

“A principal diferença entre algo que pode dar errado e algo que não pode dar errado é que quando algo que não pode dar errado dá errado, geralmente acaba sendo impossível de ser consertado” (Douglas Adams - O Guia do Mochileiro das Galáxias).

Redes Siamesas

Paulo Ricardo Lisboa de Almeida



Considerere

Considerere o problema de identificar se duas imagens pertencem à mesma pessoa.




= ???



Considerare

Ou ainda, o problema de identificar se duas assinaturas pertencem à mesma pessoa.



= ???



Resposta

O que esses problemas têm em comum?

Resposta

O que esses problemas têm em comum?

- Estamos comparando objetos.
- Possuem poucas amostras.
 - Exemplo: sua agência bancária ou cartório possui uma ou duas fotos ou assinaturas suas para comparar.
- É inviável treinar um modelo que reconheça todas as assinaturas ou fotos de pessoas possíveis.
 - Mesmo que tivéssemos amostras de todas as pessoas do mundo, quando uma nova pessoa nascer, esse modelo se tornaria obsoleto.

Ideia

Para esse tipo de problema, a ideia é que objetos que “pertencem à mesma classe” estão próximos no espaço, enquanto objetos que não pertencem à mesma classe, estão distantes.








= ???



Dissimilaridade

O conceito de dissimilaridade é projetar os objetos em um espaço em que a distância entre objetos parecidos é pequena, e a distância de objetos diferentes é grande.

Dissimilaridade





Ideia Geral



Redes de mesma estrutura para extrair a representação.

Embedding.

Entrada 1



Entrada 2



Cálculo de distância.



Resposta.

Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., & Shah, R. Signature verification using a "siamese" time delay neural network. Neurips. 1993.

Baldi, P., & Chauvin, Y. Neural networks for fingerprint recognition. Neural Computation. 1993.

Na prática

Poderíamos ter “dois braços” idênticos da rede na memória, e computar as imagens em paralelo.

Rápido, mas ocupa o dobro da memória.

Na prática

Poderíamos ter “dois braços” idênticos da rede na memória, e computar as imagens em paralelo.

Rápido, mas ocupa o dobro da memória.

Uma alternativa é manter somente um lado da rede, e mandar as imagens em sequência para ela.

Assume-se que os dois braços da rede possuem a mesma **estrutura e pesos**.

```
def __ativacao_euclidiana(self, x, y):  
    return torch.sqrt(F.relu(torch.sum((x - y) ** 2, dim=1, keepdim=True)))  
  
def forward(self, input1, input2):  
    #Mantém apenas um “braço da rede” na memória. Manda as duas imagens para ele  
    #em sequência. Depois, calcula a função de ativação que combina as respostas.  
    tower1 = self.__embedding_network(input1)  
    tower2 = self.__embedding_network(input2)  
    return self.__ativacao_euclidiana(tower1, tower2)
```

Um exemplo

O exemplo a seguir é uma rede que usa a combinação das estruturas de rede dos seguintes artigos.

CVF This CVPR2015 paper is the Open Access version, provided by the Computer Vision Foundation. The authoritative version of this paper is available in IEEE Xplore.

FaceNet: A Unified Embedding for Face Recognition and Clustering

Floriano Schroff
florian@fb.com
Google Inc.

Dmitry Kuznetsov
dmitrykuznetsov@google.com
Google Inc.

James Philbin
jphilbin@google.com
Google Inc.

Abstract
Despite significant recent advances in the field of face recognition [10, 12, 17], representing face recognition and recognition effectively of wide personal various challenges in various applications. In this paper, we present a system, called FaceNet, that directly learns a mapping from face images to a compact face embedding space which directly encodes a measure of face similarity. Once this space has been produced, tasks such as face recognition, verification and clustering can be easily implemented using standard methods with FaceNet embeddings as the face vectors.

The method uses a deep convolutional neural network to directly optimize the embedding itself, rather than an intermediate bottleneck layer or previous deep learning approaches. To train, we use triplets of roughly aligned matching face images but just perform gradient descent on novel online triplet mining method. The benefit of our approach is much greater representational efficiency, we achieve state-of-the-art face recognition performance using only 4096 bits per face.

On the widely used Labeled Faces in the Wild (LFW) dataset, our system achieves the best overall accuracy of 97.6%. On YouTube Faces DB it achieves 95.2%. Our system can be leveraged to a comparison to the best published result (57.6) by 38% on both datasets.

1. Introduction

In this paper we present a unified system for face verification in that the same person, recognition (who is this person) and clustering (what contains people among these faces). Our method is based on learning a face-level embedding per image using a deep convolutional network. The network is trained such that the squared L2 distance of the embedding space directly correspond to face similarity. One of the same person has small distances and faces of distinct people have large distances.

Once the embedding has been produced, then the aforementioned tasks become straight forward. Face verification simply involves thresholding the distance between the two embeddings; recognition becomes a k-NN classification problem, and clustering can be achieved using off-the-shelf techniques such as K-means or agglomerative clustering.

Previous face recognition approaches based on deep neural networks used a classification layer [5, 7] trained over a set of known face identities and then use an intermediate bottleneck layer to produce a representation with generalization beyond the set of identities used in training. The downside of this approach are its inflexibility and its inefficiency: one has to hope that the bottleneck representation generalizes well to new faces, and by using a bottleneck layer the representation size per face is usually very large (1000s of di-

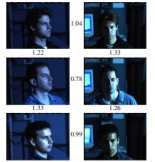


Figure 1. Illustration and Face features. Face and distance have been a long standing problem in face recognition. This figure shows the same distance of FaceNet between pairs of faces of the same and different poses in different pose and in background combination. A distance of 0.1 means the faces are identical. 0.2 corresponds to the images species, two different identities. You can see that a threshold of 1.3 would classify every pair correctly.

Schroff, F., Kalenichenko, D., & Philbin, J. Facenet: A unified embedding for face recognition and clustering. CVPR. 2015.

Quantifying Parking Dwell Time with Siamese Networks

Dwell Time

Antonio Arribas

Abstract: In smart cities, it is common practice to define a maximum length of time for a given parking space to ensure the space's mobility and allowing the flow of individual transportation vehicles. However, systematically determining individual use dwell times from images faces challenges, such as images collected from low-resolution cameras, shifting positions, and variable angles. In this work, we propose a method that combines a Siamese neural network with triplet loss to compute the dwell time of each car in a parking lot. The proposed method first defines the parking space classes between adjacent and empty using the classification network. Then, it uses the Siamese network to check if the parked car is the same as the previous image. Using an experimental protocol that focuses on a controlled scenario, we show that a perfect prediction using the proposed system can reach a mean average score (MAP) of 0.9, whereas, on average, only 0.7 of perfect prediction. Nevertheless, our experiments show a drop in the prediction quality when a real-world scenario is used to quantify the parking space classes, reaching a MAP of 0.75 instead of 0.9. In all perfect predictions, showing that the proposed Siamese network is providing the required by the quality of the detection and the assignment of the classes.

1. INTRODUCTION
Developments regarding parking space monitoring through images have been proposed in the recent past, including the detection of license to be used in techniques [1], [2]. Approaches to classify the individual parking spaces between empty or occupied [3], [4], and the automatic recognition of the parking space position [5], [6].

Following the trend of these innovations, this paper presents an approach to counting the time a car stays parked in a parking space. To accomplish this, we use image triplets and propose to train images using Siamese networks. This is an unusual task since by collecting like about the time each car stays parked, we may generate natural distributions about the amount of classes' parking slots.

Furthermore, as a common practice for parking spaces, works maximum length of time (i.e., 15 minutes) and a system that counts the time a driver stays parked may help monitoring the parking spaces. Other usage of such a system may include identifying abandoned or broken cars, or use parked for long time (e.g., 24 hours) or broken and not being repaired (e.g., 24 hours) to be used as a parking space.

Quantifying the parking time is a complex task because of the vehicles' dynamic effects collected at a distance, with low-resolution cameras, e.g., in the PKL dataset, the bounding box of a car is 75x75 pixels in size, on average. 2. Related work

In smart cities, it is common practice to define a maximum length of time for a given parking space to ensure the space's mobility and allowing the flow of individual transportation vehicles. However, systematically determining individual use dwell times from images faces challenges, such as images collected from low-resolution cameras, shifting positions, and variable angles. In this work, we propose a method that combines a Siamese neural network with triplet loss to compute the dwell time of each car in a parking lot. The proposed method first defines the parking space classes between adjacent and empty using the classification network. Then, it uses the Siamese network to check if the parked car is the same as the previous image. Using an experimental protocol that focuses on a controlled scenario, we show that a perfect prediction using the proposed system can reach a mean average score (MAP) of 0.9, whereas, on average, only 0.7 of perfect prediction. Nevertheless, our experiments show a drop in the prediction quality when a real-world scenario is used to quantify the parking space classes, reaching a MAP of 0.75 instead of 0.9. In all perfect predictions, showing that the proposed Siamese network is providing the required by the quality of the detection and the assignment of the classes.

Figure 1 shows examples from the PKL and CNPARK-EXT datasets. In this work, we focus on the following contributions:

- We define a Siamese Network to compare cars.
- We define a verification protocol to track the parking status and count the dwell time of each car.

The complete pipeline includes a classification network to identify the parking spaces, a Siamese network to compare cars and determine their position, the determination of both networks to update the dwell time of each parked car.

The structure of this work is organized as follows: we show the related works in Section II. The proposed approach, which includes the description of the classification and Siamese networks used and the algorithm used to compare the car's position, is detailed in Section III. In Section IV, we define the experimental protocol, which we follow a similar way to the one used in our train samples from the larger parking lot given for the PKL dataset. In Section V, we describe the experimental results in the PKL and CNPARK-EXT datasets, respectively. Finally, in Section VI, we present our conclusions and future works.




Fig. 2. Image examples from the PKL and CNPARK-EXT datasets.

Ribas, Benedet, Zanlorensi, Oliveira, Almeida. Quantifying Parking Dwell Time with Siamese Networks. ICMLA. 2024.


Ideias: Facenet -> Triplet Loss

Triplet Loss.


Introduzido no Artigo da Facenet.

Cada amostra de treinamento é uma tripla (triplet) contendo uma âncora, uma imagem positiva, da mesma classe da âncora, e uma imagem negativa, de uma classe diferente.

âncora



positivo




negativo




Triplet Loss

O objetivo da rede é projetar as imagens em um espaço as âncora e os exemplos positivos estão próximos, ao passo que as âncoras e os exemplos negativos estão distantes.

âncora



positivo



negativo




Triplet Loss


Desejamos que:

$$\|x_i^a - x_i^p\|_2^2 + \alpha < \|x_i^a - x_i^n\|_2^2, \forall (x_i^a, x_i^p, x_i^n) \in \mathcal{T}$$

âncora



positivo



negativo



Onde:

x^a é a âncora.

x^p é a imagem positiva.

x^n é a imagem negativa.

$\|\cdot\|_2$ é a norma L2.

\mathcal{T} é o conjunto de todas triplets, e tem tamanho N .


α é uma margem mínima entre pares positivos e negativos.

Triplet Loss


Logo, a função de Loss fica:

$$L = \sum_i^N \|x_i^a - x_i^p\|_2^2 - \|x_i^a - x_i^n\|_2^2 + \alpha$$

âncora



positivo



negativo



Onde:

x^a é a âncora.

x^p é a imagem positiva.

x^n é a imagem negativa.

$\|\cdot\|_2$ é a norma L2.

\mathcal{T} é o conjunto de todas triplets, e tem tamanho N .

α é uma margem mínima entre pares positivos e negativos.

Facenet

Existem ainda outros detalhes no artigo, como garantir que os dados residem em uma hiperesfera de raio 1, e que as triplets são “difíceis” para melhorar a convergência e evitar o colapso da rede na origem.

Veja os detalhes nos artigos.



Parking Dwell Time

Como usar redes neurais siamesas para cronometrar automaticamente o tempo de permanência de veículos parados?

ICMLA 2024 - Miami, Estados Unidos.

Aceito para publicação.



Marcelo

Heloisa

Paulo

Luiz

Luiz





Parking Dwell Time

9h40

9h45

9h50





Duas fases: 1 - Encontrar os veículos nas imagens.

Parking Dwell Time

9h40

9h45

9h50






Duas fases: 1 - Encontrar os veículos nas imagens.

Parking Dwell Time

9h40

9h45

9h50



=



≠



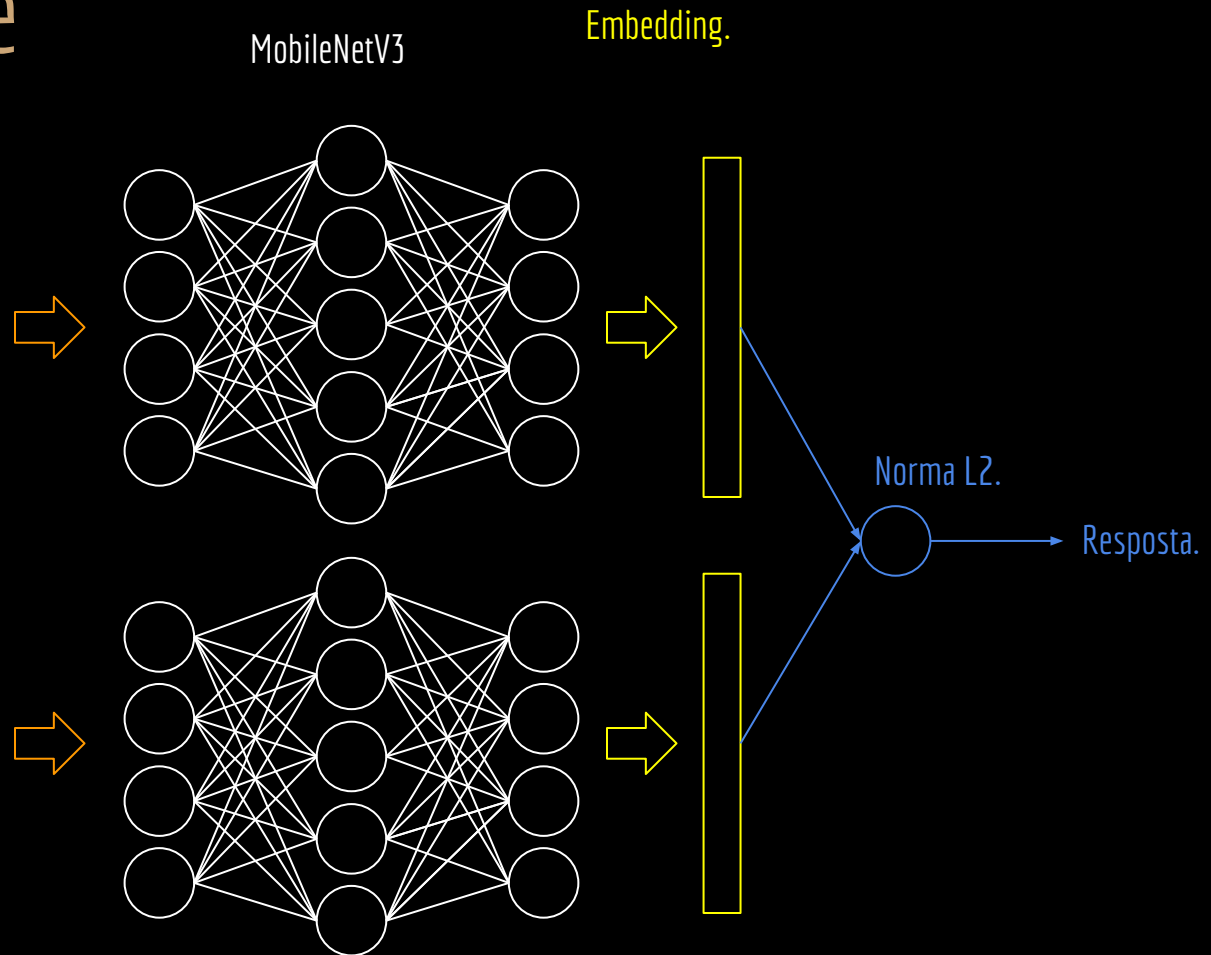
Duas fases: 2 - Comparar usando uma rede siamesa. Usar o resultado para cronometrar o tempo de permanência.

Parking Dwell Time

Rede Siamesa que usa uma MobileNetV3
pré-treinada na ImageNet.

Fine-Tuning para comparar carros.

Contrastive Loss com função de perda.




Parking Dwell Time

Rede Siamesa que usa uma MobileNetV3 pré-treinada na ImageNet.

Fine-Tuning para comparar carros.

Contrastive Loss com função de perda.

No exemplo disponibilizado, trocamos pelo triplet loss.



Faça você mesmo

Execute e entenda o exemplo disponibilizado.

Exercícios

1. A rede treinada usa um limiar de decisão arbitrário.
 - a. Estude como calcular o Equal Error Rate (EER) para problemas binários utilizando a curva RoC.
 - b. Calcule o EER nos dados de validação. Utilize o valor como novo limiar de decisão.
2. Troque a triplet loss pela contrastive loss (usada no paper original), e faça os ajustes necessários.
 - a. O algoritmo de treinamento precisará ser alterado. A seguir é dada a função que faz o cálculo do contrastive loss, e a nova função de forward.


```
def __euclidean_distance(self, x, y):  
    return torch.sqrt(F.relu(torch.sum((x - y) ** 2, dim=1, keepdim=True)))
```

```
def forward(self, input1, input2):  
    tower1 = self.__embedding_network(input1)  
    tower2 = self.__embedding_network(input2)  
    return self.__euclidean_distance(tower1, tower2)
```

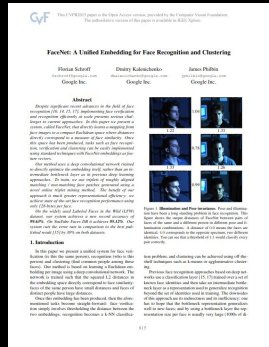
```
class ContrastiveLoss(torch.nn.Module):  
    def __init__(self, margin=1.0):  
        super(ContrastiveLoss, self).__init__()  
        self.__margin = margin  
  
    def forward(self, y_true, y_pred):  
        square_pred = y_pred ** 2  
        margin_square = torch.relu(self.__margin - y_pred) ** 2  
        return torch.mean((1 - y_true) * square_pred + y_true * margin_square)
```

Referências

Hugh Cartwright (Editor).
Artificial Neural Networks.
Springer New York. 2017.



Schroff, F., Kalenichenko, D., & Philbin, J. Facenet: A unified embedding for face recognition and clustering. CVPR. 2015.



Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., & Shah, R. Signature verification using a "siamese" time delay neural network. Neurips. 1993.

Baldi, P., & Chauvin, Y. Neural networks for fingerprint recognition. Neural Computation, 5(3), 402-418. 1993.

Ribas, Benedet, Zanlorensi, Oliveira, Almeida. Quantifying Parking Dwell Time with Siamese Networks. ICMLA. 2024.

Signature Verification using a "Siamese" Time Delay Neural Network

Jose Bromley, Isabelle Guyon, Yann LeCun, Edward Säckinger, and Rajmohan Shah

AT&T Bell Laboratories

Research Triangle Park

Research Triangle Park, NC 27709

jbromley@att.com

guyon@att.com

lecun@att.com

sackinger@att.com

shah@att.com

Abstract

The paper describes an algorithm for verification of signatures written in an arbitrary style. The algorithm is based on a neural network with a time delay element. The network is trained on a set of samples of real handwritten signatures. The network is then used to verify a signature. The network is trained on a set of samples of real handwritten signatures. The network is then used to verify a signature. The network is trained on a set of samples of real handwritten signatures. The network is then used to verify a signature.

1 INTRODUCTION

The aim of the project was to make a signature verification system based on the AT&T Bell Laboratories Time Delay Neural Network (TDNN) [1]. The system was designed to be used in a system for signature verification. The system was designed to be used in a system for signature verification. The system was designed to be used in a system for signature verification.

2 CONNECTION TO RELATED WORK

Neural Networks for Fingerprint Recognition

Plaza Ball, Ed. Handbook of Fingerprint Recognition, Cambridge University Press, 1999.

3 CONCLUSION

The time delay neural network (TDNN) architecture was used to recognize the signature images in a realistic problem in pattern recognition that has not, to the best of our knowledge, been solved before. The TDNN architecture was used to recognize the signature images in a realistic problem in pattern recognition that has not, to the best of our knowledge, been solved before.

4 REFERENCES

[1] LeCun, Y., Sackinger, E., Guyon, I., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[2] Guyon, I., LeCun, Y., Sackinger, E., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[3] LeCun, Y., Sackinger, E., Guyon, I., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[4] Guyon, I., LeCun, Y., Sackinger, E., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[5] LeCun, Y., Sackinger, E., Guyon, I., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[6] Guyon, I., LeCun, Y., Sackinger, E., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[7] LeCun, Y., Sackinger, E., Guyon, I., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[8] Guyon, I., LeCun, Y., Sackinger, E., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[9] LeCun, Y., Sackinger, E., Guyon, I., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[10] Guyon, I., LeCun, Y., Sackinger, E., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[11] LeCun, Y., Sackinger, E., Guyon, I., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[12] Guyon, I., LeCun, Y., Sackinger, E., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[13] LeCun, Y., Sackinger, E., Guyon, I., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[14] Guyon, I., LeCun, Y., Sackinger, E., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[15] LeCun, Y., Sackinger, E., Guyon, I., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[16] Guyon, I., LeCun, Y., Sackinger, E., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[17] LeCun, Y., Sackinger, E., Guyon, I., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[18] Guyon, I., LeCun, Y., Sackinger, E., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[19] LeCun, Y., Sackinger, E., Guyon, I., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

[20] Guyon, I., LeCun, Y., Sackinger, E., and Bromley, J. Signature verification using a "siamese" time delay neural network. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), 1993, pp. 789-796.

Licença

Esta obra está licenciada com uma Licença [Creative Commons Atribuição 4.0 Internacional](https://creativecommons.org/licenses/by/4.0/).