

# Persistent Evidence of Local Image Properties in Generic ConvNets

Ali Sharif Razavian<sup>(✉)</sup>, Hossein Azizpour, Atsuto Maki, Josephine Sullivan,  
Carl Henrik Ek, and Stefan Carlsson

Computer Vision and Active Perception Lab (CVAP),  
School of Computer Science and Communication (CSC),  
Royal Institute of Technology (KTH), 100 44 Stockholm, Sweden  
{razavian, azizpour, atsuto, sullivan, chek, stefanc}@csc.kth.se  
<http://www.csc.kth.se/cvap/>

**Abstract.** Supervised training of a convolutional network for object classification should make explicit any information related to the class of objects and disregard any auxiliary information associated with the capture of the image or the variation within the object class. Does this happen in practice? Although this seems to pertain to the very final layers in the network, if we look at earlier layers we find that this is not the case. In fact, strong spatial information is implicit. This paper addresses this, in particular, exploiting the image representation at the first fully connected layer, i.e. the global image descriptor which has been recently shown to be most effective in a range of visual recognition tasks. We empirically demonstrate evidences for the finding in the contexts of four different tasks: 2d landmark detection, 2d object keypoints prediction, estimation of the RGB values of input image, and recovery of semantic label of each pixel. We base our investigation on a simple framework with ridge regression commonly across these tasks, and show results which all support our insight. Such spatial information can be used for computing correspondence of landmarks to a good accuracy, but should potentially be useful for improving the training of the convolutional nets for classification purposes.

## 1 Introduction

There is at least one alchemy associated with deep convolutional networks (ConvNets). It occurs when  $\sim 100,000$  iterations of stochastic gradient descent (SGD), in tandem with  $\sim 1$  million labelled training images from ImageNet, transform the  $\sim 60$  million randomly initialized weights of a deep ConvNet into the best, by a huge margin, performing known visual image classifier [7, 10, 13, 17, 19, 20]. Alongside this high-level alchemy is another related one w.r.t. the image representations learnt by the fully connected layers of a deep ConvNet [7, 10, 17, 20]. These representations are explicitly trained to retain information relevant to semantic classes. But we show in this paper a striking fact, through various tasks, that these representations also retain *spatial* information, including the location of object parts and keypoints of object.



**Fig. 1.** How many different local image properties can be predicted from a generic ConvNet representation using a *linear* model? In this figure, given the ConvNet representation of an image, we have estimated three different local properties. Namely, semantic segmentation for background (top left) and 20 object classes (bottom left) present in PASCAL VOC dataset, 194 facial landmarks (top right) and RGB reconstruction of the original image (bottom right). One can see that a generic ConvNet representation optimized for ImageNet semantic classification has embedded high level of local information.

The notion of predicting spatial information using a ConvNet itself is not new. Recent studies have introduced several approaches to extract spatial information from an image with deep ConvNet. Some have trained a specialized ConvNet to predict specific spatial information such as body parts and facial landmarks [8, 21, 24]. Others [9, 15, 19] have shown that it is possible to extract spatial correspondences using a generic ConvNet representation. But they consider representations from the ConvNet layers that only describe sub-patches of the whole image. Then in a similar manner to a sliding window approach they have an exhaustive spatial search, in tandem with their patch descriptor, to find the locations.

Unlike those works, we show that a global image representation extracted from the first fully connected layer of a *generic* ConvNet (*i.e.* trained for predicting the semantic classes of ImageNet) is capable of predicting spatial information *without doing an explicit search*. In particular, we show that one can learn a linear regression function (with results ranging from promising to good)

from the representation to spatial properties: 2d facial landmarks, 2d object keypoints, RGB values and class labels of individual pixels, see figure 1. We chose these experiments to highlight the network’s ability to reliably extract spatial information.

Why do we concentrate on the fully connected layers? Prior work has shown that these layers correspond to the most generic and compact image representation and produce the best results, when combined with a simple linear classifier, in a range of visual recognition tasks [2]. Therefore the starting point of this work was to examine what other information, besides visual semantic information, is encoded and easily accessible from these representations.

The results we achieve for the tasks we tackle indicate that the spatial information is implicitly encoded in the ConvNet representation we consider. Remember, the network has not been explicitly encouraged to learn spatial information during the training.

The contributions of the paper are:

- For the first time we systematically demonstrate that spatial information is persistently transferred to the representation in the first fully connected layer of a generic ConvNet (section 3).
- We show that one can learn a linear regression function from the ConvNet representation to both object parts and local image properties. In particular, we demonstrate that it is possible to estimate 2d facial landmarks (section 3.1), 2d object keypoints (section 3.1), RGB values (section 3.1) and pixel level segmentations (section 3.1).
- By using a simple *look-back* method we achieved accurate predictions of facial landmarks on a par with state of the art (section 3.1).
- We qualitatively show examples where semantically meaningful directions in the ConvNet representation space can be learned and exploited to accordingly alter the appearance of a face (section 4).

Before describing our experiments and results in the next section we explain why spatial information can be ever retained and so easily accessed in the first fully-connected layer of a generic ConvNet.

## 2 Flow of Information Through a ConvNet

A generic ConvNet representation extracted from the first fully-connected layer is explicitly trained to retain information relevant to semantic class. The semantic classes in the training data are independent of spatial information and therefore this information, as it is deemed unnecessary to perform the task, should be removed or at least structured in such a manner that it does not conflict with the task.

The weights of a ConvNet’s convolutional layers encode a very large number of compositional patterns of appearance that occur in the training images. Thus, the multiple response maps output by a convolutional layer indicate which appearance patterns occur in different sub-patches (a.k.a. receptive fields) of

a fixed size in the original image. The size of these sub-patches increases as we progress through the convolutional layers. When we come to the first fully-connected layer the network must compress the set of response maps ( $13 \times 13 \times 256$  numbers assuming an AlexNet [13] ConvNet) produced by the final convolutional layer into a mere 4096 numbers. The compression performed seeks to optimize the ability of the network’s classification layer (with potentially some more intermediary fully-connected layers) to produce semantic labels as defined by the ImageNet classification task.

The weights of the first fully connected layer are, in general, not particularly sparse. Therefore the *what* and *where* explicitly encoded in the convolutional layers are aggregated, merged and conflated into the output nodes of this first fully connected layer. At this stage it is impossible to backtrack from these responses to spatial locations in the image. Nevertheless, we show it is possible to predict from this global image descriptor, using linear regressors, the spatial locations of object parts and keypoints and also pixel level descriptors such as colour and semantic class.

### 3 Experiment

We study two families of tasks to explore which spatial information resides in the ConvNet representation:

- Estimate the  $(x, y)$  coordinate of an item in an image.
- Estimate the local property of an image at  $(x, y)$ .

Given the ConvNet representation of an image for the first task, we *i*) estimate the coordinates of facial landmarks in three challenging datasets [3, 14, 18], and *ii*) predict the positions of object keypoints. We use the annotations [4] from the Pascal VOC 2011 dataset as our testbed. While for the second task, given a ConvNet representation, we *i*) predict the RGB values of every pixel in the original image (we use the ImageNet validation set as our test set), and *ii*) predict the semantic segmentation of each pixel in the original image (VOC 2012 Pascal dataset).

#### 3.1 Experimental Setup

In all our experiments we use the same ConvNet. It has the AlexNet architecture [13] and is trained on ImageNet [1] using the reference implementation provided by Caffe [11]. Our image representation then corresponds to the responses of the first fully connected layer of this network because of its compactness and ability to solve a wide range of recognition tasks [2, 7, 10, 17, 20]. We will denote this representation by  $\mathbf{f}$ . Then the only post-processing we perform on  $\mathbf{f}$  is to  $l_2$  normalize it.

For every scalar quantity  $y$  we predict from  $\mathbf{f}$ , we do so with a linear regression model:

$$y \approx \mathbf{w}^T \mathbf{f} + w_0 \tag{1}$$

**Table 1.** Evaluation of facial landmark estimation on three standard face datasets and comparison with baselines and recent state of the art methods. The error measure is the average distance between the predicted location of a landmark and its ground truth location. Each error distance is normalized by the inter-ocular distance.

	Helen [14]	LFPW [3]	IBUG [18]
Dataset Bias	0.501	0.242	0.352
RGB + ridge	0.096	0.074	0.160
STASM [16]	0.111	-	-
CompASM [14]	0.091	-	-
ConvNet + ridge	0.065	0.056	0.096
RCPR [5]	0.065	0.035	-
SDM [22]	0.059	0.035	0.075
ESR [6]	0.059	0.034	0.075
ETR [12]	0.049	0.038	0.064
ConvNet + <i>look-back</i>	0.058	0.049	0.074

We use a ridge regularised linear model because of its simplicity and for the following reason. All the class, pose and semantic information does exist in the original RGB image, as the human vision proves, but it is not easily accessible and especially not through linear models. However, we want to study if all this information is still encoded in the ConvNet representation, but in a much more accessible way and this is demonstrated by the use of a linear model compared to a much more capable prediction algorithm.

There are, of course, numerous ways we can estimate the coefficients  $(\mathbf{w}, w_0)$  from labelled training. Assume that we have labelled training data  $(y_1, \mathbf{f}_1), \dots (y_n, \mathbf{f}_n)$  where each  $y_i \in \mathbb{R}$  and  $\mathbf{f} \in \mathbb{R}^d$  ( $d = 4096$ ). The optimal values for  $(\mathbf{w}, w_0)$  are then found solving this optimization problem


$$(\mathbf{w}^*, w_0^*) = \arg \min_{\mathbf{w}, w_0} \sum_{i=1}^n (y_i - \mathbf{w}^T \mathbf{f}_i - w_0) + \lambda \|\mathbf{w}\|^2 \tag{2}$$

The closed form solution to this optimization problem is easily shown to be:


$$\mathbf{w}^* = (X^T X + \lambda I)^{-1} X^T \mathbf{y} \quad \text{and} \quad w_0^* = \frac{1}{n} \sum_{i=1}^n y_i \tag{3}$$

where

$$X = \begin{pmatrix} \leftarrow \mathbf{f}_1^T \rightarrow \\ \leftarrow \mathbf{f}_2^T \rightarrow \\ \vdots \\ \leftarrow \mathbf{f}_n^T \rightarrow \end{pmatrix} \quad \text{and} \quad \mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} \tag{4}$$



(a) Fig. 2(a) The normalized prediction error for different subsets of the landmarks after *look-back* (see the caption of table 1 for the error measure). The error is shown for three different face datasets. Since the bounding box around the chin is bigger than the bounding box around the other parts, the error for chin is higher than the rest of facial landmarks.



(b) Fig. 2(b) PCK evaluation for keypoint prediction of 20 classes of PASCAL VOC 2011

and it is assumed the columns of  $X$  have been centred. In all our experiments we set the regularization parameter  $\lambda$  with four-fold cross-validation.

**Facial Landmarks.** The first problem we address is the popular task of 2d facial landmark detection. Facial landmark detection is interesting since a large body of work has been applied to it. To train our landmark estimation model, for each landmark we estimate two separate linear regression functions, one for the  $x$ -coordinate of the landmark and one for the  $y$ -coordinate. Therefore we estimate the  $(x, y)$  coordinates of all the  $L$  landmarks from the image’s ConvNet representation,  $\mathbf{f}$ , with

$$\hat{\mathbf{x}} = W_{\text{landmarks}} \mathbf{f} + \mathbf{w}_{\text{landmarks},0} \tag{5}$$

where  $W_{\text{landmarks}} \in \mathbb{R}^{2L \times d}$  and  $\mathbf{w}_{\text{landmarks},0} \in \mathbb{R}^{2L}$ . Remember each row of  $W_{\text{landmarks}}$  is learnt independently via the ridge regression solution of equation (3). For the rest of the tasks explored in this section we use a similar formulation to the one just described so we will not introduce new notation to describe them.

**Table 2.** Quantitative evaluation of our keypoint estimation for general objects on VOC11. The performance measure is the average PCK score with  $\alpha = 0.1$ .

	airplane	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean
SIFT	17.9	16.5	15.3	15.6	25.7	21.7	22.0	12.6	11.3	7.6	6.5	12.5	18.3	15.1	15.9	21.3	14.7	15.1	9.2	19.9	15.7
SIFT+prior	33.5	36.9	22.7	23.1	44.0	42.6	39.3	22.1	18.5	23.5	11.2	20.6	32.2	33.9	26.7	30.6	25.7	26.5	21.9	32.4	28.4
ConvNet + ridge	21.3	25.1	22.7	16.4	47.3	27.2	29.9	25.4	19.7	26.3	22.0	27.1	25.5	21.8	33.8	41.0	28.2	23.0	23.9	47.3	27.8
Conv5 + sliding window [15]	38.5	37.6	29.6	25.3	54.5	52.1	28.6	31.5	8.9	30.5	24.1	23.7	35.8	29.9	39.3	38.2	30.5	24.5	41.5	42.0	33.3






**Fig. 3.** The landmarks predicted by our linear regressors for eight different images from the Helen Dataset. The leftmost image in each triplet shows the ground truth. The middle image shows the landmarks predicted by the linear regression functions from the ConvNet representation of the whole image to landmark coordinates. In the middle image the bounding boxes, defined by the initial predictions for the landmarks, used by *look-back* method are also shown. The rightmost image shows the final landmark predictions made by the *look-back* method.

Table 1 details the average errors, of our approaches and other methods, in the predicted location of the landmarks on three standard datasets: Helen [14], LFPW [3] and IBUG [18]. The reported errors are normalized by the distance between two eyes in the image according to the standard practice in the field [12]. The table reports the performance of both the baseline predictors of linear ridge regression from RGB and a random predictor and recent high performing systems [5, 6, 12, 22] which generally involve learning a complicated non-linear function from RGB to the landmarks. Our predictor, ConvNet+ridge, produces a significantly better estimate than the baselines and its performance is comparable with state-of-the-art methods specifically designed to solve this problem. Our result indicates that the locations of landmarks can be reliably extracted from the ConvNet representation.

ConvNet+ridge inherently loses around  $\pm 10$  pixel accuracy due to the pooling and strides in the first and second convolutional layers of the ConvNet. However, we can overcome this limitation in a simple manner which we term the *look-back trick*. We partition the landmarks into different subsets (such as



**Fig. 4.** RGB information linearly predicted from the ConvNet representation. For each pixel we train 3 independent linear regressors to predict the pixel’s RGB value from the image’s global ConvNet representation. We used the first 49K images from ImageNet’s cross validation set for training and visualized the result for the last 1k images. Shown above are the results for 25 random images (left) taken from the test set and their reconstructions (right).

chin, left eye and left eyebrow, right eye and right eyebrow, and then mouth and nose), figure 3 shows some examples. We let the predicted position of each set of landmarks, using equation 5, define a square bounding box containing them with some margin based on the maximum prediction error for the landmarks in the training set. We then extract the ConvNet representation for this sub-image and use linear regression, as before, to estimate the coordinates of the landmarks in the bounding-box. This simple trick significantly boosts the accuracy of the predictions, and allows us to outperform all the s.o.a. methods except for the recent work of [12]. The more sets we have in the partition the better results we get. We used six sets for each dataset. See figure 3 for qualitative examples of the result of this method on sample images from the Helen dataset. Figure 2a also shows the prediction errors for different parts after look-up for three face datasets.

**Object Keypoints.** In our next task we predict the location of object keypoints. These keypoints exhibit more variation in their spatial location than facial landmarks. We use the keypoint annotations provided by [4] for 20 classes of PASCAL VOC 2011. We make our predictions using exactly the same basic approach as for facial landmarks. The classes of PASCAL task include many deformable objects (dog, cat, human, *etc.*) and objects which have high intra-class variation (bottle, plant, *etc.*) which makes the problem of key point



**Table 3.** Evaluation of Semantic Segmentation on the validation set of VOC12 measured in mean Average Precision (mAP)

	background	airplane	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mean
ConvNet	79.12	16.01	0.02	12.93	9.26	13.69	37.29	33.75	40.01	0.01	8.62	12.24	30.89	9.43	24.94	44.03	6.22	18.77	1.64	25.33	11.05	20.73


detection extremely difficult. In order to model these keypoints a separate set of regressors is learnt for each object.

Table 2 reports the accuracy of our results for keypoint prediction, together with those achieved by other methods [15]. The accuracy is measured using mean PCK [23]. A keypoint is considered to be correctly estimated if the prediction’s Euclidean distance from the ground truth position is  $\alpha \in [0, 1]$  times the maximum of the bounding box width and height. Our simple approach outperforms SIFT by a huge margin on localizing landmarks and is only slightly below the performance of SIFT+prior [15]. Figure 2b shows the plot of PCK vs  $\alpha$  for 20 classes.

**RGB Reconstruction.** The results from the previous tasks show that our ConvNet representation does encode some levels of spatial information. The natural question is then *what does it actually remove?* To investigate this we try to evaluate if it is possible to invert the ConvNet mapping. First, we try to estimate the RGB values of the original input image from our ConvNet representation. For this, again, we simply learn  $3 \times n_p$  linear regressors where  $n_p$  is the number of pixels in the image. In other words we learn an independent regressor for each pixel and each colour channel.

We use ImageNet as our testbed. We used the first 49k images from ImageNet’s validation set for training and the last 1k images for testing. We resized each image to  $46 \times 46 \times 3$  and trained 6348 independent linear regressors. Some examples of the resulting RGB reconstruction are illustrated in figure 4. The mean absolute error of image reconstruction is 0.12. It is rather surprising that RGB values of an image can be extracted with this degree of accuracy from the ConvNet representation.

**Semantic Segmentation.** We applied the same framework which we employed for RGB reconstruction further to recover semantic labels of each pixel instead of its RGB values. The procedure is as follows: we resized each semantic segmentation map of VOC 2012 segmentation task down to a  $30 \times 30 \times 21$  image. We train a separate linear regressor to predict whether the pixel at position  $(x, y)$  belongs to class  $c$  or not encoded as 1 and 0. We have  $x \in \{1, 2, \dots, 30\}$  and similarly for  $y$  and  $c \in \{1, 2, \dots, 21\}$ . Therefore a total of 18900 linear regressors are trained with ridge regression. Solving a classification problem via regression is not ideal. But the qualitative results shown in figure 5 are visually pleasing. They show the semantic segmentations produced by our approach for some images from the VOC12 validation set.



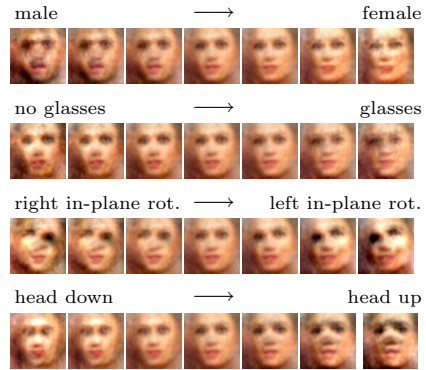
**Fig. 5.** Semantic segmentation results for images from PASCAL VOC. For each block of pictures, the top left hand picture is the original image and the one directly to its right is the *probability* map for the background class. The brighter the pixel the higher the probability. The bottom set of smaller images in the block display, in the same manner, the probabilities for the 20 classes of PASCAL VOC 2011. The probability masks are computed independently of one another though the scaling of the intensities in the displayed masks is consistent across all the masks. The learning is based on linear regression, the details of which are found in the main text.

After we apply the linear regressor for each class to each pixel, we get 21 responses. We then turn these responses into a single prediction using another linear model. We multiply the response vector by a matrix  $M \in \mathbb{R}^{21 \times 21}$  and then choose the class which corresponds to the highest response in the output vector. Ideally  $M$  should model the relations between different class responses at a single pixel. We learn  $M$ , once again with ridge regression, and during training it tries to return a binary vector of length 21 with only one non-zero entry.

As our segmentation masks only have size  $30 \times 30$  we resize the them back to their original size. We used the VOC12 training-set as the training data and augmented this set tenfold to get a better estimate and reported the result on the cross-validation set. Quantitative results for our segmentations are given in table 3. Although this result itself is not as good as s.o.a. on semantic

segmentation task (mean average precision of 20.7 compared to 47.5 of s.o.a. method), it is intriguing to see that the global ConvNet representation contains this level of information. It is easy to envisage that such a segmentation could be incorporated into an object classifier or detector.

**Fig. 6.** Semantically altering a face using its generic ConvNet representation and semantically meaningful directions in the representation space. Given an image of a face and its ConvNet representation each row above shows the effect of altering the face’s representation by moving in one direction of the representation space and then regressing from the resulting ConvNet representation back to its RGB representation. Each direction was learned from labeled training data and corresponds to a specific semantic concept: gender, glasses/no-glasses, head angle and head tilt. For gender we can see that the left-most image which corresponds to male has dark patches corresponding to a beard while the right-most image is clearly female. The glasses/no-glasses clearly alters the region around the eyes. The last two rows show variations caused by changes in head pose. Both the head angle and the tilt (last row) are clearly visible.



### 4 Semantic Directions in Representation Space

The ConvNet is trained with one objective in mind, to learn a representation where every pair of classes is linearly separable. The representation space is thus carved into different volumes corresponding to the different classes. The results of the paper so far show each class volume retains significant intra-class variations. In this section we make a first step towards understanding how these variations are structured. To proceed we learn a separate linear regressor from the representation to each of the following variates for the LFW data-set *gender*, *have-glasses*, and *pose*. Each linear regressor specifies a direction in the representation space along which a semantic concept varies.

What happens with the images if we alter the representation along this direction? Can we change the gender or add glasses or continuously change the pose of the face in the image? We can achieve this if we extrapolate along an identified direction and then regress back to the RGB image as described in section 3.1. In more detail: the ConvNet representation,  $\mathbf{f}$ , of a face can be written in terms of

its projection onto a semantic direction, such as gender  $\mathbf{w}_{\text{gender}}$ , found via linear regression and its component orthogonal to  $\mathbf{w}_{\text{gender}}$

$$\mathbf{f} = (\mathbf{w}_{\text{gender}}^T \mathbf{f}) \mathbf{w}_{\text{gender}} + \tilde{\mathbf{f}} \quad (6)$$

Then we can create a new ConvNet representation where the gender attribute of the face has been altered but not the other factors in the following simplistic manner.

$$\mathbf{f}' = \tilde{\mathbf{f}} + \lambda \mathbf{w}_{\text{gender}} \quad (7)$$

with  $\lambda \in [\lambda_{\min}, \lambda_{\max}]$ . We then regress from  $\mathbf{f}'$  back to the RGB image to visualize the result of the alteration, see figure 6 for some sample results.

What do the results of our small scale experiment convey? We can see that altering the pose direction in the representation does correspond well to the actual image transformation and that changing the gender corresponds to an altering of the face’s color composition. Similarly glasses/no-glasses alters the appearance of the region around the eyes. However, it is easy to read too much into the experiments as it is severely limited by the linear structure of the regressors. And the fear of hallucinating experimental evidence for the elephant in the room - the concept of *disentanglement* - means we will leave our speculations to these comments. However, as such a simple approach is capable of finding some structures it would be interesting to investigate if the representation factorizes the variations according to semantic factors.

## 5 Conclusion

In this paper we have shown that a generic ConvNet representation from the first fully connected layer retains significant spatial information. We demonstrated this fact by solving four different tasks, that require local spatial information, using the simple common framework of linear regression from our ConvNet representation. These tasks are *i)* 2d facial landmark prediction, *ii)* 2d object key-points prediction, *iii)* estimation of the RGB values of the original input image, and *iv)* semantic segmentation, *i.e.* recovering the semantic label of each pixel. The results demonstrated throughout all these tasks, using diverse datasets, show spatial information is implicitly encoded in the ConvNet representation and can be easily accessed. This result is surprising because the employed network was not explicitly trained to keep spatial information and also the first fully connected layer is a global image descriptor which aggregates and conflates appearance features, extracted from the convolutional layers, from all spatial locations in the image.

**Acknowledgments.** We would like to gratefully acknowledge the support of NVIDIA for the donation of multiple GPU cards for this research.

## References

1. Imagenet large scale visual recognition challenge (2013). <http://www.image-net.org/challenges/LSVRC/2013/>
2. Azizpour, H., Sharif Razavian, A., Sullivan, J., Maki, A., Carlsson, S.: From generic to specific deep representations for visual recognition (2014). [arXiv:1406.5774](https://arxiv.org/abs/1406.5774) [cs.CV]
3. Belhumeur, P.N., Jacobs, D.W., Kriegman, D.J., Kumar, N.: Localizing parts of faces using a consensus of exemplars. In: CVPR, pp. 545–552 (2011)
4. Bourdev, L., Malik, J.: Poselets: body part detectors trained using 3d human pose annotations. In: ICCV (2009)
5. Burgos-Artizzu, X.P., Perona, P., Dollár, P.: Robust face landmark estimation under occlusion. In: ICCV (2013)
6. Cao, X., Wei, Y., Wen, F., Sun, J.: Face alignment by explicit shape regression. In: CVPR, pp. 2887–2894 (2012)
7. Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., Darrell, T.: Decaf: a deep convolutional activation feature for generic visual recognition. In: ICML (2014)
8. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. In: NIPS (2014)
9. Fischer, P., Dosovitskiy, A., Brox, T.: Descriptor matching with convolutional neural networks: a comparison to sift (2014). [arXiv:1405.5769v1](https://arxiv.org/abs/1405.5769v1) [cs.CV]
10. Girshick, R.B., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: CVPR (2014)
11. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R.B., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding (2014)
12. Kazemi, V., Sullivan, J.: One millisecond face alignment with an ensemble of regression trees. In: CVPR, pp. 1867–1874 (2014)
13. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)
14. Le, V., Brandt, J., Lin, Z., Bourdev, L., Huang, T.S.: Interactive facial feature localization. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012, Part III. LNCS, vol. 7574, pp. 679–692. Springer, Heidelberg (2012)
15. Long, J., Zhang, N., Darrell, T.: Do convnets learn correspondence? (2014). [arXiv:1411.1091](https://arxiv.org/abs/1411.1091) [cs.CV]
16. Milborrow, S., Nicolls, F.: Locating facial features with an extended active shape model. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part IV. LNCS, vol. 5305, pp. 504–513. Springer, Heidelberg (2008)
17. Oquab, M., Bottou, L., Laptev, I., Sivic, J.: Learning and transferring mid-level image representations using convolutional neural networks. In: CVPR (2014)
18. Sagonas, C., Tzimiropoulos, G., Zafeiriou, S., Pantic, M.: A semi-automatic methodology for facial landmark annotation. In: CVPR Workshops, pp. 896–903 (2013)
19. Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., LeCun, Y.: Overfeat: integrated recognition, localization and detection using convolutional networks. In: ICLR (2014)
20. Sharif Razavian, A., Azizpour, H., Sullivan, J., Carlsson, S.: CNN features off-the-shelf: an astounding baseline for visual recognition. In: CVPR workshop of DeepVision (2014)

21. Toshev, A., Szegedy, C.: Deeppose: human pose estimation via deep neural networks. In: CVPR (2014)
22. Xiong, X., De la Torre, F.: Supervised descent method and its applications to face alignment. In: CVPR (2013)
23. Yang, Y., Ramanan, D.: Articulated human detection with flexible mixtures of parts. PAMI **35**(12), 2878–2890 (2013)
24. Zhang, N., Paluri, M., Ranzato, M., Darrell, T., Bourdev, L.: Panda: pose aligned networks for deep attribute modeling. In: CVPR (2014)