Facial Landmarks Detector Learned by the Structured Output SVM

Michal Uřičář¹, Vojtěch Franc¹, and Václav Hlaváč¹

Center for Machine Perception, Department of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague, Technická 2, 166 27 Prague 6, Czech Republic {uricamic, xfrancv, hlavac}@cmp.felk.cvut.cz

Keywords: Facial Landmarks Detection, Deformable Part Models, Structured Output Classification, Structured Output SVM.

Abstract. We propose a principled approach to supervised learning of facial landmarks detector based on the Deformable Part Models (DPM). We treat the task of landmarks detection as an instance of the structured output classification. To learn the parameters of the detector we use the Structured Output Support Vector Machines algorithm. The objective function of the learning algorithm is directly related to the performance of the detector and controlled by the user-defined loss function, in contrast to the previous works. Our proposed detector is real-time on a standard computer, simple to implement and easily modifiable for detector on a challenging "Labeled Faces in the Wild" (LFW) database. The empirical results show that our detector consistently outperforms two public domain implementations based on the Active Appearance Models and the DPM. We are releasing open-source code implementing our proposed detector along with the manual annotation of seven facial landmarks for nearly all images in the LFW database.

1 Introduction

The detection of facial landmarks like canthi nose and mouth corners (see Fig. 1) is an essential part of face recognition systems. The accuracy of the detection significantly influences its final performance [1,2,3]. The problem of the precise and robust detection of facial landmarks has received a lot of attention in the past decade. We briefly review only the approaches relevant to the method proposed in this paper.

Among the most popular are detectors based on the Active Appearance Models (AAM) [4] which use a joint statistical model of appearance and shape. Detectors build on AAM provide a dense set of facial features, allowing to extract whole contours of facial parts like eyes, etc. However high resolution images are required for both training and testing stage and the detection leads to solving a non-convex optimization problem susceptible to local optima unless a good initial guess of the landmark positions is available.



Fig. 1. Functionality of the facial landmark detector.

A straightforward approach to landmark detection is based on using independently trained detectors for each facial landmark. For instance the AdaBoost based detectors and its modifications have been frequently used [5]. If applied independently, the individual detectors often fail to provide a robust estimate of the landmark positions. The weakness of the local evidence can be compensated by using a prior on the geometrical configuration of landmarks. The detection is typically carried out in two consecutive steps. In the first step, the individual detectors are used to find a set of candidate positions for each landmark separately. In the second step, the best landmark configuration with the highest support from the geometrical prior is selected. The landmark detectors based on this approach were proposed for example in [6,7,8,9].

The Deformable Part Models (DPM) [10,11,12,13] go one step further by fusing the local appearance model and the geometrical constraints into a single model. The DPM is given by a set of parts along with a set of connections between certain pairs of parts arranged in a deformable configuration. A natural way how to describe the DPM is an undirected graph with vertices corresponding to the parts and edges representing the connections between the pairs of connected parts. The DPM detector estimates all landmark positions simultaneously by optimizing a single scoring function composed of a local appearance model and a deformation cost. The complexity of finding the best landmark configuration depends on the structure of underlying graph. Acyclic graph allows efficient estimation by a variant of the Dynamic Programming.

An instance of finely tuned facial landmark detector based on the DPM has been proposed in [14]. The very same detector was also used in several successful face recognition systems described in [15] and [16]. In this case, the local appearance model is learned by a multiple-instance variant of the AdaBoost algorithm with Haar-like features used as the weak classifiers. The deformation cost is expressed as a mixture of Gaussian trees whose parameters are learned from examples. This landmark detector is publicly available and we use it for comparison with our detector. ¹ Importantly, learning of the local appearance model and the deformation cost is done in two independent steps which simplifies learning, but may not be optimal in terms of detectors accuracy.

¹ There also exists a successful commercial solution OKAO Vision Facial Feature Extraction API (http://www.omron.com) which is used for example in Picasa[™] or Apple iPhoto[™] software.

We propose to learn the parameters of the DPM discriminately in one step by directly optimizing accuracy of the resulting detector. The main contributions of this paper are as follows:

- We treat the landmark detection with the DPM as an instance of the structured output classification problem whose detection accuracy is measured by a loss function natural for this application. We propose to use the Structured Output SVM (SO-SVM) [17] for supervised learning of the parameters of the landmark detector from examples. The learning objective of the SO-SVMs is directly related to the accuracy of the detector. In contrast, all existing approaches we are aware of optimize surrogate objective functions whose relation to the detector accuracy is not always clear.
- 2. We empirically evaluate accuracy of the proposed landmark detector learned by the SO-SVMs on a challenging "Labeled Faces in the Wild" (LFW) database [18].
- We provide an empirical comparison of two popular optimization algorithms the Bundle Method for Regularized Risk Minimization (BMRM) [19] and the Stochastic Gradient Descend (SGD) [20] — which are suitable for solving the convex optimization problem emerging in the SO-SVM learning.
- 4. We provide an open source library which implements the proposed detector and the algorithm for supervised learning of its parameters. In addition we provide a manual annotation of the facial landmarks for all images from the LFW database.

The paper is organized as follows. Section 2 defines the structured output classifier for facial landmark detection based on the DPM. Section 3 describes the SO-SVM algorithm for learning the parameters of the classifier from examples. Experimental results are presented in Sect. 4. Section 5 shortly describes the open source implementation of our detector and the provided manual annotation of the LFW database. Section 6 concludes the paper.

2 Structured Output Classifier for Facial Landmark Detection

We treat the landmark detection as an instance of the structured output classification problem. We assume that the input of our classifier is a still image of fixed size containing a single face. In our experiments we construct the input image by cropping a window around a bounding box found by a face detector (enlarged by a fixed margin ensuring that the whole face is contained) and normalizing its size. The classifier output are estimated locations of a set of facial landmarks. A formal definition is given next.

Let $\mathcal{I} = \mathcal{X}^{H \times W}$ be a set of input images with $H \times W$ pixels, where \mathcal{X} denotes a set of pixel values which in our experiments, dealing with 8bit gray-scale images, is $\mathcal{X} = \{0, \dots, 255\}$. We describe the configuration of M landmarks by a graph G = (V, E), where $V = \{0, \dots, M-1\}$ is a set of landmarks and $E \subset V^2$ is a set of edges defining the neighboring landmarks. Each landmark is assigned a position $s_i \in S_i \subset \{1, \dots, H\} \times \{1, \dots, W\}$, where S_i denotes a set of all admissible positions of the *i*-th landmark within the image $I \in \mathcal{I}$. The quality of a landmark configuration $\mathbf{s} = (\mathbf{s}_0, \dots, \mathbf{s}_{M-1}) \in \mathcal{S} = S_0 \times \cdots \times S_{M-1}$ given an input image $I \in \mathcal{I}$ is measured by a scoring function $f: \mathcal{I} \times S \to \mathbb{R}$

defined as

$$f(I,\mathbf{s}) = \sum_{i \in V} q_i(I,\mathbf{s}_i) + \sum_{(i,j) \in E} g_{ij}(\mathbf{s}_i,\mathbf{s}_j) \quad .$$

$$\tag{1}$$

The first term in (1) corresponds to a local appearance model evaluating the match between landmarks on positions s and the input image I. The second term in (1) is the deformation cost evaluating the relative positions of the neighboring landmarks i and j.

We assume that the costs $q_i: \mathcal{I} \times \mathcal{S}_i \to \mathbb{R}, i = 0, ..., M-1$ and $g_{ij}: \mathcal{S}_i \times \mathcal{S}_j \to \mathbb{R}, (i, j) \in E$ are linearly parametrized functions

$$q_i(I, \mathbf{s}_i) = \langle \mathbf{w}_i^q, \Psi_i^q(I, \mathbf{s}_i) \rangle \tag{2}$$

$$g_{ij}(\mathbf{s}_i, \mathbf{s}_j) = \langle \mathbf{w}_{ij}^g, \Psi_{ij}^g(\mathbf{s}_i, \mathbf{s}_j) \rangle \quad , \tag{3}$$

where $\Psi_i^q: \mathcal{J} \times \mathcal{S}_i \to \mathbb{R}^{n_{iq}}, \Psi_{ij}^g: \mathcal{S}_i \times \mathcal{S}_j \to \mathbb{R}^{n_{ig}}, i = 0, \dots, M-1$ are predefined maps and $\mathbf{w}_i^q \in \mathbb{R}^{n_{iq}}, \mathbf{w}_{ij}^g \in \mathbb{R}^{n_{ig}}, i = 0, \dots, M-1$ are parameter vectors which will be learned from examples. Let us introduce a joint map $\Psi: \mathcal{J} \times \mathcal{S} \to \mathbb{R}^n$ and a joint parameter vector $\mathbf{w} \in \mathbb{R}^n$ defined as a column-wise concatenation of the individual maps Ψ_i^q, Ψ_{ij}^g and the individual parameter vectors $\mathbf{w}_i^q, \mathbf{w}_{ij}^g$ respectively. With these definitions we see that the scoring function (1) simplifies to $f(I, \mathbf{s}) = \langle \mathbf{w}, \Psi(I, \mathbf{s}) \rangle$.



Fig. 2. Definition of (a) the underlying graph G = (V, E) for the landmark configuration and (b) the components of the proposed detector.

Given an input image *I*, the structured output classifier returns the configurations \hat{s} computed by maximizing the scoring function f(I, s), i.e.

$$\mathbf{\hat{s}} \in \arg\max_{\mathbf{s}\in\mathcal{S}} f(I,\mathbf{s}) \,. \tag{4}$$

We assume that the graph G = (V, E) is acyclic (see Fig. 2(a)), which allows efficient solving of the maximization problem (4) by dynamic programming.

A complete specification of the structured classifier (4) requires to define:

- The maps $\Psi_i^q(I, \mathbf{s}_i), i = 0, \dots, M-1$ where $\Psi_i^q(I, \mathbf{s}_i)$ defines a local feature descriptor of *i*-th landmark computed on a rectangular window centered at \mathbf{s}_i . We call the rectangular window a component (see Fig. 2(b)). The size of the component and its feature descriptor are crucial design options which have to be made carefully. In Sect. 2.1 we describe a list of feature descriptors we have considered.
- The fixed maps $\Psi_{ij}^{g}(\mathbf{s}_{i},\mathbf{s}_{j}), (i, j) \in E$ defining the parametrization of the deformation cost. Section 2.2 describes the parametrization which we have considered.
- The set $S = (S_0 \times \cdots \times S_{M-1})$ defining the search space of the landmark positions. These sets can be interpreted as hard constraints on the admissible configurations of the landmarks, i.e. the landmark positions outside these sets corresponds to $-\infty$ value of the deformation $\cos g_{ii}(\mathbf{s}_i, \mathbf{s}_i)$.

We tune the size of these search spaces experimentally — we keep track of the axis aligned bounding box (AABB) for each component throughout the whole database excluding the images whose components does not fit in the image. We set the size of components in order to keep at least 95% images of the original database. Consequently, the AABB of each component is made vertically symmetric along the center *y*-axis in order to remove bias to certain positions. Figure 3 visualizes the found search spaces.

- The joint parameter vector $\mathbf{w} \in \mathbb{R}^n$ learned from the training examples by the SO-SVM algorithm described in Sect. 3.



Fig. 3. Left: optimal search spaces for each component. Right: the same search spaces made symmetrical along the vertical magenta line.

Finally we would like to stress that the particular number of landmarks and their neighborhood structure can be arbitrary as long as the inference problem (1) can be solved efficiently. In this paper we experiment with the 8-landmarks variant of the graph G = (V, E) shown in Fig. 2(a).

2.1 Appearance Model

We have experimented with several feature descriptors Ψ_i^q for the local appearance model $q_i(I, \mathbf{s}_i)$. In particular, we considered i) normalized intensity values, ii) derivatives of image intensity values, iii) histograms of Local Binary Patterns (LBP) [21]

and iv) the LBP pyramid feature descriptor [22]. We obtained the best results with the LBP pyramid feature descriptor which is used in the experiments. The LBP pyramid descriptor is constructed by concatenating binary encoded LBP features computed in each pixel (up to boundary pixels) and in several scales. In particular, we use the LBP pyramid computed in 4 scales starting from the original image and consequently down-scaling the image 3 times by 1/2. The resulting feature vector is high dimensional but very sparse.

2.2 Deformation Cost

We have experimented with two parametrizations of the deformation cost $g_{ij}(\mathbf{s}_i, \mathbf{s}_j)$: i) a table representation and ii) a quadratic function of a displacement vector between landmark positions.

The table representation is the most generic form of the deformation cost useful when no prior knowledge is available. Table elements specify cost for each combination of \mathbf{s}_i and \mathbf{s}_j separately. $\Psi_{ij}^g(\mathbf{s}_i, \mathbf{s}_j)$ is a sparse vector with all elements zero but the element corresponding to the combinations $(\mathbf{s}_i, \mathbf{s}_j)$ which is one. Such a representation of the deformation cost is very flexible and easy to implement. Though the table representation is very flexible its main disadvantage is a very large number of parameters to be learned. In turn, a large number of training examples is required to avoid over-fitting. In fact, each combination $(\mathbf{s}_i, \mathbf{s}_j)$ should be present in training examples at least once to make the corresponding cost non-zero.

As the second option, we considered the deformation $\cos t g_{ij}(\mathbf{s}_i, \mathbf{s}_j)$ to be a quadratic function of a displacement vector $\mathbf{s}_j - \mathbf{s}_i$. Following [12], we define the deformation cost as

$$\left.\begin{array}{l}
\Psi_{ij}^{g}(\mathbf{s}_{i},\mathbf{s}_{j}) = (dx,dy,dx^{2},dy^{2})\\
(dx,dy) = (x_{i},y_{j}) - (x_{i},y_{i})
\end{array}\right\}$$
(5)

This representation accounts for the distance and the direction of the *j*-th landmark with respect to *i*-th landmark. This representation is determined only by four parameters which substantially reduces the risk of over-fitting.

We found experimentally the quadratic deformation cost to give slightly better results compared to the table representation, therefore experiments described further use only this parametrization of deformation cost.

3 Learning the Parameters of the Structured Output Classifier

We learn the joint parameter vector **w** by the SO-SVM algorithm [17]. The requirements on the classifier are specified by a user defined loss-function $\ell: S \times S \to \mathbb{R}$. The value $\ell(\mathbf{s}, \mathbf{s}^*)$ penalizes the classifier estimate **s** provided the actual configuration of the landmarks is \mathbf{s}^* . The SO-SVM requires loss function to be non-negative and zero iff the estimate is absolutely correct, i.e. $\ell(\mathbf{s}, \mathbf{s}^*) \ge 0$, $\forall \mathbf{s}, \mathbf{s}^* \in S$, and $\ell(\mathbf{s}, \mathbf{s}^*) = 0$ iff $\mathbf{s} = \mathbf{s}^*$. In particular, we use the mean normalized deviation between the estimated and the ground truth positions as the loss function, i.e.

$$\ell(\mathbf{s}, \mathbf{s}^*) = \kappa(s^*) \frac{1}{M} \sum_{j=0}^{M-1} \|\mathbf{s}_j - \mathbf{s}_j^*\| \quad .$$
(6)

The normalization factor $\kappa(s^*) = \|\frac{1}{2}(\mathbf{s_{eyeR}} + \mathbf{s_{eyeL}}) - \mathbf{s_{mouth}}\|^{-1}$ is reciprocal to the face size which we define as the length of the line connecting the midpoint between the eye centers $\mathbf{s_{eyeR}}$ and $\mathbf{s_{eyeL}}$ with the mouth center $\mathbf{s_{mouth}}$. The normalization factor is introduced in order to make the loss function scale invariant which is necessary because responses of the face detector used to construct the input images do not allow accurate estimation of the scale. Figure 4 illustrates the meaning of the loss function (6). Finally, we point out that any other loss function meeting the constraints defined above can be readily used, e.g. one can use maximal normalized deviation.

Given a set of training examples $\{(I^1, \mathbf{s}^1), \dots, (I^m, \mathbf{s}^m)\} \in (\mathcal{I} \times \mathcal{S})^m$ composed of pairs of the images and their manual annotations, the joint parameter vector **w** of the classifier (4) is obtained by solving the following convex minimization problem

$$\mathbf{w}^* = \arg\min_{\mathbf{w}\in\mathbb{R}^n} \left[\frac{\lambda}{2} \|\mathbf{w}\|^2 + R(\mathbf{w})\right] , \text{ where}$$
(7)

$$R(\mathbf{w}) = \frac{1}{m} \sum_{i=1}^{m} \max_{s \in \mathcal{S}} \left(\ell(\mathbf{s}^{i}, \mathbf{s}) + \left\langle \mathbf{w}, \Psi(I^{i}, \mathbf{s}) \right\rangle \right) - \frac{1}{m} \sum_{i=1}^{m} \left\langle \mathbf{w}, \Psi(I^{i}, \mathbf{s}^{i}) \right\rangle \quad . \tag{8}$$

The number $\lambda \in \mathbb{R}^+$ is a regularization constant whose optimal value is tuned on a validation set. $R(\mathbf{w})$ is a convex piece-wise linear upper bound on the empirical risk $\frac{1}{m}\sum_{i=1}^{m} \ell(\mathbf{s}^i, \arg\max_{\mathbf{s}\in S} f(I^i, \mathbf{s}))$. That is, the learning algorithm directly minimizes the performance of the detector assessed on the training set and at the same time it controls the risk of over-fitting via the norm of the parameter vector.

Though the problem (7) is convex its solving is hard. The hardness of the problem can be seen when it is expressed as an equivalent quadratic program with m|S| linear constraints (recall that |S| is the number of all landmark configurations). This fact rules out off-the-shelf optimization algorithms.

Thanks to its importance a considerable effort has been put to a development of efficient optimization algorithms for solving the task (7). There has been an ongoing discussion in the machine learning community trying to decide whether approximative on-line solvers like the SGD are better than the accurate slower solvers like the BMRM. No definitive consensus has been achieved so far. We contribute to this discussion by providing an empirical evaluation of both approaches on the practical large-scale problem required to learn the landmark detector. The empirical results are provided in Sect. 4.5. For the sake of self-consistency, we briefly describe the considered solvers, i.e. the BMRM and the SGD algorithm, in the following two sections.

3.1 Bundle Methods for Regularized Risk Minimization

The BMRM is a generic method for minimization of regularized convex functions [19], i.e. BMRM solves the following convex problem

$$\mathbf{w}^* = \arg\min_{\mathbf{w}\in\mathbb{R}^n} F(\mathbf{w}) := \frac{\lambda}{2} \|\mathbf{w}\|^2 + R(\mathbf{w}) \quad , \tag{9}$$

where $R: \mathbb{R}^n \to \mathbb{R}$ is an arbitrary convex function. The risk term $R(\mathbf{w})$ is usually the complex part of the objective function which makes the optimization task hard. The

core idea is to replace the original problem (9) by its reduced problem

$$\mathbf{w}_t = \arg\min_{\mathbf{w}\in\mathbb{R}^n} F_t(\mathbf{w}) := \frac{\lambda}{2} \|\mathbf{w}\|^2 + R_t(\mathbf{w}) \quad . \tag{10}$$

The objective function $F_t(\mathbf{w})$ of the reduced problem (10) is obtained after replacing the risk $R(\mathbf{w})$ in the original objective $F(\mathbf{w})$ by its cutting plane model

$$R_t(\mathbf{w}) = \max_{i=0,1,\dots,t-1} \left[R(\mathbf{w}_i + \langle R'(\mathbf{w}_i), \mathbf{w} - \mathbf{w}_i \rangle \right] \quad , \tag{11}$$

where $R'(\mathbf{w}_i) \in \partial_{\mathbf{w}_i} R(\mathbf{w}_i)$ denotes an arbitrary sub-gradient of R evaluated at the point $\mathbf{w}_i \in \mathbb{R}^n$.

Starting from an initial guess $\mathbf{w}_0 = \mathbf{0}$, the BMRM algorithm computes a new iterate \mathbf{w}_t by solving the reduced problem (10). In each iteration t, the cutting plane model (11) is updated by a new cutting plane computed at the intermediate solution \mathbf{w}_t leading to a progressively tighter approximation of $F(\mathbf{w})$. The BMRM algorithm halts if the gap between $F(\mathbf{w}_t)$ (an upper bound on $F(\mathbf{w}^*)$) and $F_t(\mathbf{w}_t)$ (a lower bound on $F(\mathbf{w}^*)$) falls below a desired ε , meaning that $F(\mathbf{w}_t) \leq F(\mathbf{w}^*) + \varepsilon$. The BMRM algorithm stops after at most $O(1/\varepsilon)$ iterations for arbitrary $\varepsilon > 0$ [19].

The reduced problem (10) can be expressed as an equivalent convex quadratic program with t variables. Because t is usually small (up to a few hundreds), off-the-shelf QP solvers can be used.

Before applied to a particular problem, the BMRM algorithm requires a procedure which for a given w returns the value of the risk $R(\mathbf{w})$ and its sub-gradient $R'(\mathbf{w})$. In our case the risk $R(\mathbf{w})$ is defined by (8) and its sub-gradient can be computed by the Danskin's theorem as

$$R'(\mathbf{w}) = \frac{1}{m} \sum_{i=1}^{m} \left(\Psi(I^i, \hat{\mathbf{s}}^i) - \Psi(I^i, \mathbf{s}^i) \right)$$
(12)

$$\hat{\mathbf{s}}^{i} = \arg\max_{\mathbf{s}\in\mathcal{S}} \left[\ell(\mathbf{s}^{i}, \mathbf{s}) + \left\langle \mathbf{w}, \Psi(I^{i}, \mathbf{s}) \right\rangle \right] . \tag{13}$$

Note that the evaluation of $R(\mathbf{w})$ and $R'(\mathbf{w})$ is dominated by the computation of the scalar products $\langle \mathbf{w}, \Psi(l^i, \mathbf{s}) \rangle$, $i = 1, ..., m, \mathbf{s} \in S$, which, fortunately, can be efficiently parallelized.

3.2 Stochastic Gradient Descent

Another popular method solving (7) is the Stochastic Gradient Descent (SGD) algorithm. We use the modification proposed in [20] which uses two neat tricks. Starting from an initial guess \mathbf{w}_0 , the SGD algorithm iteratively changes \mathbf{w} by applying the following rule:

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \frac{\lambda^{-1}}{t_0 + t} g_t, \quad g_t = \lambda \mathbf{w}_t + \mathbf{h}_t \quad . \tag{14}$$

 t_0 is a constant and t is the number of the iteration. The SGD implementation proposed in [20] tunes the optimal value of t_0 on a small portion of training examples subsampled

from training set. The sub-gradient is computed in almost the same manner as in (12), but only for one training image at a time, i.e. $\mathbf{h}_t = \Psi(I^t, \hat{\mathbf{s}}^t) - \Psi(I^t, \mathbf{s}^t)$.

In addition, [20] propose to exploit the sparsity of the data in the update step. Equation (14) can be expressed as

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha_t \mathbf{w}_t - \beta_t \mathbf{h}_t, \text{ where}$$
(15)

$$\alpha = \frac{1}{t_0 + t}, \quad \beta = \frac{\lambda^{-1}}{t_0 + t} \quad . \tag{16}$$

Note that if \mathbf{h}_t is sparse then subtracting $\beta_t \mathbf{h}_t$ involves only the nonzero coefficients of \mathbf{h}_t , but subtracting $\alpha_t \mathbf{w}_t$ involves all coefficients of \mathbf{w}_t . In turn, it is beneficial to reformulate (15) as

$$\mathbf{w}_{t+1} = (1 - \alpha_t)\mathbf{w}_t - \beta_t \mathbf{h}_t \quad . \tag{17}$$

By using this trick, the complexity O(d) corresponding to the naïve implementation of the update rule (14) reduces to the complexity $O(d_{non-zero})$ corresponding to the reformulated rule (17), where *d* is the dimension of the parameter vector and $d_{non-zero}$ is the number of the non-zero elements in \mathbf{h}_t . Typically, like in our case, $d_{non-zero}$ is much smaller than *d*.

A considerable advantage of the SGD algorithm is its simplicity. A disadvantage is a lack of any certificate of optimality and thus also of the theoretically grounded stopping condition.

4 Experiments

In this section, we present experimental evaluation of the proposed facial landmark detector and its comparison against three different approaches. We considered the detector estimating positions of the eight landmarks: the canthi of the left and right eye, the corners of the mouth, the tip of the nose and the center of the face. The corresponding graph (V, E) is shown in Fig. 2(a).

In Sect. 4.1, we describe the face database and the testing protocol used in the experiments. The competing methods are summarized in Sect. 4.2. The results of the comparison in terms of detection accuracy and basic timing statistics are presented in Sect. 4.4. Finally, in Sect. 4.5 we compare two algorithms for solving the large-scale optimization problems emerging in the SO-SVM learning, namely, the BMRM and the SGD algorithm.

4.1 Database and Testing Protocol

We use the LFW database [18] for evaluation as well as for training of our detector. This database consists of 13,233 images each of size 250×250 pixels. The LFW database contains a great ethnicity variance and the images have challenging background clutter. We augmented the original LFW database by adding manual annotation of the eight considered landmarks.

We randomly split the LFW database into training, testing and validation sets. Table 1 describes this partitioning. The experimental evaluation of all competing detectors



Fig. 4. The illustration of two accuracy statistics used to benchmark the detectors. The green and the red crosses denote the manually annotated landmarks and the detected landmarks, respectively. The deviations $\varepsilon_0, \ldots, \varepsilon_7$ correspond to radii of the dashed circles.

Table 1. The partitioning of the LFW database into training, validation and testing set.

Data set	Training	Validation	Testing
Percentage	60%	20%	20%
# of examples	6,919	2,307	2,316

was made on the same testing set. The training and the validation parts were used for learning of the proposed detector and the base line SVM detector. The other competing detectors had their own training databases.

In order to evaluate the detectors, we use two accuracy measures: i) the mean normalized deviation $\ell(\mathbf{s}, \mathbf{s}')$ defined by equation (6) and ii) the maximal normalized deviation

$$\ell^{\max}(\mathbf{s}, \hat{\mathbf{s}}) = \kappa(\mathbf{s}) \max_{j=0,\dots,M-1} \|\mathbf{s}_j - \hat{\mathbf{s}}_j\| \quad , \tag{18}$$

where $\mathbf{s} = (\mathbf{s}_0, \dots, \mathbf{s}_{M-1})$ are the manually annotated landmark positions and $\hat{\mathbf{s}} = (\hat{\mathbf{s}}_0, \dots, \hat{\mathbf{s}}_{M-1})$ are the landmark positions estimated by the tested detector. Figure 4 illustrates both accuracy measures.

4.2 Competing Methods

In this section, we outline all detectors that were used in the experimental evaluation.

Proposed Detector The proposed detector estimates the landmark positions according to (4). As the feature descriptor $\Psi_i^q(I, \mathbf{s}_i)$ defining the local appearance model $q_i(I, \mathbf{s}_i)$, we use the LBP pyramid described in Sect. 2.1. As the parametrization $\Psi_{ij}^g(\mathbf{s}_i, \mathbf{s}_j)$ of the deformation cost $g_{ij}(\mathbf{s}_i, \mathbf{s}_j)$, we use the quadratic function described in Sect. 2.2. The parameter vector \mathbf{w} of the classifier (4) is trained from the training part of the LFW database using the BMRM algorithm (c.f. Sect. 3.1). The regularization constant λ

appearing in the learning problem (7) was selected from the set $\{10, 1, 0.1, 0.01, 0.001\}$ to minimize the average mean normalized deviation R_{VAL} computed on the validation part of the LFW database.

Independently Trained SVM Detector This detector is formed by standard two-class linear SVM classifiers trained independently for each landmark. For training, we use the SVM solver implemented in LIBOCAS [22]. For each individual landmark we created a different training set containing examples of the positive and negative class. The positive class is formed by images cropped around the ground truth positions of the respective component. The negative class contains images cropped outside the ground truth regions. Specifically, the negative class images satisfy the following condition

$$|P_{-}^{x} - P_{\text{GT}}^{x}| > \frac{1}{2} \text{width}_{\text{GT}}, |P_{-}^{y} - P_{\text{GT}}^{y}| > \frac{1}{2} \text{height}_{\text{GT}},$$
 (19)

where P_{-}^{x} and P_{GT}^{x} is the *x*-coordinate of the negative and the ground truth component respectively. height_{GT} and width_{GT} denote the width and the height of the component.

We use the LBP-pyramid descriptor (see Sect. 2.1) as the features. The parameters of the linear SVM classifier are learned from the training part of the LFW database. The SVM regularization constant *C* was selected from the set $\{10, 1, 0.1, 0.01, 0.001\}$ to minimize the classification error computed on the validation part of the LFW database.

Having the binary SVM classifiers trained for all components, the landmark position is estimated by selecting the place with the maximal response of the classifier scoring function, evaluated in the search regions defined for each component differently. The search regions as well as the sizes of the components are exactly the same as we use for the proposed SO-SVM detector.

Note that the independently trained SVM detector is a simple instance of the DPM where the deformation cost $g_{ij}(\mathbf{s}_i, \mathbf{s}_j)$ is zero for all positions inside the search region and $-\infty$ outside. We compare this baseline detector with the proposed SO-SVM detector to show that one can improve the accuracy by learning the deformation cost from data.

Active Appearance Models We use a slightly modified version of a publicly available implementation of the AAM [23]. As the initial guess of the face position required by the AAM, we use the center of the bounding box obtained from a face detector. The initial scale is also computed from this bounding box. The AAM estimates a dense set of feature points which are distributed around important face contours like the contour of mouth, eyes, nose, chin and eyebrows. The AAM requires a different training database which contains high resolution images along with annotation of all contour points.

For training the AAM model we use a publicly available IIM Face database [24]. The IIM database consists of 240 annotated images (6 images per person). Each image is 640×480 pixel in size and comes with 58 manually annotated points which are distributed along the main face contours. Note that the creation of training examples for the AAM put much higher demands on the annotator who has to click a large number of uniformly distributed points. In contrast, our method requires annotation of only a small number of well defined points. Specifically, the whole IIM database requires to annotate

13920 points, carefully distributed along each contour, while the LFW database requires to annotate 48433 points, which are well defined and easy to annotate.

To compare the AAM based detector with our detector, we have to transform the output of the AAM, i.e. the points on contours around important face parts, to the land-mark positions returned by our detector. We simply select the relevant points on the corresponding contours.

4.3 Detector of Everingham et al.

The last competing detector is the DPM based detector of [25]. This detector was trained on a collection of consumer images which, however, are not available. This detector returns canthi of both eyes (4 landmarks), corners of the mouth (2 landmarks) and 3 landmarks on the nose. To compare this detector, we consider only the relevant landmarks for our detector. Note that unlike the proposed SO-SVM detector, this detector learns the local appearance model and the deformation cost of the DPM independently.

4.4 Results

In this section, we describe results of the experimental evaluation of the detection accuracy of all competing detectors. We have measured the mean and the maximal normalized deviation computed on the test part of the LFW database.

Table 2 shows the average mean normalized deviation R_{TST} and the the average maximal normalized deviation $R_{\text{TST}}^{\text{max}}$ for each individual detector. The results show that the proposed detector consistently outperforms all other competing methods irrespective to the accuracy measure. Surprisingly, the independently trained SVM detector is comparable with the DPM based detector of [25]. The far worst results were obtained for the AAM based detector which can be attributed to a relatively low resolution input images.

In Fig. 5 we show the cumulative histograms of the mean and maximal normalized deviation. Table 3 shows the percentage of examples from the test part of the LFW database with the mean/maximal normalized deviation less or equal to 10% (this corresponds to the line at 10% of x-axis taken from Fig. 5. It is seen that the proposed detector estimates around 97% of images with the mean normalized deviation less than 10%. This results is far better than was achieved for all other competing methods. In Fig. 6, we show examples of images with the mean normalized deviation equal to 10% for better understanding of these statistics.

We have also measured the average time required by the proposed detector to process a single image. The measurements were done on a notebook with Intel Core 2 Duo T9300 2.50 GHz. The average detection time was 7 ms per face.

4.5 Comparison of BMRM and SGD

In this section, we compare performance of the BMRM and the SGD algorithm on the problem emerging when learning the proposed detector.

The task of the solvers is to minimize the problem stated in (7). Besides the value of the objective function $F(\mathbf{w})$ of the task (7) we also measured the validation risk

Table 2. Average mean normalized deviation and the average maximal normalized deviation computed on the test part of the LFW database.



Fig. 5. Cumulative histograms for the mean and the maximal normalized deviation shown for all competing detectors.

 $R_{VAL}(\mathbf{w})$ being another important criterion characterizing convergence of the learning algorithm.

To make the iterations of both algorithms comparable, we define one iteration of the SGD as a sequence of single update steps equal to the number of training examples. This makes the computational time of both solvers approximately proportional to the number of iterations. The optimal value of parameter t_0 for SGD was selected to minimize the objective function $F(\mathbf{w})$ computed on 10% of the training examples after one pass of the SGD algorithm throughout the data. The parameter t_0 have to be tuned for each value of λ separately. We fixed the total number of iterations of the SGD algorithm to 50.

We run both solver on the problem (7) with the parameters $\lambda \in \{0.001, 0.01, 0.1, 1\}$ recording both $F(\mathbf{w})$ and $R_{VAL}(\mathbf{w})$. Results of the experiment are summarized in Table 4.

It can be seen that the SGD converges quickly at the beginning and it stalls as it approaches the minimum of the objective *F*. Similarly for the validation risk. The optimal value of λ minimizing the validation error was 0.01 for both SGD and BMRM. The test errors computed for the optimal λ were $R_{\text{TST}} = 5.44$ for the SGD and $R_{\text{TST}} =$ 5.54 for the BMRM, i.e., the difference is not negligible. The results for the SGD could be improved by using more than 50 iterations, however, in that case both algorithms

Table 3. The percentage of images from the test part of the LFW database where the mean/maximal normalized deviation of the estimated landmark positions was less or equal to 10%.

	Mean	Maximal
AAM	8.98%	0.62%
Everingham et al.	85.28%	22.93%
binary SVM	85.66%	34.50%
proposed detector	96.59%	53.23%



Fig. 6. Sample images where the estimated landmark positions have the mean normalized deviation equal to 10%. The green and red points denote the manually annotated and estimated landmarks, respectively.

would require the comparable time. Moreover, without the reference solution provided by the BMRM one would not know how to set the optimal number of iterations for the SGD. We conclude that for the tested problem the BMRM produced more accurate solution, but the SGD algorithm was significantly faster. This suggests that the SGD is useful in the cases when using the precise but slower BMRM algorithm is prohibited. In the opposite case the BMRM algorithm returning a solution with the guaranteed optimality certificate is preferable.

5 Open-source library and LFW annotation

We provide an open-source library implementing the proposed DPM detector including the BMRM algorithm for supervised learning of its parameters. The detector is implemented in C and we also provide the MEX interface to MATLAB. The library comes with several toy example applications, e.g. running the detector on still images or a video stream and displaying the landmarks. The library is licensed under the GNU/GPL version 3 and was tested on GNU/Linux and Windows platform.

In addition, we provide a manual annotation of the LFW database for noncommercial use. The following set of landmarks is annotated for each face: canthi for both eyes, the tip of the nose and the corners of mouth, i.e. 7 annotated landmarks in total for each image. Both the library and the annotation can be downloaded from: http://cmp.felk.cvut.cz/~uricamic/flandmark

Table 4. Comparison of the BMRM and the SGD. We show the value of primal objective function $F(\mathbf{w})$ and validation risk R_{VAL} for the 50th iteration (assuming termination of SGD after this iteration) as well as for the number of iterations needed by the BMRM algorithm to find the ε -precise solution.

	λ	$\lambda = 1$		$\lambda = 0.1$		$\lambda = 0.01$		$\lambda = 0.001$	
	# of iterations	50	106	50	201	50	462	50	1200
BMRM	$F(\mathbf{w})$	77.48	62.19	45.13	29.68	35.33	14.62	34.35	7.459
	$R_{\rm VAL}({\bf w})$	23.24	10.48	9.054	6.067	9.054	5.475	9.054	5.876
SGD	$F(\mathbf{w})$	50.88	50.44	20.62	20.52	13.72	10.80	12.86	6.309
	$R_{\rm VAL}({\bf w})$	9.719	9.627	6.156	6.142	5.577	5.496	5.544	5.818

6 Conclusions

In this paper, we have formulated the detection of facial landmarks as an instance of the structured output classification problem. Our structured output classifier is based on the DPM and its parameters can be learned from examples by the SO-SVM algorithm. In contrast to the previous works, the learning objective is directly related to the accuracy of the resulting detector. Experiments on the LFW database show that the proposed detector consistently outperforms a baseline independently trained SVM detector and two public domain detectors based on the AAM and DPM.

Acknowledgements

The first two authors were supported by EC projects FP7-ICT-247525 HUMAVIPS and PERG04-GA-2008-239455 SEMISOL. The third author was supported by the Grant Agency of the Czech Republic under Project P103/10/0783. All authors were supported by The Grant Agency of the CTU Prague project SGS12/187/OHK3/3T/13. Any opinions expressed in this paper do not necessarily reflect the views of the European Community. The Community is not liable for any use that may be made of the information contained herein.

References

- Beumer, G., Veldhuis, R.: On the accuracy of EERs in face recognition and the importance of reliable registration. In: 5th IEEE Benelux Signal Processing Symposium (SPS-2005), IEEE Benelux Signal Processing Chapter (2005) 85–88
- Cristinacce, D., Cootes, T., Scott, I.: A multi-stage approach to facial feature detection. In: 15th British Machine Vision Conference (BMVC-2004). (2004) 277–286
- 3. Riopka, T., Boult, T.: The eyes have it. In: Proceedings of ACM SIGMM Multimedia Biometrics Methods and Applications Workshop. (2003) 9–16
- Cootes, T., Edwards, G.J., Taylor, C.J.: Active appearance models. IEEE Trans. Pattern Analysis and Machine Intelligence 23(6) (June 2001) 681–685
- Viola, P., Jones, M.: Robust real-time face detection. International Journal of Computer Vision 57(2) (2004) 137–154

- Beumer, G., Tao, Q., Bazen, A., Veldhuis, R.: A landmark paper in face recognition. In: 7th International Conference on Automatic Face and Gesture Recognition (FGR-2006), IEEE Computer Society Press (2006)
- Cristinacce, D., Cootes, T.: Facial feature detection using AdaBoost with shape constraints. In: 14th Proceedings British Machine Vision Conference (BMVC-2003). (2003) 231–240
- Erukhimov, V., Lee, K.: A bottom-up framework for robust facial feature detection. In: 8th IEEE International Conference on Automatic Face and Gesture Recognition (FG2008). (2008) 1–6
- Wu, J., Trivedi, M.: Robust facial landmark detection for intelligent vehicle system. In: IEEE International Workshop on Analysis and Modeling of Faces and Gestures. (2005)
- Crandall, D., Felzenszwalb, P., Huttenlocher, D.: Spatial priors for part-based recognition using statistical models. In: In CVPR. (2005) 10–17
- Felzenszwalb, P.F., Huttenlocher, D.P.: Pictorial structures for object recognition. Internatinal Journal of Computer Vision 61 (January 2005) 55–79
- Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D.: Object detection with discriminatively trained part based models. IEEE Transactions on Pattern Analysis and Machine Intelligence 99(1) (2009)
- Fischler, M.A., Elschlager, R.A.: The representation and matching of pictorial structures. IEEE Transactions on Computers C-22(1) (1973) 67–92
- 14. Everingham, M., Sivic, J., Zisserman, A.: "Hello! My name is... Buffy" automatic naming of characters in TV video. In: Proceedings of the British Machine Vision Conference. (2006)
- Everingham, M., Sivic, J., Zisserman, A.: Taking the bite out of automatic naming of characters in TV video. Image and Vision Computing 27(5) (2009)
- Sivic, J., Everingham, M., Zisserman, A.: "Who are you?" learning person specific classifiers from video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2009)
- Tsochantaridis, I., Joachims, T., Hofmann, T., Altun, Y., Singer, Y.: Large margin methods for structured and interdependent output variables. Journal of Machine Learning Research 6 (2005) 1453–1484
- Huang, G.B., Ramesh, M., Berg, T., Learned-Miller, E.: Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst (October 2007)
- Teo, C.H., Vishwanthan, S., Smola, A.J., Le, Q.V.: Bundle methods for regularized risk minimization. J. Mach. Learn. Res. 11 (2010) 311–365
- Bordes, A., Bottou, L., Gallinari, P.: SGD-QN: Careful quasi-newton stochastic gradient descent. Journal of Machine Learning Research 10 (July 2009) 1737–1754
- Heikkilä, M., Pietikäinen, M., Schmid, C.: Description of interest regions with local binary patterns. Pattern Recognition 42(3) (2009) 425–436
- Franc, V., Sonnenburg, S.: LIBOCAS library implementing OCAS solver for training linear svm classifiers from large-scale data. www: http://cmp.felk.cvut.cz/~xfrancv/ocas/ html/index.html (2010)
- 23. Kroon, D.J.: Active shape model (ASM) and active appearance model (AAM). MATLAB Central, www: http://www.mathworks.com/matlabcentral/fileexchange/26706active-shape-model-asm-and-active-appearance-model-aam (February 2010)
- Nordstrøm, M.M., Larsen, M., Sierakowski, J., Stegmann, M.B.: The IMM face database

 an annotated dataset of 240 face images. Technical report, Informatics and Mathematical Modelling, Technical University of Denmark, DTU (may 2004)
- Everingham, M., Sivic, J., Zisserman, A.: Willow project, automatic naming of characters in tv video. MATLAB implementation, www: http://www.robots.ox.ac.uk/~vgg/research/ nface/index.html (2008)