Computer Vision for Fraud detection

Bucharest, 6th June, 2019

Katarina Milosevic
Ioana Gherman

How to cook good image similarity?

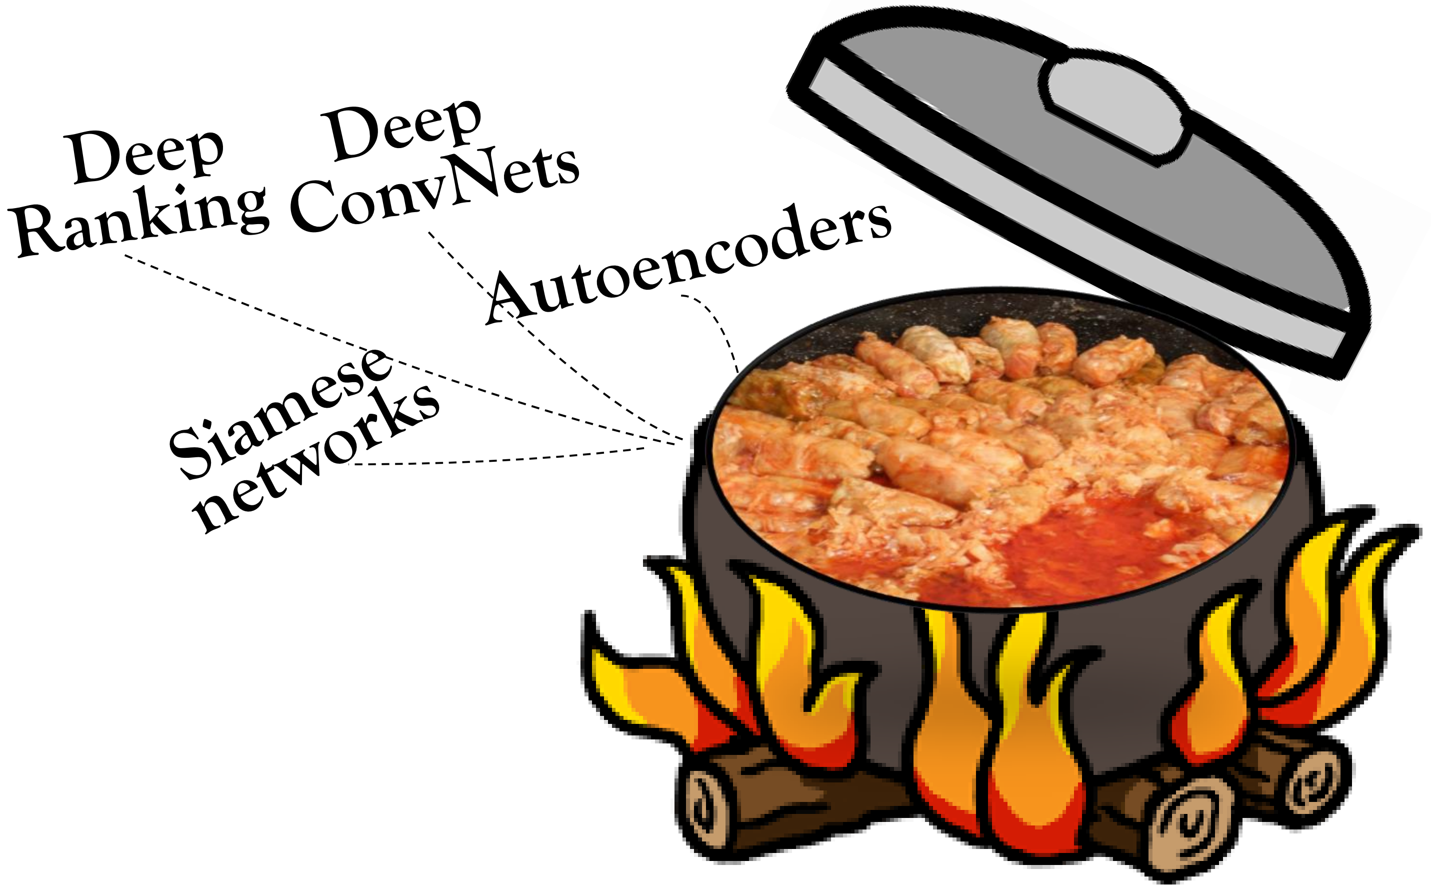

Autoencoders – Siamese networks – Deep Ranking

Autoencoders – One Man Show

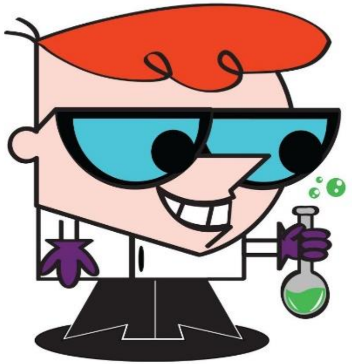

Siamese networks – Powerful Duo

Deep Ranking – Triplets power

Autoencoders – One Man Show

Autoencoders

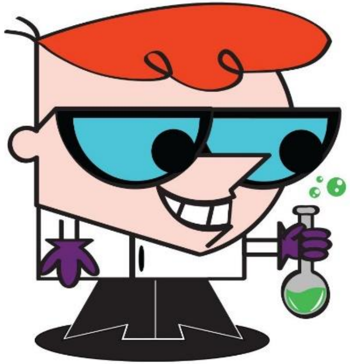

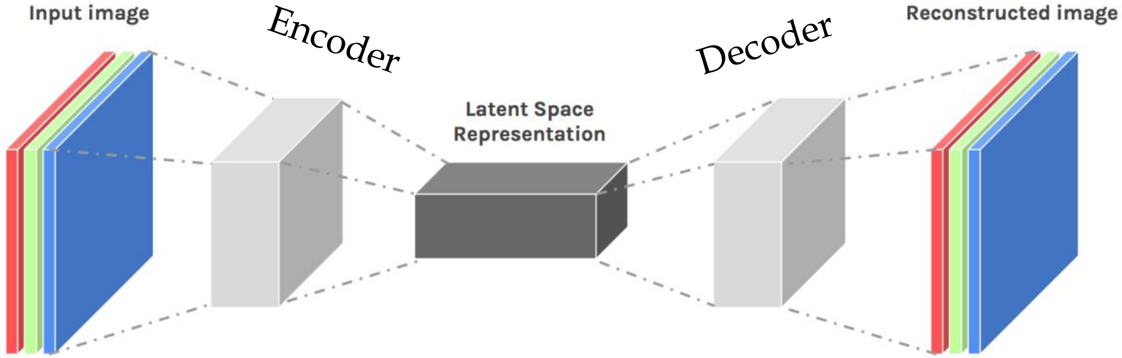

Pros: We don't need to label the data

Cons: Autoencoders learn "blindly" - do not focus on what we are interested in

Autoencoders – One Man Show

Autoencoders

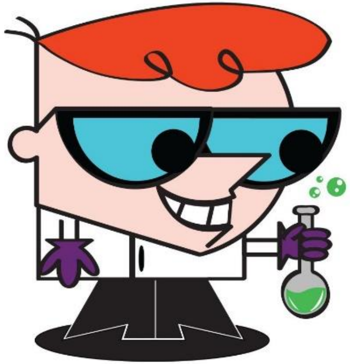

Input images

Output images

Similarity of *query image* calculated on *encoded images*

Siamese networks – Is it the same person?

Siamese network

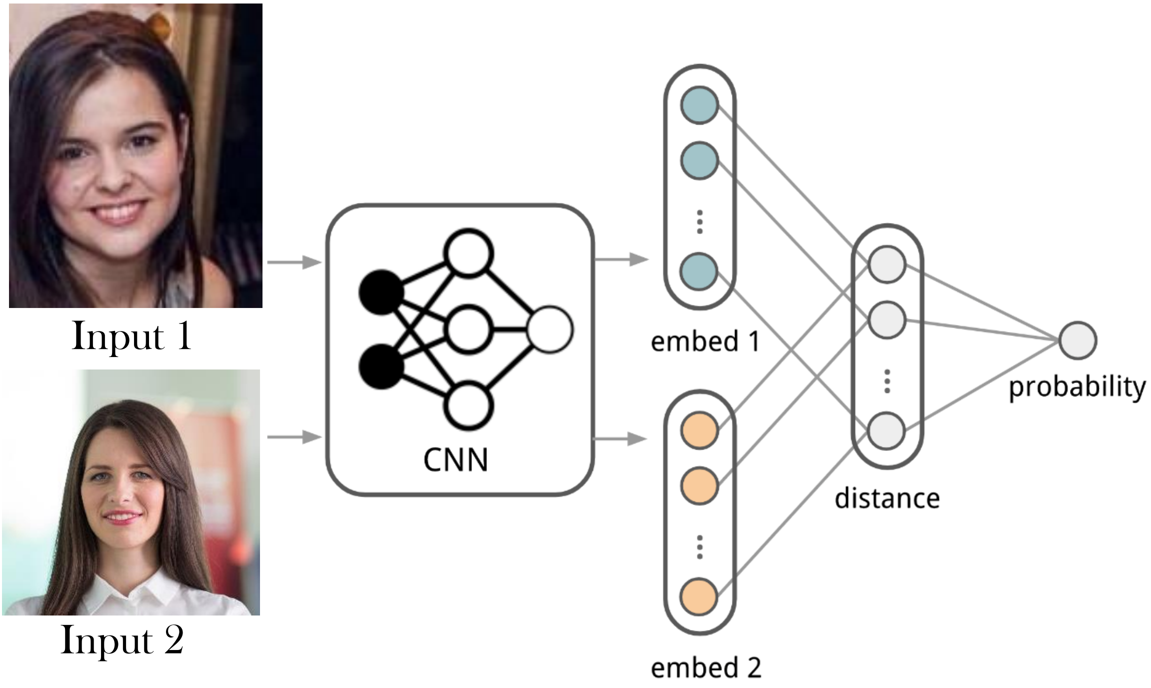

Koch, Gregory, Richard Zemel, and Ruslan Salakhutdinov. "Siamese neural networks for one-shot image recognition." ICML deep learning workshop. Vol. 2. 2015.

Siamese networks – Damage similarity

Siamese network

1 – same damage

0 – not the same damage

Pros: Recognizes if the instances are from the same collection

Cons: Hard to teach the model to focus on the damage itself

Deep Ranking – Distinguish between positive and negative example

Deep Ranking

Query

Positive

Negative

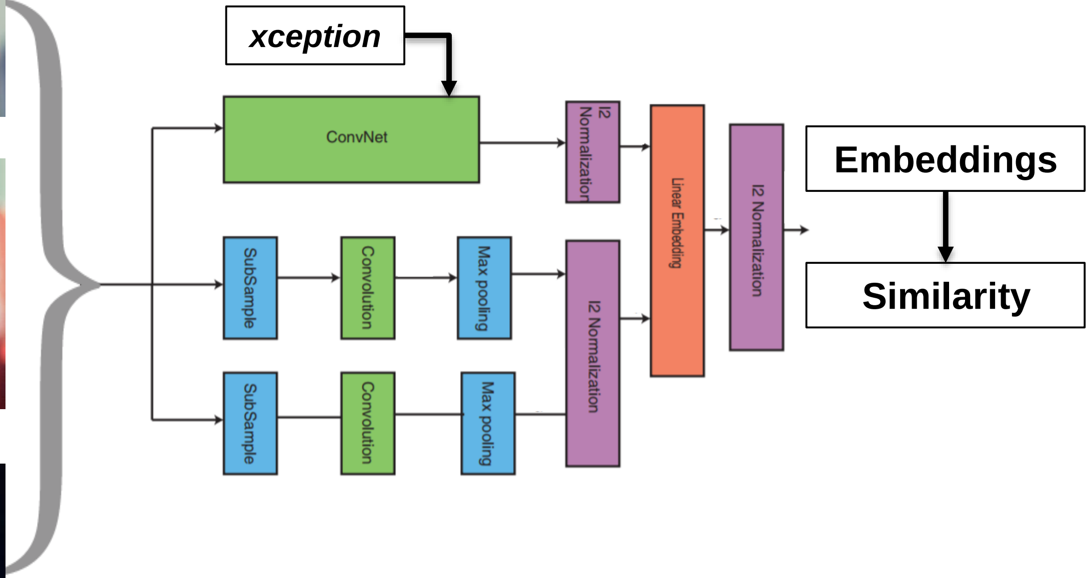

Wang, Jiang, et al. "Learning fine-grained image similarity with deep ranking." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2014.

Deep Ranking–Distinguish between same/different damage

Query

Positive

Negative

Pros: Loss is calculated taking into account both positive and negative instance for each input
=> Model learns how to recognize positive and negative examples for the given input image

Cons: Construction of the training dataset, Evaluation of the model

Deep Ranking – Distinguish between same / different damage

Deep Ranking

1 – same damage

0 – not the same damage

How to cook good image similarity? Brand new recipe!

Deep ConvNet – Xception

Multi-labelling

Clustering

Similarity metrics

Extracting features with Xception

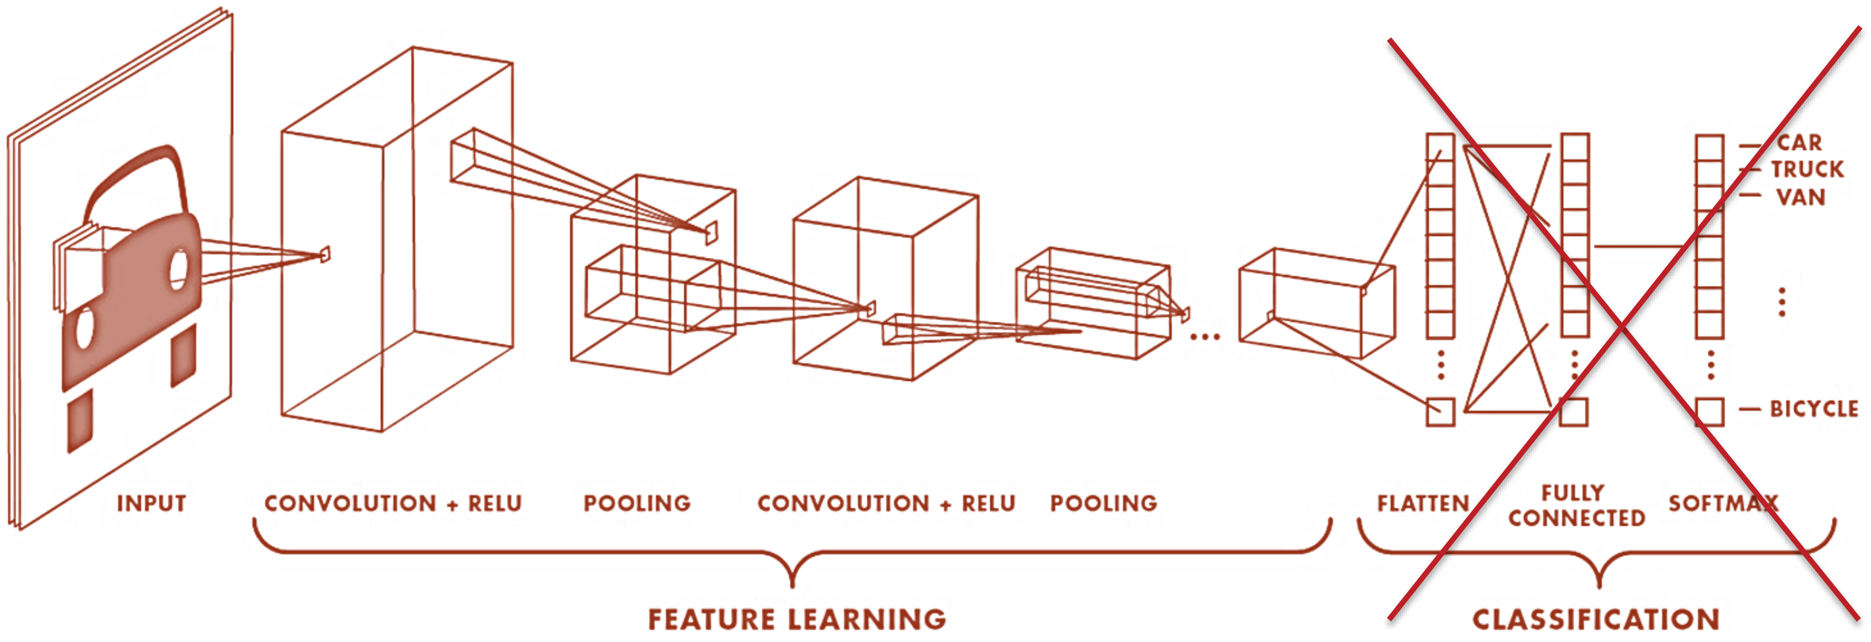

Predicting multi-labels based on xception features

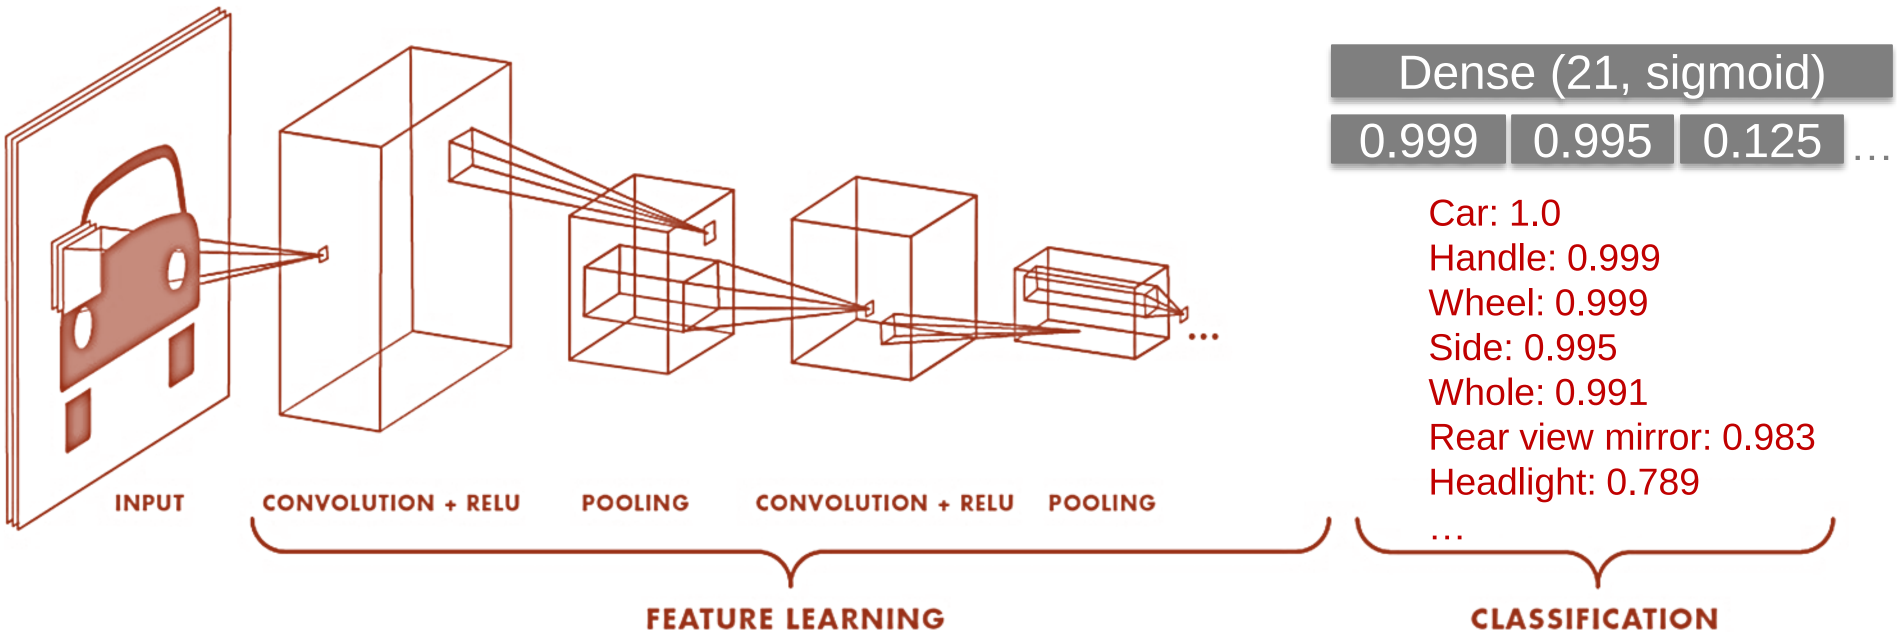

Which car parts do you see in this image?

Car : 0.999
Side : 0.995
Handle : 0.894
Rear view mirror : 0.582
Wheel : 0.267

Cluster images based on multi-label probabilities

Multi-labels

Car: 1.0
Handle: 0.999
Wheel: 0.999
Side: 0.995
Whole: 0.991
Rear view mirror: 0.983
Headlight: 0.789
…

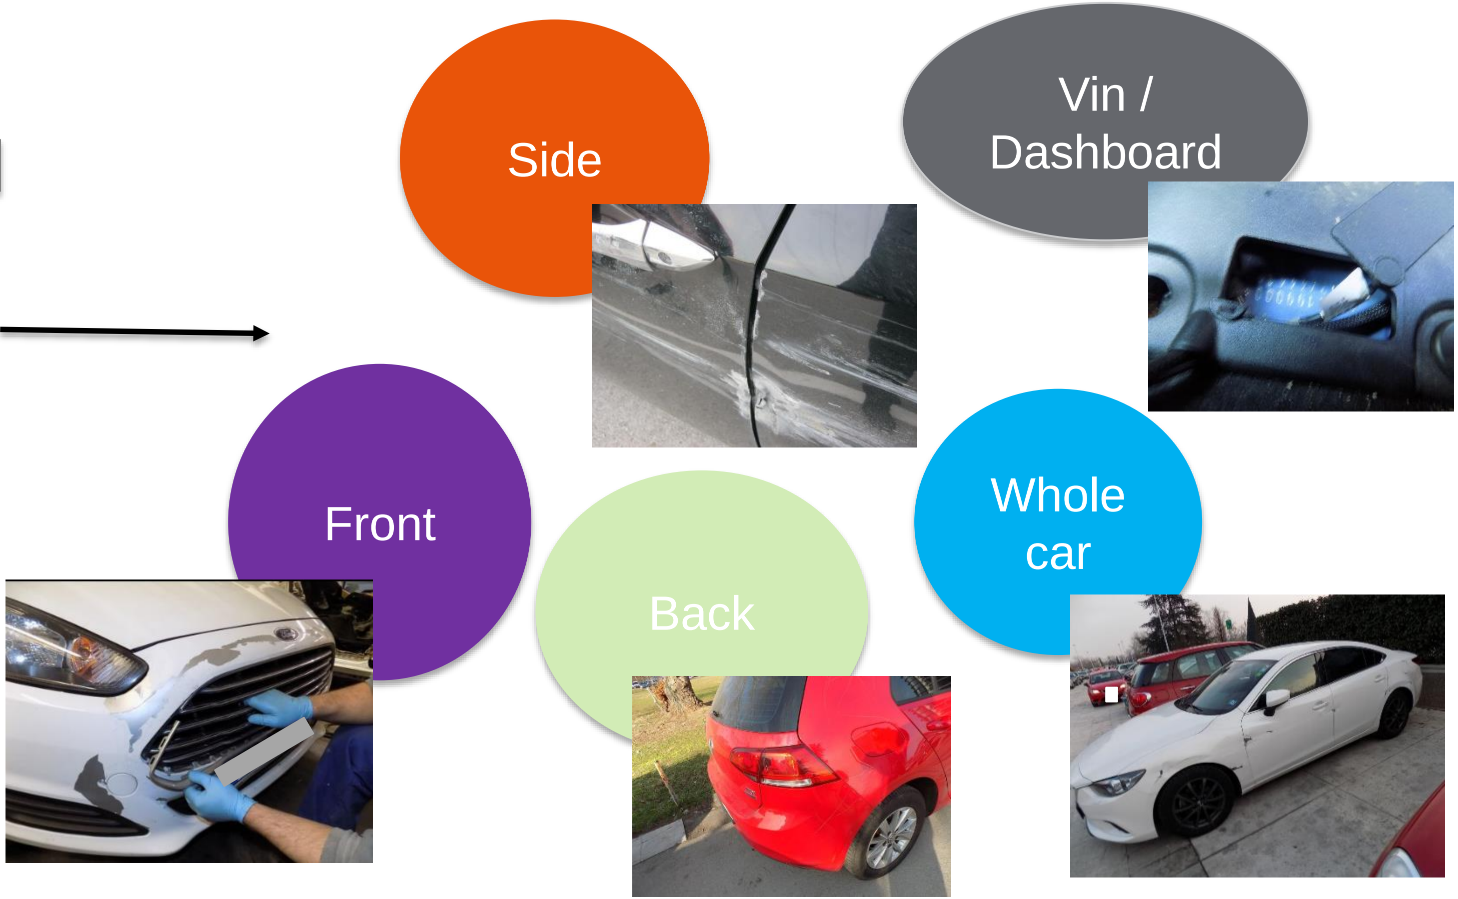

Find the most similar images inside the cluster

Front

Cosine similarity of xception embeddings of 2 photos: 0.81

Claim 1

Claim 2

Eliminate symmetric pairs

Front

Cosine similarity of xception embeddings of 2 photos: 0.71

Claim 1

Claim 2

Eliminate symmetric pairs – Neural network on Gist features

Xception was trained with image augmentation

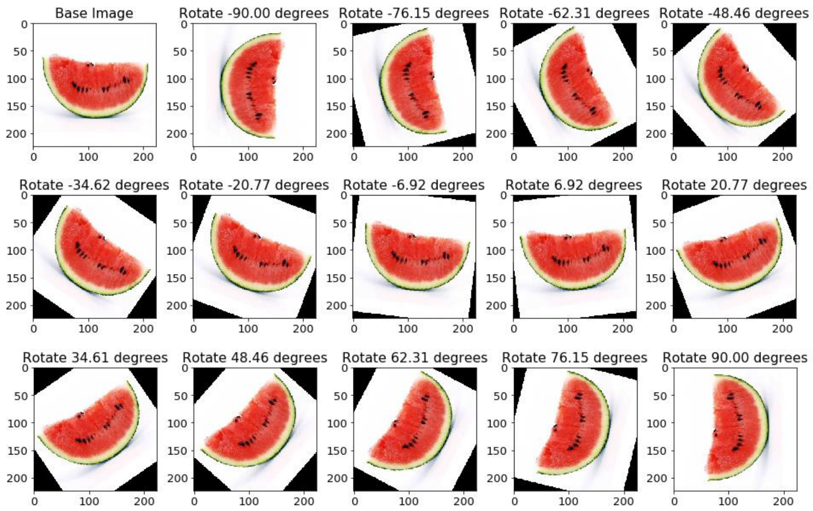

Extracting GIST features which are invariant to images symmetricity

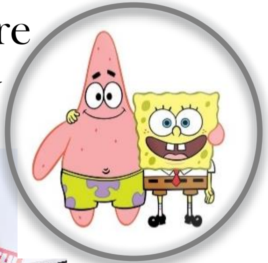

1 – symmetric

0 – not symmetric

How to cook good image similarity? Winning recipe!

Deep ConvNet - Xception

Multi-labelling

Clustering

Similarity metrics

Symmetricity

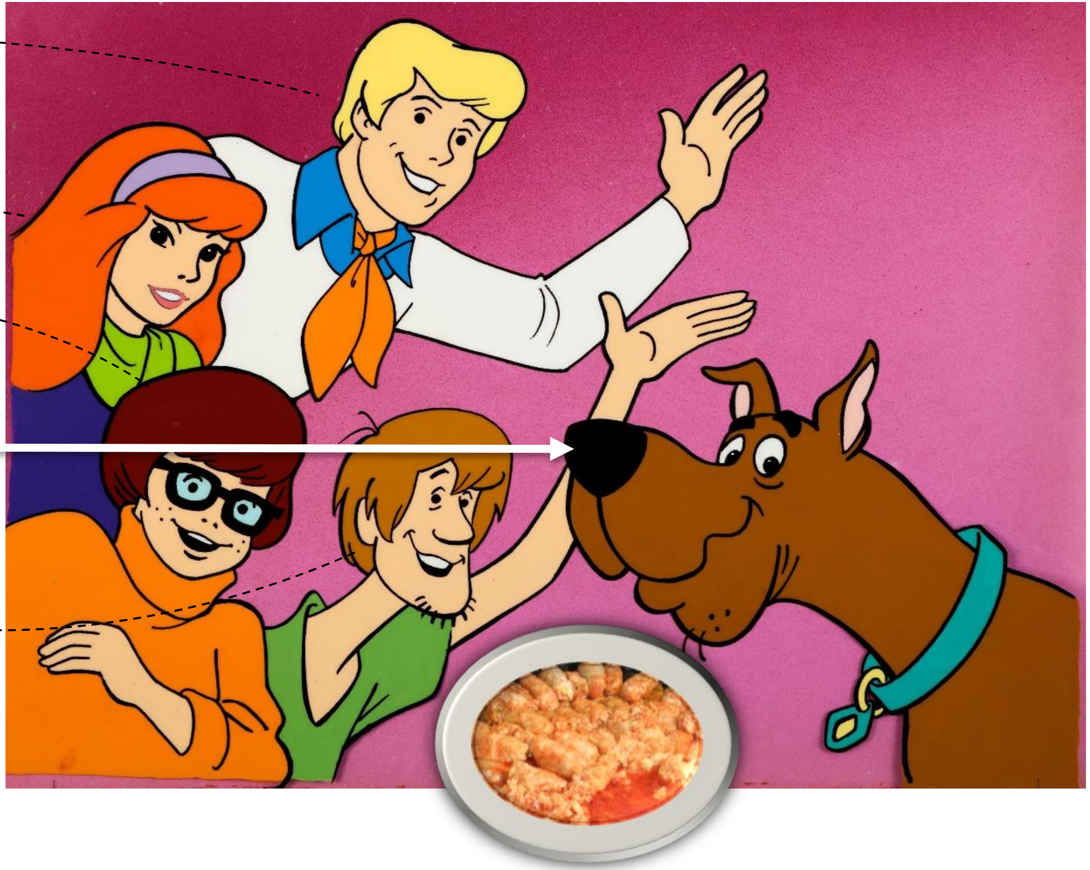

Thank You.

Contacts:

Ioana.Gherman@generali.com
Katarina.Milosevic2@generali.com