<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper published in *IET Biometrics*. This paper has been peer-reviewed but does not include the final publisher proof-corrections or journal pagination.

Citation for the original published paper (version of record):

Ranftl, A., Alonso-Fernandez, F., Karlsson, S., Bigun, J. (2017)

A Real-Time AdaBoost Cascade Face Tracker Based on Likelihood Map and Optical Flow.

IET Biometrics, 6(6): 468-477

<https://doi.org/10.1049/iet-bmt.2016.0202>

Access to the published version may require subscription.

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:hh:diva-33836>

the inference. In short, it is a static algorithm which restarts face detection in each new frame. Apart from lacking temporal consistency, it does not save information about near-positives either (candidate windows that pass a significant number of stages of the classification cascade). Not only positions of detected faces may behave erratically due to restarting detection in every frame, but if a face becomes temporarily distorted or occluded, so that the very last part of the cascade fails, the detection fails abruptly.

This paper investigates a way to extend the Viola-Jones method to incorporate information about near-positives, as well as to include tracking capabilities. We extend our previous study presented in a conference paper [6]. The developed tracker takes advantage of optical flow calculations between frames. The results of the optical flow computations are used to interpolate a likelihood map with respect to time. The likelihood map describes the probability for a pixel being part of a face. It is built-up by summing up the number of classification stages that a window has passed during evaluation by the Viola-Jones algorithm. This way, information about near-positives is preserved, even if they do not pass all stages of the cascade of classifiers. Our face tracker calls this modified version of the Viola-Jones algorithm every n frames, and in the frames in between, face detection and tracking is done by interpolating the likelihood map from frame to frame with optical flow calculations. The proposed face tracking system shows a higher detection rate w.r.t. Viola-Jones, while keeping a good accuracy. More importantly, it enables face detection and tracking under partial or complete occlusion, and our system has less false negatives (i.e. face losses) due to the use of the proposed likelihood map. We also achieve less erratic movements in the bounding box surrounding the detected faces. Our system also compares well in terms of accuracy to two recently published algorithms based on Multi-task Cascaded Convolutional Networks [7] and Deformable Part Models [8, 9]. With an accuracy comparable to these newer approaches, our system has a detection speed one order of magnitude smaller.

The remaining of this section presents a literature review on the topics of face detection and tracking, as well as the contributions of this paper. Section 2 describes the proposed face detection and tracking system. The experimental framework, including baseline detection methods, database, performance metrics and results, is given in Section 3, followed by conclusions in Section 4. A supporting video can be found at <https://youtu.be/U7PiQu11Auo>.

1.1. Literature Review

1.1.1. Face Detection: Several algorithms have been developed for face detection [5, 10], mostly based in scanning the image with sub-windows in a raster-like fashion, and classifying the pattern in the sub-window as face/non-face. Scale change is overcome with image pyramids, a time consuming approach. The face/non-face classifier is learned from face/non-face training sub-windows using statistical learning methods. The rationale behind this is that pixels on a face are highly correlated, whereas those in a non-face sub-window present much less regularity [5]. Variations caused by lighting or illumination, face expression, pose or appearance makes necessary the use of non-linear classifiers. Examples are neural networks [11, 12] or Support Vector Machines [13]. But the breakthrough in face detection happened with Viola-Jones detector [14]. This algorithm has evolved as a de-facto standard [5], with an implementation included in OpenCV.

The Viola-Jones detector uses a set of Haar features for image scanning, which resemble features such as center-surround and directional responses. A variant of the AdaBoost learning algorithm is used to select strong features and to combine weak classifiers in cascade into strong ones, in a way that background regions are quickly discarded while spending more time (i.e. classifiers) on candidate, object-like windows. If a window passes all the classification stages, then it is con-

sidered to contain a face. Accuracy and speed are balanced with the number of classifiers and its individual performance, so a non-face candidate window is rejected in the early stages of the classification cascade. Since it is expected that a face is detected by several overlapped windows that differ slightly in position and/or scale, the number of false positives is further reduced by imposing a minimum number of windows that must overlap in order to have a face. Position of the detected face is thus refined by combining all the overlapped windows involved. The Viola-Jones algorithm is able to work in real time, and it has proven to work very well with near-frontal faces, but its performance is significantly decreased for faces at arbitrary poses and other degradations [15]. Multi-view face detection is usually dealt by training separate classifiers for each face view, but if pose estimation fails, e.g. a face is misclassified as frontal, it may never be detected by the frontal face cascade, not to mention the increased computational complexity [5]. Subsequent improvements of the Viola-Jones algorithm have included the use of more complex features to obtain a more accurate classification, such as diagonal Haar-like feature sets [16] or Multi-Block Local Binary Patterns (MB-LBP) [17]. The latter, which is regarded as faster and more discriminative than Haar-like features by its authors, is the implementation that we will employ in this paper.

More recently, the cascade framework (i.e. quick rejection of background regions) has been applied to features learned by Convolutional Neural Networks (CNN) [18], and the amount of research works on face detection making use of CNNs is exploding, e.g. [7, 19, 20, 21], inspired by the remarkable recent success of CNNs in many computer vision tasks. A drawback of these approaches, and of many approaches for unconstrained face detection, is that they usually need a considerable amount of annotated training data, apart from being computationally expensive [7]. Unconstrained face detection has gained attention recently also due to benchmarks for the development of unconstrained face detection algorithms, such as the Face Detection Data Set and Benchmark (FDDDB) [22], or the Annotated Facial Landmarks in the Wild (AFLW) [23]. Proposed approaches for unconstrained face detection include for example SURF features in an Adaboost cascade [24], Probabilistic Elastic Parts (PEP) [25], Deformable Parts Models (DPM) [9], or Normalized Pixel Differences (NPD) in a deep quadratic tree [26]. However, accuracy of unconstrained face detection is not yet satisfactory, especially when the detector is required to have low false alarms [26].

1.1.2. Face Tracking: A number of algorithms to track detected faces in videos have been proposed in the literature. Face tracking algorithms can be classified in several ways [27], for example considering if the whole face is tracked as a single entity (head tracking), if individual face features are tracked instead (feature tracking), or whether the tracking is done in the 2D or in the 3D space (with the latter facilitating estimation of the 3D pose). To track facial features, it is useful to know the head pose in order to have a rough estimate of the location of the features, but on the other hand it is also possible to estimate head pose by using the detected/tracked facial features. It is usual then that detection and tracking of the whole face and individual features is done simultaneously in a two-levels hierarchy, so one approach complements the other, as for example if Active Appearance Models are used [28, 29, 30]. Among the main challenges that a face tracking system has to solve, they include variations of pose and lightning, facial deformation, occlusions, or real-time capabilities [31]. Most of the methods only address one or two of these challenges, and many of them need to be trained or initialized manually [32]. Several authors have proposed the use of optical flow estimation to aid in the track of faces between frames [33, 34, 35, 36, 37, 32]. Other works, not discussed here, make use for example of mean-shift [38, 39, 40, 41, 42, 43], particle filters [44, 45, 46, 47], Kalman filters [48, 49, 50], Active Contour Models [34, 35, 51],

wavelet transform [52] or Bayesian Networks [53], or Markov Random Fields [54]. This list is not exhaustive but only gives works published after 2000. We refer the reader to [55, 56] for additional information. Recently, unconstrained correlation filters based method show a promising performance on benchmark datasets [57, 58], with works on face tracking making use of this technique too [59, 60, 61, 62, 63].

The optical flow describes the displacement of features from one image to another image within an image sequence [64], giving the spatial offset. Two main types of optical flow methods exist: sparse and dense. Dense flow methods operate on all pixels of the image, whereas sparse methods operate only on a subset of them. Those using sparse data rely heavily upon existence of discriminatory and identifiable local features. An example is the algorithm of Shi and Tomasi [65], which selects points possessing significant image gradients in two orthogonal directions, which are considered good features for tracking. A neighborhood containing two orthogonal directions occurring equally frequently is known as a perfectly balanced neighborhood [66]. It is called so, because translation of the neighborhood in an arbitrary direction will produce the same error (measured in L2 norm), which also happens to be the maximal error, the translational error signaling isotropically the inflicted displacement maximally. Such neighborhoods are also known as lacking linear-symmetry [66] or corner-like [67], and can be conveniently found via eigenvalues of the structure tensor in image neighborhoods of any dimension, not just two, and even in non-cartesian coordinates [64].

Other inherent problem of sparse methods is correspondence, i.e. which sparse point corresponds to which between two frames, along with the occlusion of points by moving objects. In the work of [33], face detection is done by applying 3-D deformable face models. Optical flow information (computed from a sparse set of points detected via [65]) is combined with edge information using a Kalman filter, with the aim of minimizing the noise which is present in the optical flow calculations. The algorithm, however, is not suitable for real-time processing. The method in [37] is based on the principle of Tracking-Learning-Detection (TLD). Each frame is processed by a frontal face detector and a tracker, and their outputs are passed to an integrator which estimates the object location. According to the TLD principle, the face under tracking is simultaneously learned, so the detector is updated with information of the object being tracked. Displacement of sparse candidate patches for tracking is estimated with the Lucas-Kanade method [68]. The work in [32] tracks mouth and eyebrow patches after the head is detected, with the aim of capturing facial expressions. Other works have dealt with the issue of tracking the mouth for lip reading purposes by optical flow using Active Contours [34, 35] or shape-from-motion features [69]. Tracking the mouth of a talking person is useful for many applications such as identity recognition or speechless human-computer interaction. Lastly, the method in [36] tracks pupils and nostrils as well by optical flow, in addition to mouth corners.

1.2. Contributions

This paper extends our previous study [6] with additional experimental results. The face detection and tracking system is described more formally and in greater detail. The developed method utilizes a modified version of the well-known Viola-Jones algorithm for face detection, which is complemented with the addition of tracking capabilities via optical flow computation. We also include two recently published algorithms based on Multi-task Cascaded Convolutional Networks [7] and Deformable Part Models [8, 9] to our comparison framework.

Within the Viola-Jones method, classification is done with a cascade of 20 stages, and a candidate window has to pass all the stages to be considered a face window. It does not keep information

Fig. 1: Flow chart of the developed optical flow based face tracker. The likelihood map L_t of the input frame and the detected faces are also shown. Grey intensity in the likelihood map indicates probability (re-scaled to 0-255 for visualization purposes).

about near-positives, and does not consider detection information from previous frames. Instead, Viola-Jones restarts face detection in each new frame. In contrast, the developed approach calls Viola-Jones method only every $n > 1$ frames for refreshing the likelihood map. We modify the OpenCV implementation of the Viola-Jones face detector in the way that a likelihood map is built. The likelihood values within this map are dependent on the numbers of classification stages of the cascade that each detection window passes. In order to track the faces and preserve information from previous frames, position of the likelihood map in the next frame is interpolated via dense flow map. When Viola-Jones is triggered again, the likelihood map is refreshed by combination of the likelihood map interpolated from the previous frame, and the newly generated likelihood map from the modified Viola-Jones cascade. This way, information from previous features is not dismissed when Viola-Jones is called again.

Compared to the original Viola-Jones implementation, our approach enables faces to be detected even when they do not pass all of the stages of the cascade classifier. An image region gets also a high value in the likelihood map if the respective window passes a certain number of classification stages, and the likelihood map is populated in proportion to the number of stages passed by the window. This way, information about the reliability of detection of each window is preserved, in contrast to the *all-or-nothing* approach of Viola-Jones. Another advantage of the developed face tracker is that it can also track faces during partial occlusion and even under complete occlusion, due to interpolation of the likelihood map from frame to frame by optical flow. Our approach is also one order of magnitude faster than the two other algorithms evaluated in this paper, while keeping a similar detection accuracy, something that can be an advantage in environments with limited computing resources.

2. Optical Flow Enhanced AdaBoost Cascade Face Detector and Tracker

2.1. Overview of the Proposed Algorithm

A flowchart of the proposed system can be found in Figure 1. In the proposed algorithm, the Viola-Jones algorithm is executed every n frames (called *refresh* frame), but instead of using its binary output of detected/non-detected faces, we produce a likelihood map which computes for every pixel the probability of having a face located at that respective position. The mentioned probability is derived from the number of classifiers of the cascade that every image window has passed, and face position is then extracted from the likelihood map by thresholding. In addition, the optical flow between two consecutive frames is computed. In intermediate frames where Viola-Jones is not executed, position of the new likelihood map is interpolated according to optical flow

computations, and it is refreshed by recursive filtering after n frames, when Viola-Jones is re-executed. If no face can be detected in a refresh frame, the modified Viola-Jones is executed again in the next frame, with this procedure repeated until a frame with a face is obtained. In the meantime, the likelihood map is updated frame by frame using optical flow information.

The proposed setup allows to detect a face window as long as it passes a sufficient number of classifiers, and prevents that it is lost between frames by means of dense optical flow information even in case of occlusion. Additionally, size of the face window does not change erratically as in the original Viola-Jones algorithm, which handles face detection frame by frame separately, so that the face window may appear unstable even if the face does not move.

2.2. Likelihood Map Computation

When Viola-Jones is executed at time t , a refresh likelihood map L is built by summing up at each pixel the number of stages of the Viola-Jones classification cascade passed by each detection window. In order to preserve temporal information from the previous frames, the likelihood map L_t at time t is then computed recursively as:

$$L_t = (1 - \alpha) \times L + \alpha \times L'_t \quad (1)$$

where L'_t is the likelihood map L_{t-1} at time $t - 1$, interpolated using the optical flow between frames at $t - 1$ and t . If Viola-Jones does not need to be called at time t , the likelihood map is just updated as $L_t = L'_t$.

Time savings can be achieved by computing L using only windows that have passed a certain number τ of classifiers. To compute L , we set every pixel of the detection window to the corresponding number of classification stages that have been passed by that window. Figure 1 (right) gives an example of a likelihood map from an image with two faces. All detection windows having passed a number τ of classifiers are summed up, but before a window is added, it is shrunk by a factor s . This is because Viola-Jones detection windows are usually bigger than the faces outlined by them, including pixels that do not contain a face. Since pixels not belonging to the face usually will not contain motion that is concurrent with the head movement, boundaries of the likelihood map would not be correctly interpolated with the optical flow if these pixels are kept. At the same time, this allows to separate windows from faces that appear adjacent to each other.

2.3. Likelihood Map Interpolation by Optical Flow

We use the Farneback dense flow method [70], which is suitable for real-time applications and exists as ready-to-use function in OpenCV. Dense flow methods compute displacement of all pixels of the image. We employ dense flow as we seek pixel-wise interpolation of the likelihood map in order to account for non-rigidity of the face. Sparse flow is faster to compute, but it would not be suitable for the proposed interpolation between frames, since we want to interpolate the set of adjacent pixels that represent a face in the likelihood map (Figure 1, right). Geometric 2-D interpolation of L_{t-1} to L'_t is done using the OpenCV function `remap()`.

2.4. Face Detection from the Likelihood Map

To extract the face area from L_t , we first apply a threshold c to binarize the likelihood map. This is to avoid false positives at smaller peaks. We then label individual adjacent areas as different faces, and their centers are used as the associated face centers. To compute the bounding box outlining the face, the individual areas are stretched by the inverse of the shrinking factor s applied

to the windows before they were summed up to build the likelihood map. This procedure outputs a rectangle instead of an square (as it is the case of Viola-Jones), allowing a more accurate estimation of the face region. Figure 1 shows the likelihood map of an image containing two adjacent faces, as well as the detected faces using the procedure indicated in this section. Areas containing the two faces show a high probability, which decreases towards the boundaries of the white regions.

3. Experimental Framework

3.1. Baseline Face Detection Methods

We also conduct face detection experiments using two recently published algorithms [7, 8]. The method in [7] is based on a Multi-task Cascaded Convolutional Network (MTCNN) with three stages of deep convolutional networks (CNN) which predict face and landmark location in a coarse-to-fine manner. The complexity of the CNN increases from stage to stage, in a way that quick rejection of non-face windows is done at the beginning, then a more complex CNN is used in the last stage to refine the results. The training databases employed include WIDER FACE [71] and CelebA [19] databases, while evaluation results reported in [7] were obtained using Face Detection Data Set and Benchmark (FDDB) [22], Annotated Facial Landmarks in the Wild (AFLW) [23] and WIDER FACE databases. We use the code released by the authors¹ for our experiments, which is implemented using Caffe. Regarding the method of [8], it is based on the Deformable Part Models (DPM) algorithm of [9], which is included in the release of the VGG-Face CNN recognition descriptor² described in [8], that we also use for our experiments here. The implementation used is run using MatConvNet. The method originally described in [9] was trained using images from AFLW [23] and Pascal Faces [72] databases, while evaluation was done using Annotated Faces in the Wild (AFW) [73] and FDDB [22] databases.

3.2. Database and Evaluation Metrics

Experiments with the Viola-Jones algorithm have been done in a Samsung 5 Ultra 530U3C A0A laptop (i3-3217U processor, 8 Gb DDR3 RAM, built-in Intel HD Graphics 4000) with MS Windows 7 Pro. The proposed face tracking algorithm has been implemented in C++ using OpenCV 2.4.8.0 with activated parallelization. The Viola-Jones algorithm (with Multi-Block Local Binary Patterns, or MB-LBP features) and the Farnebäck dense optical flow method are from OpenCV as well. The utilized implementation of Viola-Jones uses a cascade of 20 classifiers. Experiments with MTCNN and VGG face detectors have been done in a Dell E7240 laptop (i7-4600 processor, 16 Gb DDR3 RAM, built-in Intel HD Graphics 4400) with MS Windows 8.1 Pro. MTCNN is implemented in Caffe, while VGG is implemented in MatConvNet, and both are called from Matlab r2014b x64. Regarding the experimental dataset, we have used the Boston Head Tracking Database [74], having 45 videos of resolution 320×240 recorded at 30fps under uniform light conditions. Each video has 199 frames and contains one person moving the head around. Ground truth information is available through the UVAEYES annotations which have the positions of the eyes in each frame (see Figure 2 for an example) [75]. This database is also employed in some previous face tracking studies, e.g. [32, 76, 77].

The developed face tracker has some parameters that have been set empirically. Viola-Jones is called every $n=20$ frames and when this is done, the updated likelihood map is computed using

¹https://kpzhang93.github.io/MTCNN_face_detection_alignment

²http://www.robots.ox.ac.uk/~vgg/software/vgg_face

Fig. 2: Example frames (23, 92, 100, 150 and 188) of the video jam1.avi from the Boston Head Tracking Database. Red dots in the first row indicate ground-truth positions of the eyes, and green dots represent face centers computed from the ground-truth. The blue rectangles and blue dots represent the output given by the respective detector (note that MTCNN also outputs facial landmarks, while the other detectors only output the face bounding box). In frame 23, a refresh is done by the Viola-Jones algorithm (denoted by the white bounding box in the top left image). In frames 150 and 188, the Viola-Jones algorithm resulted in a false negative (no face detected).

$\alpha=0.5$. The reject level threshold for accepting windows to contribute to the likelihood map has been set to $\tau=15$, the shrinking factor used is $s=1/3$, and the binarization threshold is $c=65$. The latter value of c is set by knowing that a window containing a face increments corresponding pixels by at least 15 (due to τ), and that Viola-Jones usually output several windows of different sizes centered around the same face. We compare our method with the original Viola-Jones face detector, and with MTCNN and VGG detectors. Accuracy of the algorithms is measured by the Euclidean distance between the computed face center, and the ground-truth face center $[x, y]$, which is extracted via Equations 2 and 3:

$$x = x_1 + \frac{x_2 - x_1}{2} \quad (2)$$

$$y = y_1 + \frac{y_2 - y_1}{2} + 10 \quad (3)$$

Fig. 3: Percentage of correctly detected face centers per video. Better seen in colour.

These calculate the middle point between the eyes from the ground-truth information (x_1, y_1 and x_2, y_2 respectively), and add 10 pixels in y (vertical) direction to estimate the face center. Origin of coordinates is in the upper left corner of the image. This works quite well as there is not much difference in the sizes of faces in our database.

The evaluated algorithms output a bounding box with the detected face in a frame (see Figure 2). The center $[x_t, y_t]$ of the bounding box is used as the detected face center, and it is counted as a valid detection if its distance to the ground-truth is lower than 20 pixels, otherwise it is counted as a false positive. In case of multiple detections in a frame, all the distances to the ground-truth are computed, and the nearest one is chosen as valid, as long as it is within the threshold defined above. All other detections are counted as false positives. If no bounding boxes are detected in a frame, it is counted as a false negative. It should be noted that only one false negative per frame can occur, while the number of false positives per frame can be higher than one.

3.3. Detection Rate

The detection rate r of a video is calculated by summing up the valid detections d_i at each frame, and dividing by the number of frames n_{frames} (ideally $d_i = 1$ in each frame):

$$r = \frac{1}{n_{frames}} \times \sum_{i=1}^{n_{frames}} d_i \quad (4)$$

Figure 3 shows the average detection rate per video of the Boston Head Tracking Database. The average detection rate (across all videos) of the proposed optical flow face tracker is of 93.68%, superior to the Viola-Jones detector, which has an average detection rate of 86.81%. The other two detectors have an average detection rate of 99.17% (MTCNN) and 95.16% (VGG). Therefore, the modifications proposed to the Viola-Jones face tracker brings the detection rate just one step behind newer methods based on Convolutional Neural Networks or Deformable Part Models.

(a) Average detection accuracy per video.

(b) Detection accuracy in video jam1.avi

Fig. 4: Detection accuracy. Better seen in colour.

3.4. Detection Accuracy and False Detections

Figure 4a shows the average detection accuracy per video of the face detectors. The false negatives and false positives per video are also given in Figure 5. The average accuracy over the whole database is of 5.99 pixels (our algorithm), 2.56 pixels (Viola-Jones), 5.45 (MTCNN) and 11.07 (VGG). Our detector shows an accuracy comparable to MTCNN, with both algorithms having very few false negatives (face loss). We do not consider, however, these numbers to be representatively different than the accuracy of Viola-Jones, given the fact that we are estimating the face center from the ground-truth by adding a fixed amount of 10 pixels in vertical direction to the middle point between the two eyes, but this offset may be slightly different for each frame. These numbers, however, are computed considering only the frames where a face is detected, so frames with false negatives are not counting towards these values, but Viola-Jones and VGG show false negatives (face loss) in a number of frames (Figure 5, top).

In addition to a number of false negatives and positives in some videos, it is worth noting the worse accuracy of the VGG algorithm. A further analysis of the output images given by this

Fig. 5: False negatives and positives per video. Better seen in colour.

detector reveals that it tends to overestimate the face region (see Figure 2), resulting in bounding boxes that are slightly bigger than the face (even in frontal upright faces, as can be seen in the bottom left image). Because of that, the center of the bounding box will be displaced, resulting in a worse accuracy, or in a false positive if the distance to the ground-truth exceeds our threshold.

As mentioned above, the Viola-Jones detector produces a number of false negatives in some of the videos. This is because a window has to pass the 20 classifiers of the cascade to be considered a face, which may not be the case in certain poses. This condition, on the other hand, causes Viola-Jones to show few false positives. The number of false negatives in the optical flow detector is kept low due to: *i*) the likelihood map, which considers if a window has passed most of the classifiers of the cascade (even if it does not pass the 20 classifiers), and *ii*) the interpolation of the likelihood map via optical flow, which allows to track the face across frames where Viola-Jones may not detect a face. By design, the optical flow face tracker will always detect a face as long as it is correctly initialized with one. On the other hand, if Viola-Jones produces false negatives, the likelihood map is not refreshed with new face information, causing that the tracker becomes inaccurate and produces false positives over time. For this reason, the videos where our tracker has a high number of false positives are videos where Viola-Jones has a high number of false negatives too (Figure 5). Figure 4b shows the accuracy of the detectors across frames in the video jam1.avi (Figure 2). Frames with false negatives are marked with a value of -5. When Viola-Jones loses the face, the optical flow detector is able to track its position, although the distance to the ground-truth center starts to increase. The latter eventually would cause false positives after a period of time. However, when Viola-Jones detects the face again, the optical flow tracker regains a better accuracy. It can also be observed in Figure 4b the similar accuracy of the Optical Flow and MTCNN detectors, as indicated above.

(a) Average detection stability per video.

(b) Detection stability in video jam1.avi

Fig. 6: Detection stability. Better seen in colour.

3.5. Detection Stability

Erratic movements of the face across frames is not desirable in applications, as faces move relatively slow in practice. The stability e is measured as the Euclidean distance between the face centers of two successive frames with valid detections:

$$e = \sqrt{(x_t - x_{t-1})^2 + (y_t - y_{t-1})^2} \quad (5)$$

Figure 6a shows the average stability per video. In general the optical flow face tracker has less erratic movements. The average on all videos equals 0.9 pixel for the optical flow face tracker, 1.5 pixels for the Viola-Jones detector, 1.8 for MTCNN, and 2.8 for VGG. By increasing α , which determines the weight of the interpolated likelihood map in Equation 1, the peaks of the optical flow face tracker could be lowered. However, it is likely that the accuracy and the detection rate be worsened. An improvement to further smooth the bounding box coordinates would be to incorporate longer sequences in the computation of the likelihood map via multi-frame optical flow computation [78], although this would imply higher computational cost. Figure 6b shows the sta-

(a) Average computation time per frame per video.

(b) Average computation time per frame per video of the optical flow tracker. Time is split into computation time of optical flow (gray) and of the remaining operations (blue).

Fig. 7: Computation time.

bility across frames in the video jam1.avi, further illustrating the higher smoothness of the optical flow tracker in the detection. The values of Viola-Jones algorithm are partly discontinuous due to false negatives in the respective frames.

3.6. Detection Speed

Speed of the algorithms is measured in milliseconds. The average computation time per frame in each video is shown in Figure 7a. The Viola-Jones algorithm is faster (6 ms in average vs. 13 ms), although both are capable of working in real time. The average computation time of the other two algorithms is 137.18 ms (MTCNN) and 253.74 ms (VGG), which is considerably higher despite being executed in a more powerful machine (Section 3.2). The latter two systems could be benefited of GPU acceleration (not used here), although it may not be feasible in some applications involving limited computing resources, e.g. smart-phones. It has to be mentioned too that the average CPU usage of Viola-Jones is higher due to the parallelization of the OpenCV implementation utilized, resulting in lower computation time. In contrast, the Farnebäck optical flow implementation employed only uses one core, reducing the average CPU usage but increasing the execution time. Figure 7b shows the split between computation of the optical flow (gray) and the rest of our tracking algorithm (blue), where it can be observed that the optical flow takes the majority of the computation time (11.33 ms in average vs. 1.67 ms of the remaining operations). Based on the above considerations, however, significant time savings could be achieved by parallelizing the optical flow computation, since it is the most time-consuming operation. The proposed algorithm

Fig. 8: The optical flow face tracker is able to track faces under partial occlusion (left) and also under complete occlusion (right).

could be also ideally used in applications demanding face detection, and where the optical flow should be computed anyway, as lip-motion analysis [69].

3.7. Tracking During Occlusions

The developed face tracking algorithm is able to track faces during partial and complete occlusion under certain circumstances, such as fast moving objects passing (like cars). This is because the Farnebäck dense flow technique employed is invariant to very fast motion. Figure 8 illustrates two examples where the developed face tracker is able to track faces under occlusion. In these two cases, the Viola-Jones algorithm is not able to detect the face.

4. Conclusion

We present here a novel real-time face tracking algorithm which utilizes a modified version of the Viola-Jones algorithm for face detection. In contrast to a pure Viola-Jones face detector, the developed approach calls a modified Viola-Jones method only at every $n=20$ frames for refreshing a likelihood map. The likelihood values within this map are dependant on the numbers of classification stages which each detection window passes, thus preserving information about near-positives. In order to track faces, the likelihood map is interpolated frame by frame with a flow map computed by the Farnebäck dense optical flow method. The resulting likelihood map of the modified Viola-Jones algorithm contributes to the system's likelihood map by recursive filtering every n frames.

The developed face tracking algorithm and the original Viola-Jones face detector have been evaluated on the Boston Head Tracking Database. We also use in our benchmark two recently published face detection algorithms based on Multi-Task Cascaded Convolutional Networks (MTCNN) and Deformable Part Models, the latter included in the release of the VGG-Face CNN recognition descriptor (VGG). The developed face tracker achieves a higher detection rate than the Viola-Jones face detector, and just one step behind the other two methods. The valid detections of our method are less accurate than Viola-Jones, but on average the distance to ground truth face center is below 6 pixels. This is a very good value considering that the coordinates for the face center are derived from ground truth information of eye coordinates. The accuracy of our method is comparable to MTCNN, while VGG tends to overestimate the size of the face bounding box, resulting in worse accuracy and a higher number of false negatives. Furthermore, the optical flow face tracker shows less erratic movements of detections than any other detector.

In terms of false detections, the Viola-Jones detector shows more false negatives, because it demands that candidate windows have to pass the whole cascade of classifiers to be considered a face, which on the other hand results in less false positives. Our system shows zero false negatives since it allows candidate windows to pass only a number of classifiers of the cascade, and once

a face window is detected, it is interpolated to subsequent frames via optical flow. Compared to the original Viola-Jones implementation, the likelihood map approach enables faces to be detected even when they do not pass all of the stages of the cascade classifier. Due to the fact that the likelihood map is never discarded completely, a region gets also a high value in the likelihood map if the respective window passes for example 17, 18 or 19 stages within several executions of the modified Viola-Jones method, allowing to reduce false negatives. Another advantage of the developed face tracker is that it can also track faces under partial and complete occlusion. On the other hand, our system is dependant of the success of Viola-Jones in finding a face, so in videos where Viola-Jones has many false negatives, our system has a higher number of false positives since it cannot receive a refresh frame.

In comparison to the original Viola-Jones algorithm, our face tracker needs more computation time. The Farnebäck flow computation takes a majority of computation time as this method has been used ‘as is’ in OpenCV. The demand in computational power is caused by the fact that Farnebäck flow is not parallelized in OpenCV, whereas the Viola-Jones algorithm is. Subtracting the computation time of Farnebäck flow from the face tracker’s computation time, our system is about four times faster than the Viola-Jones detector. This fact makes the face tracker an appropriate tool where in applications where dense flow computation has to be done anyway and where face tracking is needed, e.g. lip analysis [69]. By optimizing the Farnebäck flow computation method, the tracker execution time could be lowered drastically as well. The MTCNN and VGG detectors, on the other hand, have an execution time one order of magnitude higher than Viola-Jones and our face tracker. These two systems could be benefitted of GPU acceleration, since they are implemented using CNN frameworks, although it may not be feasible in some applications with limited resources.

The developed tracker relies on the Viola-Jones algorithm, which means that errors or inaccuracies of the Viola-Jones algorithm during initialization or refresh of the likelihood map is propagated to our system. This may result in false positives after a period of time if Viola-Jones is not able to detect a face when it is called every $n=20$ frames. This effect could be minimized for example by triggering other methods for face detection with higher accuracy if Viola-Jones fails, or by using a detection algorithm with lower computation time in each frame to quick check if a face is lost when interpolating the likelihood map and if so, call Viola-Jones earlier than every 20 frames. Another possibility is to execute a modified version of Viola-Jones only on the extracted face areas, in order to check if the tracker has lost the faces or not. A complementary strategy to improve the accuracy of our method would be to make the weighting parameter α in Equation 1 dependant on the detection confidence of the algorithm used in a given frame, so the likelihood map is refreshed selectively in proportion to the confidence on the detection. A good algorithm with good detection accuracy and few false negatives could be the MTCNN method employed in this paper, although it should be considered that its execution time is considerably higher than Viola-Jones. To further improve the accuracy and stability of the detection, we are also exploring the use of multi-frame optical flow computation in order to consider longer sequences in the computation of the likelihood map [78].

The Viola-Jones framework is not perfectly suited for the implementation of our interpolated likelihood map, as some extra amount of tweaking was necessary to build up an initial map from the detection results. Viola-Jones was chosen for this work as it is the most well known and widely used approach to static face detection today, and it is able to work in real time [26]. However, this is rapidly changing as deep learning approaches are steadily improving on performance, with hardware dedicated to them being developed jointly (embedded GPUs). Although Viola-Jones

works very well with near-frontal faces, its performance significantly suffers with faces having arbitrary poses or other degradations. It could be expected that accuracy of MTCNN, VGG or other approaches for unconstrained face detection would be improved in the same way that Viola-Jones has been improved here. It is an interesting future approach to replace both Farnebäck flow and Viola Jones with deep learning methods, such as the MTCNN approach evaluated here [7], although some tweaking would be also necessary. DeepFace [79] is very promising as it improves performance greatly, it is even better adapted to detecting identity (not just face occurrence) and it outputs a likelihood map to use. Together with recent advances in optical flow computation using deep learning approach, especially FlowNet [80], we have a promising framework that is much better adapted to our likelihood map generation.

5. Acknowledgments

Author F. A.-F. thanks the Swedish Research Council for funding his research. Authors acknowledge the CAISR program and the SIDUS-AIR project of the Swedish Knowledge Foundation.

6. References

- [1] S. Z. Li and A. K. Jain, Eds., *Handbook of Face Recognition*, Springer Verlag, 2nd edition, 2011.
- [2] A. K. Jain and A. Kumar, *Second Generation Biometrics*, chapter Biometrics of Next Generation: An Overview, Springer, 2010.
- [3] J. Stallkamp, H. K. Ekenel, and R. Stiefelhagen, "Video-based face recognition on real-world data," in *Proc. Intl Conf on Computer Vision, ICCV*, Oct 2007, pp. 1–8.
- [4] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," *Proc Computer Vision and Pattern Recognition Conf, CVPR*, vol. 1, pp. 511–518, Dec 2001.
- [5] S. Z. Li, *Handbook of Face Recognition*, chapter Face Detection, Springer Verlag, 2nd edition, 2011.
- [6] A. Ranftl, F. Alonso-Fernandez, and S. Karlsson, "Face tracking using optical flow," *Proc Intl Conf of the Biometrics Special Interest Group, BIOSIG (Best Paper Award)*, pp. 1–5, Sep 2015.
- [7] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, Oct 2016.
- [8] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," *Proc British Machine Vision Conference, BMVC*, 2015.
- [9] Markus Mathias, Rodrigo Benenson, Marco Pedersoli, and Luc Van Gool, "Face detection without bells and whistles," *Proc European Conf Computer Vision, ECCV*, pp. 720–735, 2014.

- [10] M.-H. Yang, D. Kriegman, and N. Ahuja, "Detecting faces in images: A survey," *IEEE Trans Pattern Analysis and Machine Intelligence*, vol. 24, no. 1, pp. 34–58, Jan 2002.
- [11] H.A. Rowley, S. Baluja, and T. Kanade, "Neural network-based face detection," *IEEE Trans Pattern Analysis and Machine Intelligence*, vol. 20, no. 1, pp. 23–38, 1998.
- [12] K.-K. Sung and T. Poggio, "Example-based learning for view-based human face detection," *IEEE Trans Pattern Analysis and Machine Intelligence*, vol. 20, no. 1, pp. 39–51, 1998.
- [13] Y. M. Li, S. G. Gong, and H. Liddell, "Support vector regression and classification based multi-view face detection and recognition," *Proc Intl Conf Automatic Face and Gesture Recognition, FG*, 2000.
- [14] P. Viola and M. Jones, "Robust real-time face detection," *Intl Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2004.
- [15] D. Teferi and J. Bigun, "Evaluation protocol for the dxm2vts database and performance comparison of face detection and face tracking on video," in *Intl Conf Pattern Recognition, ICPR*, Dec 2008, pp. 1–4.
- [16] R. Lienhart, A. Kuranov, and V. Pisarevsky, "Empirical analysis of detection cascades of boosted classifiers for rapid object detection," *Proc DAGM 25th Pattern Recognition Symposium*, 2003.
- [17] L. Zhang, R. Chu, S. Xiang, S. Liao, and S. Li, "Face detection based on multi-block lbp representation," *Proc Intl Conf Biometrics, ICB*, 2007.
- [18] H. Li, Z. Lin, X. Shen, J. Brandt, and G. Hua, "A convolutional neural network cascade for face detection," in *Proc Intl Conf on Computer Vision and Pattern Recognition, CVPR*, Jun 2015, pp. 5325–5334.
- [19] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proc Intl Conf on Computer Vision, ICCV*, 2015, pp. 3730–3738.
- [20] S. Yang, P. Luo, C. C. Loy, and X. Tang, "From facial parts responses to face detection: A deep learning approach," in *Proc Intl Conf on Computer Vision, ICCV*, 2015, pp. 3676–3684.
- [21] C. Zhang and Z. Zhang, "Improving multiview face detection with multi-task deep convolutional neural networks," in *Proc IEEE Winter Conference on Applications of Computer Vision*, March 2014, pp. 1036–1041.
- [22] J. V. Jain and E. G. Learned-Miller, "FDDB: A benchmark for face detection in unconstrained settings," *Technical Report UMCS-2010-009, University of Massachusetts, Amherst*, 2010.
- [23] M. Köstinger, P. Wohlhart, P. M. Roth, and H. Bischof, "Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization," in *Proc. Intl Conf on Computer Vision Workshops, ICCV*, Nov 2011, pp. 2144–2151.
- [24] Jianguo Li, Tao Wang, and Yimin Zhang, "Face detection using surf cascade," in *Proc. Intl Conf on Computer Vision Workshops, ICCV*, Nov 2011, pp. 2183–2190.
- [25] H. Li, G. Hua, Z. Lin, J. Brandt, and J. Yang, "Probabilistic elastic part model for unsupervised face detector adaptation," in *Proc. Intl Conf on Computer Vision, ICCV*, Dec 2013, pp. 793–800.

- [26] S. Liao, A. K. Jain, and S. Z. Li, "A fast and accurate unconstrained face detector," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 2, pp. 211–223, Feb 2016.
- [27] W. Zhao, R. Chellappa, P. Phillips, and A. Rosenfeld, "Face recognition: A literature survey," *ACM Comput. Surv.*, vol. 35, no. 4, pp. 399–458, 2003.
- [28] F. Dornaika and J. Ahlberg, "Fast and reliable active appearance model search for 3-d face tracking," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 34, no. 4, pp. 1838–1853, Aug 2004.
- [29] F. Dornaika and J. Ahlberg, "Fitting 3d face models for tracking and active appearance model training," *Image and Vision Computing*, vol. 24, no. 9, pp. 1010 – 1024, 2006.
- [30] P. Maurel, A. McGonigal, R. Keriven, and P. Chauvel, "3d model fitting for facial expression analysis under uncontrolled imaging conditions," *Proc Intl Conf on Pattern Recognition, ICPR*, pp. 1–4, Dec 2008.
- [31] A.-K. Roy-Chowdhury and Y. Xu, *Encyclopedia of Biometrics*, chapter Face Tracking, Springer, 2nd edition, 2015.
- [32] E. R. Gast, "A framework for real-time face and facial feature tracking using optical flow pre-estimation and template tracking," M.S. thesis, Leiden University, 2010.
- [33] D. DeCarlo and D. Metaxas, "Optical flow constraints on deformable models with applications to face tracking," *Intl Journal of Computer Vision*, vol. 38, no. 2, pp. 99–127, 2000.
- [34] P. Delmas, N. Eveno, and M. Lievin, "Towards robust lip tracking," *Proc Intl Conf on Pattern Recognition, ICPR*, vol. 2, pp. 528–531 vol.2, 2002.
- [35] Z. Wu, P.S. Aleksic, and A.K. Katsaggelos, "Lip tracking for mpeg-4 facial animation," *Proc IEEE Intl Conf on Multimodal Interfaces*, pp. 293–298, 2002.
- [36] J. Chen and B. Tiddeman, "Multi-cue facial feature detection and tracking," *Proc Intl Conf Image and Signal Processing, ICISP*, pp. 356–367, 2008.
- [37] Z. Kalal, K. Mikolajczyk, and J. Matas, "Face-tld: Tracking-learning-detection applied to faces," *Proc Intl Conf Image Processing, ICIP*, pp. 3789–3792, 2010.
- [38] D. Comaniciu, V. Ramesh, and P. Meer, "Kernel-based object tracking," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 5, pp. 564–577, May 2003.
- [39] J. Tao and Y.-P. Tan, "A unified probabilistic approach to face detection and tracking," *IEEE International Symposium on Circuits and Systems, SCAS*, pp. 3797–3800 Vol. 4, May 2005.
- [40] H.-P. Huang and C.-T. Lin, "Multi-camshift for multi-view faces tracking and recognition," *Proc IEEE International Conference on Robotics and Biomimetics, ROBIO*, pp. 1334–1339, Dec 2006.
- [41] Q.-A. Nguyen, A. Robles-Kelly, and C. Shen, "Enhanced kernel-based tracking for monochromatic and thermographic video," *IEEE Intl Conf on Video and Signal Based Surveillance, AVSS*, pp. 28–28, Nov 2006.

- [42] P. Vadakkepat, P. Lim, L.C. De-Silva, Liu Jing, and Li Li Ling, "Multimodal approach to human-face detection and tracking," *IEEE Transactions on Industrial Electronics*, vol. 55, no. 3, pp. 1385–1393, Mar 2008.
- [43] P. S. Hiremath, Manjunath Hiremath, and Mahesh R., "Face detection and tracking in video sequence using fuzzy geometric face model and mean shift," *International Journal of Advanced Trends in Computer Science and Engineering*, vol. 2, no. 1, pp. 41–46, 2013.
- [44] D. Grest and R Koch, "Realtime multi-camera person tracking for immersive environments," *Proc IEEE Workshop on Multimedia Signal Processing*, pp. 387–390, Sep 2004.
- [45] C. Su and L. Huang, "Spatio-temporal graphical-model-based multiple facial feature tracking," *EURASIP Journal on Applied Signal Processing*, vol. 2005, pp. 2091–2100, 2005.
- [46] T. Yun and L. Guan, "Fiducial point tracking for facial expression using multiple particle filters with kernel correlation analysis," *Proc IEEE Intl Conf on Image Processing, ICIP*, pp. 373–376, Sep 2010.
- [47] Y. N. Chae, J. Ha, and H. S. Yang, "Development of an efficient face detection and tracking system for mobile devices," *Proc IEEE Intl Conf on Virtual Systems and Multimedia, VSMM*, pp. 192–196, Oct 2010.
- [48] Z. Zhu and Q. Ji, "3d face pose tracking from an uncalibrated monocular camera," *Proceedings Intl Conf on Pattern Recognition, ICPR*, vol. 4, pp. 400–403 Vol.4, Aug 2004.
- [49] V. Girondel, A. Caplier, and L. Bonnaud, "Real time tracking of multiple persons by kalman filtering and face pursuit for multimedia applications," *Proc IEEE Southwest Symposium on Image Analysis and Interpretation*, pp. 201–205, Mar 2004.
- [50] A. Destrero, F. Odone, and A. Verri, "A system for face detection and tracking in unconstrained environments," *Proc IEEE Conference on Advanced Video and Signal Based Surveillance, AVSS*, pp. 499–504, Sep 2007.
- [51] X. Bing, Y. Wei, and C. Chareonsak, "Automatic focusing technique for face detection and face contour tracking," *IEEE International Workshop on Biomedical Circuits and Systems*, pp. S3/2–9–12, Dec 2004.
- [52] B. Mohabbati and S. Kasaei, "Face localization and versatile tracking in wavelet domain," *Proc IEEE Conf on Information and Communication Technologies, ICTTA*, vol. 1, pp. 1552–1556, 2006.
- [53] E. Loutas, C. Nikou, and I. Pitas, "An information theoretic approach to joint probabilistic face detection and tracking," *Proc Intl Conf on Image Processing, ICIP*, vol. 1, pp. I–505–I–508 vol.1, 2002.
- [54] Baoyuan Wu, Bao-Gang Hu, and Qiang Ji, "A coupled hidden markov random field model for simultaneous face clustering and tracking in videos," *Pattern Recognition*, vol. 64, pp. 361 – 373, 2017.
- [55] R. Subban, S. Muthukumar, P. Pasupathi, A. Benedict, and J. Jayapal, "Face tracking techniques in color images : A study and review," *International Journal of Engineering Research & Technology, IJERT*, vol. 2, no. 12, 2013.

- [56] R. Jagathishwaran, K.S. Ravichandran, and Premaladha Jayaraman, “A survey on face detection and tracking,” *World Applied Sciences Journal*, vol. 29, pp. 140–145, 2014.
- [57] Matej Kristan, Roman Pflugfelder, Aleš Leonardis, Jiri Matas, Luka Čehovin, Georg Nebelhay, Tomáš Vojtř, Gustavo Fernández, Alan Lukežič, Aleksandar Dimitriev, Alfredo Petrosino, Amir Saffari, Bo Li, Bohyung Han, CherKeng Heng, Christophe Garcia, Dominik Pangeršič, Gustav Häger, Fahad Shahbaz Khan, Franci Oven, Horst Possegger, Horst Bischof, Hyeonseob Nam, Jianke Zhu, JiJia Li, Jin Young Choi, Jin-Woo Choi, João F. Henriques, Joost van de Weijer, Jorge Batista, Karel Lebeda, Kristoffer Öfjäll, Kwang Moo Yi, Lei Qin, Longyin Wen, Mario Edoardo Maresca, Martin Danelljan, Michael Felsberg, Ming-Ming Cheng, Philip Torr, Qingming Huang, Richard Bowden, Sam Hare, Samantha YueYing Lim, Seunghoon Hong, Shengcai Liao, Simon Hadfield, Stan Z. Li, Stefan Duffner, Stuart Golodetz, Thomas Mauthner, Vibhav Vineet, Weiyao Lin, Yang Li, Yuankai Qi, Zhen Lei, and Zhi Heng Niu, *The Visual Object Tracking VOT2014 Challenge Results*, pp. 191–217, Springer International Publishing, Cham, 2015.
- [58] Y. Wu, J. Lim, and M. H. Yang, “Object tracking benchmark,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 9, pp. 1834–1848, Sept 2015.
- [59] V. D. My and A. Zell, “Real time face tracking and pose estimation using an adaptive correlation filter for human-robot interaction,” in *Proc European Conference on Mobile Robots*, Sept 2013, pp. 119–124.
- [60] L. B. Zhou, H. Wang, W. Mou, and Z. C. Hu, “Robust face detection and tracking under natural conditions,” in *Proc Intl Conf on Robotics and Biomimetics, ROBIO*, Dec 2013, pp. 934–939.
- [61] Leopoldo N. Gaxiola, Víctor Hugo Díaz-Ramírez, Juan J. Tapia, Arnoldo Diaz-Ramirez, and Vitaly Kober, *Robust Face Tracking with Locally-Adaptive Correlation Filtering*, pp. 925–932, Springer International Publishing, Cham, 2014.
- [62] H. K. Galoogahi, T. Sim, and S. Lucey, “Correlation filters with limited boundaries,” in *Proc. Conf on Computer Vision and Pattern Recognition, CVPR*, June 2015, pp. 4630–4638.
- [63] Wei Chen, Kaihua Zhang, and Qingshan Liu, “Robust visual tracking via patch based kernel correlation filters with adaptive multiple feature ensemble,” *Neurocomputing*, vol. 214, pp. 607 – 617, 2016.
- [64] J. Bigun, *Vision with Direction*, Springer, 2006.
- [65] J. Shi and C. Tomasi, “Good features to track,” *Proc Intl Conf Computer Vision and Pattern Recogn, CVPR*, 1994.
- [66] J. Bigun and G.H. Granlund, “Optimal orientation detection of linear symmetry,” *Proc Intl Conf Computer Vision, ICCV*, pp. 433–438, Jun 1987.
- [67] C. Harris and M. Stephens, “A combined corner and edge detector,” in *Proc. of Fourth Alvey Vision Conference*, 1988, pp. 147–151.
- [68] B. D. Lucas and T. Kanade, “An iterative image registration technique with an application to stereo vision,” *Proc Intl Joint Conf on Artificial Intelligence, IJCAI*, pp. 674–679, Oct 1981.

- [69] S. M. Karlsson and J. Bigun, “Lip-motion events analysis and lip segmentation using optical flow,” *Proc IEEE Computer Vision and Pattern Recognition Biometrics Workshop, CVPRW*, 2012.
- [70] G. Farneback, “Two-frame motion estimation based on polynomial expansion,” *Proc Swedish Symposium on Image Analysis, SSBA*, 2003.
- [71] Shuo Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang, “Wider face: A face detection benchmark,” in *Proc. Conf on Computer Vision and Pattern Recognition, CVPR*, 2016.
- [72] Junjie Yan, Xuzong Zhang, Zhen Lei, and Stan Z. Li, “Face detection by structural models,” *Image and Vision Computing*, vol. 32, no. 10, pp. 790 – 799, 2014, Best of Automatic Face and Gesture Recognition 2013.
- [73] X. Zhu and D. Ramanan, “Face detection, pose estimation, and landmark localization in the wild,” in *Proc. Conf on Computer Vision and Pattern Recognition, CVPR*, June 2012, pp. 2879–2886.
- [74] M. La Cascia, S. Sclaroff, and V. Athitsos, “Fast, reliable head tracking under varying illumination: An approach based on registration of texture-mapped 3d models,” *IEEE Trans Pattern Analysis and Machine Intelligence*, 2000.
- [75] R. Valenti and T. Gevers, “Robustifying eye center localization by head pose cues,” *Proc Intl Conf Computer Vision and Pattern Recogn, CVPR*, 2009.
- [76] S. U. Jung and M. S. Nixon, “Model-based feature refinement by ellipsoidal face tracking,” in *Proc. Intl Conf on Pattern Recognition, ICPR*, Nov 2012, pp. 1209–1212.
- [77] C. N. Duong, T. C. P. Dinh, T. D. Ngo, D. D. Le, D. A. Duong, B. H. Le, and S. Satoh, “Robust eye localization in video by combining eye detector and eye tracker,” in *Proc. Intl Conf on Pattern Recognition, ICPR*, Nov 2012, pp. 242–245.
- [78] Michal Irani, “Multi-frame correspondence estimation using subspace constraints,” *International Journal of Computer Vision*, vol. 48, no. 3, pp. 173–194, 2002.
- [79] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, “Deepface: Closing the gap to human-level performance in face verification,” in *Proc. Conf on Computer Vision and Pattern Recognition, CVPR*, Jun 2014, pp. 1701–1708.
- [80] P. Fischer, A. Dosovitskiy, E. Ilg, P. Häusser, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox, “FlowNet: Learning optical flow with convolutional networks,” in *Intl Conf on Computer Vision, ICCV*, 2015.