



Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr

Eye pupil localization with an ensemble of randomized trees

Nenad Markuš^{a,*}, Miroslav Frliak^a, Igor S. Pandžić^a, Jörgen Ahlberg^b, Robert Forchheimer^b^a University of Zagreb, Faculty of Electrical Engineering and Computing, Unska 3, 10000 Zagreb, Croatia^b Linköping University, Dept. of Electrical Engineering, Information Coding Group, SE-581 83 Linköping, Sweden

ARTICLE INFO

Article history:

Received 3 December 2012

Received in revised form

17 April 2013

Accepted 1 August 2013

Keywords:

Eye pupil localization

Boosting

Randomized trees

ABSTRACT

We describe a method for eye pupil localization based on an ensemble of randomized regression trees and use several publicly available datasets for its quantitative and qualitative evaluation. The method compares well with reported state-of-the-art and runs in real-time on hardware with limited processing power, such as mobile devices.

© 2013 Published by Elsevier Ltd.

1. Introduction

The human eyes play an essential role in everyday interaction, communication and other routine activities. Thus, tracking the eyes and eye pupils opens a wide range of applications. Eye pupil localization can be used in biometric systems, gaze/focus estimation, human-machine interfaces, character animation, etc. Using the eyes as a means of human-computer interaction is helping disabled people to improve their daily lives, and may become a hands-free alternative in other applications or an entertaining element in innovative games. Eye tracking is increasingly finding use in safety applications to detect situations such as sleepiness or lack of attention while driving or using hazardous machinery.

High quality gaze estimation systems most commonly use head-mounted or multiple near-infrared cameras. Such systems can achieve high accuracy but they are expensive, intrusive or both. Furthermore, they often require calibration. Therefore, they are not well suited for user friendly applications working on off-the-shelf hardware where eye information is useful. There is a need for more versatile and simple systems and our research is directed this way. We are interested in the case when the images are supplied from a monocular video stream on consumer hardware, especially on mobile devices. These images may have low resolution or are often noisy and have bad contrast. Furthermore, we are interested in supporting a wide range of PC and mobile devices with limited

processing power. To satisfy such requirements, a robust yet high-performance system is needed.

In this paper, we describe a framework for eye pupil localization. Small computational cost of the developed framework makes it ideal for achieving high performance using off-the-shelf hardware. Only one uncalibrated camera is needed and this makes the framework completely non-intrusive. The implemented system runs in real-time on mobile devices.

1.1. Related work

Our work is related to a large body of research on eye and gaze tracking. An extensive overview has been done by Hansen et al. [18].

Of particular interest is the research done on eye center localization. Methods that report state-of-the-art results have been described in [29,31]. The basic idea of both approaches is to find the face within the image and then locate eye centers using geometric information (symmetry or curvature of the iris). Timm et al. [29] propose an approach based on analysis of image gradients. The idea is to define an objective function which obtains its maximum in the location where most gradient vectors intersect. Since the iris has a circular appearance in most frontal images, this also corresponds to the position of the pupil. Valenti et al. [31] use the curvature of isophotes (curves connecting points of equal intensity) to design a voting scheme for pupil localization. Additionally, they extend their approach by extracting a SIFT [24] vector for each candidate location and match it with examples in a database to obtain the final decision. In some cases this significantly increases the accuracy. Both methods are designed to deal with frontal faces and are accurate at solving this problem.

* Corresponding author. Tel.: +385 16129738; fax: +385 16129832.

E-mail addresses: nenad.markus@fer.hr, neno.markus@gmail.com (N. Markuš), miroslav.frliak@fer.hr (M. Frliak), igor.pandzic@fer.hr (I.S. Pandžić), jorgen.ahlberg@isy.liu.se (J. Ahlberg), robert@isy.liu.se (R. Forchheimer).

Although they reportedly work in real-time, the authors have not demonstrated their performance on mobile devices.

The methods are fundamentally different from our approach since we use a machine learning algorithm to discover the structure in the problem rather than assuming one *a priori*. We formulate the problem of pupil localization as regression based on low-level image features. The crucial component is based on randomized trees. These were introduced to the computer vision and image processing community by Amit and Geman [1], and used since to solve classification [15,22,23,26,27,32] and regression [7–12,17,28,36] problems. Randomized trees were a natural choice in our case due to the fact that they handle large training sets with multiple inputs and outputs, and their high generalization power and fast computation.

A similar approach has been taken in [35] to estimate the gaze of the user. The authors decided to select a subset of Haar-like features and train a neural network to estimate the point of gaze on the screen.

1.2. Our contributions

We describe a method for pupil location estimation based on an ensemble of randomized regression trees. The method compares well with the current state-of-the-art in terms of accuracy and runs in real-time on hardware with limited computational resources, such as mobile devices. Additionally, we perform experiments that could prove valuable to other researchers and engineers in the development of image based regression algorithms for similar problems.

2. Method

We assume that the approximate locations of the eyes are known and that we can extract rectangular regions that contain them. This does not pose a problem today since very efficient and accurate face detectors exist. In the next step we use a regression tree ensemble to estimate the coordinates of eye pupils within these regions.

For the purpose of our experiments, we implemented a framework for regression tree construction and evaluation. Furthermore, we use a face detector available in OpenCV to find approximate locations of the eyes in videos and still images.

2.1. Regression trees

Regression trees [5] are tools for function approximation. The basic idea is to split recursively the original problem into two simpler ones, each solvable by a model of reduced complexity. The splits are performed at internal nodes of the tree, based on problem-specific binary tests. Terminal nodes contain simple models that approximate the desired output. In practice, we treat the tree construction process as a supervised learning problem. In other words, we have to choose the binary tests in internal nodes and output models at terminal nodes based on a finite set of input–output pairs in a training set.

Let us briefly describe the configuration of a single tree in our implementation. The binary tests at internal nodes of the tree are based on pixel intensity differences (this is motivated by good results obtained, for example, by [6,27]). We define a binary test on image I as

$$\text{bintest}(I; \mathbf{l}_1, \mathbf{l}_2, t) = \begin{cases} 0, & I(\mathbf{l}_1) - I(\mathbf{l}_2) \leq t \\ 1 & \text{otherwise} \end{cases}$$

where $I(\mathbf{l}_i)$ is the pixel intensity at location \mathbf{l}_i and t is the threshold. Locations \mathbf{l}_1 and \mathbf{l}_2 are in normalized coordinates. This results in

their resolution independence and means that the binary tests can easily be resized based on the data obtained from depth cameras or trackers/detectors that provide information about the scale (resizing is needed when the user moves around and changes his/her distance from the camera as this affects the scale of eyes within the video frame). Each terminal node of the tree contains a constant vector that models the output.

The construction of the tree is supervised. The training set consists of images annotated with values in \mathbb{R}^2 . In our case, these values represent the location of eye pupil in normalized coordinates. The parameters of each binary test in internal nodes of the tree are optimized in a way to maximize clustering quality obtained when the incoming training data is split by the test. This is performed by minimizing the following quantity:

$$Q = \sum_{\mathbf{x} \in C_0} \|\mathbf{x} - \bar{\mathbf{x}}_0\|_2^2 + \sum_{\mathbf{x} \in C_1} \|\mathbf{x} - \bar{\mathbf{x}}_1\|_2^2, \quad (1)$$

where C_0 and C_1 are clusters that contain pupil coordinates $\mathbf{x} \in \mathbb{R}^2$ of all eye patches for which the outputs of binary test were 0 and 1, respectively. The vector $\bar{\mathbf{x}}_0$ is the mean of C_0 and vector $\bar{\mathbf{x}}_1$ is the mean of C_1 .

As the set of all pixel intensity differences is prohibitively large, we generate only a small subset during optimization and compute a threshold for each of them. Although there are many strategies to generate this subset [6], in our experiments we select each pixel difference by sampling two locations from a uniform distribution on a square $[-1, +1] \times [-1, +1]$.

For each pixel intensity difference feature, an optimal threshold can be computed in the following way. Sort the training images by the value of their feature response and then perform a linear search over the array for the best threshold. The test that achieves the smallest error according to Eq. (1) is selected. Thus, the computational complexity of binary test optimization is $O(F \cdot S \log S)$, where S is the number of training samples at the node and F is the number of generated features. The training data is recursively clustered in this fashion until some termination condition is met. In our experiments, we limit the depth of our trees to reduce training times and memory requirements. The output at each leaf node is obtained as the mean of pupil coordinates arriving at that node.

2.2. Tree ensembles

It is well known that a single tree will most likely overfit the training data. On the other hand, an ensemble of trees can achieve impressive results. The most common ways of combining multiple trees are random forests [4] and gradient boosting [14]. The former grows deep trees independently and then averages the result of prediction of each individual tree in the ensemble. The latter grows trees sequentially. Each new one added to the ensemble is learned to reduce the remaining training error further. Its output is shrunk by a scalar factor (called shrinkage, $\nu \in [0, 1]$) that plays a similar role as does learning rate when training neural networks.

Both methods have as parameters the maximum depth of each tree, d , and the number of trees in an ensemble, T . Memory required to store the regression ensemble scales as $O(2^d \cdot T)$ and time for its evaluation as $O(d \cdot T)$. Additionally, gradient boosting has shrinkage ν as a parameter. This parameter does not affect the speed or memory requirements during runtime.

One of the reasons behind the success of tree ensembles is the randomness present in the learning process. Randomness is injected in the induction of each single tree by bootstrap sampling the training set and randomly choosing features as split candidates at each internal node. This procedure reduces the correlation between different trees in the ensemble and leads to better generalization [19]. The other

benefit of this approach is the reduction of training time needed to learn each individual tree.

Both methods show good results in practice, as evidenced in numerous experiments conducted by other researchers and engineers. Thus, we have decided to use numerical measures in order to determine which method to use.

2.3. Random perturbations

We have observed that the output of regression trees is noisy and can be unreliable in some frames, especially if the video stream is supplied from a low quality camera. This can be attributed to variance of the regressor as well as to the simplicity of binary test at internal nodes of the trees: pixel footprint size changes significantly with variations in scale of the eyes and we can expect problems with aliasing. These problems can be reduced during runtime with random perturbations. The idea is to sample multiple rectangular patches at different positions and scales around the eye and estimate the pupil location in each of them. The resulting pupil coordinate is the median over the detections for each spatial dimension. Similar solutions were also proposed by Dollár et al. [10] and Cao et al. [7].

The runtime complexity of the method scales as $O(p \cdot T \cdot d)$, where p is the number of sampled rectangles (random perturbations). Notice that this complexity does not depend on the resolution of the patch.

2.4. A chain of multi-scale ensembles

We have observed that accuracy and robustness of the method critically depend on the scale of the rectangle within which we perform detection. If the rectangle is too small, we risk that it will not contain the eye at all due to the uncertainty introduced by face tracker/detector. If the rectangle is big, the detection is more robust but accuracy suffers. Thus, we propose a solution similar to the one provided by Ong et al. [25]. The approach also resembles the standard pyramidal optical flow algorithm [3].

The idea is to learn multiple tree ensembles, each for estimation at different scale. The method proceeds in a recursive manner, starting with an ensemble learned for largest scale. The obtained intermediate result is used to position the rectangle for the next ensemble in the chain. The process continues until the last one is reached. Its output is accepted as the final result. Images illustrating the typical size of rectangles for multi-scale estimation can be seen in Fig. 1.

We provide numerical results in the experimental section to back our arguments.

3. Training the method

3.1. Annotated data

For our experiments we use the UULM [34] and Gi4E [33] publicly available datasets. Additionally, we constructed our own database for the purposes of the paper. It consists of various students, people from our lab and images acquired on the internet.

All databases consist of images of a number of different subjects with varying head rotation and gaze direction. The images are of varying resolution (800×600 and 640×480 pixels) and each is annotated with locations of eye pupils. All datasets together contain around 5000 face images.

3.2. Training set preparation

Under real world conditions, we must expect that face trackers and detectors, on which our method relies for rough eye region localization, will introduce errors in image patch extraction. Thus, in order to make our system more robust, we intentionally introduce position and scale perturbations in the training data.

For each annotated face in available databases we extract a number of eye patches and corresponding pupil coordinates. All these patches are generated from the *basic rectangle* using random perturbations. The basic rectangle has a predetermined aspect ratio and is obtained from annotated eye corner coordinates which provide both eye region location and a scale factor. Each position perturbation is sampled from a uniform distribution on a rectangle determined by -25% to $+25\%$ of the width and height of the basic rectangle. Similarly, each scale perturbation is sampled from a uniform distribution on an interval determined by -25% to $+25\%$ of the basic rectangle scale.

3.3. Generalization ability of learned tree ensembles

In order to establish the generalization ability of tree ensembles for different parameter values we have performed numerous experiments and it is not possible to show all results due to limited space. Thus, we have decided to present only the graphs that demonstrate the reduction of mean squared error (in image coordinates) achieved by different ensembles for various tree depths, d .

For each annotated image of the eye, we generate 100 training samples according to the method described in the previous section. The UULM database and our annotated images are used for training and Gi4E is used for validation. Thus, the training set consists of around 700 000 images and the validation set of around 200 000 images. We grow each tree on a 50% bootstrap sample of the prepared training data. The number of binary tests



Fig. 1. Some typical rectangle sizes for multi-scale pupil location estimation.




Fig. 2. Mean squared errors on validation data plotted against the number of evaluated binary tests, $d \cdot T$, for different maximum depth d of trees in boosted ensembles and random forests. Number of binary test evaluations is proportional to the amount of computation needed at runtime.

considered at each internal node is $4 + 4d'$, where d' is the depth of the node. For boosting, the shrinkage parameter is fixed to 0.4 as we have noticed that this value gives good results.

The whole evaluation procedure was repeated 5 times and the results were averaged. In Fig. 2 we plot the mean squared errors against the number of binary test evaluations, $d \cdot T$, for different ensemble parameters. We can observe that boosting outperforms random forests in our setup. Note that the errors are not normalized, i.e., they are in image coordinates. Thus, as the validation images vary in resolution, these errors do not have an intuitive relation to the accuracy of pupil localization. In later sections we will use a more appropriate measure to compare our method with other approaches.

All tree ensembles mentioned in the rest of the paper were trained using data from all mentioned databases. To generate the training set, we used 400 random perturbations per annotated image of the eye.

3.4. Comments on training time

We implemented our framework for regression tree construction in standard C. The training runs in one thread although it could be trivially parallelized (in our case it was just not necessary). For comparison, training time for an ensemble of one hundred trees of depth equal to ten takes about 6 h on a training set with six million examples, which is acceptable.

4. Experiments

We intend to make a system capable of real-time performance on hardware of limited processing power, such as mobile devices. Thus, we need to take care of both memory and computation requirements. The complexity of the estimation structure can be regulated by varying the number of trees and the maximum depth of each tree. We fix the run-time complexity by imposing a constraint on the number of binary tests that can be evaluated during estimation to 1000. Furthermore, to limit the memory requirements, we fix the maximum depth of each tree to $d=10$. This limits the number of trees to $T=100$. For learning boosted ensembles the shrinkage was set to $\nu=0.4$. In our preliminary experiments we have observed that these parameters give good tradeoff between accuracy and performance. It is left to determine

the appropriate number of stages for multi-scale estimation process and the method for training each ensemble. We address these questions through experimental analysis.

We are interested in evaluating the usefulness of the method in relevant applications such as pupil localization in still images and pupil tracking in video streams. In our implementation, eye pupil localization in still images is summarized by the following steps:

1. Obtain a face bounding box using a face detector.
2. Estimate eye regions using simple anthropometric relations.
3. Estimate pupil location for each eye region using a chain of multi-scale tree ensembles.

Pupil tracking uses the previously described method for initialization and proceeds by repeating the following steps for each frame:

1. Estimate positions of the eyes based on pupil locations in the previous frame.
2. Estimate scale of each eye region based on distance between the eyes.
3. Estimate pupil location for each eye region using a chain of multi-scale tree ensembles.

In the rest of this section we discuss quantitative and qualitative results obtained by the method. Additionally, we measure the required processing times on different hardware.

4.1. Quantitative results

The normalized error is adopted as the accuracy measure for the estimated eye locations. It was proposed by Jesorsky et al. [21] and is defined as follows:

$$e = \frac{\max\{D_L, D_R\}}{D}, \quad (2)$$

where D_L and D_R are the Euclidean distances between the found left and right eye centers and the ones in the ground truth, and D is the Euclidean distance between the eyes in the ground truth. Roughly, an error of $e \leq 0.25$ corresponds to the distance between the eye center and the eye corners, $e \leq 0.1$ corresponds to the range of the iris, and $e \leq 0.05$ corresponds to the range of the pupil.

We plot regression error characteristic (REC) curves [2] to evaluate the accuracy of our method for different parameter choices on different datasets. The first experiment examines the accuracy on a set of still images. The other evaluates real-time pupil localization in a video stream.

4.1.1. Experimentation with still images

The BioID database (<http://www.bioid.com>) is used for testing. It consists of 1521 frontal face grayscale images with significant variation in illumination, scale and pose. In several samples of the database the subjects are wearing glasses. In some instances the eyes are closed, turned away from the camera, or completely hidden by strong highlights on the glasses. Due to these conditions, the database is considered difficult and realistic. Some snapshots can be seen in Fig. 3. Our approach yields inaccurate estimations if the eyes are closed or strong reflections on the glasses occur as we did not include such examples in the training set.

One important thing to note before presenting numerical results is that we omitted the images where the detector did not find the face, following [31]. This did not significantly alter the results since approximately 1% images were discarded this way.

Let us first examine the benefits of multi-scale estimation (we do not perform random perturbations in this experiment,




Fig. 3. First row illustrates some successful detections on the BioID database. The second and third one illustrate some typical failures.

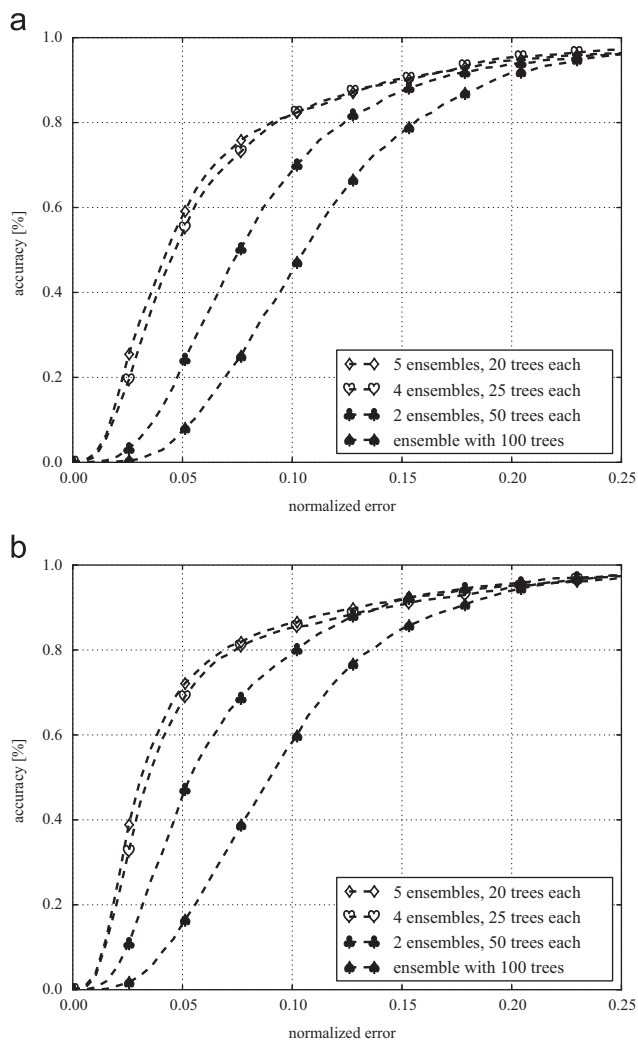


Fig. 4. Increase in accuracy when using a multi-scale estimation approach. We observe that boosting gives better results: (a) Estimation using random forests and (b) estimation using boosted ensembles.

i.e., we set $p=1$). Quantitative results can be seen in Fig. 4. We can observe that this procedure significantly improves the accuracy for $e \leq 0.15$. As we are mostly interested in accurate pupil localization, we can conclude that the multi-scale approach provides significant benefits.

The results can be improved even further by using random perturbations, as explained in Section 2.3. Quantitative results can be seen in Fig. 5. We observe a significant increase in accuracy when p increases. Of course, the number of computations increases linearly with p . Thus, in order to justify the use of this improvement we demonstrate in Section 4.3 that real-time frame rates can be achieved on mobile devices with moderate values of p .

We compare our results in Table 1 to the two state-of-the-art methods for eye center localization. We can observe that our method compares well to the competition in terms of accuracy. Of course, there is always the problem of dataset bias [30]: We cannot conclude which method is superior based just on accuracy results on one dataset.

We noticed a few sources of discrepancies with annotated data (examples can be seen in Figs. 3 and 6):

- *Presence of eyeglasses*: In some images the eyes are completely hidden by strong highlights on the glasses. Our method fails under these conditions. Fig. 6a illustrates this case.
- *Closed eyes*: Eye center localization can not work on closed eyes when the iris is not visible. Thus, the results obtained from our system are unpredictable during eye closure. An example can be seen in Fig. 6b.
- *Eye region extraction failure*: Some bad results are due to the limitations of the used face detection procedure. An example can be seen in Fig. 6c.

4.1.2. Experiment on video

We use the Talking Face video database (http://www-prima.inrialpes.fr/FGnet/data/01-TalkingFace/talking_face.html) to evaluate our system quantitatively in real-time applications. The video contains 5000 frames taken from a video of a person engaged in conversation. Each frame was annotated semi-automatically using an active appearance model [16] trained specifically for the person in

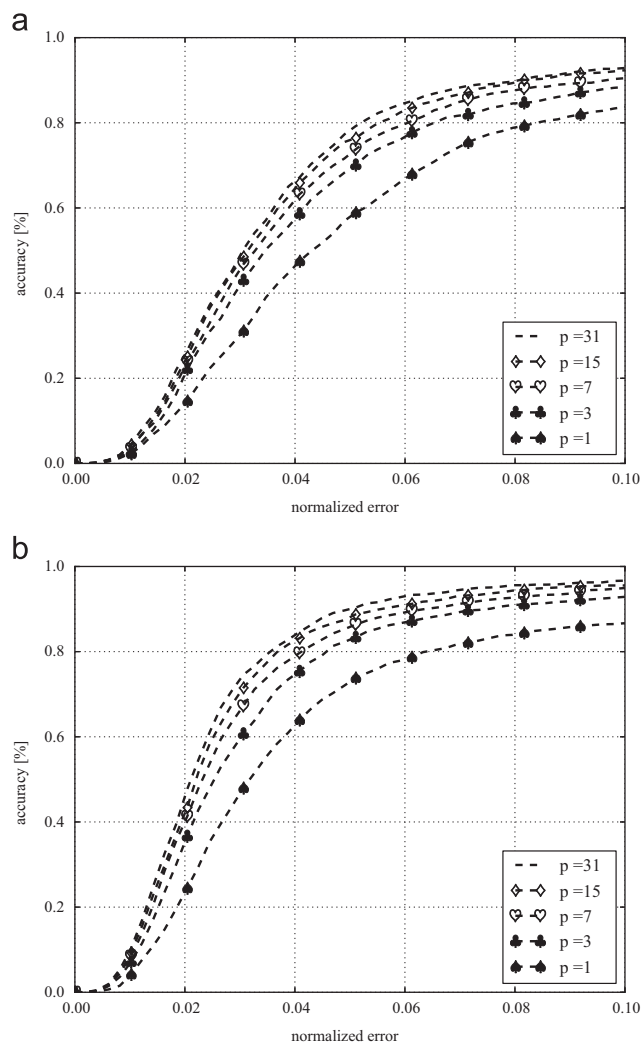


Fig. 5. Increase in accuracy when using different number of random perturbations, p . We can observe that boosting gives significantly better results: (a) Estimation using random forests at 5 different scales, 20 trees per scale and (b) estimation using boosted ensembles at 5 different scales, 20 trees per scale.

Table 1

Comparison of normalized error scores on the BioID database. The values in the table are percentages.

Method	$e \leq 0.05$	$e \leq 0.1$	$e \leq 0.25$
Our results, $p=7$	85.7	95.3	99.7
Our results, $p=31$	89.9	97.1	99.7
Timm et al. [29]	82.5	93.4	98.0
Valenti et al. [31], basic method	80.6	85.8	96.6
Valenti et al. [31], extended method	86.1	91.7	97.9

the video. The annotations are sufficiently accurate to represent the facial movements during the sequence. They contain the locations of eye centers, among other characteristic points on the face.

We used an estimation chain consisting of 5 boosted ensembles with 20 trees of depth equal to 10. The number of random perturbations was set to $p=15$. The video showing tracking results is available as supplementary material.

The normalized error (Eq. (2)) averaged over the video sequence was equal to 0.0236 and its median 0.0201. Accuracy




Fig. 7. Accuracy curve for the Talking Face video pupil tracking.




Fig. 8. The distribution of normalized errors over the frames of the Talking Face video.




Fig. 6. Snapshots from the BioID database illustrating some typical failures of our method: (a) Failure due to strong highlights on the glasses; (b) failure due to eye closure; and (c) failure of face detection procedure.

curve can be seen in Fig. 7 and error distribution over the frames in Fig. 8. These show that most of the time our results were in the range of pupil radius from annotated data. Some snapshots from the video can be seen in Fig. 9.

Although the average error is sufficiently low, we analyze the error spikes present in the graph. We noticed a few sources of these discrepancies with the annotated data:

- *Closed eyes:* Eye center localization can not work on closed eyes when the iris is not visible. Thus, the results obtained from our system are unpredictable during eye closure. The proper solution is to include eye closure detection, which we plan in the future. An example can be seen in Fig. 10a.
- *Error in annotated data:* Because the annotated data was obtained semi-automatically, it contains occasional errors. For example, this happens around the 4000th frame (a big spike in the error graph). Fig. 10b illustrates the case. We can observe that our system is more accurate than the annotated data in this frame.
- *Failure of our method:* In some frames our method outputs incorrect eye pupil locations even when all appropriate conditions are met. An example of such a case can be seen in Fig. 10c.

4.2. Qualitative results

Besides the quantitative analysis on datasets that provide ground truth data, we performed less formal qualitative experiments on still images and video input, both on PCs with webcams and on mobile devices. A video showcasing these results is available at <https://www.youtube.com/watch?v=7J30yNHIXIQ> (which is also available as supplementary material). Furthermore, for readers who wish to test the method themselves, the windows-based demo application is available (http://hotlab.fer.hr/_download/repository/puploc.zip).

Fig. 11 illustrates some examples of eye pupil localization in still images. The system correctly estimates the pupil coordinates even when eye rotations are present. Fig. 12 illustrates some results for eye regions extracted during tracking in a video stream. We can observe that the method provides accurate results for large variations in position, scale and rotation. This demonstrates that our approach can successfully deal with real-world conditions.

We have observed occasional failures that can be attributed to the reasons discussed in Sections 4.1.1 and 4.1.2.

4.3. Analysis of performance and required memory

We measure the average time to process one eye-region rectangle and locate the pupil using our method with standard parameters ($T=100$ and $d=10$). The system is implemented in C and runs in a single thread. Most computations are performed in integer arithmetic to further speedup the processing. The obtained results are reported in Table 2.

Considering that typical face tracking performance may be around 30–50 frames per second on most devices, this gives a frame processing time of 20 ms, and considerably more in case of full face detection. This means that detection of both eye pupils presents a negligible performance impact within a face detection or tracking system. Also note that the system can track eyes on its own at these computational requirements (Section 4.1.2).

A structure of $T=100$ trees of depth $d=10$ requires approximately 1.3 megabytes of storage. We believe that this is acceptable both at runtime and distribution with an application.

5. Conclusion and future directions

We have described an eye pupil localization framework based on an ensemble of randomized regression trees. The implementation



Fig. 9. Snapshots from the Talking Face video in which our system outputs the correct locations of eye pupils.




Fig. 10. Snapshots from the Talking Face video illustrating some typical discrepancies with annotated data. Annotations are represented with a cross and the output of our system with a circular target symbol: (a) Failure due to eye closure; (b) error in annotated data; and (c) our method outputs wrong results.




Fig. 11. Pupil localization in still images.



Fig. 12. Example eye patches with estimated pupil locations.

is capable of achieving real-time performance on mobile devices. The developed software provides a low-cost alternative to complex and costly commercial eye tracking systems.

The described method is very general and can be applied to solve similar problems, such as tracking of learned 2D templates. An interesting application would be to track multiple points

Table 2

Average time required to process one image region for different number of random perturbations, p . The ensemble consists of 100 trees, each of depth 10.

Device	CPU	Time (ms)			
		$p=3$	$p=7$	$p=15$	$p=31$
PC1	3.4 GHz Core i7-2600	0.02	0.05	0.12	0.24
PC2	2.53 GHz Core 2 Duo P8700	0.06	0.13	0.29	0.60
iPhone 5	1.3 GHz Apple A6	0.12	0.27	0.58	1.20
iPad 2	1 GHz ARM Cortex-A9	0.22	0.53	1.13	2.32
iPhone 4S	800 MHz ARM Cortex-A9	0.28	0.66	1.42	2.93

scattered on a rigid object and use RANSAC algorithm [13] to recover the motion, similar to the work done by Zimmermann et al. [37].

Future work includes the possibility of using randomized trees for eye closure estimation. It would be interesting to see if we could obtain good results by synthesizing a training set using methods from computer graphics. Another interesting direction of research is to use our method for estimating the point of gaze of the user. This would require the integration of a 3D model to our framework (similar to [20]).

Conflict of interest

None declared.

Acknowledgments

This research is partially supported by Visage Technologies AB, Linköping, Sweden, by the Ministry of Science, Education and Sports of the Republic of Croatia, grant number 036-0362027-2028 “Embodied Conversational Agents for Services in Networked and Mobile Environments” and by the European Union through ACROSS project, 285939 FP7-REGPOT-2011-1.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version of <http://dx.doi.org/10.1016/j.patcog.2013.08.008>.

References

- [1] Y. Amit, D. Geman, Shape quantization and recognition with randomized trees, *Neural Computation* 9 (1997) 1545–1588.
- [2] J. Bi, K.P. Bennett, Regression error characteristic curves, in: *Proceedings of the 20th International Conference on Machine Learning*, 2003, pp. 43–50.
- [3] J.Y. Bouquet, Pyramidal implementation of the lucas kanade feature tracker description of the algorithm, Technical Report, Intel Microprocessor Research Labs., 1999.
- [4] L. Breiman, Random forests, *Journal of Machine Learning Research* 45 (2001) 5–32.
- [5] L. Breiman, J. Friedman, C.J. Stone, R.A. Olshen, *Classification and Regression Trees*, Chapman and Hall, 1984.
- [6] M. Calonder, V. Lepetit, M. Ozuysal, T. Trzinski, C. Strecha, P. Fua, Brief: computing a local binary descriptor very fast, *Transactions on Pattern Analysis and Machine Intelligence* 34 (2011) 1281–1298.
- [7] X. Cao, Y. Wei, F. Wen, J. Sun, Face alignment by explicit shape regression, in: *CVPR*, 2012, pp. 2887–2894.
- [8] A. Criminisi, E. Konukoglu, J. Shotton, Regression forests for efficient anatomy detection and localization in CT studies, in: *Medical Computer Vision 2010: Recognition Techniques and Applications in Medical Imaging, MICCAI Workshop*, 2010, pp. 106–117.
- [9] M. Dantone, J. Gall, G. Fanelli, L.V. Gool, Real-time facial feature detection using conditional regression forests, in: *CVPR*, 2012, pp. 2578–2585.
- [10] P. Dollár, P. Welinder, P. Perona, Cascaded pose regression, in: *CVPR*, 2010, pp. 1078–1085.
- [11] G. Fanelli, M. Dantone, J. Gall, A. Fossati, L.V. Gool, Random forests for real time 3d face analysis, *International Journal of Computer Vision* 101 (2012) 437–458.
- [12] G. Fanelli, J. Gall, L.V. Gool, Real time head pose estimation with random regression forests, in: *CVPR*, 2011, pp. 617–624.
- [13] M.A. Fischler, R.C. Bolles, Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography, *Communications of the ACM* 24 (1981) 381–395.
- [14] J.H. Friedman, Greedy function approximation: a gradient boosting machine, *Annals of Statistics* 29 (2001) 1189–1232.
- [15] J. Gall, A. Yao, N.R.L.V. Gool, V. Lempitsky, Hough forests for object detection, tracking, and action recognition, *Transactions on Pattern Analysis and Machine Intelligence* 33 (2011) 2188–2202.
- [16] X. Gao, Y. Su, X. Li, D. Tao, A review of active appearance models, *IEEE Systems, Man, and Cybernetics, Part C: Applications and Reviews* 40 (2010) 145–158.
- [17] B. Glocker, O. Pauly, E. Konukoglu, A. Criminisi, Joint classification-regression forests for spatially structured multi-object segmentation, in: *ECCV*, 2012, pp. 870–881.
- [18] D.W. Hansen, Q. Ji, In the eye of the beholder: a survey of models for eyes and gaze, *Transactions on Pattern Analysis and Machine Intelligence* 32 (2010) 478–500.
- [19] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning*, Springer, 2009.
- [20] T. Heyman, V. Spruyt, A. Ledda, 3d Face Tracking and Gaze Estimation Using a Monocular Camera, Technical Report, 2012.
- [21] O. Jesorsky, K.J. Kirchbergand, R. Frischholz, Robust face detection using the hausdorff distance, in: *Audio and Video Biom. Pers. Auth.*, 1992, pp. 90–95.
- [22] Z. Kalal, K. Mikolajczyk, J. Matas, Tracking-learning-detection, *Transactions on Pattern Analysis and Machine Intelligence* 34 (2012) 1409–1422.
- [23] V. Lepetit, P. Fua, Keypoint recognition using randomized trees, *Transactions on Pattern Analysis and Machine Intelligence* 28 (2006) 1465–1479.
- [24] D.G. Lowe, Object recognition from local scale-invariant features, in: *ICCV*, vol. 2, 1999, pp. 1150–1157.
- [25] E.-J. Ong, Y. Lan, B. Theobald, R. Bowden, Robust facial feature tracking using selected multi-resolution linear predictors, *Transactions on Pattern Analysis and Machine Intelligence* 33 (2011) 1844–1859.
- [26] M. Ozuysal, P. Fua, V. Lepetit, Fast keypoint recognition in ten lines of code, in: *CVPR*, 2007, pp. 1–8.
- [27] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, A. Blake, Real-time Human Pose Recognition in Parts from Single Depth Images, Technical Report, Microsoft Research, 2011.
- [28] M. Sun, P. Kohli, J. Shotton, Conditional regression forests for human pose estimation, in: *CVPR*, 2012, pp. 3394–3401.
- [29] F. Timm, E. Barth, Accurate eye centre localisation by means of gradients, in: *International Conference on Computer Vision Theory and Applications*, vol. 1, 2011, pp. 125–130.
- [30] A. Torralba, A.A. Efros, Unbiased look at dataset bias, in: *CVPR*, 2011, pp. 1521–1528.
- [31] R. Valenti, T. Gevers, Accurate eye center location through invariant isocentric patterns, *Transactions on Pattern Analysis and Machine Intelligence* 34 (2012) 1785–1798.
- [32] M.J. Villamizar, A. Andrade-Cetto, F.M.-N. Sanfeliu, Bootstrapping boosted random ferns for discriminative and efficient object classification, *Pattern Recognition* 45 (2012) 3141–3153.
- [33] V. Ponz, A. Villanueva, R. Cabeza, Dataset for the evaluation of eye detector for gaze estimation, in: *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, 2012, pp. 681–684. <http://dx.doi.org/10.1145/2370216.2370364>.
- [34] U. Weidenbacher, G. Layher, P.-M. Strauss, H. Neumann, A comprehensive head pose and gaze database, in: *3rd IET International Conference on Intelligent Environments*, Ulm (Germany), 2007, pp. 455–458.
- [35] Y. Zhang, A. Bulling, H. Gellersen, Towards pervasive eye tracking using low-level image features, in: *Symposium on Eye Tracking Research and Applications*, 2012, pp. 261–264.
- [36] S.K. Zhou, B. Georgescu, X.S. Zhou, D. Comaniciu, Image based regression using boosting method, in: *ICCV*, 2005, pp. 541–548.
- [37] K. Zimmermann, J. Matas, T. Svoboda, Tracking by an optimal sequence of linear predictors, *Transactions on Pattern Analysis and Machine Intelligence* 31 (2009) 677–692.

Nenad Markuš is a research and teaching assistant at the Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia. He graduated at the same institution in July 2011 with a thesis in the area of sparse signal processing.

Miroslav Frljak is a research associate at the Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia. He graduated at the same institution in July 2011, with a thesis in the area of face tracking on mobile devices.

Igor S. Pandžić is a Professor at the Department of Telecommunications, Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia. His main research interests are in the field of computer graphics and virtual environments, with particular focus on character animation, embodied conversational agents, and their applications in networked and mobile environments.

Jörgen Ahlberg received his Ph.D. in Electrical Engineering from Linköping University in 2002. 2002–2011 he worked at the Swedish Defense Research Agency, and currently he is the R&D manager at Termisk Systemteknik AB. He is also a visiting senior lecturer at Linköping University, and a co-founder of Visage Technologies AB.

Robert Forchheimer is a professor at the dept. of Electrical Engineering, Linköping University, Sweden. His main research area for nearly three decades has been video coding. In parallel he has been involved in packet radio networking and optical computing. More recently he has published works in areas such as optical networks, organic electronics and systems biology.