ACCURACY CONSEQUENCES OF RECOGNIZING MULTIPLE FACES IN A SINGLE IMAGE

Image 1

Image 2

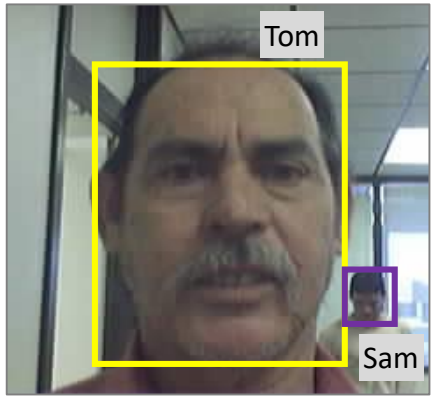

NON-MATE COMPARISON	Patrick	Yao
Tom	0.3	0.3
Sam	0.2	0.8

$\max(s_{11}, s_{21}, s_{12}, s_{22})$ → 0.8

Used in FMR computation

MATE COMPARISON	Mei	Patrick
Patrick	0.1	0.9
Yao	0.3	0.2

$\max(s_{11}, s_{21}, s_{12}, s_{22})$ → 0.9

Used in FNMR computation

Image 3

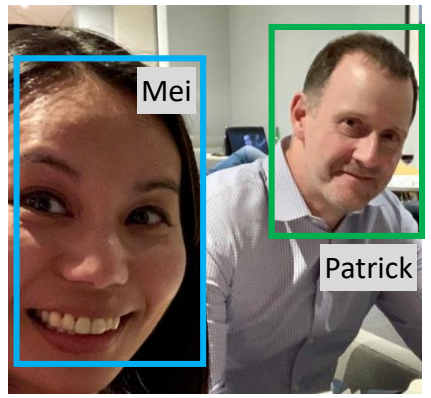

1. If an algorithm outputs N templates from a first image, and M templates from a second image, the NIST test software will call the 1:1 comparison function NM times and will retain only the maximum of the NM scores. For a mate (genuine) comparison, this will report the true mate score (0.9 above).
2. In the FRVT API prior to February 14, 2022 the algorithm under test may have found the correct face in image 1 (Patrick) but the incorrect face in image 3 (Mei) and reported a low score (0.1).
3. In the new API, the maximum score approach will be effective at finding the correct mate (0.9) and, over the entire mated test set, will reduce FNMR.
4. However, in a non-mate (impostor) comparison the maximum score approach will, over the set of non-mated pairs, increase FMR.
5. Further, some algorithms give high impostor scores when comparing a low resolution face with another low resolution face. The scores are similar to mate scores (0.8 is near 0.9, in the example). Over the test set, to achieve a low FMR (say 10^{-6}) a high threshold is needed and this greatly increases FNMR.
6. It is generally an undesirable property of an algorithm to produce high non-mate scores from face images with low information content. The software should either not produce templates from low resolution faces, or should correctly discount the lack of information in the photo.