

Face verification vs Face Recognition

Verification -

- input image, name / ID
- Output whether the input image is that of the claimed person

recognition -

- has a database of K persons
- get an input image
- output ID if the image is of the K persons (or "not recognized")


Learning a "similarity" function

$d(\text{img 1}, \text{img 2})$ = degree of difference between images

If $d(\text{img1}, \text{img2}) \leq \tau$ "same" } verification
 $> \tau$ "different" }
 ↑
 Some threshold

→ we try to predict if the two images have the same person by comparing it to others

Siamese Network



$$d(x^{(i)}, x^{(j)}) = \|f(x^{(i)}) - f(x^{(j)})\|_2^2$$

Parameters of NN define an encoding $f(x^{(i)})$

Learn parameters so that:

if $x^{(i)}, x^{(j)}$ are the same person, $\|f(x^{(i)}) - f(x^{(j)})\|_2^2$ is small

if $x^{(i)}, x^{(j)}$ are different persons, $\|f(x^{(i)}) - f(x^{(j)})\|_2^2$ is large

Triplet Loss

- Comparing images to anchors

- we want -

$$\underbrace{\|f(A) - f(P)\|_2^2 + \alpha}_{d(A, P)} \leq \underbrace{\|f(A) - f(N)\|_2^2}_{d(A, N)}$$

$$\rightarrow \|f(A) - f(P)\|_2^2 - \|f(A) - f(N)\|_2^2 + \alpha \leq 0$$

\nwarrow margin

- Loss function -

given 3 images A, P, N :

$$L(A, P, N) = \max (\|f(A) - f(P)\|^2 - \|f(A) - f(N)\|^2 + \alpha, 0)$$





$$J = \sum_{i=1}^m L(A^{(i)}, P^{(i)}, N^{(i)})$$



we want to make this zero or less than zero
 so that we train $f(A), f(P), f(N)$ to make
 the similarity function + margin equal to zero
 so that the distances are at least α apart
 from each other

→ we try to choose $d(A, P)$ and $d(A, N)$ such that
 they are close and the learning algorithm works
 hard to push them apart


Face Verification & Binary Classification



$$\hat{y} = \sigma \left(\sum_{k=1}^{128} |f(x^{(i)})_k - f(x^{(j)})_k| + b \right)$$


$$\frac{(f(x^{(i)})_k - f(x^{(j)})_k)^2}{f(x^{(i)})_k + f(x^{(j)})_k}$$

χ^2 similarity





→ here we just compare 2 images and output the result

Neural Style Transfer



Visualizing Deep Layers



$$J(G) = \alpha J_{\text{content}}(c, G) + \beta J_{\text{style}}(s, G)$$

→ initialize G randomly

$G: 100 \times 100 \times 3$

→ use gradient descent to minimize $J(G)$

$$G := G - \frac{\partial J(G)}{\partial G}$$

Content Cost Function

- say you use layer l to compute content cost
- use pre-trained ConvNet (VGG)
- let $a^{[l](c)}$ and $a^{[l](G)}$ be the activation of layer l on the images
- if $a^{[l](c)}$ and $a^{[l](G)}$ are similar, both images have similar content

$$J_{\text{content}}(c, G) = \frac{1}{2} \| a^{[l](c)} - a^{[l](G)} \|^2$$

↑

learn so overall cost is low
to find an image G so the
hidden layer activations are
similar to the content image

Style Cost Function

Let $a_{i,j,k}^{[l]}$ = activation at (i, j, k) . $G^{[l]}$ is $n_c^{[l]} \times n_c^{[l]}$

Style Matrix : $G_{k k'}^{[l](s)} = \sum_i \sum_j a_{ijk}^{[l](s)} \cdot a_{ijk'}^{[l](s)}$ $\left\{ \begin{array}{l} k \text{ and } k' \text{ are correlated if } G_{kk'} \text{ is large} \\ k \text{ and } k' \text{ are un-correlated if } G_{kk'} \text{ is small} \end{array} \right.$
"gram" $G_{k k'}^{[l](G)} = \sum_i \sum_j a_{ijk}^{[l](G)} \cdot a_{ijk'}^{[l](G)}$

Cost Function : $J_{\text{style}}^{[l]}(s, G) = \| G^{[l](s)} - G^{[l](G)} \|_F^2$
 $= \frac{1}{(2 \cdot n_h^{[l]} \cdot n_w^{[l]} \cdot n_c^{[l]})^2} \sum_k \sum_{k'} (G_{kk'}^{[l](s)} - G_{kk'}^{[l](G)})^2$

$$J_{\text{style}}(s, G) = \sum_l \lambda^{[l]} J_{\text{style}}^{[l]}(s, G) \quad \leftarrow \text{applying to all layers}$$

finally,

$$J(G) = \alpha J_{\text{content}}(c, G) + \beta J_{\text{style}}(s, G)$$