IEEE Transactions on Image Processing*, in 2017.

In the past two years, other methods which work on toy examples were used to increase the resolution of facial images but these could not guarantee that the restored face was that of the person being captured.

Existing forensic software packages have a function called 'super-resolution' which tries to register a number of video frames and then combines them together.

"This method works fairly well if the object is static, for example, a licence plate of a stationary car, but it will not work when the object is dynamic and non-rigid, as for example a face," Farrugia points out.

"Another limitation of these methods is that while the operator knows what he is restoring (a face, licence plate, etc) – the method used to restore is agnostic of what it is restoring. Therefore, the super-resolution method adopted has to generalise for all the objects."

These software packages are also very computationally intensive and time-consuming as the operator needs to mark a set of points on a face. With Deep-FIR, the operator will just have to mark the region where there is a face and the algorithm will try to restore that part, Farrugia says. The user will also be able to restore the whole head, including the hair region, which is important for person identification.

The advanced AI techniques used involve a data-driven approach. The Deep-FIR system will be made up of a large number of filters with millions of parameters that have to be learned to restore facial images.

"This will make the training phase very expensive in terms of computational time," he says. "It might also take weeks to train, but once these parameters are trained properly, then the restoration part is very fast – it should take less than one second to restore a face."

These restored images should have more than 40 pixels, which is considered high-resolution. Camilleri says that, however, this depends on the compression used.

"Since images and videos take up a large amount of space [in terms of memory requirements], these are processed by special algorithms that compress or reduce the amount of data, the image or video. When the image or video needs to be seen, the image data is restored from the compressed file using other algorithms. However, to obtain high levels of compression, very often algorithms throw away some of the original data, so when the images/videos are restored they can be corrupted.

"The loss of data may be small or large, depending on the compression algorithm used and how much compression was made. Therefore, this project will not only increase the resolution of images that were small to start with but will also mitigate the lost data."

Ascent Software, a local premier software development house with a vast experience in applied AI in various ICT-related sectors, is actively involved in the Deep-FIR project and is expected to develop the first working prototype that will be evaluated on real-world cases by 2019. This project is also supported by Amped Software, which is one of the leading international companies in forensic video enhancement software. Besides, the University team has contacted several forensic scientists across Europe to find out what they need in their software, such as how much time they are willing to wait to get a restored face.

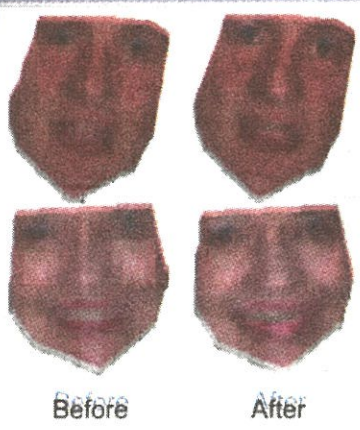

Deep-FIR has been granted around €200,000 by the Malta Council for Science & Technology, for and on behalf the Foundation for Science and Technology, through the Fusion: R&I Technology Development Programme.

More information about this project can be found on www.um.edu.mt/projects/deep-fir/.