

Memory-efficient Global Refinement of Decision-Tree Ensembles and its Application to Face Alignment

Nenad Markuš¹
nenad.markus@fer.hr

Ivan Gogić¹
ivan.gogic@fer.hr

Igor S. Pandžić¹
igor.pandzic@fer.hr

Jörgen Ahlberg²
jorgen.ahlberg@liu.se

¹ Faculty of Electrical Engineering and Computing
University of Zagreb
Unska 3, 10000 Zagreb, Croatia

² Computer Vision Laboratory
Dept. of Electrical Engineering
Linköping University
SE-581 83 Linköping, Sweden

Abstract

Ren et al. [1] recently introduced a method for aggregating multiple decision trees into a strong predictor by interpreting a path taken by a sample down each tree as a binary vector and performing linear regression on top of these vectors stacked together. They provided experimental evidence that the method offers advantages over the usual approaches for combining decision trees (random forests and boosting). The method truly shines when the regression target is a large vector with correlated dimensions, such as a 2D face shape represented with the positions of several facial landmarks. However, we argue that their basic method is not applicable in many practical scenarios due to large memory requirements. This paper shows how this issue can be solved through the use of quantization and architectural changes of the predictor that maps decision tree-derived encodings to the desired output.

1 Introduction

Decision trees [2] are a tried-and-true machine learning method with a long tradition. They are especially powerful when combined in an ensemble [3] (the outputs of multiple trees are usually summed together). On certain problems, these approaches are still competitive [4, 5].

A nice property of decision-tree ensembles is that the method easily deals with multi-dimensional prediction (e.g., in multi-class classification). This is achieved by placing a vector in the leaf node of each tree. This means that the multidimensional output $\Delta(x)$ for the input sample x is computed as $\Delta(x) = \sum_i \mathbf{w}_i$, where the i th vector is output by the i th tree: $\mathbf{w}_i = \text{Tree}_i(x)$. Ren et al. [1, 6] interpret this computation as a linear projection step:

$$\Delta(x) = \sum_i \text{Tree}_i(x) = \mathbf{W} \cdot \Phi(x), \quad (1)$$

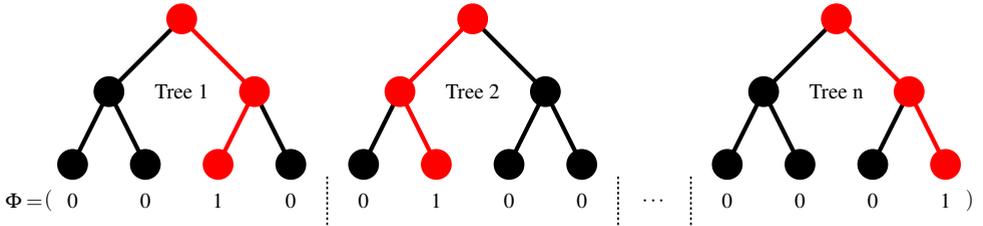


Figure 1: The process of generating a sparse feature vector Φ with an ensemble of n decision trees. The path that a sample takes through each tree is drawn in red. Exactly n components of Φ are set to 1 based on these paths. The rest are set to 0.

where \mathbf{W} is a large matrix that contains as columns the leaf-node vectors of the trees and $\Phi(x)$ is a sparse vector that indicates which columns of \mathbf{W} should be summed together in order to obtain the prediction for the sample x . See Figure 1 for an illustration that shows how to obtain $\Phi(x)$. This interpretation enabled Ren et al. to learn an efficient method for face alignment by jointly refining the outputs of multiple decision trees [17] (more details later).

If the ensemble has n trees of depth equal to d and the dimension of the output is o , then the matrix \mathbf{W} has $n \cdot 2^d \cdot o$ parameters. This number can be quite large in a practical setting. We show how to replace Equation (1) with a more memory-friendly computation: first, we investigate two different methods with a reduced number of coefficients; second, we show that the remaining coefficients can be further compressed with quantization. We apply our ideas to a face-alignment problem and compare with several recently published methods.

To summarize, the goal of this paper is twofold:

- we improve on the previous work of Ren et al. [17, 19] by significantly reducing the memory requirements of their methods with no loss in accuracy;
- we describe a new face-alignment method that is reliable, fast and memory-efficient, which makes it especially suitable for use on mobile devices and other hardware with limited resources.

2 Memory-efficient global refinement of tree ensembles

A way to compress $\mathbf{W} \in \mathbb{R}^{o \times (n \cdot 2^d)}$ is to express it as a product of $\mathbf{W}_2 \in \mathbb{R}^{o \times r}$ and $\mathbf{W}_1 \in \mathbb{R}^{r \times (n \cdot 2^d)}$: $\mathbf{W} = \mathbf{W}_2 \cdot \mathbf{W}_1$. Of course, r has to be smaller than o . For classification, we can attempt to learn these matrices using gradient descent. For regression, besides gradient descent, we can also use the reduced-rank regression (RRR) framework [16].

Another possible path to improving memory issues is to replace linear regression for computing Δ from Φ (Equation (1)) with a neural network (NN). One architecture that we found to work well in our experiments is

$$\Delta(x) = \mathbf{W}_3 \cdot \tanh(\mathbf{W}_2 \cdot \tanh(\mathbf{W}_1 \cdot \Phi(x))), \quad (2)$$

where $\mathbf{W}_3 \in \mathbb{R}^{o \times 2r}$, $\mathbf{W}_2 \in \mathbb{R}^{2r \times r}$, $\mathbf{W}_1 \in \mathbb{R}^{r \times (n \cdot 2^d)}$ and \tanh is the elementwise hyperbolic-tangent nonlinearity. See Figure 2 for an illustration. We agree that there might be better

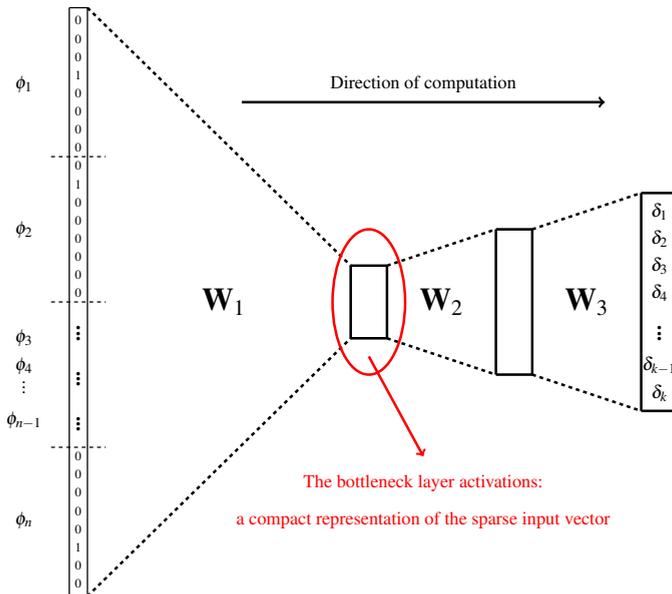


Figure 2: An illustration of the neural-net structure used in the experiments. First, the input sparse vector is compressed into a compact representation with the projection matrix \mathbf{W}_1 . Next, this representation is gradually expanded with matrices \mathbf{W}_2 and \mathbf{W}_3 to obtain the result. Hyperbolic-tangent nonlinearities (\tanh) are applied in the two inner layers.

design choices. However, we leave these question to practitioners (as the proposed framework is generic) and experimentally test the performance of the architecture specified with Equation 2. The matrices \mathbf{W}_1 , \mathbf{W}_2 and \mathbf{W}_3 can be learned with gradient descent through the use of backpropagation. The presented NN architecture improves memory issues only if r can be made reasonably small. Also, note that modern tricks that we do not use here (e.g., highway connections [23] and batch normalization [10]) could further boost performance.

Instead of a standard tree-based prediction, which is additive in nature (Equation (1)), we now have a two-step procedure:

- produce a compact encoding $\mathbf{e}(x)$, $\mathbf{e}(x) = \sum_i \text{Tree}_i(x)$;
- decode $\Delta(x)$ from $\mathbf{e}(x)$ (this can be a nonlinear, non-additive computation, such as a neural network).

Our claim is that r can be significantly smaller than o in some interesting problems and that, consequently, this leads to large reduction in storage-related issues for computing Δ from Φ . We later back these claims by experimental evidence.

2.1 Related work

For the history and applications of reduced-rank regression please see the book by Reinsel and Velu [16]. As this is a well-known and widely-used approach in machine learning and statistics, we will not repeat the details here.

Our research is related to a large body of work in neural-network compression and speed improvement (e.g., [6, 9, 21, 24, 29]). Specifically, the work of Sainath et al. [21] and other similar approaches suggest to reduce the memory-related issues of machine learning models by introducing a low-dimensional "bottleneck" before the output, as we do in this paper. However, to the best of our knowledge, we are the first to investigate the use of these techniques in the context of decision-tree ensembles. We can say that our work is a blend of the interpretation introduced by Ren et al. [17] and recent network compression techniques.

In the next section, we apply the proposed techniques to the problem of face alignment. The literature in this area is large and we give more information only on most related methods [6, 17, 19, 21] there as our work builds on the same core ideas.

3 Application to face alignment

Face alignment is the process of finding the shape of the face given its location and size in the image. Today's most popular methods (e.g., [6, 17, 21]) achieve this through $t = 1, 2, \dots, T$ iterations: $\mathbf{S}_t = \mathbf{S}_{t-1} + \Delta_t$, where $\Delta_t = \Delta(\mathbf{S}_{t-1}, F)$ is the shape update computed from the previous shape \mathbf{S}_{t-1} and image F . The shape update is usually computed in two steps [17, 21]. The first step is a feature extraction process $\Phi(\mathbf{S}_{t-1}, F)$ that encodes the image appearance around the shape \mathbf{S}_{t-1} . The second step computes the shape increment via a linear projection: $\Delta_t = \mathbf{W}_t \cdot \Phi(\mathbf{S}_{t-1}, F)$.

For the supervised descent method of Xiong and De la Torre [21], the features Φ are stacked SIFT descriptors extracted around landmark points \mathbf{S}_{t-1} . The matrices $\{\mathbf{W}_t\}$ are learned from a training set by ordinary least squares regression. A more complicated method of Ren et al. [17] extracts a sparse feature vector around each landmark point from \mathbf{S}_{t-1} with a decision tree ensemble and concatenates them to obtain Φ . The binary tests in internal tree nodes are thresholded pixel intensity differences, $F_{x_1, y_1} - F_{x_2, y_2}$. The simplicity of these tests makes the method very fast in runtime. The sparse feature vector Φ computed in this manner is very large (tens of thousands of dimensions in practical settings). Consequently, the matrices $\{\mathbf{W}_t\}$ have a lot of parameters (weights). Our goal is to experimentally show that we can replace linear projection $\Delta_t = \mathbf{W}_t \cdot \Phi_t$ with a more memory-friendly computation. For a more in-depth review on recent advances in face alignment, we refer the reader to [17].

We base all our experiments on the 300 Faces in-the-Wild framework (300W for short, [17]), which is a standard procedure for benchmarking face-alignment systems. Each face is represented with a 136-dimensional vector (68 landmark points, two coordinates each). The dataset was constructed to examine the ability of face-alignment systems to handle naturalistic, unconstrained face images. It covers many different variations of unseen subjects, pose, expression, illumination, background, occlusion and image quality. These features have made the 300W dataset a *de facto* standard for benchmarking face-alignment systems. The dataset is partitioned into the training and testing subsets according to a predefined protocol. We follow these guidelines closely.

3.1 Learning setup

To augment the training data, we apply the following transformations 20 times for each face: flip the image with probability equal to 0.5 (faces are symmetric), rotate the face for a random angle between -20 and 20 degrees, randomly perturb the scale and the position of the face bounding box. Overall, this results in a training set of around 50 000 samples.

The number of trees per landmark point is set to 5 and the depth of each individual tree is 5. It can be seen from the setup that the dimensionality of vector Φ is $68 \cdot 5 \cdot 2^5 = 10\,880$. Initially, each tree is learned to predict the position of its landmark point from a local patch around it. Next, the values in the leaf nodes are discarded and we learn the function $\Delta(x)$ using the method proposed by Ren et al. [14] (Equation (1)) or one of the methods described in Section 2. The total number of stages is also set to 5. These hyper-parameters follow the ones proposed in the original LBF (Local Binary Features) paper [14]. The initial shape is set to be the mean face computed on the training data.

For reduced-rank regression and the neural-net approach, our preliminary experiments showed that it is beneficial to increase the size of the bottleneck as the shape estimation progresses through the stages. The following values lead to good results: $r = 16, 24, 32, 40, 48$. This is based on the heuristic that the transformations/deformations in the first iterations are more rigid and, thus, require less parameters to be modeled reliably. On the other hand, the later iterations require a considerable number of parameters to fine-tune the landmark positions. These values were tuned on the training/validation set partitions (i.e., the testing set was left out of this experimentation). However, the search was not exhaustive and there is a possibility that some other values might lead to better results.

The projection matrix for each stage of the original LBF method is learned with LIBLINEAR [2] (we closely follow their recommendations [14]). To get the optimal reduced-rank regression matrices, we use BLAS and LAPACK since the process involves eigenvalue problems, matrix multiplications and inversions [16]. For the backpropagation-based methods, we use the Torch7 framework as it implements a very efficient `nn.SparseLinear()` module for computing products of sparse vectors and matrices. The SGD learning rate is initially set to $\eta = 1$ and halved each time the validation loss saturates. The minibatch size is 128. The learning for all methods is quite fast in the described setting (a few hours).

To make the exposition more compact and reduce the clutter in graphs/plots, we introduce the following abbreviations: LL for the basic LBF method [14], RRR for reduced-rank regression [16], RRR-BP for reduced-rank regression learned with backpropagation and NN for the neural networks.

3.2 First results and the effects of random initializations

The average point-to-point Euclidean error for $L = 68$ landmark points normalized by the inter-ocular distance (IOD, measured as the Euclidean distance between the outer corners of the eyes) is used to compare different methods:

$$\text{err} = 100 \times \frac{1}{L} \sum_{i=1}^L \frac{\sqrt{(x_i - x_i^*)^2 + (y_i - y_i^*)^2}}{\text{IOD}} [\%], \quad (3)$$

where (x_i, y_i) and (x_i^*, y_i^*) are the estimated and ground-truth locations of the i th landmark.

It is known that averaging the results of several initializations (slight perturbations in position and scale of the face bounding box) can improve face alignment [8]. We did not notice high gains in alignment accuracy (Equation (3)). However, the method reduces the jitter when processing videos or a webcam stream. This significantly affects the subjective quality of the face-alignment system. Thus, we also report the results for different numbers of random initializations, p .

Table 1 contains normalized errors averaged over the common and the challenging testing subsets of the 300W dataset (as the name implies, these subsets differ in difficulty). It can be

p	1	7	15	p	1	7	15
LL	4.39	4.19	4.19	LL	11.69	11.42	11.19
NN	4.39	3.96	3.96	NN	11.59	10.59	10.19
RRR	4.39	4.03	3.95	RRR	11.18	10.21	10.06
RRR-BP	4.49	4.12	4.11	RRR-BP	12.72	12.02	11.59

Table 1: Average errors for the common (left) and the challenging (right) subsets of the 300W dataset for different number of random initializations, p .

q	no quantization	8	6	4	2
LL	27901	8332	5960	3980	2098
NN	6927	2367	1808	1343	900
RRR	6790	2229	1671	1205	762

Table 2: Memory requirements for the packed and quantized regression structures. The value q shows the number of bits per weight and the rest of table entries are in kilobytes.

seen that random perturbations slightly increase the accuracy for all methods. All competing methods perform approximately the same, i.e., achieve similar average errors.

These results lead to the conclusion that NN- and RRR-based methods are viable alternatives to the basic method of Ren et al. [17]. Moreover, these methods reduce storage/bandwidth requirements by approximately three times (from 27 to 9 megabytes). However, this reduction may not be sufficient for some applications (such as mobile-phone and web programs with bandwidth constraints). Thus, the possibility of quantizing the \mathbf{W}_1 matrix is explored next.

3.3 Effects of coefficient quantization

Quantization is a straightforward way to improve storage requirements. In all the discussed methods, the matrices that multiply the (very large and sparse) input vector Φ contain the most parameters. Thus, these are most interesting to compress with quantization. In the case of face alignment, these matrices are "long and flat", i.e., they have much more columns than rows. We found empirically that the following scheme gives good results: each row of the matrix is quantized to q bits per weight/coefficient with a nonlinear quantizer. The centroids are computed with the standard k -means clustering algorithm where $k = 2^q$. Separate quantizers are used for each row to increase the expressive power without requiring a too-large quantization table.

Parameter quantization leads to significant memory savings. The exact numerical values can be seen in Table 2 (the RRR-BP method is left out since it has the same memory usage as the ordinary RRR). However, it is not clear how much accuracy degradation is introduced with quantization. Figure 3 provides more insight into this issue. The effects are similar for all values of applied randomly perturbed initializations, p . We can see that all the architectures introduced in this paper outperform the basic LBF method (LL, [17]) when quantization is used. We hypothesize that this is due to the fact that only the projection matrices (the \mathbf{W}_1 matrices; see Section 2) are quantized in the implemented scheme. In other words, parts of the NN- and RRR-based frameworks remain unquantized in all scenarios and it seems that this leads to the preservation of important information.

Overall, quantized RRR- and NN-based methods with $q = 4$ and $p = 7$ significantly improve over the basic LBF-based shape regression approach introduced by Ren et al. [17].

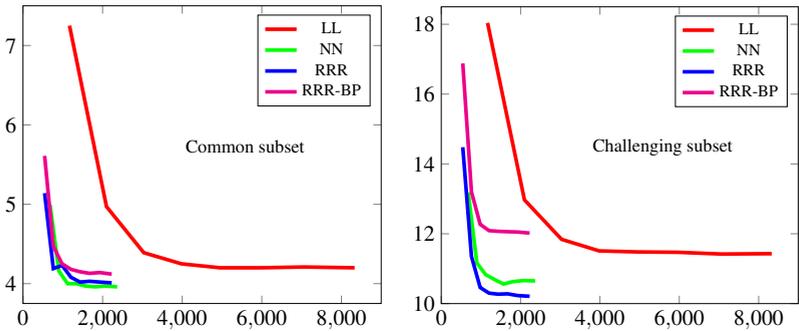


Figure 3: Storage-vs-error tradeoff for the common and challenging subsets of 300W: abscissae are in kilobytes and ordinates are normalized errors. The results on the "common" subset are in the left graph and the results on the "challenging" subset are in the right one.

p	1	3	7	15
LL	0.26	0.79	1.88	4.01
NN	0.26	0.75	1.73	3.57
RRR	0.22	0.60	1.43	2.99

Table 3: Average times in milliseconds required to process one face region for different number of random initializations, p . CPU: 3.4GHz Core i7-2600.

The storage requirements are reduced by approximately 21 times (from 27MB to 1.3MB) without loss of accuracy. Note that the gain would be even more significant in the case of a larger face model (more trees, more facial landmarks).

3.4 Processing speed analysis

Processing speed results can be seen in Table 3. We can see that all tested methods are extremely fast. For example, with $p = 1$ over 3000 FPS processing speeds can be achieved.

The measurements also show a slight advantage of RRR-based method over the other two. This can be attributed to fewer floating point operations when transforming the sparse feature vector to the face shape. However, this advantage can be beneficial only to a certain point since all methods use the same mechanism to produce the sparse feature vector through pixel-intensity comparison binary tests, which amounts for the majority of computations.

3.5 Comparison to the other methods

Table 4 compares the RRR ($q = 4$) with recent methods [10, 11, 5, 12, 14, 15, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33]. The processing speeds are based on the results from [19]. The unknown/missing values are denoted with a dash: "-". Other configurations discussed earlier in the experimental sections, such as those produced with NN- and LL-based approaches and different quantization levels are not included for clarity of presentation due to the fact that all achieve similar error rates. It can be seen that the RRR method achieves competitive results, even when compared to recent deep-learning approaches [24]. However, it is important to note that different error normalizations are used by researchers making the exact comparison difficult. We used the inter-ocular normalization (distance between the outer corner of the

Method	Common	Challenging	Full set	Proc. speed [FPS]
Zhu et al. [13]	8.22	18.33	10.20	–
DRMF [10]	6.65	19.79	9.22	–
ESR [6]	5.28	17.00	7.58	120
RCPR [9]	6.18	17.26	8.35	80
SDM [12]	5.57	15.40	7.50	70
Smith et al. [22]	–	13.30	–	–
Zhao et al. [30]	–	–	6.31	–
GN-DPM [28]	5.78	–	–	–
CFAN [28]	5.50	–	–	44
ERT [14]	–	–	6.40	1000
LBF [14]	4.95	11.98	6.32	320
LBF (fast) [14]	5.38	15.50	7.37	3100
TDCNN [30]	4.80	8.60	5.54	56
cGPRT [15]	–	–	5.71	–
CFSS [12]	4.73	9.98	5.76	25
CFSS (practical) [12]	4.79	10.92	5.99	50
DCR [12]	4.19	8.42	5.02	–
DAN [13]	3.19	5.24	3.59	5
RRR ($p = 1$)	4.52	11.34	5.86	3500+
RRR ($p = 7$)	4.08	10.30	5.26	500+

Table 4: Comparing other approaches to RRR with $q = 4$. Normalized errors are in the middle three columns.

eyes) since there are no pupil annotations on the 300-W data set while some researchers estimated the pupil position for the inter-pupil normalization.

The conclusions from [19] apply. Overall, the regression-based approaches are better than the ASM-based methods. Methods based on pixel-intensity comparisons ([6, 12, 14], RRR) and SIFT-based methods [17, 32] achieve very good results. The methods based on convolutional neural nets [13, 12, 28, 30] achieve the best results. However, the CNN-based methods do not fare so well in terms of processing speed. Specifically, the DAN [13] method achieves the state-of-the-art accuracy, however it can achieve only 5 FPS while running on a CPU (speed measured using the open source implementation from the authors). In order to achieve real-time processing speeds, CNN-based methods depend on the availability of GPUs.

The proposed RRR- and NN-based methods provide an excellent accuracy-to-speed trade-off. Also, their memory requirements are quite modest: less than 2MB of data files. These characteristics make them an excellent candidate for face alignment on mobile devices and embedded hardware.

Some typical success and failure cases for RRR ($q = 4$ and $p = 7$) can be seen in Figure 4. The method sometimes fails for points that lie on the face contour. These images demonstrate this issue in a qualitative way.

4 Conclusion

We investigated the possibility of reducing the storage requirements of tree ensembles that predict high-dimensional targets. The obtained results on the face-alignment task are encouraging. We recommend to use the backpropagation-based learning of neural architectures that transform the large and sparse tree-derived encodings of the input into the desired

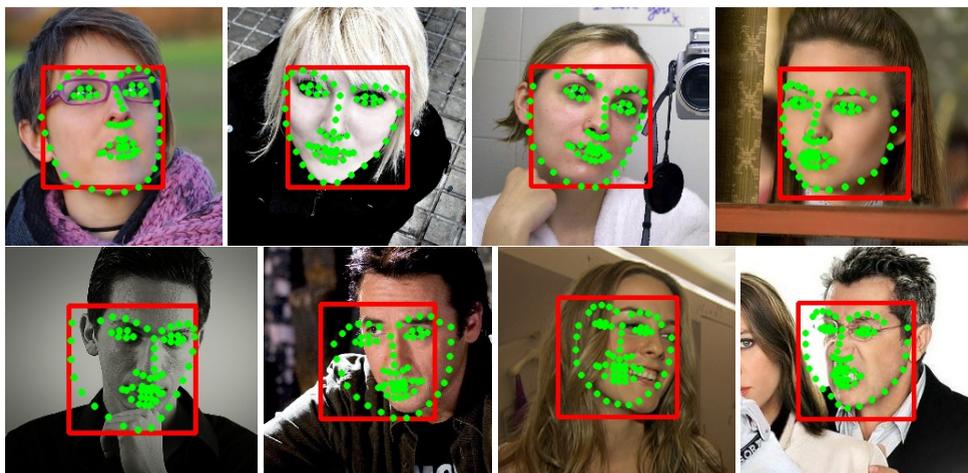


Figure 4: Some success (top row) and failure (bottom row) cases.

output. This is due to the scalability of the method and its simplicity. Also, the approach can naturally handle both classification and regression tasks, as well as multi-objective problems.

References

- [1] A. Asthana, S. Zafeiriou, S. Cheng, and M. Pantic. Robust discriminative response map fitting with constrained local models. In *CVPR*, 2013.
- [2] L. Breiman. Random forests. *Journal of Machine Learning Research*, 2001.
- [3] L. Breiman, J. Friedman, C. J. Stone, and R. A. Olshen. *Classification and Regression Trees*. Chapman and Hall, 1984.
- [4] P. Burgos-Artizzu, P. Perona, and P. Dollar. Robust face landmark estimation under occlusion. In *CVPR*, 2013.
- [5] X. Cao, Y. Wei, F. Wen, and J. Sun. Face alignment by explicit shape regression. In *CVPR*, 2012.
- [6] E. Denton, W. Zaremba, J. Bruna, Y. LeCun, and R. Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *NIPS*, 2014.
- [7] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. Liblinear: A library for large linear classification. *JMLR*, 2008.
- [8] M. Fernández-Delgado, E. Cernadas, S. Barro, and D. Amorim. Do we need hundreds of classifiers to solve real world classification problems? *Journal of Machine Learning Research*, 2014.
- [9] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. In *ICLR*, 2016.

- [10] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [11] Xin Jin and Xiaoyang Tan. Face alignment in-the-wild: A survey. *Computer Vision and Image Understanding*, 162:1–22, 2017.
- [12] V. Kazemi and J. Sullivan. One millisecond face alignment with an ensemble of regression trees. In *CVPR*, 2014.
- [13] Marek Kowalski, Jacek Naruniec, and Tomasz Trzcinski. Deep alignment network: A convolutional neural network for robust face alignment. In *Proceedings of the International Conference on Computer Vision & Pattern Recognition (CVPRW), Faces-in-the-wild Workshop/Challenge*, volume 3, page 6, 2017.
- [14] H. Lai, S. Xiao, Z. Cui, Y. Pan, C. Xu, and S. Yan. Deep cascaded regression for face alignment. *CoRR*, abs/1510.09083, 2015.
- [15] D. Lee, H. Park, and C. D. Yoo. Face alignment using cascade gaussian process regression trees. In *CVPR*, 2015.
- [16] G. C. Reinsel and R. P. Velu. *Multivariate Reduced-Rank Regression: Theory and Applications*. Springer, 1998.
- [17] S. Ren, X. Cao, Y. Wei, and J. Sun. Face alignment at 3000 fps via regressing local binary features. In *CVPR*, 2014.
- [18] S. Ren, X. Cao, Y. Wei, and J. Sun. Global refinement of random forest. In *CVPR*, 2015.
- [19] S. Ren, X. Cao, Y. Wei, and J. Sun. Face alignment via regressing local binary features. In *IEEE Transactions on Image Processing*, 2016.
- [20] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. 300 Faces in-the-Wild Challenge: The first facial landmark localization Challenge. In *ICCV workshop*, 2013.
- [21] T. N. Sainath, B. Kingsbury, V. Sindhvani, E. Arisoy, and B. Ramabhadran. Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In *ICASSP*, 2013.
- [22] B. M. Smith, J. Brandt, Z. Lin, and L. Zhang. Nonparametric context modeling of local appearance for pose-and expression-robust facial landmark localization. In *CVPR*, 2014.
- [23] R. K. Srivastava, K. Greff, and J. Schmidhuber. Training very deep networks. In *NIPS*, 2015.
- [24] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [25] G. Tzimiropoulos and M. Pantic. Gauss-newton deformable part models for face alignment in-the-wild. In *CVPR*, 2014.
- [26] M. Wainberg, B. Alipanahi, and B. J. Frey. Are random forests truly the best classifiers? *Journal of Machine Learning Research*, 2016.

- [27] X. Xiong and F. De la Torre. Supervised descent method and its applications to face alignment. In *CVPR*, 2013.
- [28] J. Zhang, S. Shan, M. Kan, and X. Chen. Coarse-to-fine auto-encoder networks (cfan) for real-time face alignment. In *ECCV*, 2014.
- [29] Y. Zhang, E. Chuangsuwanich, and J. R. Glass. Extracting deep neural network bottleneck features using low-rank matrix factorization. In *ICASSP*, 2014.
- [30] Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Facial landmark detection by deep multi-task learning. In *ECCV*, 2014.
- [31] X. Zhao, T.-K. Kim, and W. Luo. Unified face analysis by iterative multi-output random forests. In *CVPR*, 2014.
- [32] S. Zhu, C. Li, C. C. Loy, and X. Tang. Face alignment by coarse-to-fine shape searching. In *CVPR*, 2015.
- [33] X. Zhu and D. Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *CVPR*, 2012.