

## Face Recognition (Recunoasterea Faciala)

FR := Aplicatie biometrica pentru stabilirea identitatii absolute a unei persoane.

Repository (articole/algoritmi, cod sursa, seturi de date):

### FACE RECOGNITION HOMEPAGE

<http://www.face-rec.org/algorithms/>

<http://www.face-rec.org/databases/>

<http://www.face-rec.org/source-codes/>

Metode traditionale machine learning: [PCA](#) | [ICA](#) | [LDA](#) | [EP](#) | [EBGM](#) | [Kernel Methods](#) | [Trace Transform](#) | [AAM](#) | [3-D Morphable Model](#) | [3-D Face Recognition](#) | [Bayesian Framework](#) | [SVM](#) | [HMM](#) | [Boosting & Ensemble](#) | [Algorithms Comparisons](#)

Implementari:

Philipp Wagner, **FaceRecognizer - Face Recognition with OpenCV**  
(Eigenfaces, Fisherfaces, LBP)

<http://docs.opencv.org/trunk/modules/contrib/doc/facerec/>

[http://docs.opencv.org/trunk/modules/contrib/doc/facerec/facerec\\_tutorial.html](http://docs.opencv.org/trunk/modules/contrib/doc/facerec/facerec_tutorial.html)

Ver. MTALAB (Eigenfaces, Fisherfaces):

[https://github.com/bytefish/facerecognition\\_guide/raw/master/facerec\\_octave.pdf](https://github.com/bytefish/facerecognition_guide/raw/master/facerec_octave.pdf)

PCA / Eigenfaces (MATLAB): <http://www.cs.ait.ac.th/~mdailey/matlab/>

PCA on the FERET Dataset (EigenFaces): <http://www.face-rec.org/source-codes/pca.zip>

## Face recognition using Eigenfaces

Based on the Principal component analysis (PCA) theory:

[http://en.wikipedia.org/wiki/Principal\\_component\\_analysis](http://en.wikipedia.org/wiki/Principal_component_analysis)

### Eigenfaces

<http://en.wikipedia.org/wiki/Eigenface>

**Eigenfaces** is the name given to a set of [eigenvectors](#) ([https://en.wikipedia.org/wiki/Eigenvalues\\_and\\_eigenvectors](https://en.wikipedia.org/wiki/Eigenvalues_and_eigenvectors)) when they are used in the [computer vision](#) problem of human [face recognition](#).<sup>[1]</sup> The approach of using eigenfaces for [recognition](#) was developed by Sirovich and Kirby (1987) and used by [Matthew Turk](#) and [Alex Pentland](#) in face classification.<sup>[2]</sup> The [eigenvectors](#) are derived from the [covariance matrix](#) of the [probability distribution](#) over the high-[dimensional vector space](#) of face images. The eigenfaces themselves form a basis set of all images used to construct the covariance matrix. This produces dimension reduction by allowing the smaller set of basis images to represent the original training images. Classification can be achieved by comparing how faces are represented by the basis set.

### History

The Eigenface approach began with a **search for a low-dimensional representation of face images**. Sirovich and Kirby (1987) showed that [Principal Component Analysis](#) could be used on a collection of face images to form a **set of basis features**. These basis images known as **Eigenpictures**, could be **linearly combined** to **reconstruct images in the original training set**. If the **training set** consists of  $M$  **images**, principal component analysis could form a **basis set of  $N$  images**, where  $N < M$ . The reconstruction error is reduced by increasing the number of eigenpictures, however the number needed is always less than  $M$ . For example if you need to generate a number of  $N$  eigenfaces for a training set of  $M$  face images, **you can say**

**that each face image can be made up of "proportions" of all this K "features" or eigenfaces :  $\text{Face image}_1 = (23\% \text{ of } E_1) + (2\% \text{ of } E_2) + (51\% \text{ of } E_3) + \dots + (1\% \text{ of } E_n)$ .**

In 1991 M. Turk and A. Pentland expanded these results and presented the **Eigenface method of face recognition**.<sup>[3]</sup> As well as designing a system for automated face recognition using eigenfaces, they showed a way of calculating the [eigenvectors](#) of a [covariance matrix](#) in such a way as to make it possible for computers at that time to perform eigen-decomposition on a large number of face images. Face images usually occupy a high dimensional space and conventional principal component analysis was intractable on such data sets. Turk and Pentlands paper demonstrated ways to extract the eigenvectors based on matrices sized by the number of images rather than the number of pixels.

Once established, the eigenface method was expanded to include methods of preprocessing to improve accuracy.<sup>[4]</sup> Multiple manifold approaches were also used to build sets of eigenfaces for different subjects<sup>[5][6]</sup> and different features, such as the eyes.<sup>[7]</sup>

### **Eigenface generation**

A **set of eigenfaces** can be generated by performing a mathematical process called [principal component analysis](#) (PCA) on a large set of images depicting different human faces. Informally, **eigenfaces can be considered a set of "standardized face ingredients"**, derived from [statistical analysis](#) of many pictures of faces. Any human face can be considered to be a combination of these standard faces. **For example, one's face might be composed of the average face plus 10% from eigenface 1, 55% from eigenface 2, and even -3% from eigenface 3.** Remarkably, it does not take many eigenfaces combined together to achieve a fair approximation of most faces. Also, because a person's face is not recorded by a [digital photograph](#), but instead as just a list of values (one value for each eigenface in the database used), **much less space is taken for each person's face.**


The eigenfaces that are created will appear as light and dark areas that are arranged in a **specific pattern**. This pattern is how different features of a face are singled out to be evaluated and scored. There will be a **pattern** to evaluate [symmetry](#), if there is any style of **facial hair**, **where the hairline is**, or **evaluate the size of the nose or mouth**. Other eigenfaces have patterns that are less simple to identify, and the image of the eigenface may look very little like a face.

The technique used in creating eigenfaces and using them for recognition is also used outside of facial recognition. **This technique is also used for** [handwriting analysis](#), [lip reading](#), [voice recognition](#), [sign language](#)/hand [gestures](#) interpretation and [medical imaging](#) analysis. Therefore, some do not use the term eigenface, but prefer to use 'eigenimage'.

## **Practical implementation**

To create a set of eigenfaces, one must:

1. **Prepare a training set of face images.** The pictures constituting the training set should have been taken under **the same lighting conditions, and must be normalized to have the eyes and mouths aligned across all images**. They must also be all resampled to a **common [pixel](#) resolution** ( $N \times N$ ).
2. Each image is treated as **one vector**, simply by [concatenating](#) the rows of pixels in the original image, resulting in a **single row with  $N \times N$  elements** ( $\Gamma_i$ ). For this implementation, it is assumed that all images of the training set are stored in a single [matrix](#), where **each column of the matrix is an image**.



3. The **average image**  $\Psi$  has to be calculated:

$$\Psi = \frac{1}{M} \sum_{i=1}^M \Gamma_i$$

4. The **average image**  $\Psi$  has to be subtracted from each original image ( $\Gamma_i$ ).

$$\Phi_i = \Gamma_i - \Psi$$

5. Calculate the **eigenvectors** ( $u_i$ ) and eigenvalues of the [covariance matrix](#)  $C$  of  $\mathbf{A}=[\Phi_1 \Phi_2 \dots \Phi_M]$ . Each eigenvector has the same dimensionality (number of components) as the original images, and thus can itself be seen as an image. The **eigenvectors** of this covariance matrix are **therefore called eigenfaces**. *They are the directions in which the images differ from the mean image.* Usually this will be a computationally expensive step (if at all possible), but the practical applicability of eigenfaces stems from the possibility to compute the eigenvectors of  $C$  efficiently, without ever computing  $C$  explicitly, as detailed below.

$$C = \frac{1}{M} \sum_{n=1}^M \Phi_n \Phi_n^T = AA^T \quad (N^2 \times N^2 \text{ matrix})$$

$$\text{where } A = [\Phi_1 \Phi_2 \dots \Phi_M] \quad (N^2 \times M \text{ matrix})$$

## 6. Computing the eigenvectors

Performing PCA directly on the covariance matrix of the images is often computationally infeasible. If small, say  $100 \times 100$ , greyscale images are used, each image is a point in a 10,000-dimensional space and the covariance matrix **C** is a matrix of  $10,000 \times 10,000 = 10^8$  elements. However the [rank](#) of the covariance matrix is limited by the number of training examples: if there are  $N$  training examples, there will be at most  $N - 1$  eigenvectors with non-zero eigenvalues. **If the number of training examples is smaller than the dimensionality of the images**, the principal components can be computed more easily as follows:

Let  $A$  be the matrix of preprocessed training examples, where each column contains one mean-subtracted image. The **covariance matrix** can then be computed as  $C = AA^T$  and the eigenvector decomposition of  $C$  is given by:

$$Cu_i = AA^T u_i = \mu_i u_i$$

However  $AA^T$  is a large matrix ( $N^2 \times N^2$ ), and if instead we take the eigenvalue decomposition of:

$$A^T A v_i = \mu_i v_i$$

then we notice that by pre-multiplying both sides of the equation with  $A$ , we obtain:

$$AA^T A v_i = C A v_i = \mu_i A v_i$$

$$u_i = A v_i \quad \Rightarrow \quad C u_i = \mu_i u_i$$

Meaning that, if  $v_i$  is an eigenvector of  $A^T A$ , then  $u_i = A v_i$  is an eigenvector of  $C$ .

The **M eigenvalues** (eigenvectors) of  $A^T A$  correspond to the **M largest eigenvalues** (eigenvectors) of  $C = AA^T$

$\Rightarrow$  Compute the **M** best eigenvectors of  $AA^T$ :  $u_i = A v_i$

**Note1:** If we have a training set of **M = 300** images of  $100 \times 100$  ( $N \times N$ ) pixels, the matrix  $A^T A$  is a  $300 \times 300$  matrix (it can have up to **M = 300** eigenvectors /

eigenvalues), which is much more manageable than the ( $N^2 \times N^2 = 10,000 \times 10,000$ ) covariance matrix  $C$  (it can have up to  $N^2=10,000$  **eigenvectors** / eigenvalues). Notice however that the resulting vectors  $u_i$  are not normalised; if normalisation is required it should be applied as an extra step such as  $\| u_i \| = 1$ .

**Note 2:** Each eigenvector represents a **direction** in the ( $N \times N$ ) image space

7. From the  $M$  best eigenvectors keep only  $K$  eigenvectors (corresponding to the  $K$  largest **eigenvalues**  $\approx$  **Eigenfaces** / **Eigenpictures**)

These **eigenfaces** can now be used to represent both existing and new faces: we can project a new (mean-subtracted) image on the eigenfaces and thereby record how that new face differs from the mean face. **The eigenvalues associated with each eigenface represent how much the images in the training set vary from the mean image in that direction.** We lose information by projecting the image on a subset of the eigenvectors, but we minimize this loss by keeping those eigenfaces with the largest eigenvalues. For instance, if we are working with a 100 x 100 image, then we will obtain 10,000 eigenvectors. In practical applications, most faces can typically be identified using a projection on between **100 and 150 eigenfaces**, so that most of the 10,000 eigenvectors can be discarded.

### Representing faces onto this basis

Each face (minus the mean)  $\Phi_i$  in the training set can be represented as a linear combination of the best  $K$  eigenvectors:

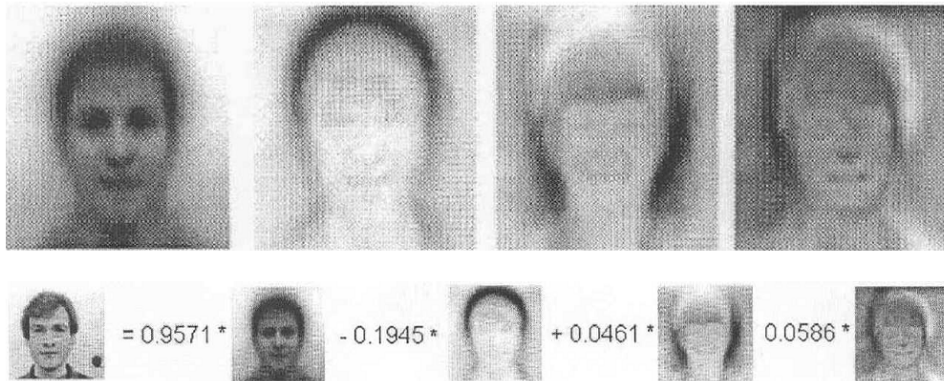
$$\hat{\Phi}_i - mean = \sum_{j=1}^K w_j u_j, \quad (w_j = u_j^T \Phi_i)$$

$u_j$  - eigenfaces / eigenvectors (vectors of size  $N^2$ )

Each normalized training face  $\Phi_i$  is represented in this basis by a

$$\Omega_i = \begin{bmatrix} w_1^i \\ w_2^i \\ \dots \\ w_K^i \end{bmatrix}, \quad i = 1, 2, \dots, M$$

vector:



### Use in facial recognition

Facial recognition was the source of motivation behind the creation of eigenfaces. For this use, eigenfaces have advantages over other techniques available, such as the system's speed and efficiency. As eigenface is primarily a dimension reduction method, a system can represent many subjects with a relatively small set of data. As a face recognition system it is also **fairly invariant to large reductions in image sizing**, however it begins to **fail considerably when the variation between the seen images and probe image is large**.

To recognize faces, gallery images, those seen by the system, are saved as collections of weights describing the contribution each eigenface has to that image. When a **new face is presented to the system for classification**, its **own weights are found by projecting the image onto the collection of eigenfaces**. This provides a **set of weights describing the probe face**. These weights are then classified against all weights in the gallery set to find the closest match. A **nearest neighbour method** is a simple approach for finding the [Euclidean Distance](#) between two vectors, where the minimum can be classified as the closest subject. ([Turk & Pentland 1991](#), p. 590)

The **weights of each gallery image only convey information describing that image, not that subject**. An image of one subject under **frontal lighting** may have very different weights to those of the same subject under strong **left lighting**. This limits the application of such a system. Experiments in the original Eigenface paper



presented the following results: an average of 96% with light variation, 85% with orientation variation, and 64% with size variation. ([Turk & Pentland 1991](#), p. 590).

### **Other methods**

Various extensions have been made to the eigenface method such [eigenfeatures](#). This method combines [facial metrics](#) (measuring distance between facial features) with the eigenface representation. Another method similar to the eigenface technique is '[fisherfaces](#)' which uses (LDA) [Linear discriminant analysis](#).<sup>[8]</sup> This method for facial recognition **is less sensitive to variation in lighting and pose** of the face than using eigenfaces. Fisherface utilizes labeled data to retain more of the class specific information during the dimension reduction stage.

A further alternative to eigenfaces and fisherfaces is the [active appearance model](#). This approach use an [Active Shape Model](#) to describe the outline of a face. By collecting many face outlines, [Principal Component Analysis](#) can be used to form a basis set of models which, encapsulate the variation of different faces.

[http://www.vision.jhu.edu/teaching/vision08/Handouts/case\\_study\\_pca1.pdf](http://www.vision.jhu.edu/teaching/vision08/Handouts/case_study_pca1.pdf)

<http://www.face-rec.org/algorithms/PCA/jcn.pdf>

<http://www.cs.columbia.edu/~belhumeur/journal/fisherface-pami97.pdf>

## Deep learning based face recognition

Colectie de referinte la metode reprezentative

<https://machinelearningmastery.com/introduction-to-deep-learning-for-face-recognition/>

Articol de survey: **Deep Face Recognition: A Survey**

<https://arxiv.org/pdf/1804.06655.pdf>

Schroff, F., Kalenichenko, D., & Philbin, J. (2015). **Facenet: A unified embedding for face recognition and clustering**. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 815-823). (6313 citari – Google Scholar)

[https://www.cv-foundation.org/openaccess/content\\_cvpr\\_2015/html/Schroff\\_FaceNet\\_A\\_Unified\\_2015\\_CVPR\\_paper.html](https://www.cv-foundation.org/openaccess/content_cvpr_2015/html/Schroff_FaceNet_A_Unified_2015_CVPR_paper.html)

<https://medium.com/@ageitgey/machine-learning-is-fun-part-4-modern-face-recognition-with-deep-learning-c3cffc121d78>



Figure 2. **Model structure.** Our network consists of a batch input layer and a deep CNN followed by  $L_2$  normalization, which results in the face embedding. This is followed by the triplet loss during training.

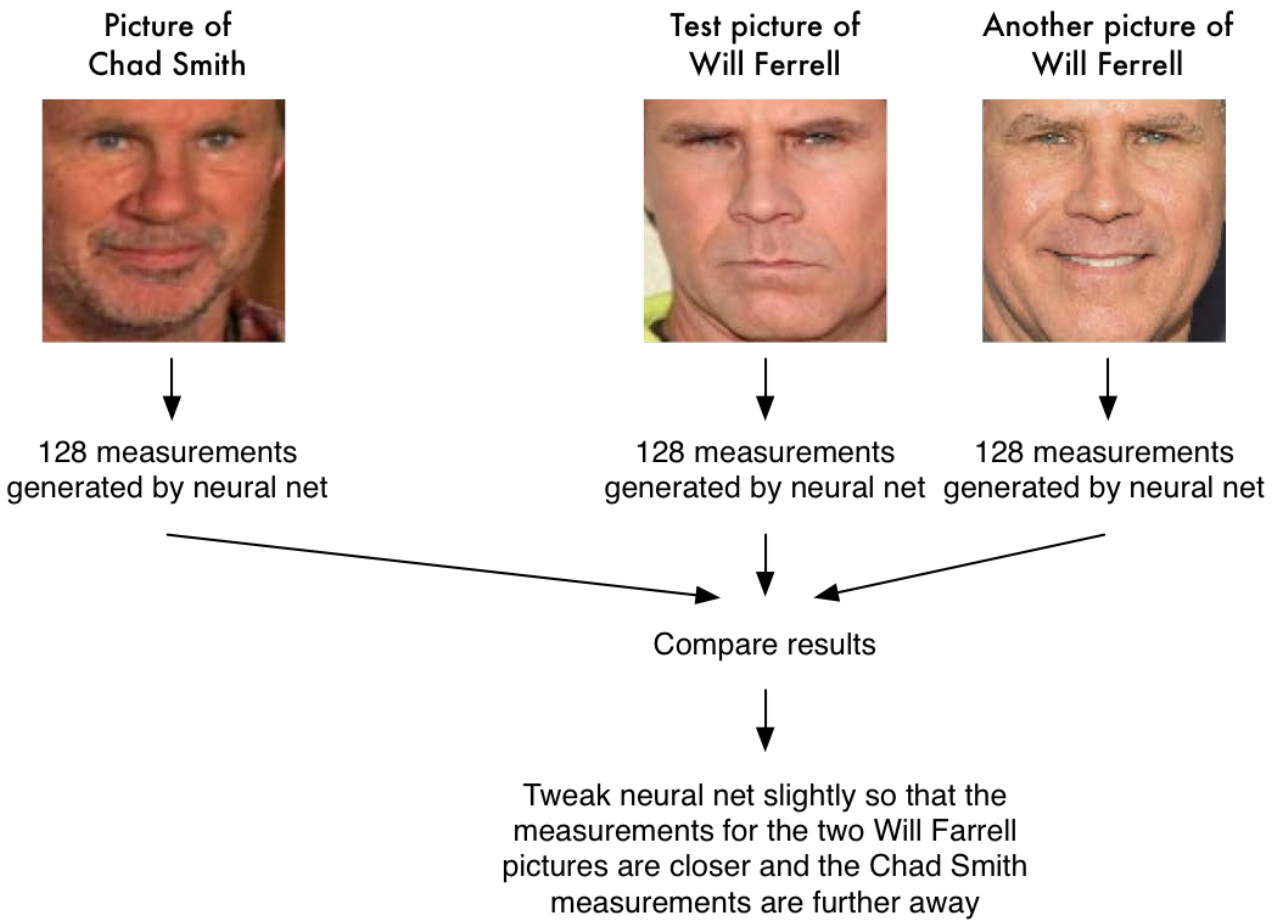


Figure 3. The **Triplet Loss** minimizes the distance between an *anchor* and a *positive*, both of which have the same identity, and maximizes the distance between the *anchor* and a *negative* of a different identity.

$$\sum_i^N \left[ \|f(x_i^a) - f(x_i^p)\|_2^2 - \|f(x_i^a) - f(x_i^n)\|_2^2 + \alpha \right]$$

[Schroff2015]

# A single 'triplet' training step:



Input Image

128 Measurements Generated from Image

0.097496084868908	0.045223236083984	-0.1281466782093	0.032084941864014
0.12529824674129	0.060309179127216	0.17521631717682	0.020976085215807
0.030809439718723	-0.01981477253139	0.10801389068365	-0.00052163278451189
0.036050599068403	0.065554238855839	0.0731306001544	-0.1318951100111
-0.097486883401871	0.1226262897253	-0.029626874253154	-0.0059557510539889
-0.0066401711665094	0.036750309169292	-0.15958009660244	0.043374512344599
-0.14131525158882	0.14114324748516	-0.031351584941149	-0.053343612700701
-0.048540540039539	-0.061901587992907	-0.15042643249035	0.078198105096817
-0.12567175924778	-0.10568545013666	-0.12728653848171	-0.076289616525173
-0.061418771743774	-0.074287034571171	-0.065365232527256	0.12369467318058
0.046741496771574	0.0061761881224811	0.14746543765068	0.056418422609568
-0.12113650143147	-0.21055991947651	0.0041091227903962	0.089727647602558
0.061606746166945	0.11345765739679	0.021352224051952	-0.0085843298584223
0.061989940702915	0.19372203946114	-0.086726233363152	-0.022388197481632
0.10904195904732	0.084853030741215	0.09463594855878	0.020696049556136
-0.019414527341723	0.0064811296761036	0.21180312335491	-0.050584398210049
0.15245945751667	-0.16582328081131	-0.035577941685915	-0.072376452386379
-0.12216668576002	-0.007277755558491	-0.036901291459799	-0.034365277737379
0.083934605121613	-0.059730969369411	-0.070026844739914	-0.045013956725597
0.087945111095905	0.11478432267904	-0.089621491730213	-0.013955107890069
-0.021407851949334	0.14841195940971	0.078333757817745	-0.17898085713387
-0.018298890441656	0.049525424838066	0.13227833807468	-0.072600327432156
-0.011014151386917	-0.051016297191381	-0.14132921397686	0.0050511928275228
0.0093679334968328	-0.062812767922878	-0.13407498598099	-0.014829395338893
0.058139257133007	0.0048638740554452	-0.039491076022387	-0.043765489012003
-0.024210374802351	-0.11443792283535	0.071997955441475	-0.012062266469002
-0.057223934680223	0.014683869667351	0.05228154733777	0.012774495407939
0.023535015061498	-0.081752359867096	-0.031709920614958	0.069833360612392
-0.0098039731383324	0.037022035568953	0.11009479314089	0.11638788878918
0.020220354199409	0.12788131833076	0.18632389605045	-0.015336792916059
0.0040337680839002	-0.094398014247417	-0.11768248677254	0.10281457751989
0.051597066223621	-0.10034311562777	-0.040977258235216	-0.082041338086128

## 2 archituri CNN ([Schroff2015])

layer	size-in	size-out	kernel	param	FLPS
conv1	220×220×3	110×110×64	7×7×3, 2	9K	115M
pool1	110×110×64	55×55×64	3×3×64, 2	0	
rnorm1	55×55×64	55×55×64		0	
conv2a	55×55×64	55×55×64	1×1×64, 1	4K	13M
conv2	55×55×64	55×55×192	3×3×64, 1	111K	335M
rnorm2	55×55×192	55×55×192		0	
pool2	55×55×192	28×28×192	3×3×192, 2	0	
conv3a	28×28×192	28×28×192	1×1×192, 1	37K	29M
conv3	28×28×192	28×28×384	3×3×192, 1	664K	521M
pool3	28×28×384	14×14×384	3×3×384, 2	0	
conv4a	14×14×384	14×14×384	1×1×384, 1	148K	29M
conv4	14×14×384	14×14×256	3×3×384, 1	885K	173M
conv5a	14×14×256	14×14×256	1×1×256, 1	66K	13M
conv5	14×14×256	14×14×256	3×3×256, 1	590K	116M
conv6a	14×14×256	14×14×256	1×1×256, 1	66K	13M
conv6	14×14×256	14×14×256	3×3×256, 1	590K	116M
pool4	14×14×256	7×7×256	3×3×256, 2	0	
concat	7×7×256	7×7×256		0	
fc1	7×7×256	1×32×128	maxout p=2	103M	103M
fc2	1×32×128	1×32×128	maxout p=2	34M	34M
fc7128	1×32×128	1×1×128		524K	0.5M
L2	1×1×128	1×1×128		0	
total				140M	1.6B

### Zeiler&Fergus [22] architecture

type	output size	depth	#1×1	#3×3 reduce	#3×3	#5×5 reduce	#5×5	pool proj (p)	params	FLOPS
conv1 (7×7×3, 2)	112×112×64	1							9K	119M
max pool + norm	56×56×64	0						m 3×3, 2		
inception (2)	56×56×192	2		64	192				115K	360M
norm + max pool	28×28×192	0						m 3×3, 2		
inception (3a)	28×28×256	2	64	96	128	16	32	m, 32p	164K	128M
inception (3b)	28×28×320	2	64	96	128	32	64	L <sub>2</sub> , 64p	228K	179M
inception (3c)	14×14×640	2	0	128	256,2	32	64,2	m 3×3,2	398K	108M
inception (4a)	14×14×640	2	256	96	192	32	64	L <sub>2</sub> , 128p	545K	107M
inception (4b)	14×14×640	2	224	112	224	32	64	L <sub>2</sub> , 128p	595K	117M
inception (4c)	14×14×640	2	192	128	256	32	64	L <sub>2</sub> , 128p	654K	128M
inception (4d)	14×14×640	2	160	144	288	32	64	L <sub>2</sub> , 128p	722K	142M
inception (4e)	7×7×1024	2	0	160	256,2	64	128,2	m 3×3,2	717K	56M
inception (5a)	7×7×1024	2	384	192	384	48	128	L <sub>2</sub> , 128p	1.6M	78M
inception (5b)	7×7×1024	2	384	192	384	48	128	m, 128p	1.6M	78M
avg pool	1×1×1024	0								
fully conn	1×1×128	1							131K	0.1M
L2 normalization	1×1×128	0								
total									7.5M	1.6B

### GoogLeNet style Inception models

architecture	VAL
NN1 (Zeiler&Fergus 220×220)	87.9% ± 1.9
NN2 (Inception 224×224)	89.4% ± 1.6
NN3 (Inception 160×160)	88.3% ± 1.7
NN4 (Inception 96×96)	82.0% ± 2.3
NNS1 (mini Inception 165×165)	82.4% ± 2.4
NNS2 (tiny Inception 140×116)	51.9% ± 2.9

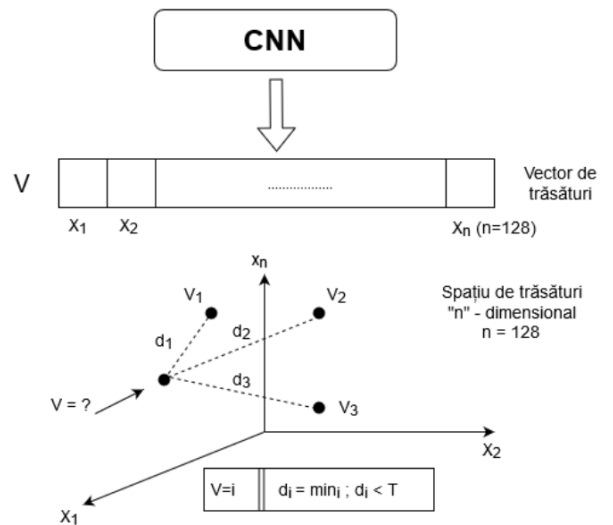
#training images	VAL
2,600,000	76.3%
26,000,000	85.1%
52,000,000	85.1%
260,000,000	86.2%

NN4 (96 x 96)

**Pas final:** Identificarea persoanei - clasificator traditional (NN, dbscan, linear SVM etc):


*NN (Nearest Neighbor)*

```
function classify(persons_data, names, unknown_frame) {
  for i=1 to len(persons_data)
    Calculate distance d(persons_data[i], unknown_frame)
  end
  set I index of min(d)
  return names[i] if (d(i) = I and d(i) < T)
}
```



*DBSCAN*

```
function dbscan(encodings) {
  for i=1 to len(encodings)
    set cID=1
    if encodings[i] is not classified
      doExpandCluster(encodings[i], cID)
      if doExpandCluster is true
        cID += 1
    end
  end
}
```



**Alte metode relevante:**

Taigman, Y., Yang, M., Ranzato, M. A., & Wolf, L. (2014). Deepface: Closing the gap to human-level performance in face verification. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 1701-1708). - (6313 citari – Google Scholar)

[https://www.cv-foundation.org/openaccess/content\\_cvpr\\_2014/html/Taigman\\_DeepFace\\_Closing\\_the\\_2014\\_CVPR\\_paper.html](https://www.cv-foundation.org/openaccess/content_cvpr_2014/html/Taigman_DeepFace_Closing_the_2014_CVPR_paper.html)