Jan, 13th - 14th 2020

Neosperience Empathy in Technology

Computer Vision Applications



www.neosperience.com | blog.neosperience.com | info@neosperience.com

Who am I?

Luca Bianchi

Chief Technology Officer @ Neosperience

Chief Technology Officer @ WizKey

Serverless Meetup Italy manager

ServerlessDays co-organizer



github.com/aletheia



https://it.linkedin.com/in/lucabianchipavia



@bianchiluca



https://speakerdeck.com/aletheia





Who am I?

Alessandro Bitetto

PhD student @ FinTech Lab Pavia

PhD intern @ Neosperience



https://www.linkedin.com/in/alessandrobitetto





Neosperience Cloud

A cloud platform built on AWS to deliver DCX projects

Neosperience Cloud is the technology platform that allows creating personalized experiences for your customers that drive loyalty faster paths and to purchase. Unlike existing technologies that rely only on demographics data, we use proprietary models, developed with Al, to personalize your offering to the right segment. A compelling experience for each customer at the right time, place, and situational context.

- Deeply understand their customers and be more useful to them by delivering relevant digital experiences.
- Delight customers by delivering relevant experiences across mobile, web, in-store.
- Maintain their Brand identity and increase value as platforms like Amazon, Google and Facebook drive up disintermediation and make companies unintentional utilities.
- associated with the alignment of apps, web apps, social media and conversational interfaces.

...which means fast time to market, machine learning and scalability by design.





Keep pace with the variety of devices and interaction models available to customers to overcome complexity and costs

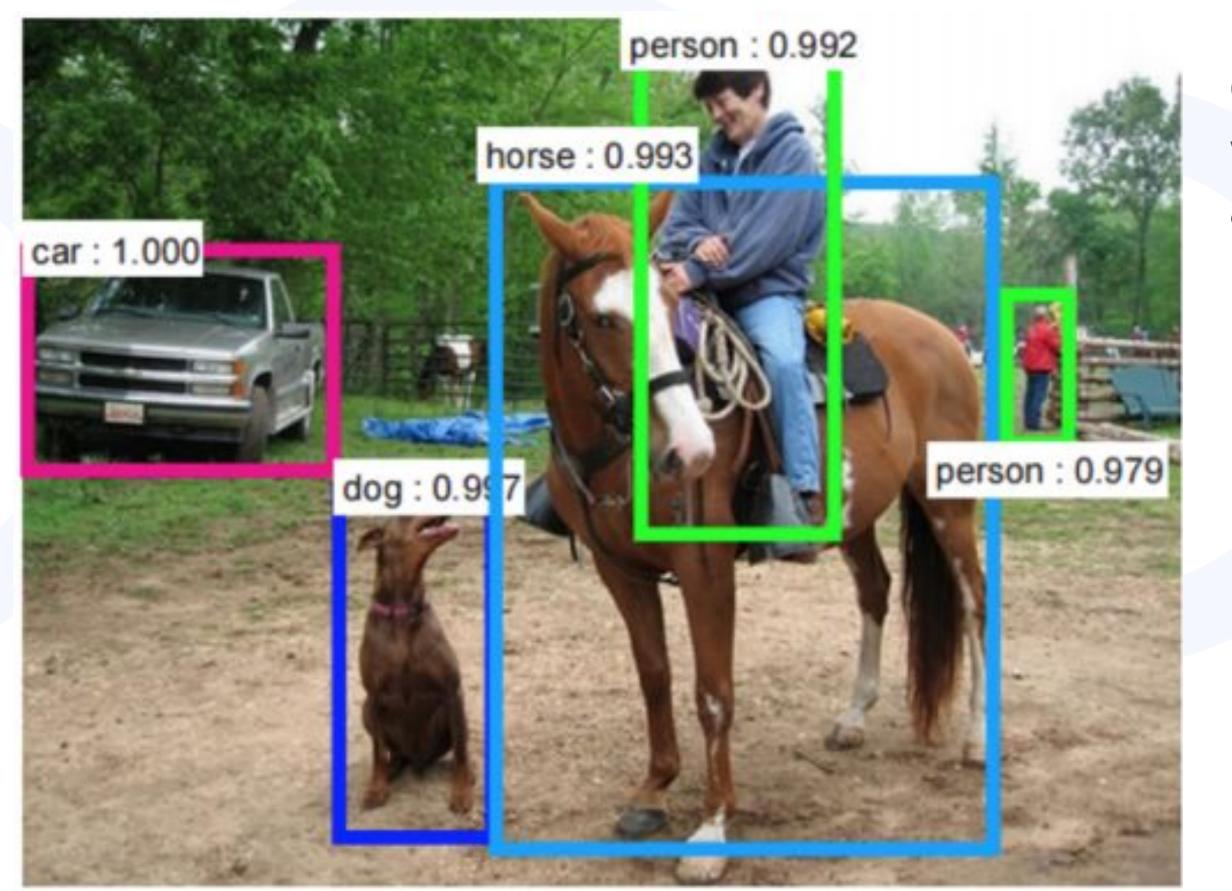


Part 1 Convolutional Neural Networks (CNNs)



Convolutional Neural Network

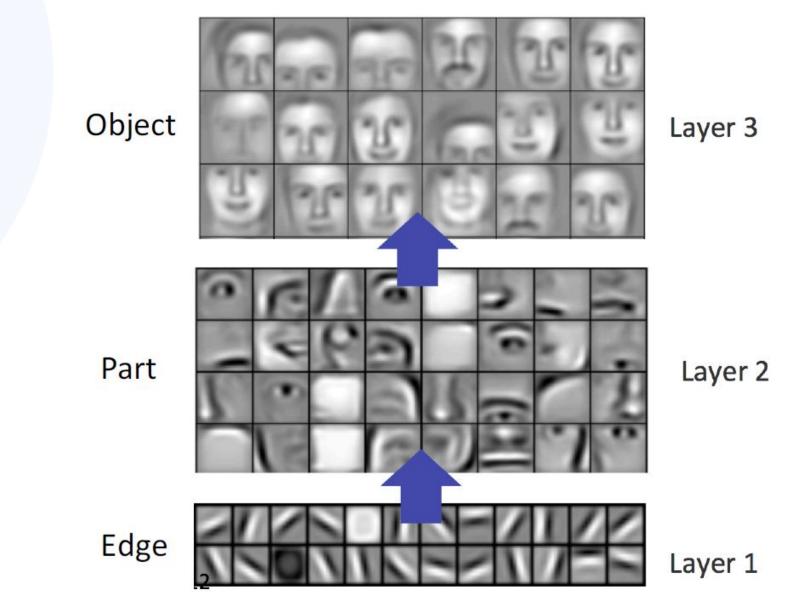
proven very effective in areas such as image recognition and classification.





Convolutional Neural Networks (ConvNets or CNNs) are a category of Neural Networks that have

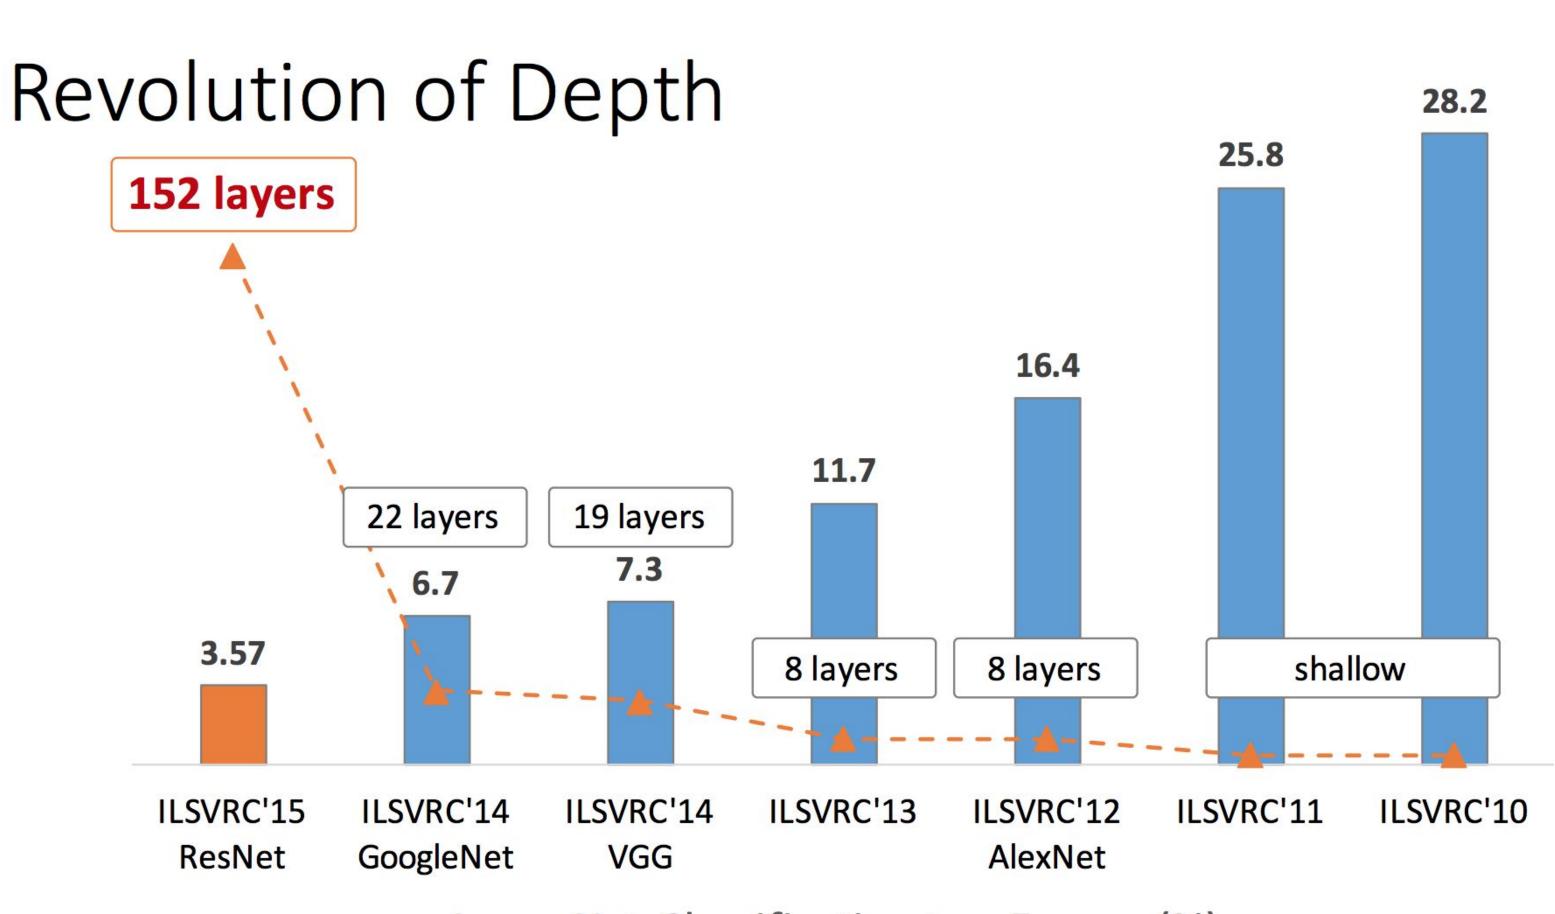
CNNs are based on **Hierarchical Compositionality**: we start from a low level input (pixel) and then we aggregate informations up to an higher interpretation level.

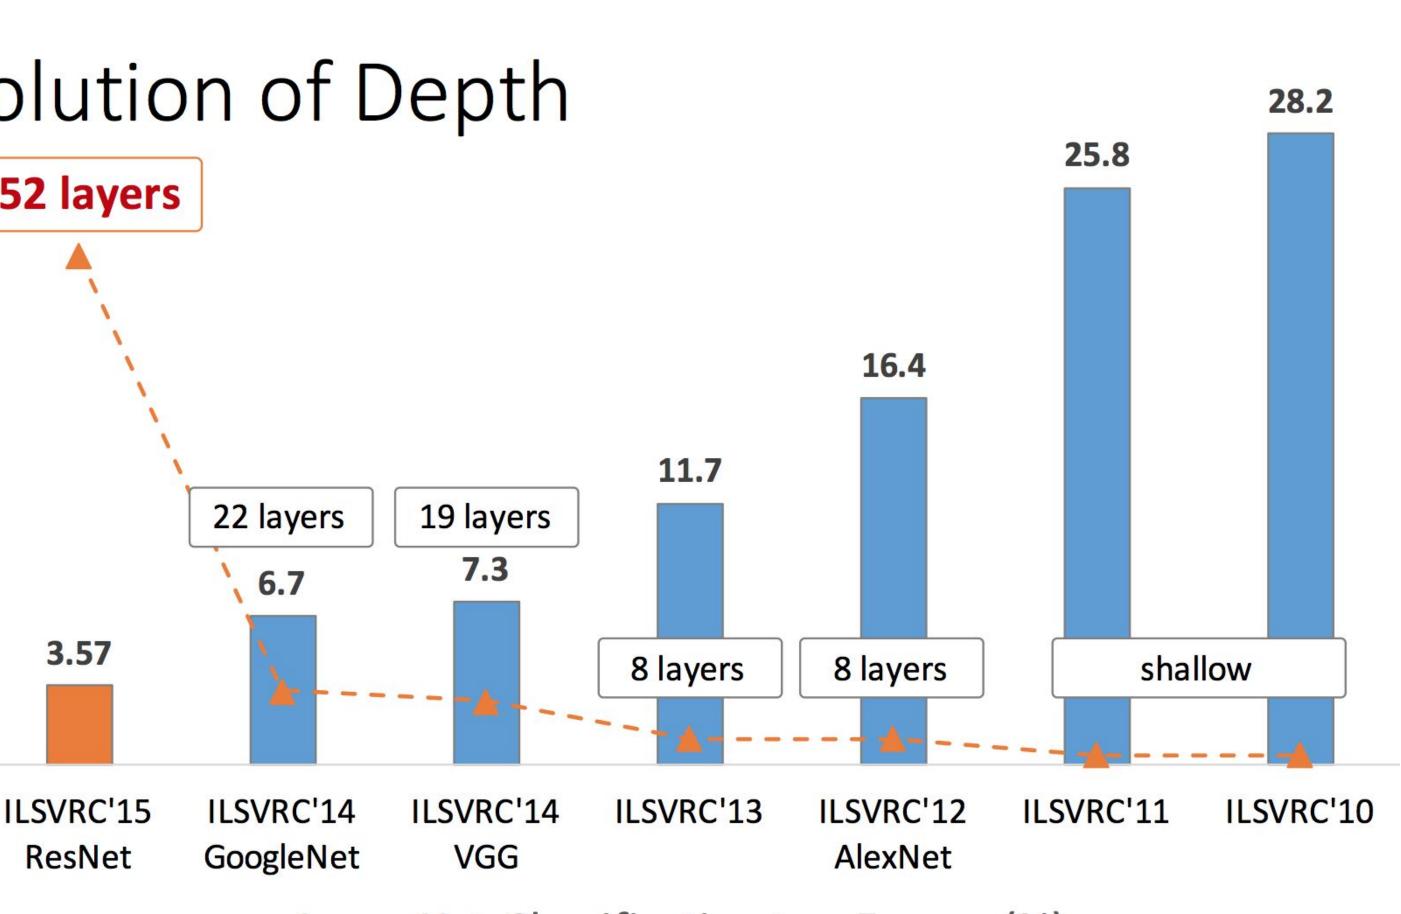




Convolutional Neural Network

First CNN was developed by Yann LeCun on 1988, called LeNet, but CNNs became popular when in 2012 AlexNet was the first CNN to win the ImageNet Large Scale Visual Recognition Challenge (ILSVCR). Since then, only DNN model where used (and won) the following editions.







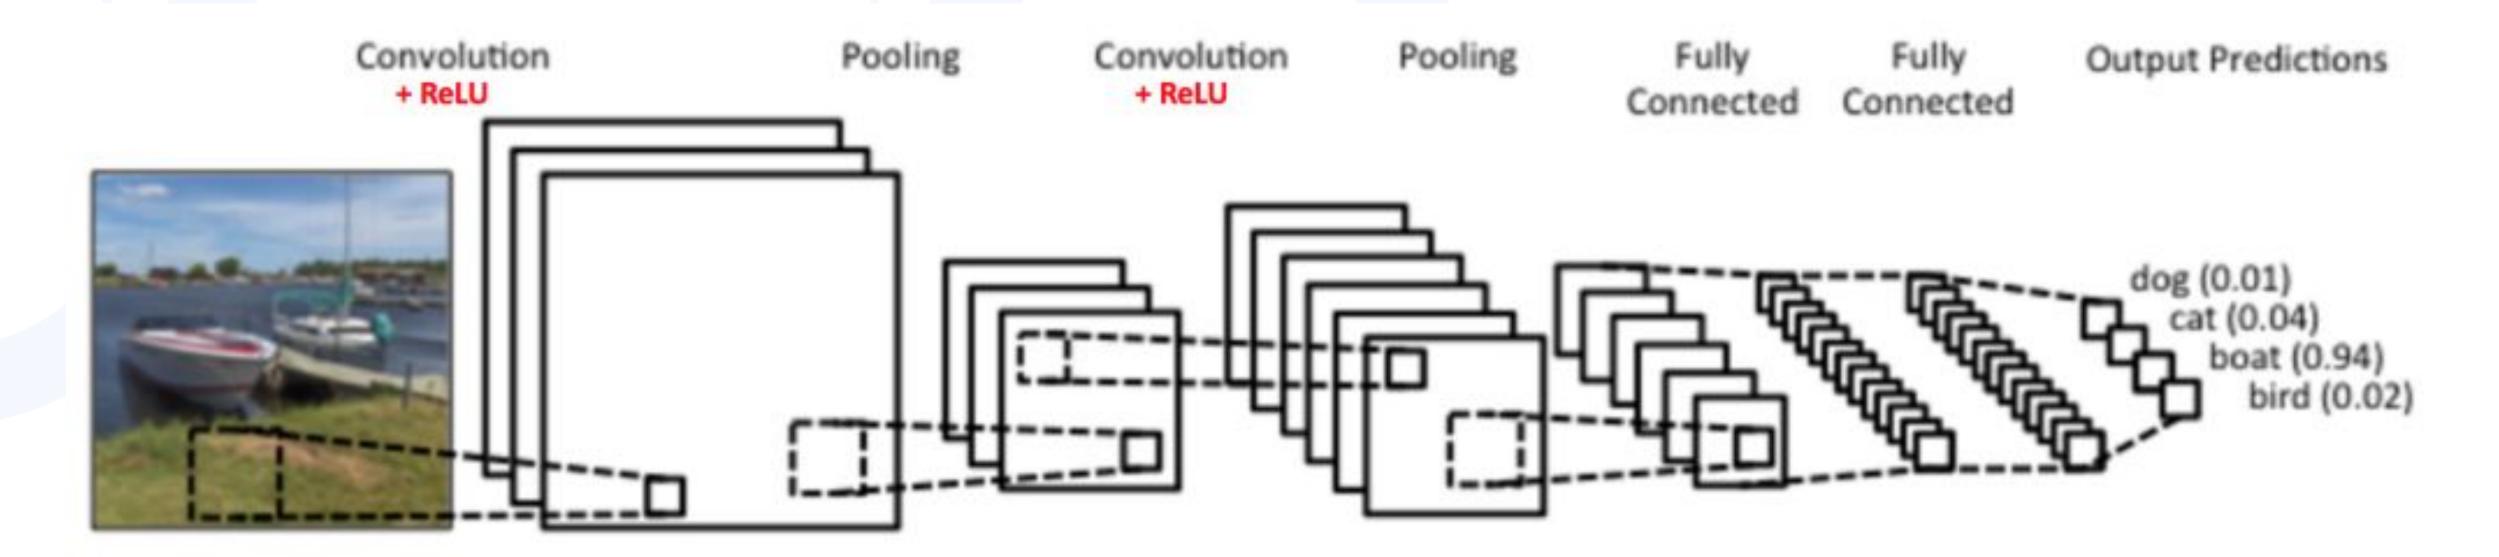


ImageNet Classification top-5 error (%)

Convolutional Neural Network

Key components of a CNN are the following:

- Convolution
- Non Linearity (activation function)
- Pooling or Sub-sampling
- Classification (fully connected layer) and training

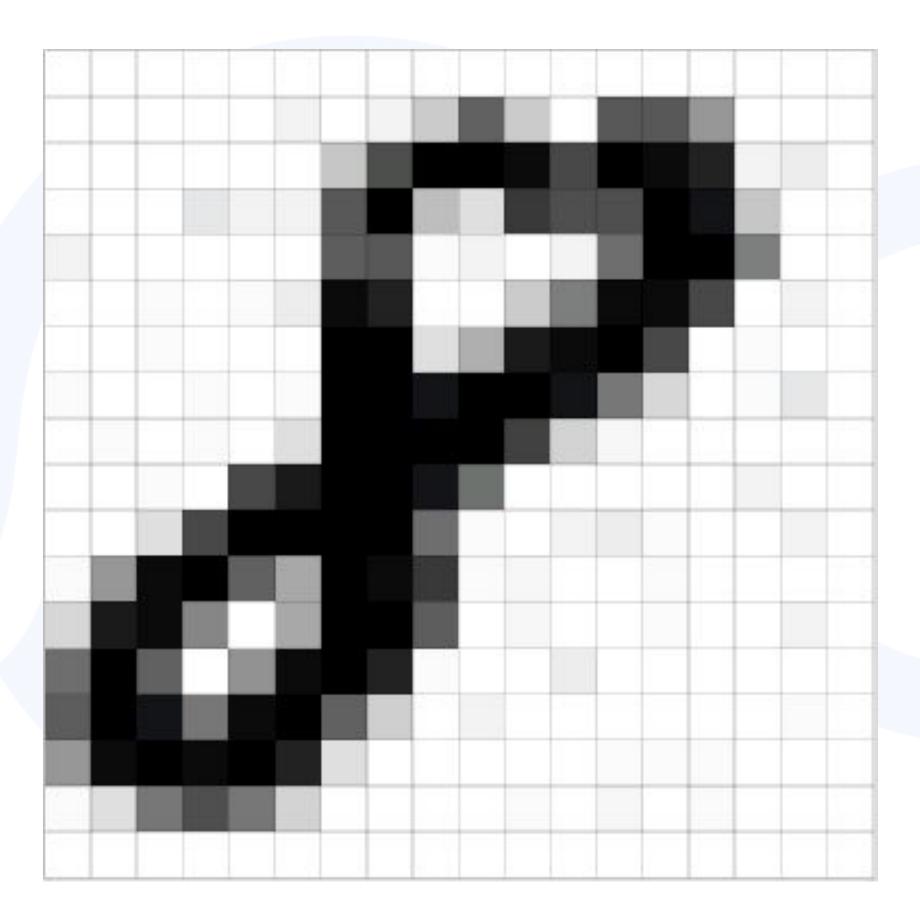




LeNet



Every image can be represented as matrices of pixels, one for each channel (RGB, HSV, etc)





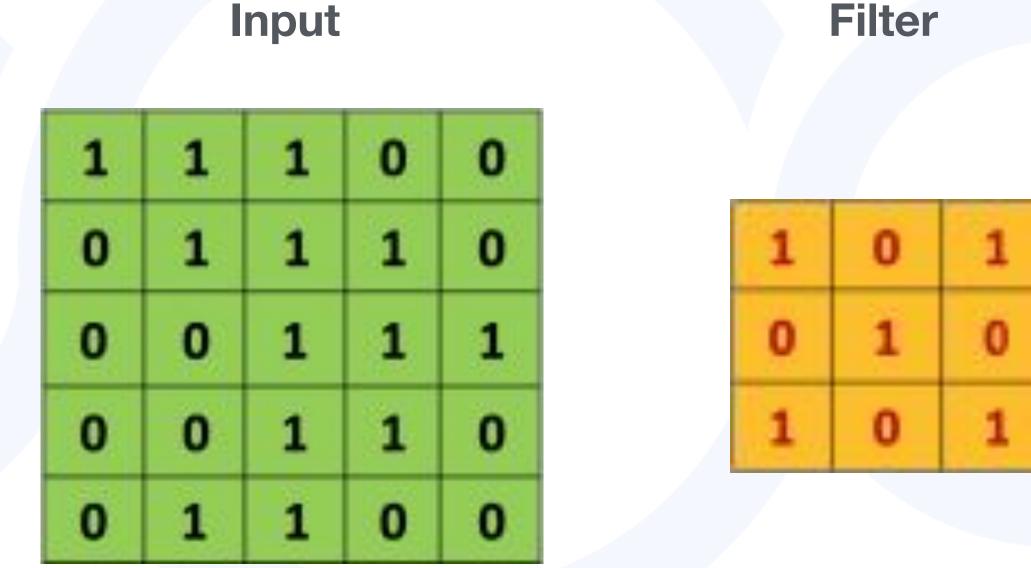
						<u>_</u>]
4	9	2	5	8	3	\mathbf{H}
5	6	2	4	0	3	
2	4	5	4	5	2	
5	6	5	4	7	8	
5	7	7	9	2	1	
5	8	5	3	8	4	

Input

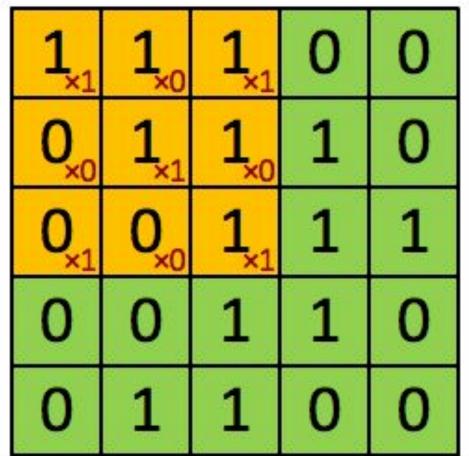
6 x 6 x 3

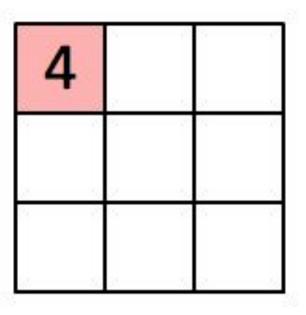
Convolution filter

We chose a **filter** (or **Kernel**) to be passed on the image. Every cell of the filter is multiplied elementwise with the corresponding area of each channel and then summed up. Outcome is called **Convolved Feature** or **Feature Map Filter Convolution filter** Input





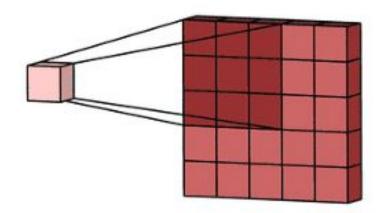


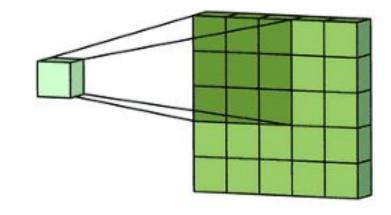


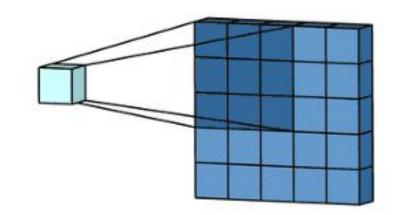
Image

Convolved Feature

Convolution filter - 3 channel example







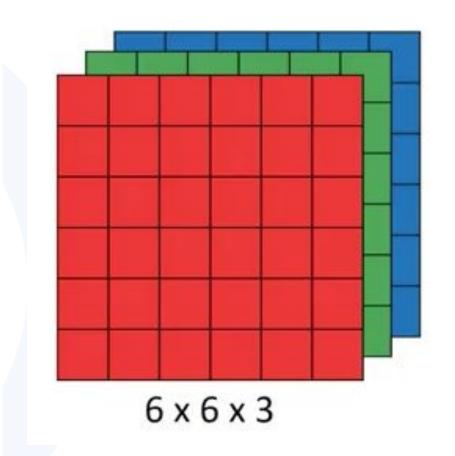
_		
_		

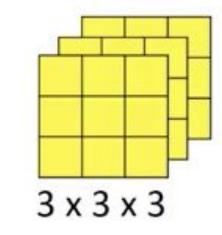
			٢
_	-	0	-
			1

-

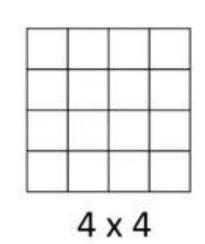








*

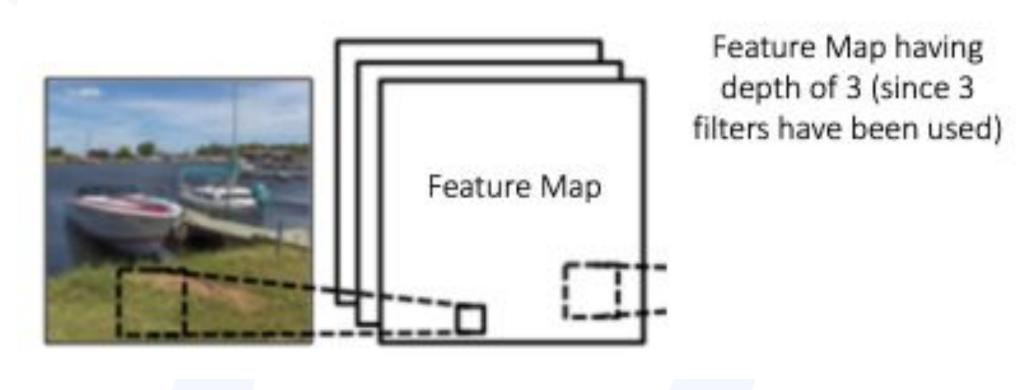


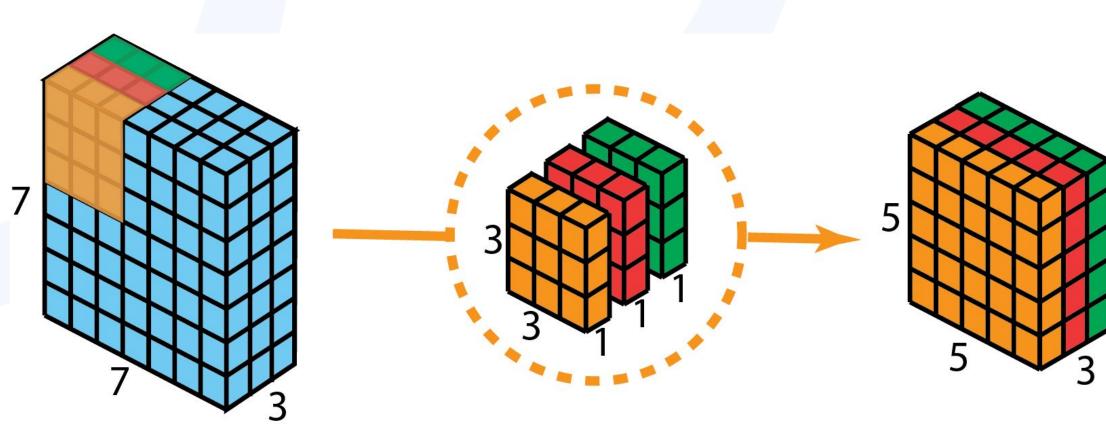
=

Convolution filter parameters

Each filter is characterized by the following parameters:

"features" of the images





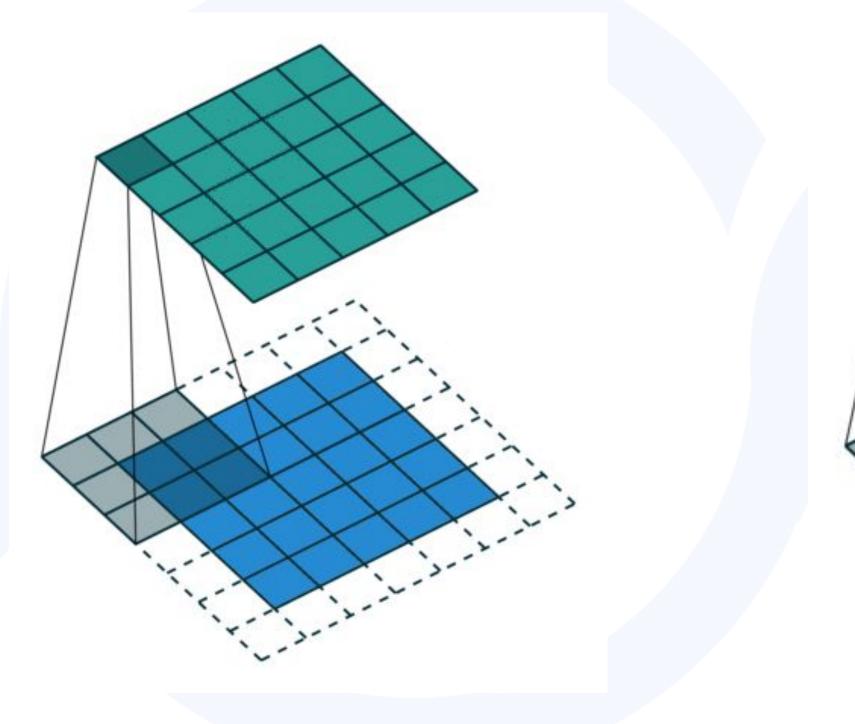




• **Depth:** number of distinct filters we use for the convolution operation. Multiple filters are used to detect different

Convolution filter parameters

maps

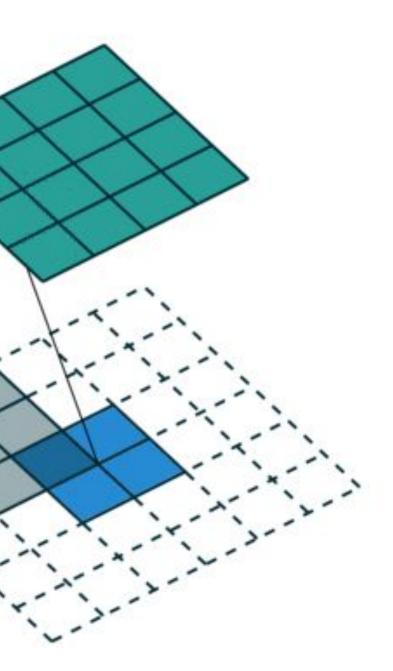


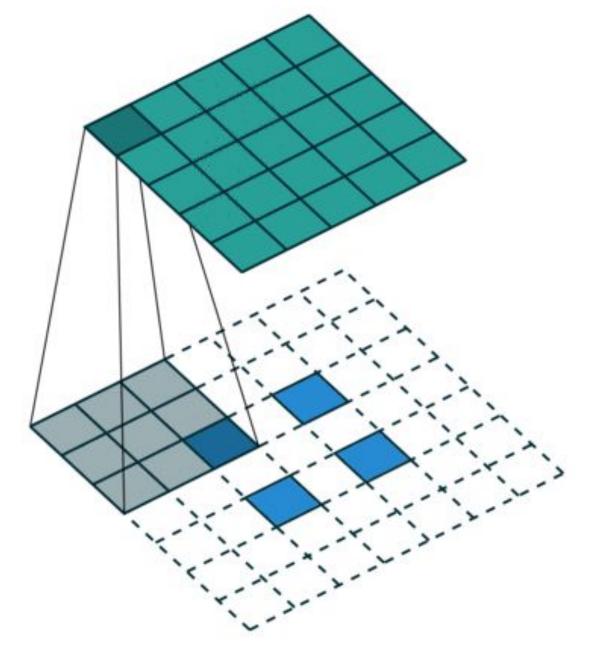
1-padding





• Zero-Padding: pad the input matrix with zeros around the border. it allows us to control the size of the feature





2-padding

2-padding with up-sampling

Convolution filter parameters

maps

0	0	0	0	0	0	
0	105	102	100	97	96	
0	103	99	103	101	102	2
0	101	98	104	102	100	
0	99	101	106	104	99	7
0	104	104	104	100	98	

Ke	rnel Ma	atrix
0	-1	0
-1	5	-1
0	-1	0

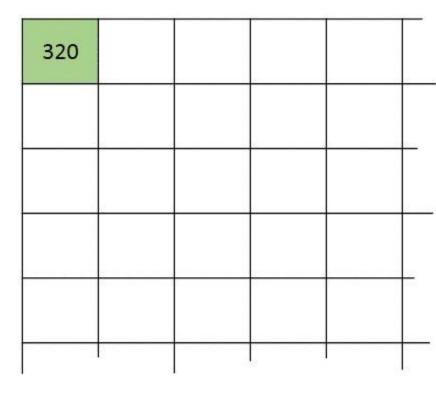


Image Matrix

0 * 0 + 0 * -1 + 0 * 0+0 * -1 + 105 * 5 + 102 * -1+0 * 0 + 103 * -1 + 99 * 0 = 320

Output Matrix

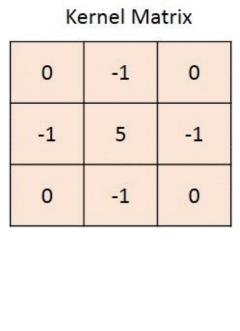
Convolution with horizontal and vertical strides = 1





• Stride: number of pixels by which we slide our filter matrix. Having a larger stride will produce smaller feature

0	0	0	0	0	0	0
0	105	102	100	97	96	
0	103	99	103	101	102	10
0	101	98	104	102	100	ì
0	99	101	106	104	99	ŕ
0	104	104	104	100	98	ĩ



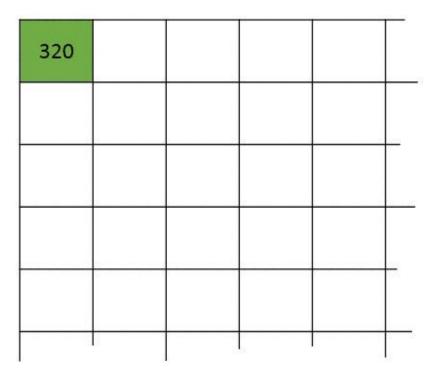


Image Matrix

0 * 0 + 0 * -1 + 0 * 0+0 * -1 + 105 * 5 + 102 * -1+0 * 0 + 103 * -1 + 99 * 0 = 320

Output Matrix

Convolution with horizontal and vertical strides = 2

Classic Computer Vision filters

Edge

Classic CV filters are set by the model designer and are "experience based", depending on the context of the images and the task to be achieved.



Input

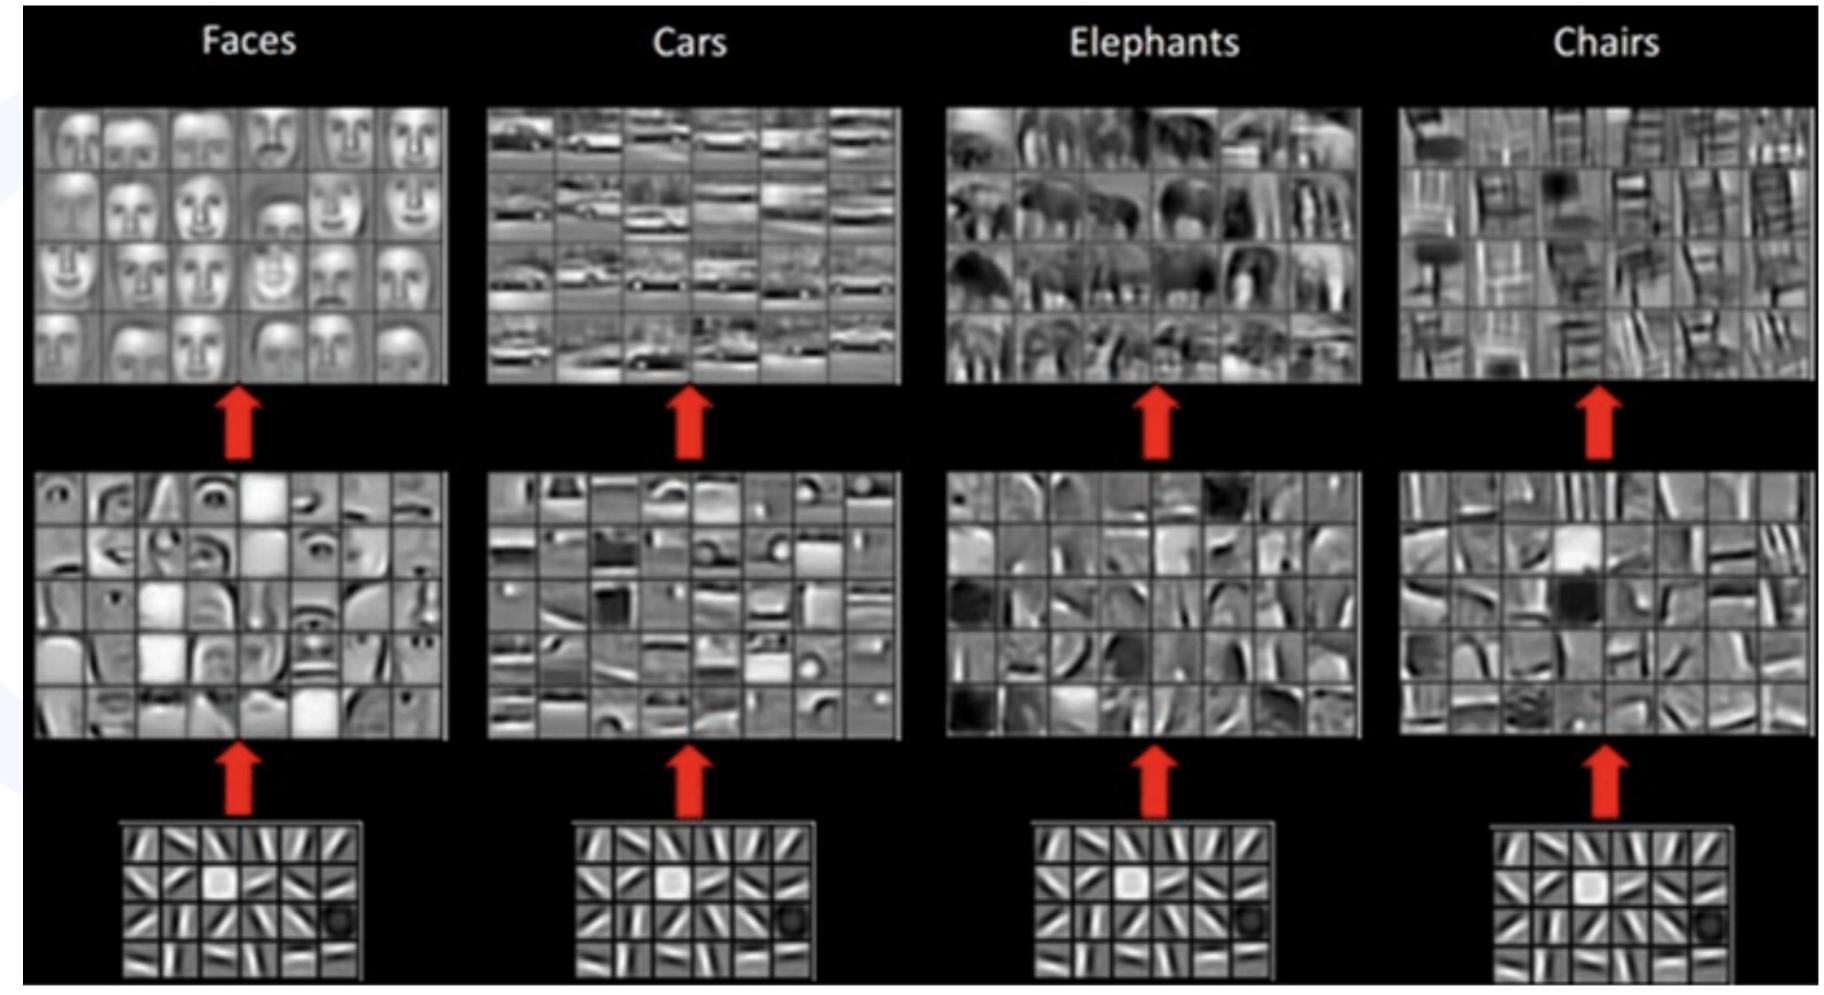


Operation	Filter	Convolved Image	Operation	Filter	Convol Imag
dge detection	$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{bmatrix}$		Sharpen	$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$	
	$\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$		Box blur (normalized)	$\frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$	
	$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$		Gaussian blur (approximation)	$\frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$	



CNN learned filters

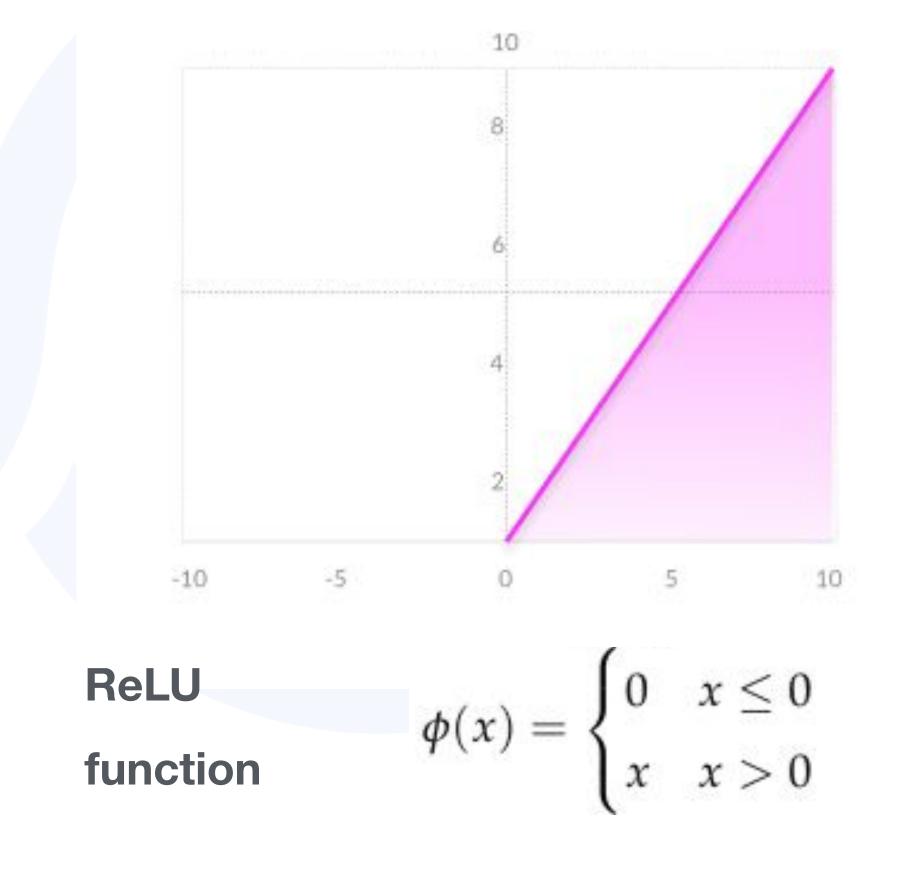
CNN filters are learned by the network itself, surprisingly identifying understandable context features



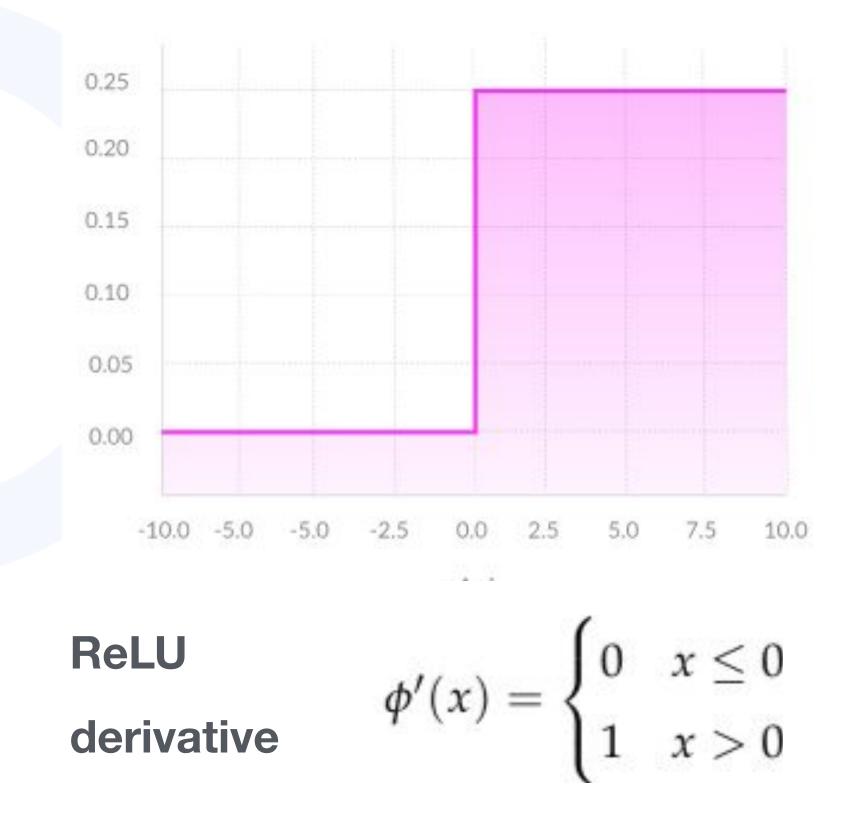


Non linearity

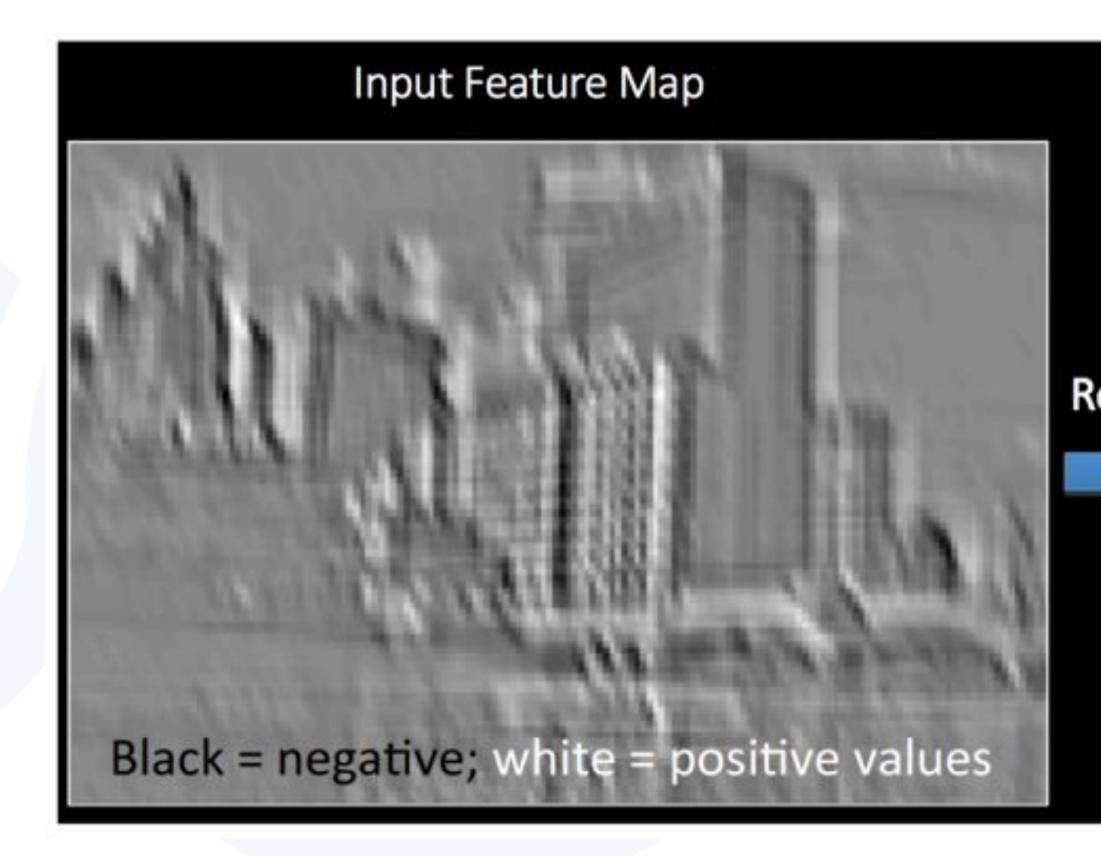
A commonly used activation function is the **Rectified Linear Unit** (ReLU), a non-linear function and element wise operation (applied per pixel) that replaces all negative pixel values in the feature map by zero.







Non linearity



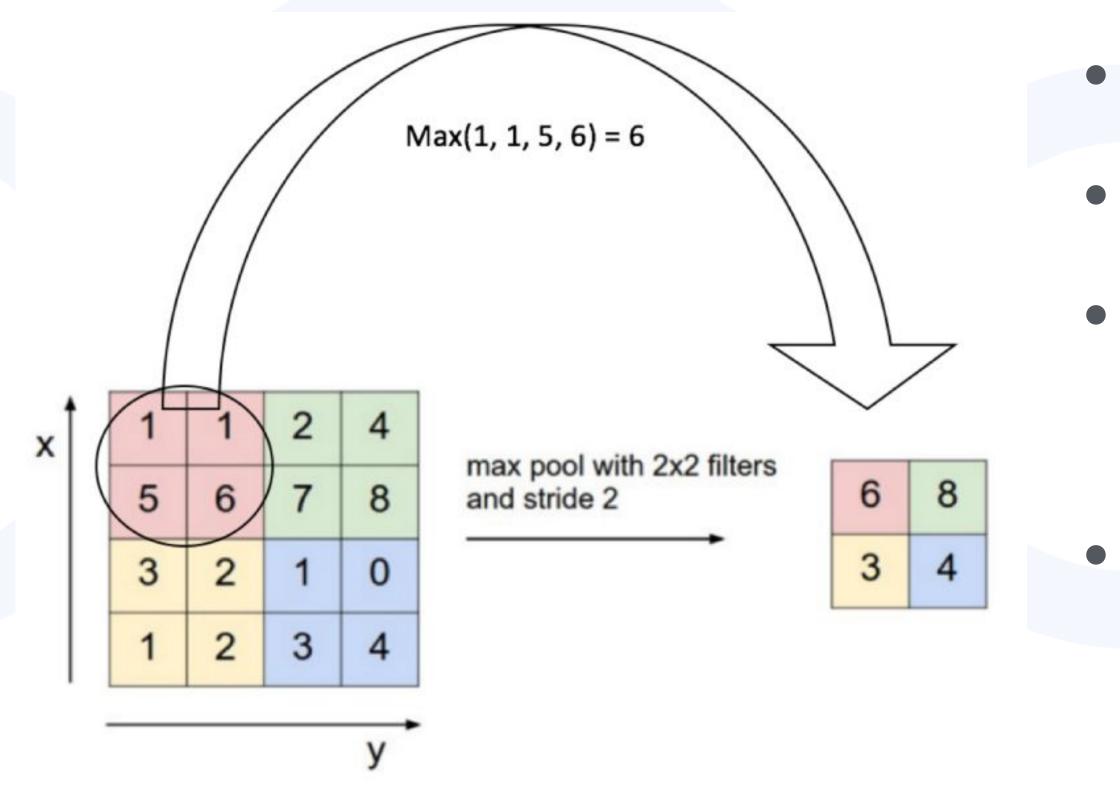


Rectified Feature Map



Pooling

Spatial Pooling (also called subsampling or downsampling) reduces the dimensionality of each feature map but retains the most important information. Spatial Pooling can be of different types: Max, Average, Sum etc.



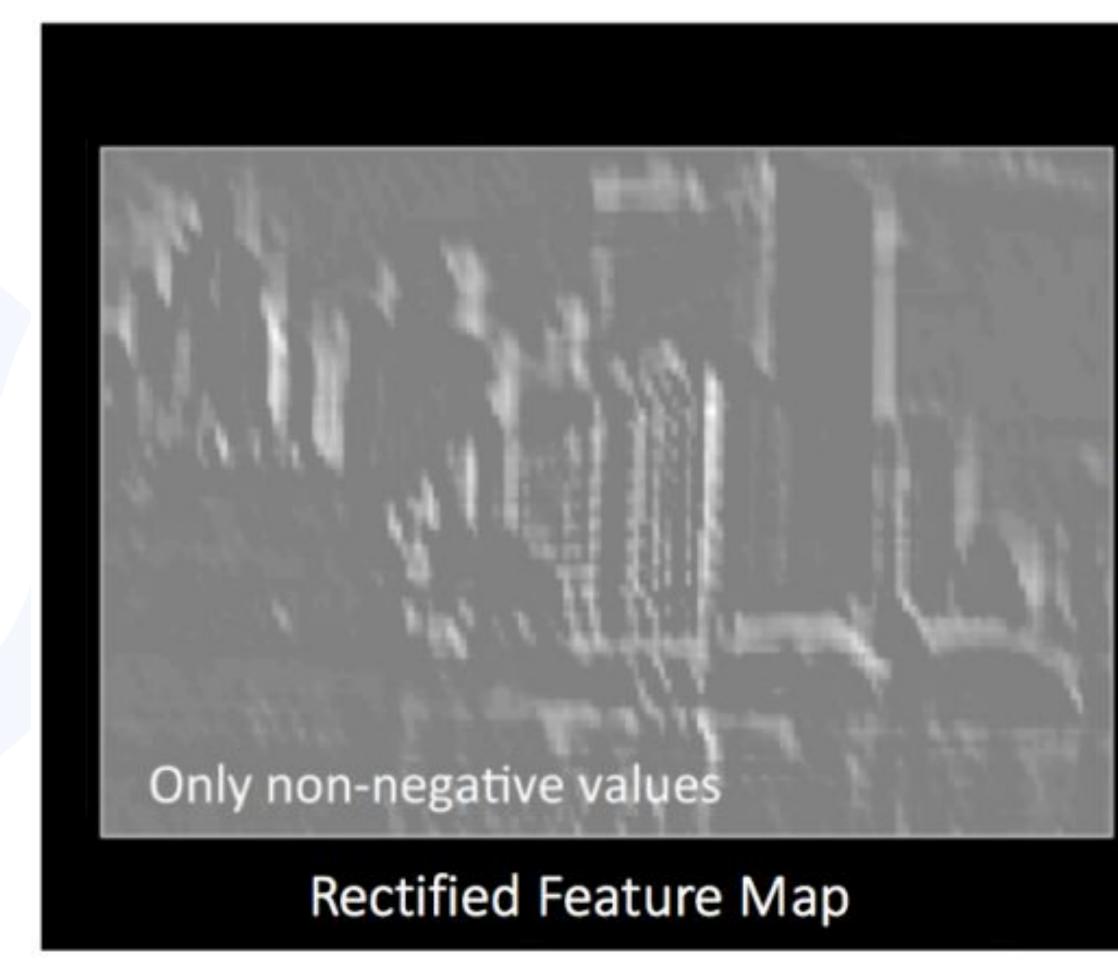
Rectified Feature Map



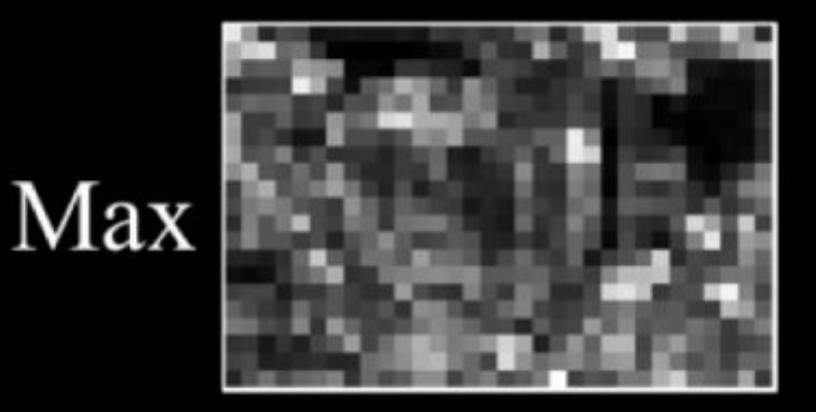
- makes the **input representations** (feature dimension) **smaller** and more manageable
- reduces the number of parameters and computations in the network
- makes the network invariant to small transformations, distortions and translations in the input image (a small distortion) in input will not change the output of Pooling – since we take the maximum / average value in a local neighborhood) • helps to arrive at an almost scale invariant (equivariant) **representation** of our image. This is very powerful since we can detect objects in an image no matter where they are located



Pooling



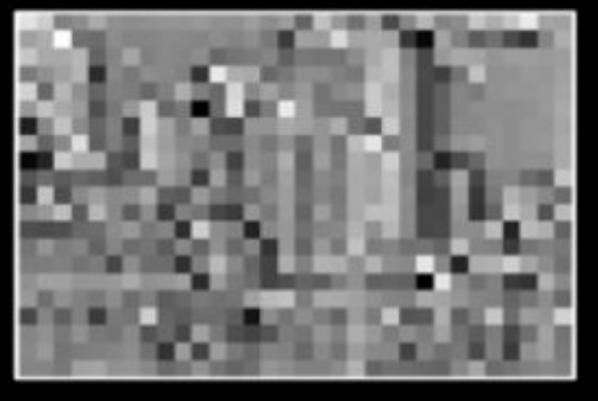




Pooling

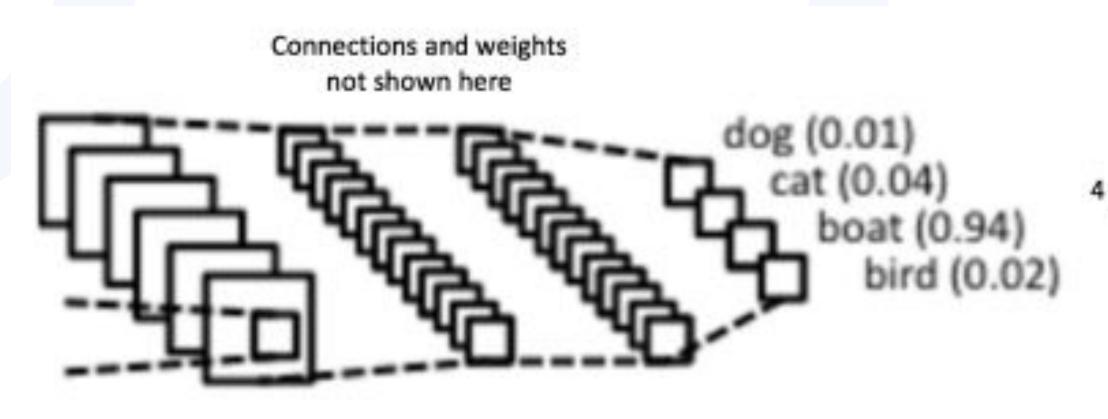


Sum



Training and loss function

- The Fully Connected layer is a traditional Multi Layer Perceptron that uses a **Softmax** activation function in the output layer, flattening the output of convolutional and pooling layers
- The output from the convolutional and pooling layers represent high-level features of the input image
- The purpose of the Fully Connected layer is to use these features for classifying the input image into various classes based on the training dataset.
- This is also a cheap way of learning non-linear combinations of these features. Most of the features from might be even better





convolutional and pooling layers may be good for the classification task, but combinations of those features

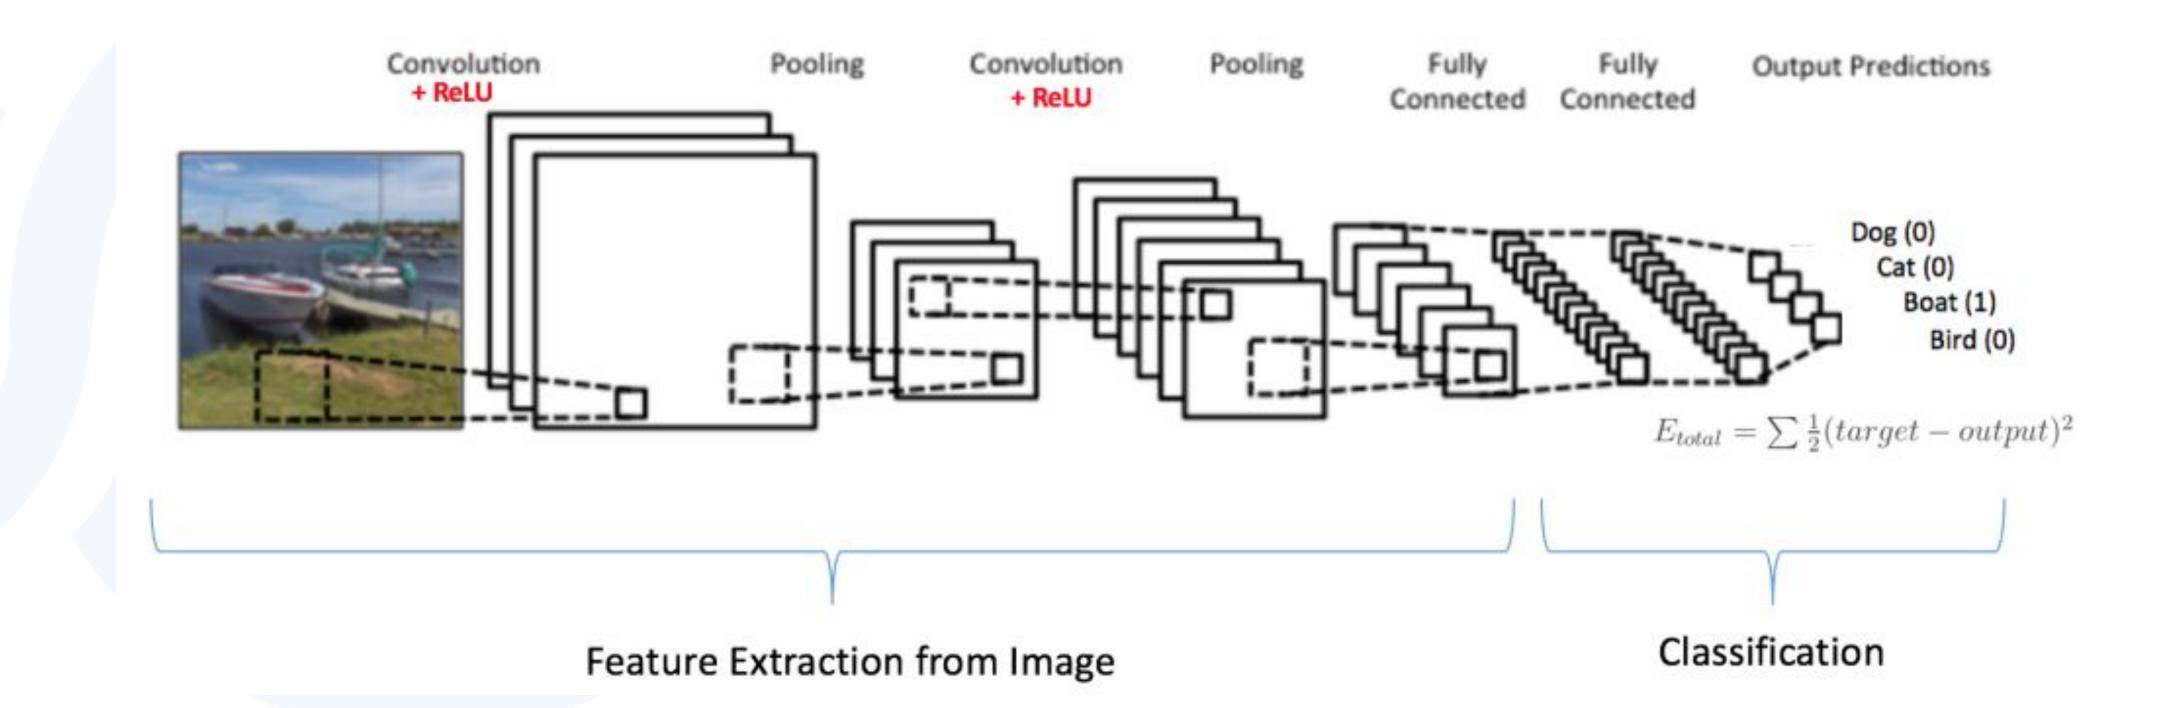
4 possible outputs

 $exp(x_i)$ $Softmax(x_i)$



Training and loss function

Now we have all the building blocks to train our neural network





Training and loss function

Training (tuning of the weights) consist of the following steps:

- 1) initialize all filters and parameters (weights) with random values
- 2) OP_i for each class (normalized with the softmax)
- output ones. Two commonly used metrics are:

Mean Squared Error

$$\frac{1}{n}\sum_{i=1}^{n}(TP_i - OP_i)^2$$

4) Use Backpropagation to calculate the gradients of the error with respect to all weights in the network and use gradient descent to update all weights and parameter values to minimize the output error

Repeat steps 2-4 with all images in the training set 5)



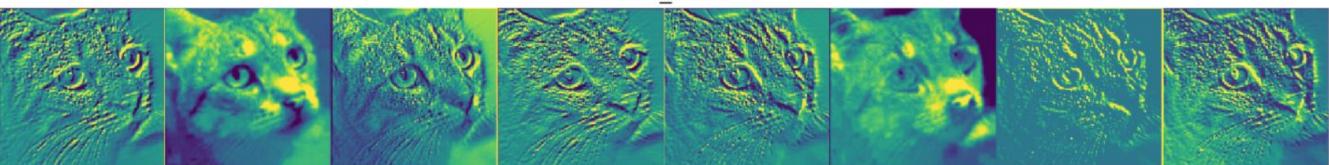
The network takes a training image as input, goes through the forward propagation step (convolution, ReLU and pooling operations along with forward propagation in the Fully Connected layer) and finds the output probabilities

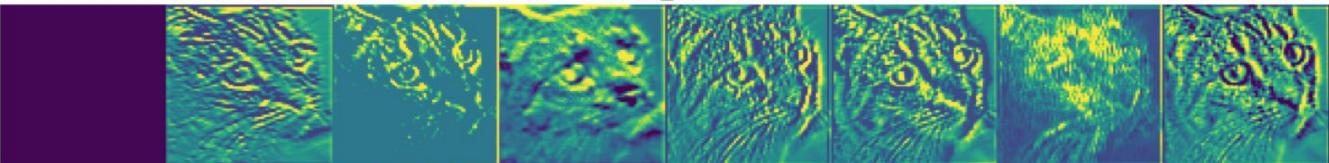
3) Calculate the total error (Loss Function) at the output layer comparing the target probabilities TP_i with the

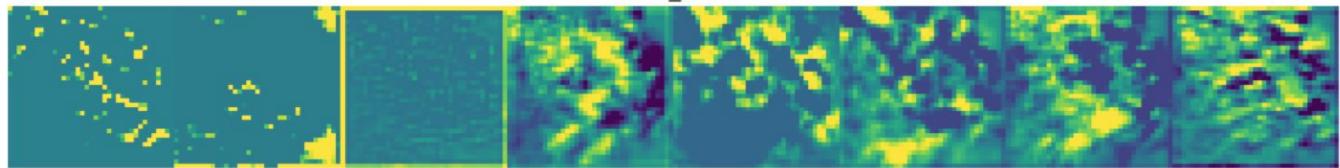
Cross-Entropy
$$-\sum_{i=1}^{n} TP_i \log(OP_i)$$

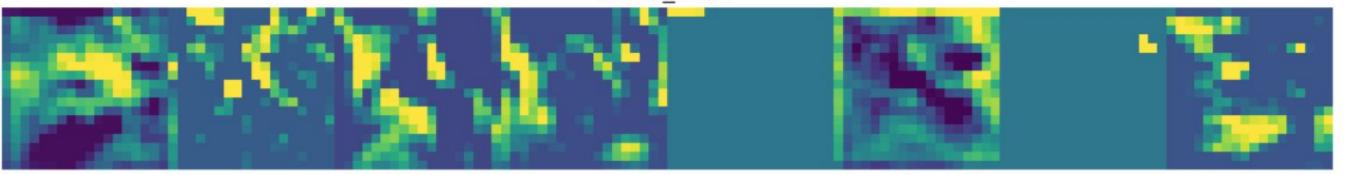


Visualizing CNN













block1_conv1

block2_conv1

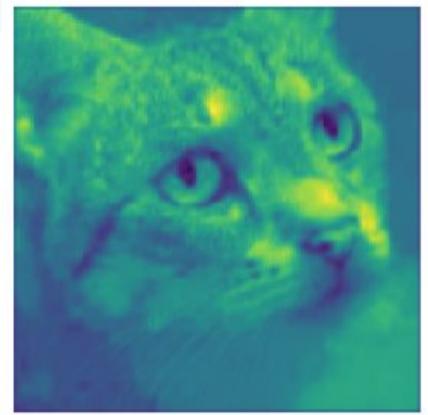
block3_conv1

block4_conv1

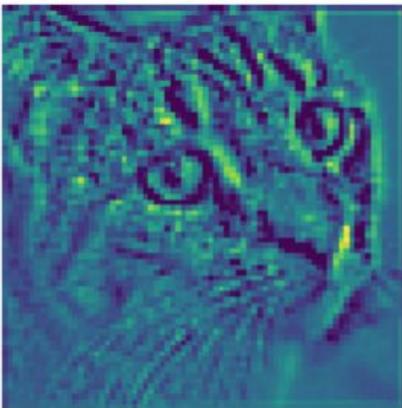
block5_conv1

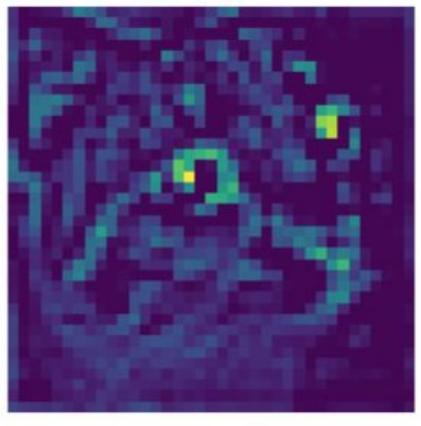
Visualizing CNN

block1_conv1



block2_conv1

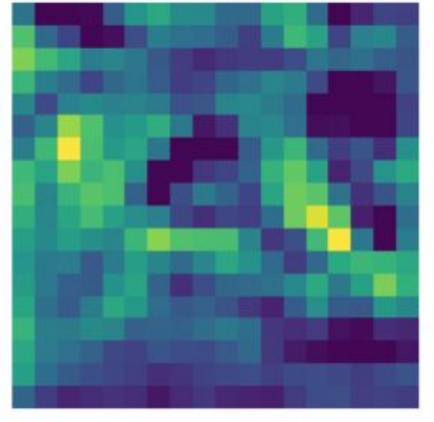


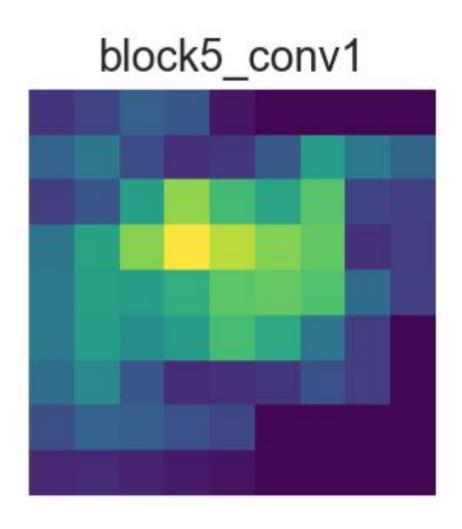




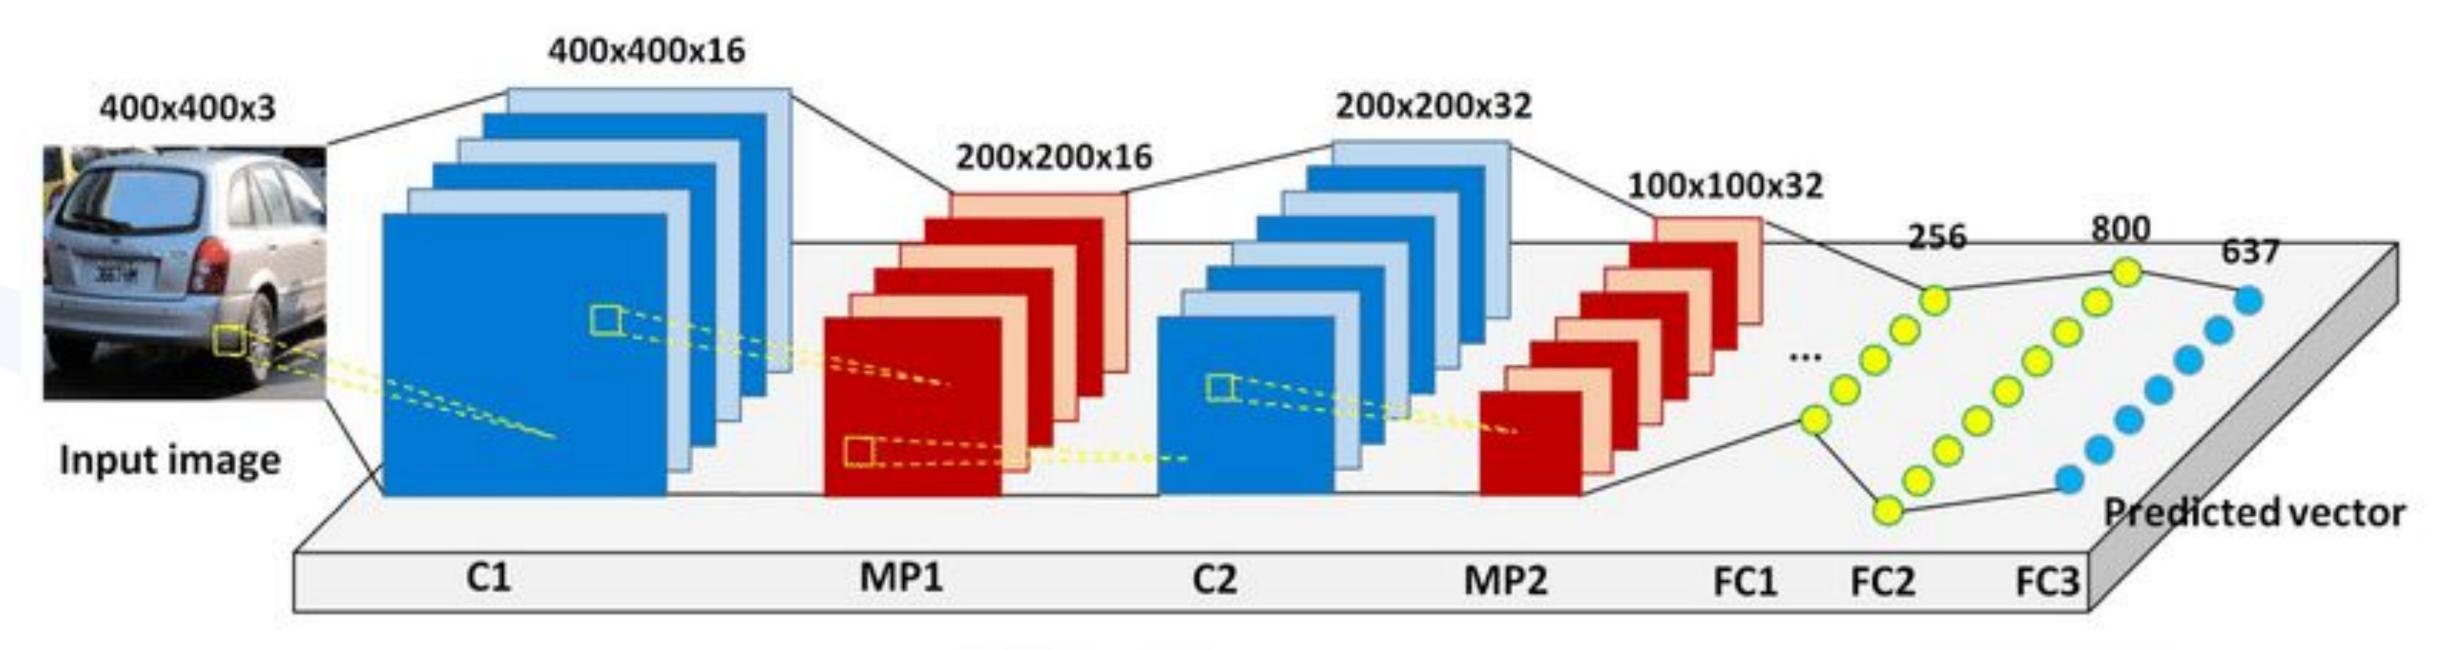
block3_conv1

block4_conv1





Visualizing CNN





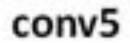
conv1

conv2

conv3



conv4





Useful implementation tips

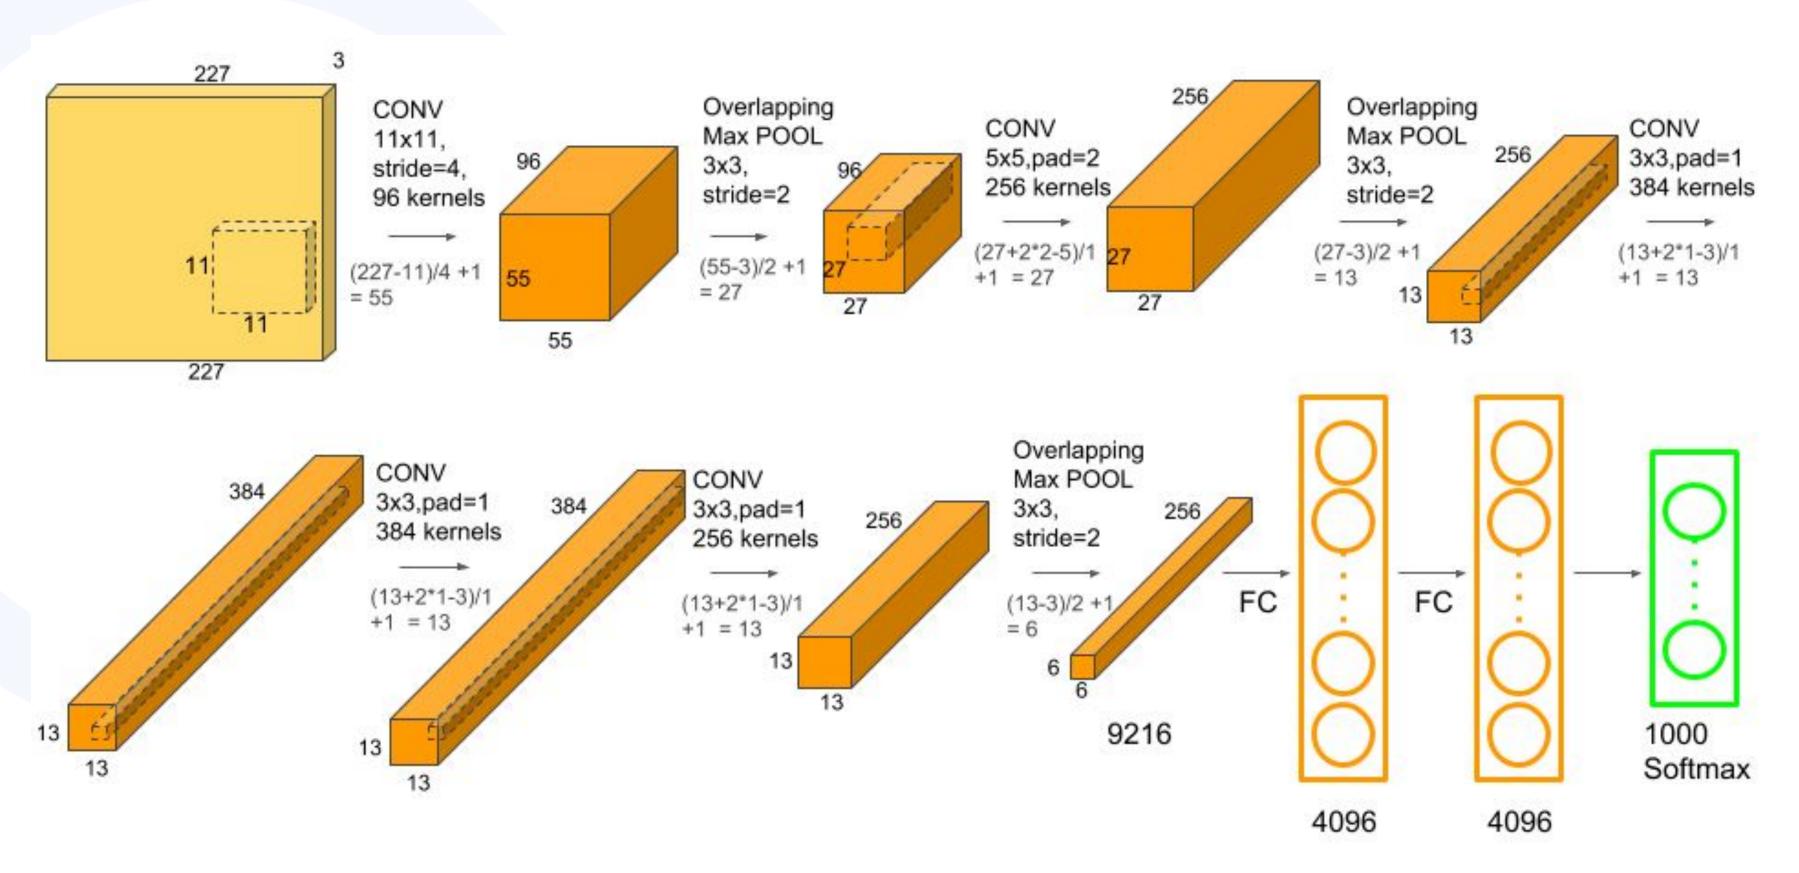
Overfitting is a common problem for all Neural Network. Following tips may prevent the problem and speed up training time:

- Preprocessing:
 - Local Mean Subtraction
 - Normalization
- Better optimization methods:
 - (Batch) Stochastic Gradient Descent (SGD)
 - ADAptive Moment (ADAM) that combines momentum and RMSprop
- Regularization
 - Weight decay
 - Dropout
- Data augmentation



CNN Architectures: AlexNet (Alex Krizhevsky - 2012)

- six days to train on two GTX 580 3GB GPUs.
- consists of 5 Convolutional Layers and 3 Fully Connected Layers





• AlexNet was much larger than previous CNNs. It has 60 million parameters and 650,000 neurons and took five to



CNN Architectures: ZFnet (Zeiler & Fergus - 2013)

- Before this model CNN were black boxes. This mod internal representations
- Main idea is to improve AlexNet introducing Deconvectors
 convolution and Unpooling (inverse of pooling)

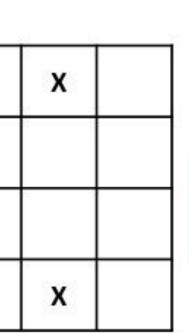
Unpooling

0.1	0.5	1.2	-0.7					
0.8	-0.2	-0.5	0.3	max-pooling	0.8	1.2		
0.4	0.9	-0.1	-0.2		0.9	0.5		
-0.6	0.1	0.5	0.3				x	
								х
0	0	0.5	0					
1.3	0	0	0		1.3	0.5	ma	x lo
0	0.4	0	0	unpooling	0.4	0.1		
0	0	0.1	0					



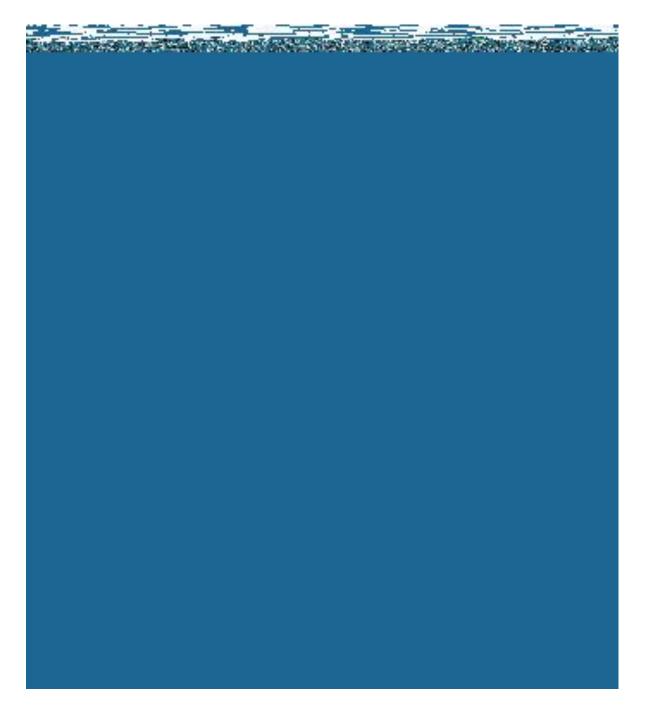
• Before this model CNN were black boxes. This model provides insights into how CNNI networks are learning

• Main idea is to improve AlexNet introducing **DeconvNet**, a deconvolutional net that acts as the opposite of



ocations

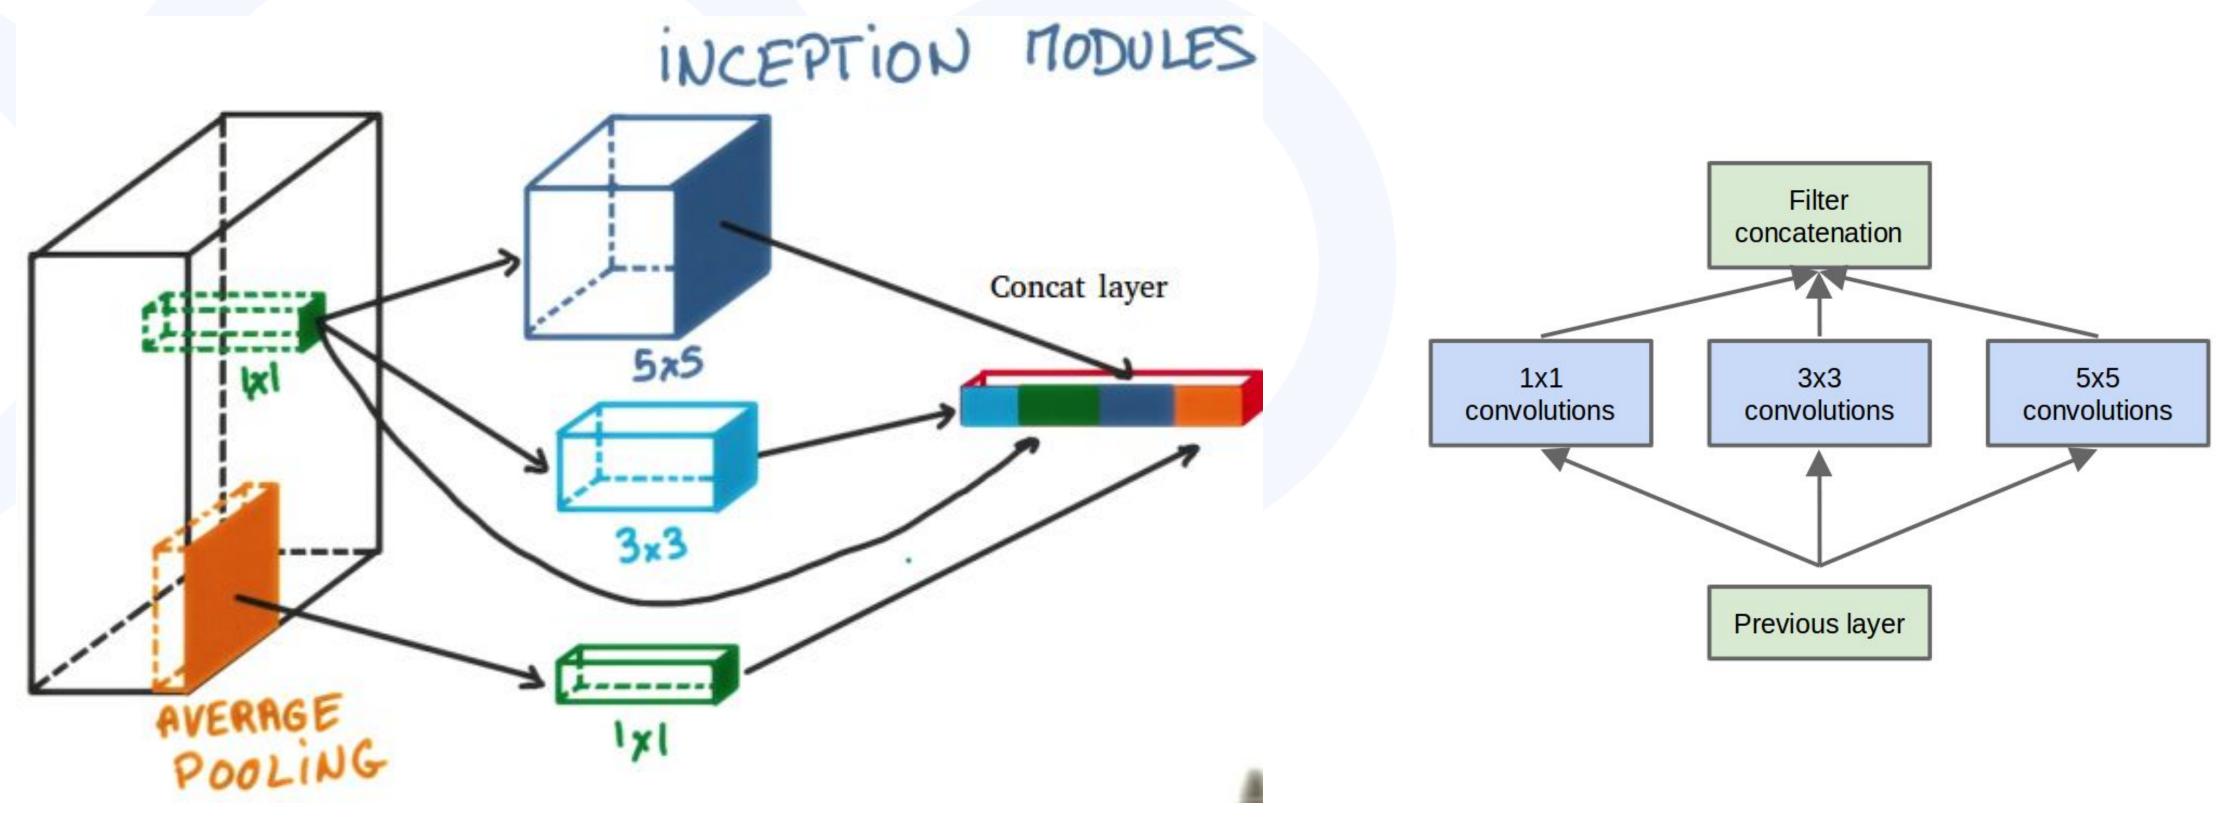
Deconvolution



Blue is input, cyan is output

CNN Architectures: GoogLeNet (2014)

- Introduced Inception layer, convolving in parallel dif one (5x5)
- The idea is that a series of filters with different sizes, that all filters on the inception layer are learnable.



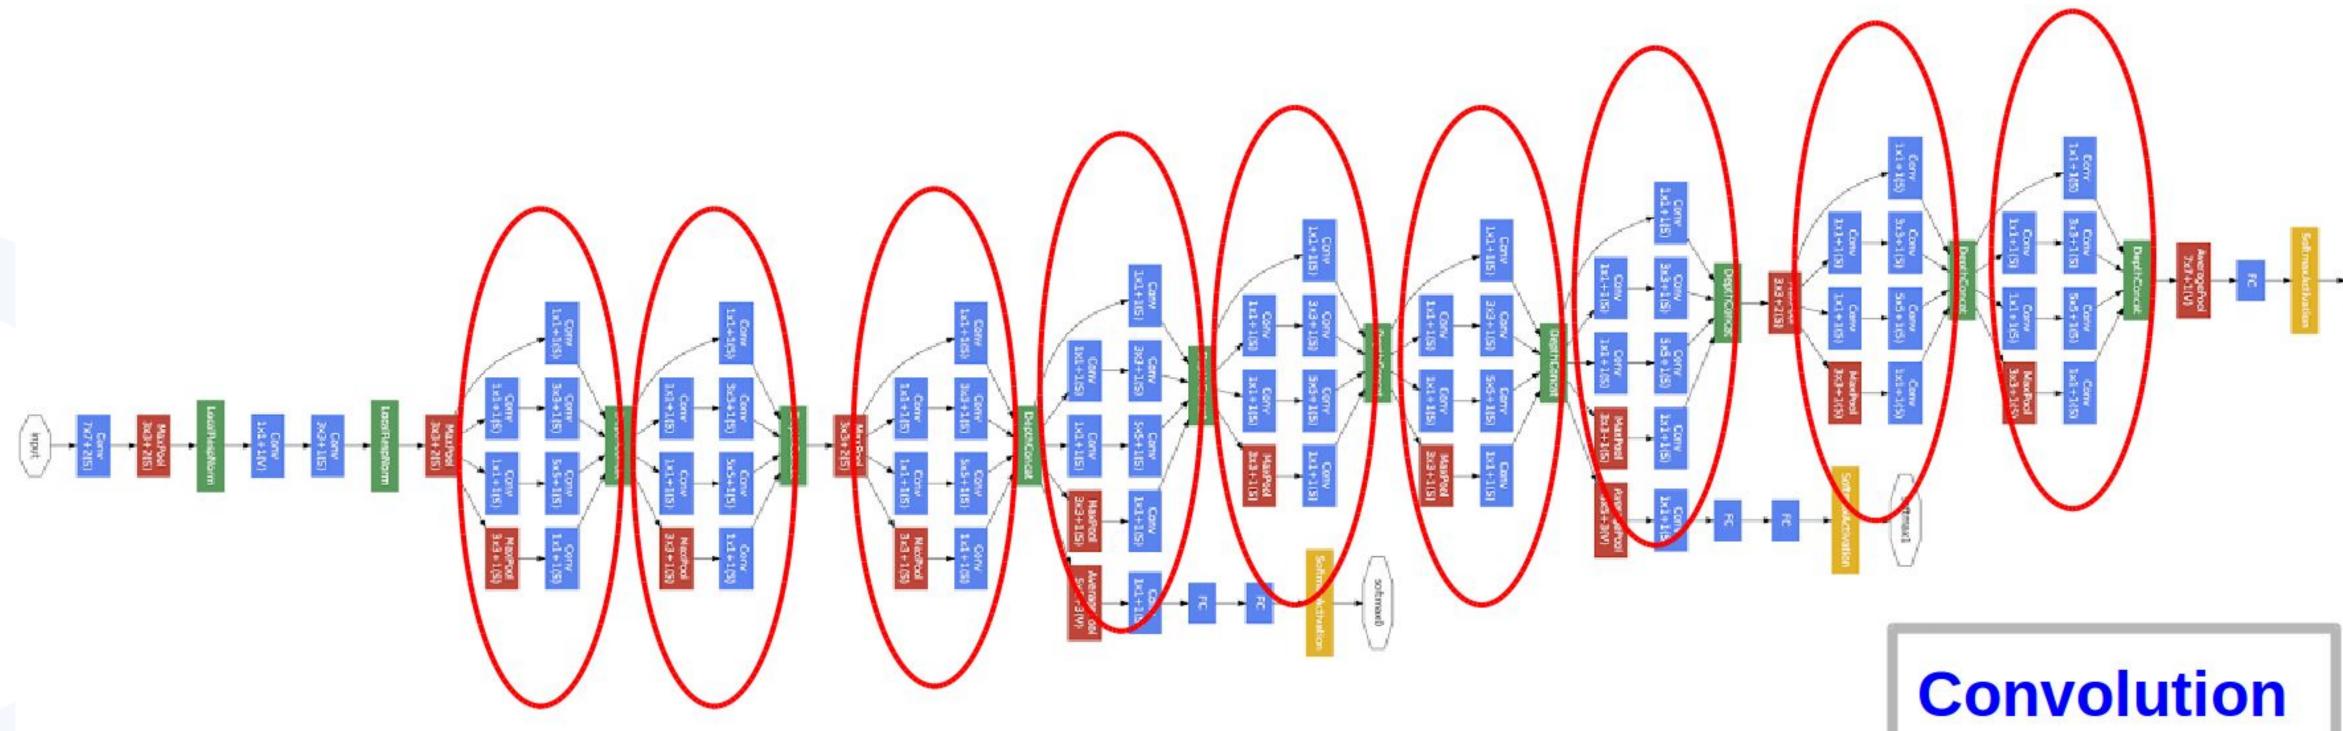


• Introduced Inception layer, convolving in parallel different sizes from the most accurate detailing (1x1) to a bigger

• The idea is that a series of filters with different sizes, will handle better multiple objects scales with the advantage



CNN Architectures: GoogLeNet (2014)





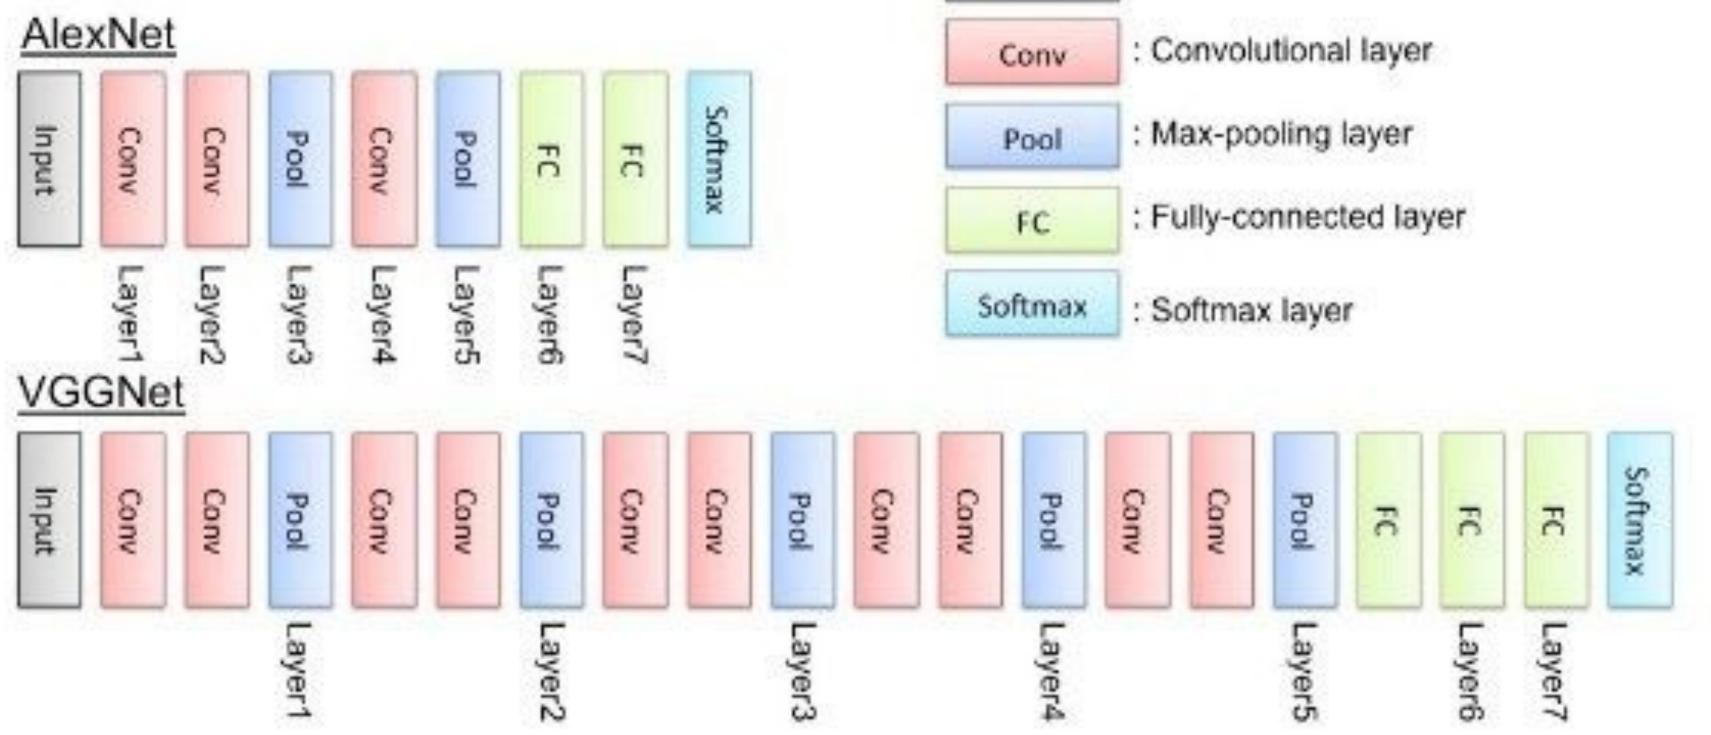
Convolution Pooling Softmax

Concat/Normalize



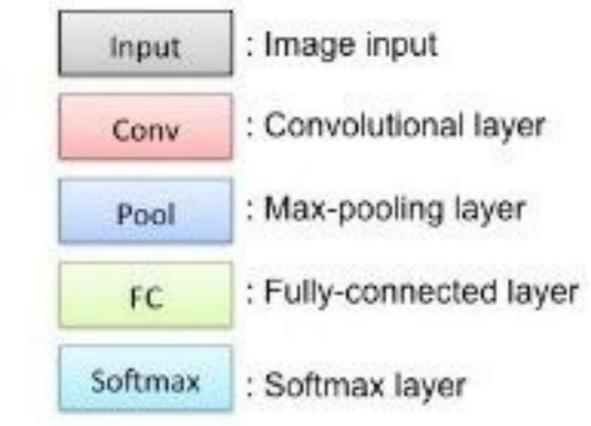
CNN Architectures: VGGNet (2014)

- Improved AlexNet using more convolutional filter blocks but with smaller size
- good performance



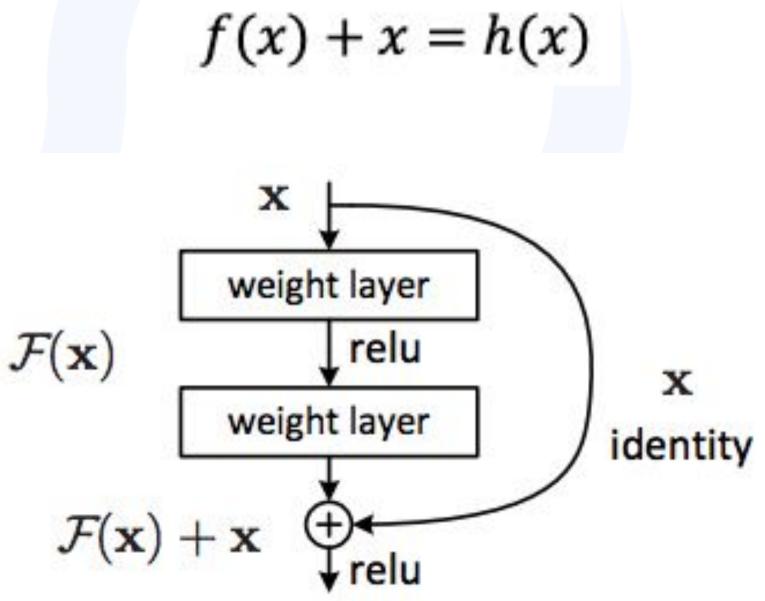


Main contribution was in showing that the depth of the network (number of layers) is a critical component for



CNN Architectures: ResNets (2015)

- Faces the **vanishing gradient** problem, allowing to increase the number of layers
- the output of a function becomes the input itself
- it





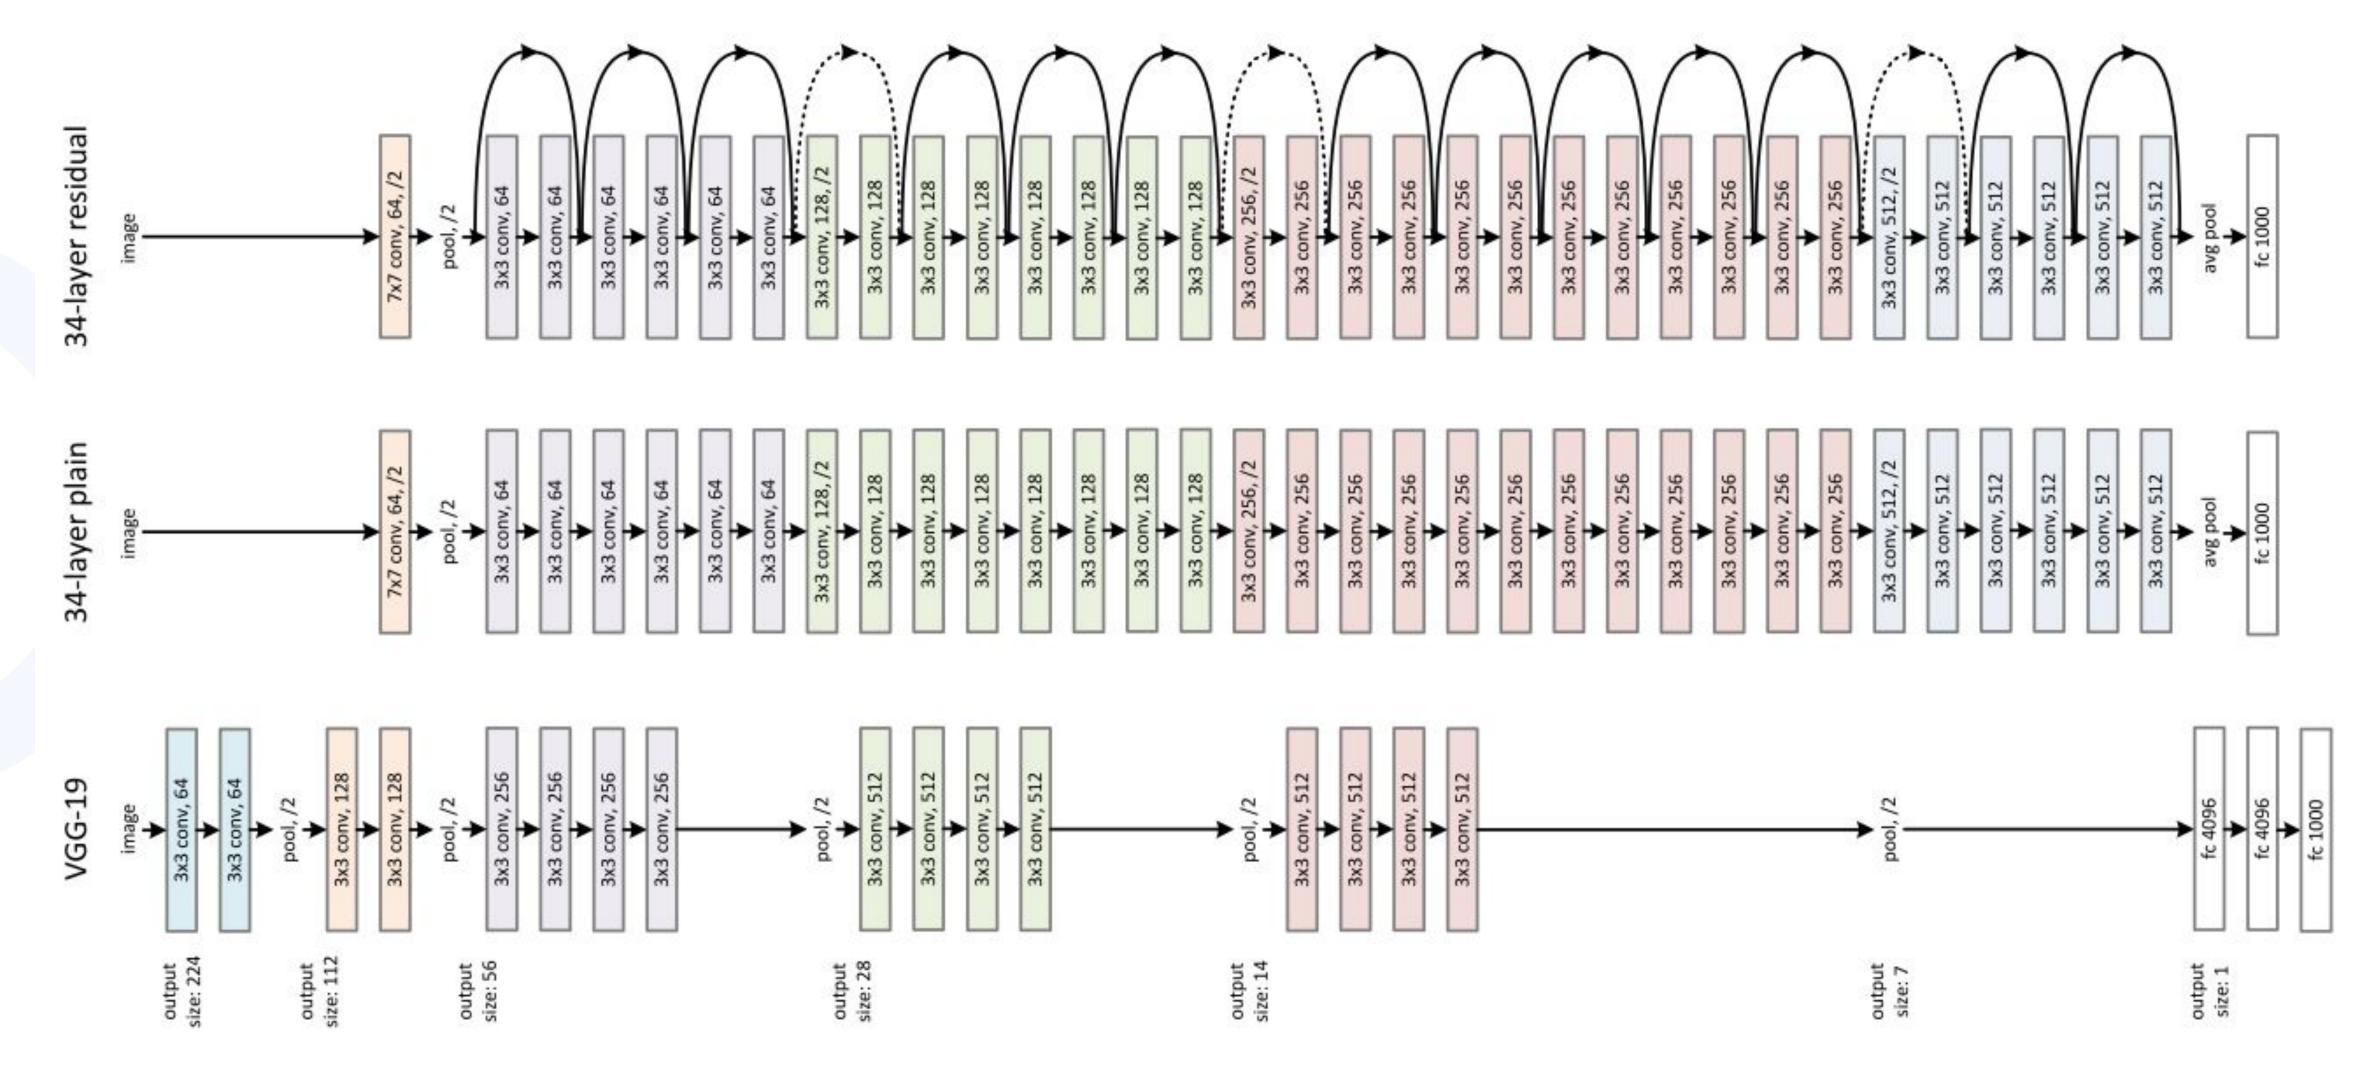
• Neural networks are good function approximators, they should be able to easily solve the identify function, where

f(x) = x

• Following the same logic, if we bypass the input to the first layer of the model to be the output of the last layer of the model, the network should be able to predict whatever function it was learning before with the input added to



CNN Architectures: ResNets (2015)

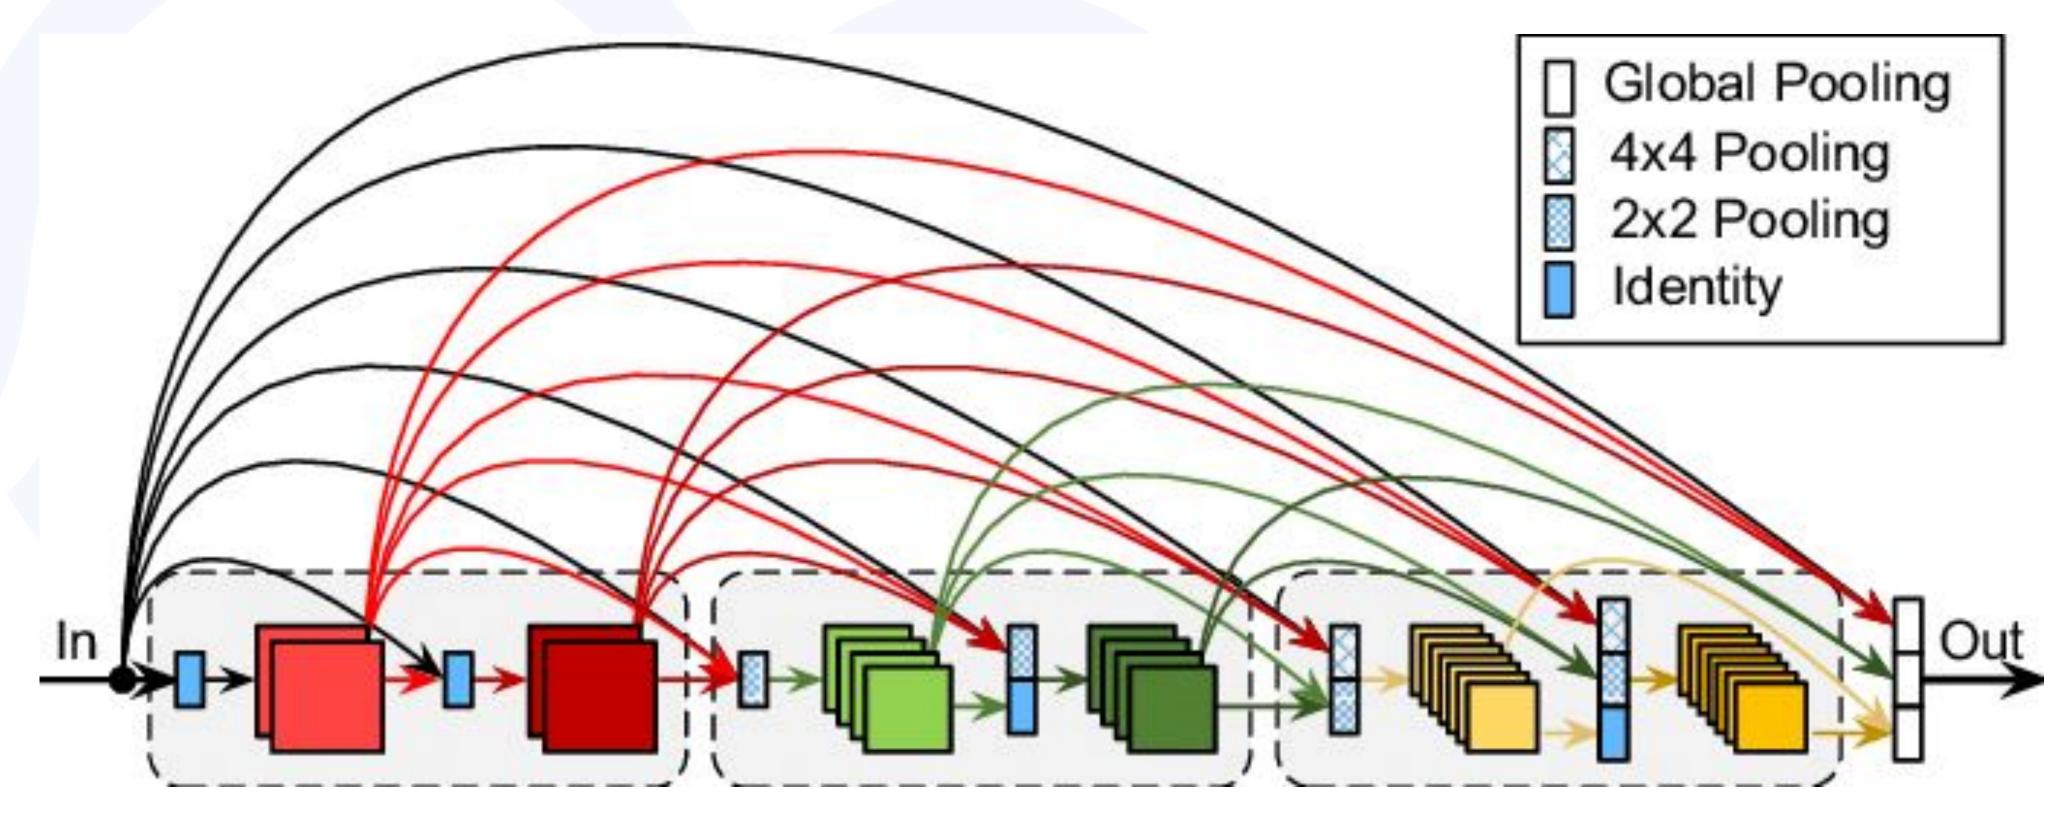


Neosperience



CNN Architectures: DenseNet (2016)

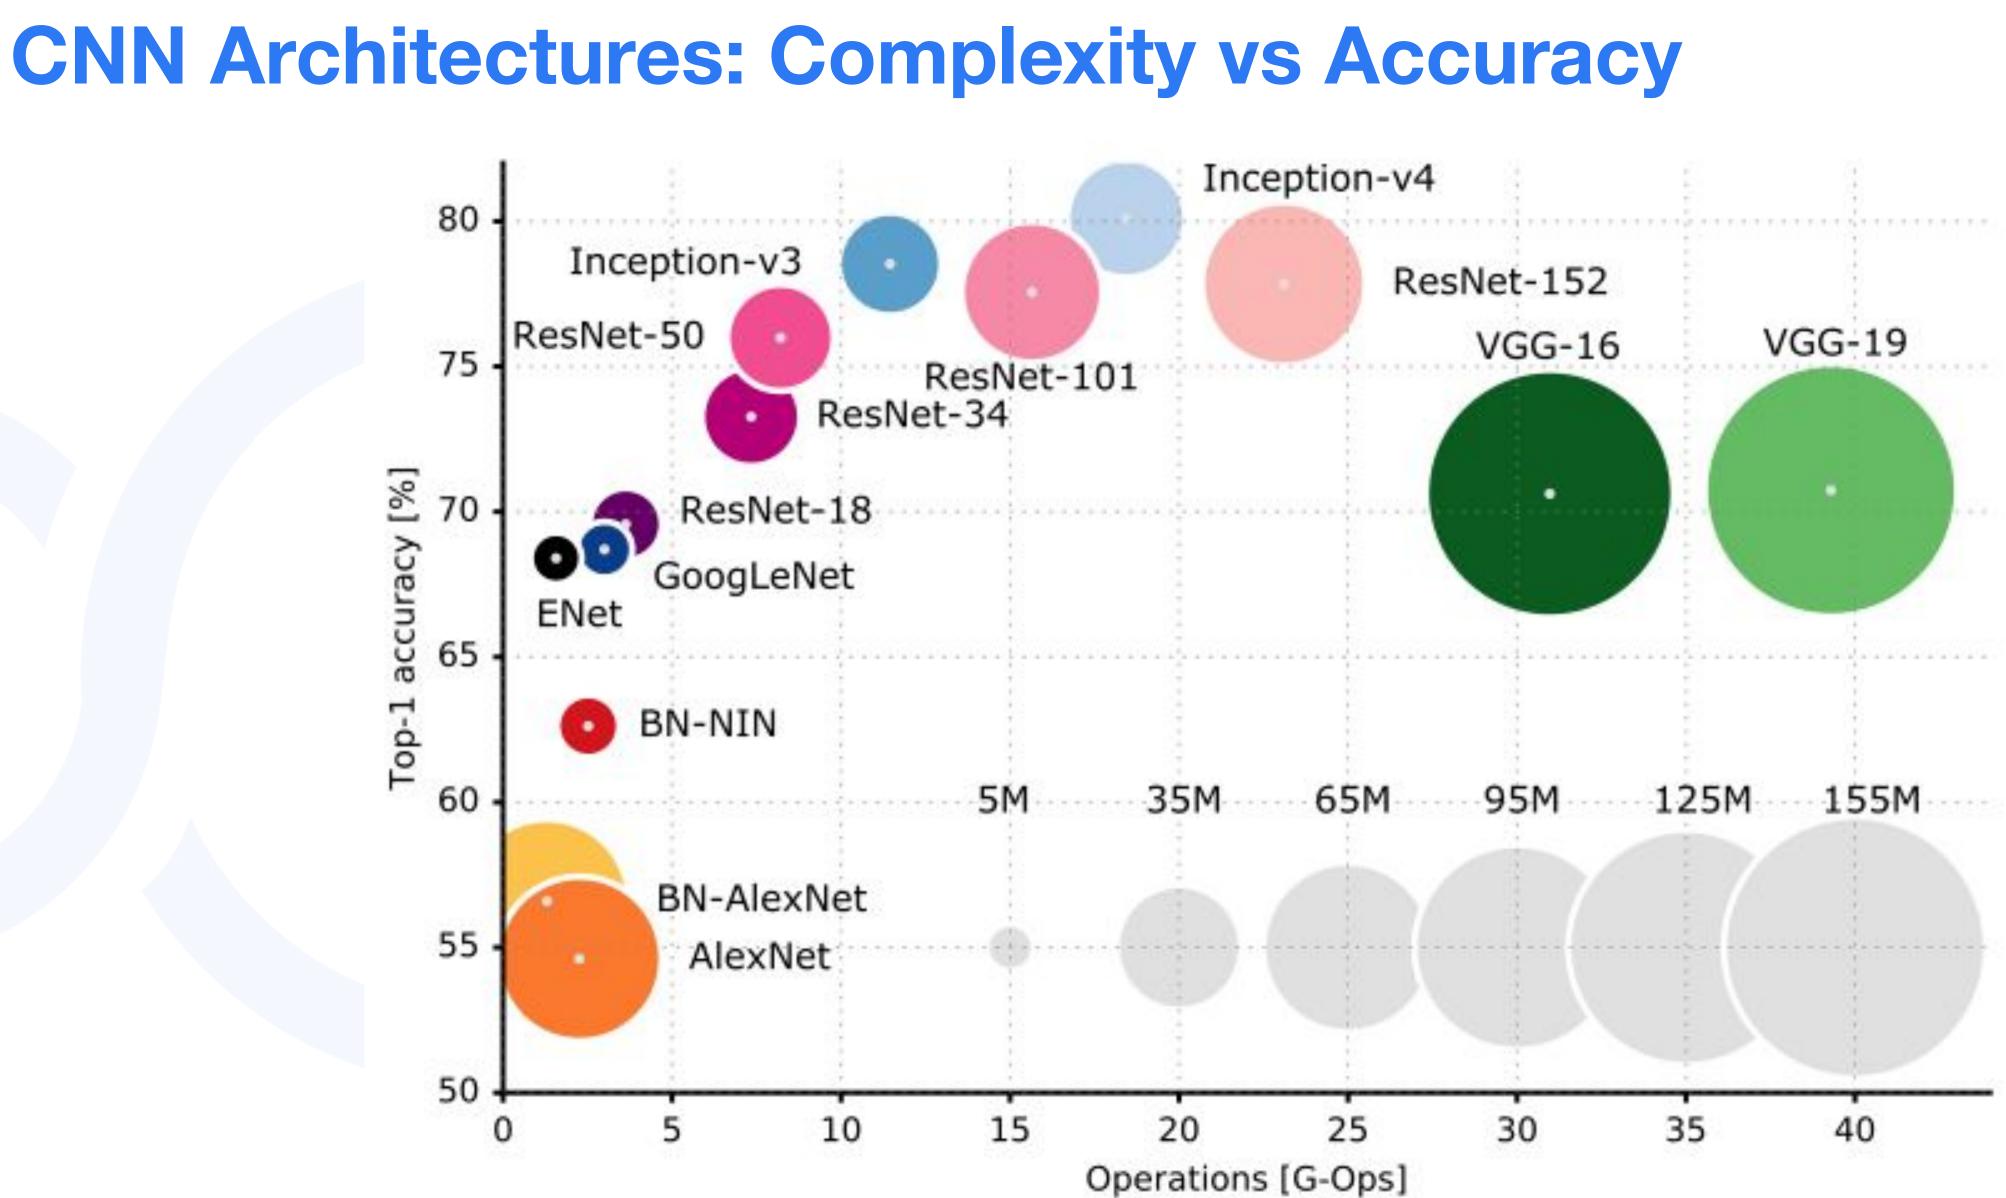
- receive in input all previous layers output feature maps
- loss function thanks to the shorter connections





• DenseNet is composed of **Dense blocks**. In those blocks, the layers are densely connected together: each layer

• This extreme use of residual creates a deep supervision because each layer receive more supervision from the







Building Neural Networks is not easy.



Deep Learning Frameworks



Tensorflow + Keras

Built by Google TensorFlow is the foundation of many DeepLearning services and one of the most used frameworks, even through its high-level library Keras.

Why TensorFlow

TensorFlow is an end-to-end open source platform for machine learning. It has a comprehensive, flexible ecosystem of tools, libraries and community resources that lets researchers push the state-of-theart in ML and developers easily build and deploy ML powered applications.

About →



Easy model building

Build and train ML models easily using intuitive high-level APIs like Keras with eager execution, which makes for immediate model iteration and easy debugging.

Robust ML production anywhere

Easily train and deploy models in the cloud, on-prem, in the browser, or on-device no matter what language you use.



Powerful experimentation for research

A simple and flexible architecture to take new ideas from concept to code, to stateof-the-art models, and to publication faster.

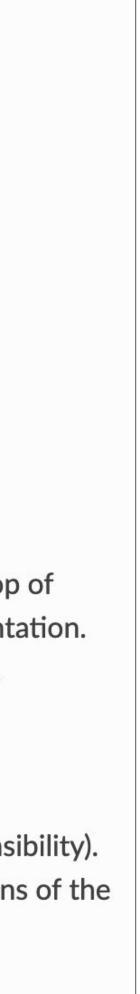
Keras: The Python Deep Learning library Keras

You have just found Keras.

Keras is a high-level neural networks API, written in Python and capable of running on top of TensorFlow, CNTK, or Theano. It was developed with a focus on enabling fast experimentation. Being able to go from idea to result with the least possible delay is key to doing good research.

Use Keras if you need a deep learning library that:

- Allows for easy and fast prototyping (through user friendliness, modularity, and extensibility).
- Supports both convolutional networks and recurrent networks, as well as combinations of the two.
- Runs seamlessly on CPU and GPU.



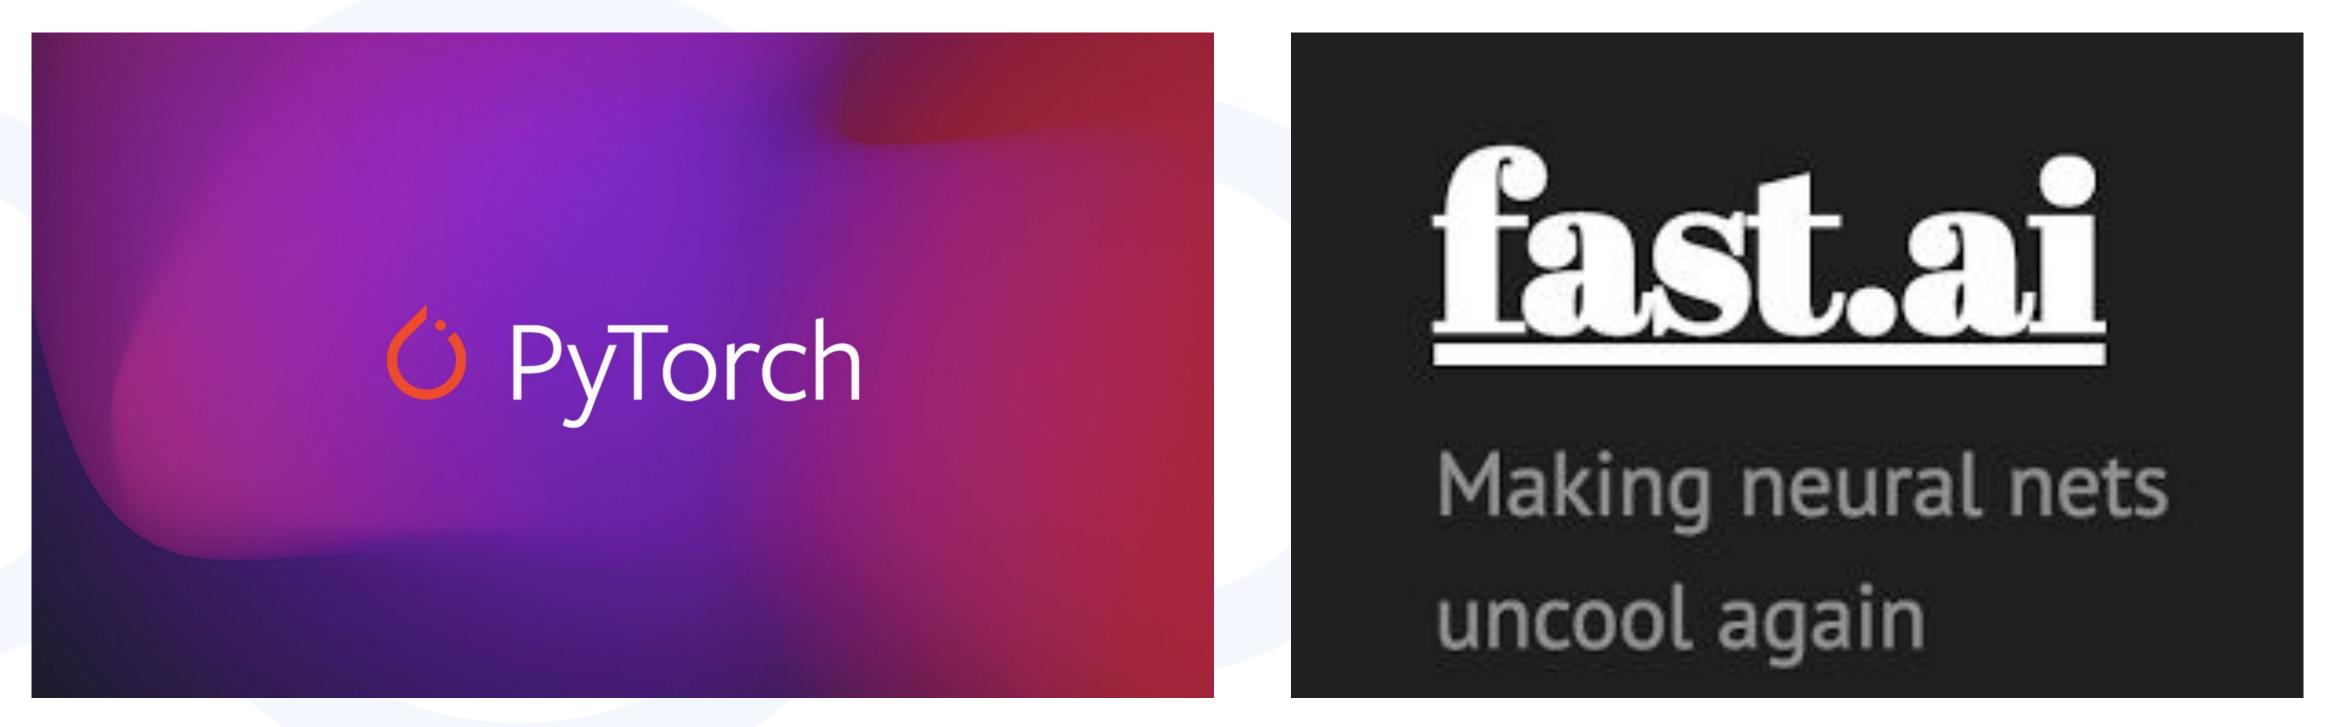
ResNet50

https://github.com/tensorflow/models/blob/master/official/vision/image_classification/resnet_model.py



PyTorch + FastAl

front-end, distributed training, and ecosystem of tools and libraries.





PyTorch enables fast, flexible experimentation and efficient production through a user-friendly

Setting up a Deep Learning environment



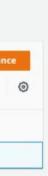
Jupyter Notebooks: an IDE for Machine Learning

Amazon SageMaker -- with Deep Learning instances

AWS Services - Res	source Groups 👻 🔯 Amazon SageMaker 丨 🛠 🗘	Oregon +	Support 👻	Amazon SageMaker offers fully manage
Amazon SageMaker $ imes$	Amazon SageMaker > Dashboard			Amazon bagemaker oners rung manage
Dashboard Search ▼ Ground Truth Labeling jobs Labeling datasets Labeling workforces ▼ Notebook Notebook instances Lifecycle configurations Git repositories ▼ Training Algorithms Training jobs Hyperparameter tuning jobs Hyperparameter tuning jobs Model packages Models Endpoint configurations	Ground Truth Notebook Training Set up and manage Availability of AWS and Train and labeling jobs for highly Availability of AWS and Train and accurate training SageMaker SDKs and Train and labeling. SageMaker SDKs and Train and labeling. Notebook Train and Labeling. Notebook The Recent activity Notebook The	ng Inference Ind tune models scale. Leverage erformance AWS hms or bring your raining jobs Create models from training jobs or import external models for hosting to run inferences on new data. Models Recent activity within the Las	Hide Amazon SageMaker > Notebook instances > Create notebook instance Amazon SageMaker > Notebook instances > Create notebook instance Create notebook instance Create notebook instance Create notebook instances Create notebook instances Cheate notebook instances that run Jupyter notebooks. The notebook instances include example code for common model training and hosting exercises. Learn more [2] Notebook instance settings Netebook instance name MySageMakerInstance Mutepook instance tops Mutepook instance type Mut	Jupyter Notebook (Jupyter Lab) running on virtual instances with on board GPUs t easily allow model training and mode inference on a variety of trained model
Endpoints Batch transform jobs AWS Marketplace	Learning Content Amazon SageMaker 10-Minute Tutorial Step-by-step guide to getting started faster Learn more ? Amazon SageMaker for Telecommunications Deployment guide for machine learning in telecommunications Learn more ?	Feature Spotlight Amazon SageMaker How-to Blog AWS machine learning experts showcase hor Amazon SageMaker Learn more [2] Amazon SageMaker Automatic Model To Skip manual trial-and-error experimentation	Elastic Inference Learn more none Additional configuration Additional configuration Permissions and encryption IAM role Notebook instances require permissions to call other services including SageMaker and S3. Choose a role or let us create a role with the AmazonSageMakerFullAccess IAM policy attached. Choose an IAM role Create a new role Enter a custom IAM role ARN Use existing role Encryption key - optional Encrypt your notebook data. Choose an existing KMS key or enter a key's ARN.	
Neosper	' ience		No Custom Encryption Amazon SageMaker Dashboard Search Ground Truth Labeling jobs Labeling datasets Labeling workforces	X Amazon SageMaker > Notebook instances Notebook instances Actions ▼ Create notebook instances Q Search notebook instances < 1 > Name ▼ Instance Creation time ▼ Status ▼ Actions MySageMakerInstance ml.t2.medium Aug 13, 2019 20:51 UTC © InService Open Jupyter Open JupyterLab







Jupyter Notebooks: an IDE for Machine Learning

Google Colab -- free managed Notebooks

File Edit View Insert Runtime Tools Help Lasts	saved at 4:15 PM	(2432)		
- Code + Text		V RAM Disk	•	🖍 Editing
 Lesson 1 - What's your pet 			∱ ↓ ⊕ E	1/11
Welcome to lesson 1! For those of you who are using prepared specially for you; click File -> Open now an		arn about this use	ful tool in a t	utorial we
In this lesson we will build our first image classifier fr	rom scratch, and see if we can achieve world-clas	ss results. Let's di	ive in!	
Every notebook starts with the following three lines; to that any charts or images displayed are shown in this		re reloaded here a	automatically	∕, and also
<pre>[] %reload_ext autoreload %autoreload 2 %matplotlib inline</pre>				
We import all the necessary packages. We are going t provides many useful functions that enable us to quic			The fastai lib	orary
<pre>[] from fastai.vision import * from fastai.metrics import error_rate</pre>				
If you're using a computer with an unusually small GF click Kernel->Restart, uncomment the 2nd line below and try again.		-		
<pre>[] bs = 64 # bs = 16 # uncomment this line if you</pre>		Doctort		

Looking at the data

We are going to use the Oxford-IIIT Pet Dataset by O. M. Parkhi et al., 2012 which features 12 cat breeds and 25 dogs breeds. Our model will



- Google Colab is a completely free managed Jupyter
 Notebook instance, where can select runtime type
 between GPU and TPU.
 - When runtime is released, or instance gets disconnected, the environment is reset and local storage is removed

Problem: build a breed detector

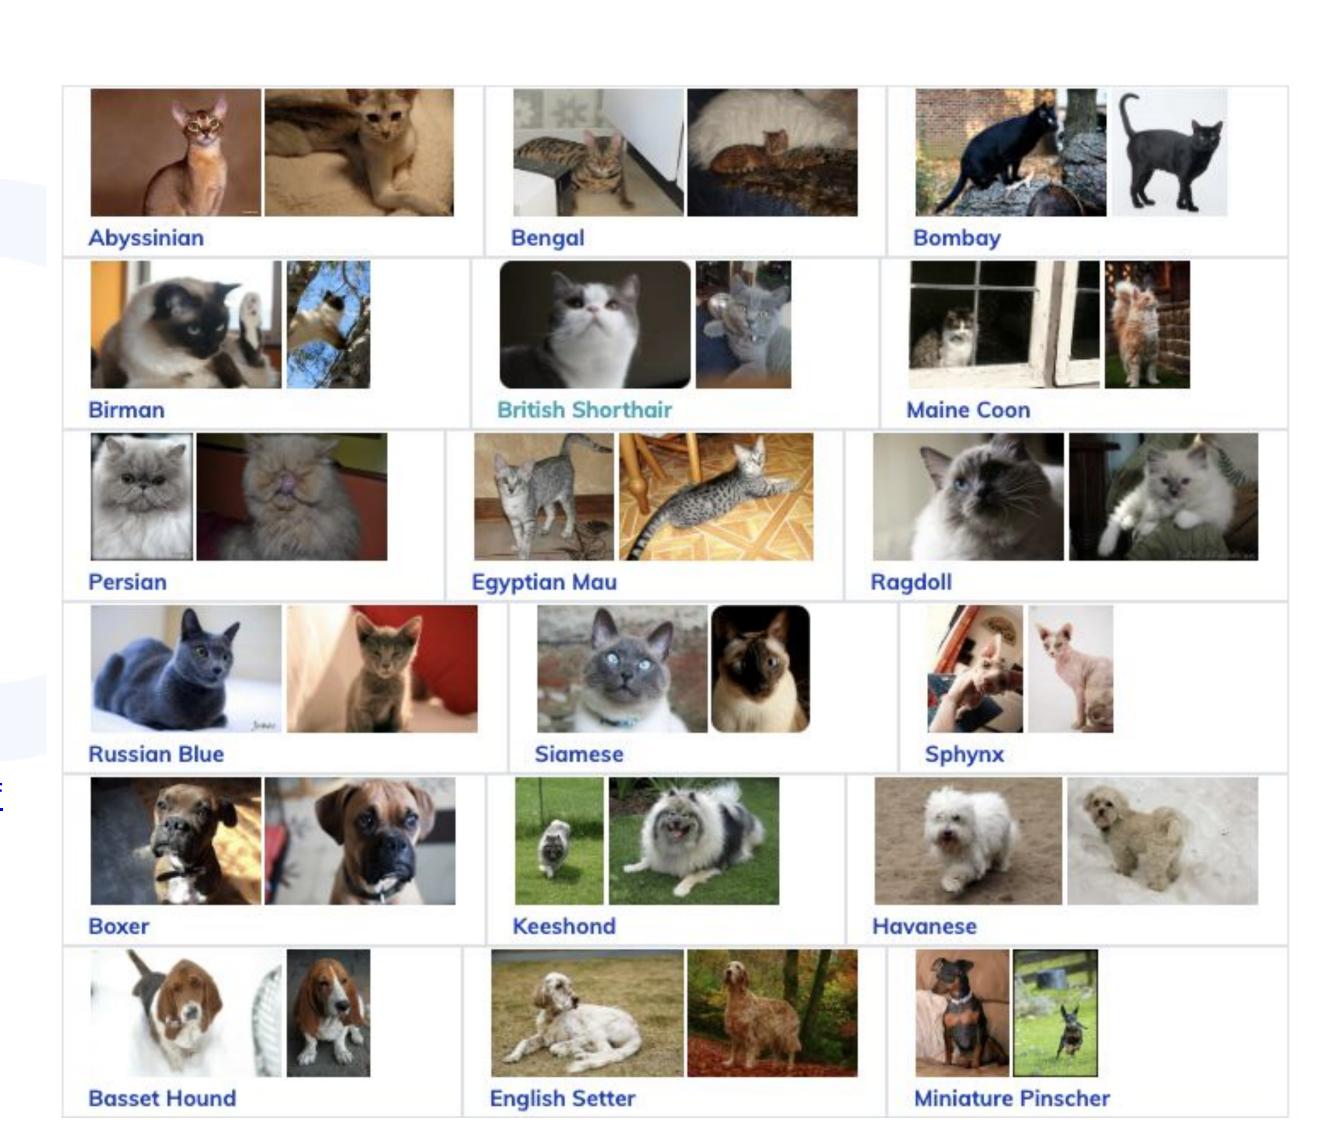
We want to detect not only whether an image contains a cat or a dog, but also which breed is the pet pictured.

One of the most difficult tasks in computer vision was, until 2013 image classification: telling the difference between a dog and a cat has been one of the best benchmarks for a CNN.

Since 2016 the computing power of GPUs makes this problem too naive to be used as benchmark, so we moved to detecting the breed of the pet in a picture

http://www.robots.ox.ac.uk/~vgg/publications/2012/parkhi12a/parkhi12a.pdf





Step 1: Data Exploration

Never under estimate your intuition looking at the data. This phase is usually named data exploration and involves extracting some statistical figures.

miniature_pinscher

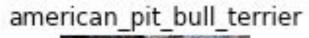


Egyptian_Mau



merican_pit_bull_terrier





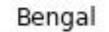


leonberger



Bengal









newfoundland



The first thing we do when we approach a problem is to take a look at the data. We always need to understand very well what the problem is and what the data looks like before we can figure out how to solve it. Taking a look at the data means understanding how the data directories are structured, what the labels are and what some sample images look like.

Labels:

'Abyssinian', 'Bengal', 'Birman', 'Bombay', 'British_Shorthair', 'Egyptian_Mau', 'Maine_Coon', 'Persian', 'Ragdoll', 'Russian_Blue', 'Siamese', 'Sphynx', 'american_bulldog', 'american_pit_bull_terrier', 'basset_hound', 'beagle', 'boxer', 'chihuahua', 'english_cocker_spaniel', 'english_setter', 'german_shorthaired', 'great_pyrenees', 'havanese', 'japanese_chin', 'keeshond', 'leonberger', 'miniature_pinscher', 'newfoundland', 'pomeranian', 'pug', 'saint_bernard', 'samoyed', 'scottish_terrier', 'shiba_inu', 'staffordshire_bull_terrier', 'wheaten_terrier', 'yorkshire_terrier'





Step 2: Data Cleaning

Remove outliers or unwanted data.

In a real-life scenario data has not been prepared into a dataset for your convenience, but needs to be converted, normalized and cleaned. Often datasets contain images that are blurred, too dark or simply wrong. Finding the right amount of data needed for a classificator

- how different are the classes that you're trying to separate?
- how aggressively can you augment the training data?
- can you use pre-trained weights to initialise the lower layers of your net?
- do you plan to use batch normalisation?
- is dataset balanced or unbalanced?

A thumb rule would be starting with thousands of images, then extending your dataset as soon as more data is required (i.e. error stops going down)

Neosperience

Step 3: Data Augmentation

If your model needs to be able to work with provide the set with rotations, skews and different sizes.

All modern frameworks allow for dataset creation with augmentation techniques zooming, flipping and rotating images. This makes your model robust to these transforms: the network learns how to classify a pet also if the image is not perfectly captured or gets distorted for any reason.

More transforms your add, more images and training time you need.

??get_transforms

Signature: get_transforms(do_flip:bool=True, flip_vert:bool=False, max_rotate:float=10.0, max_zoom:float=1.1, max_lighting:float=0.2, max_
warp:float=0.2, p_affine:float=0.75, p_lighting:float=0.75, xtra_tfms:Union[Collection[fastai.vision.image.Transform], NoneType]=None) ->
Collection[fastai.vision.image.Transform]



If your model needs to be able to work with practical images, you need to "augment" the batch

<_

Step 4: Training

Choose your network architecture, a loss function and an error metric

Many CNN models come already pre-trained into Pytorch or Keras. Using a pre-trained model and specializing the network on our dataset is often called **transfer learning**. Finding a good metric is important to tell whether our model is overfitting a dataset (loss functions goes down, error goes up).

Some metrics are already built in, such as MSE, RMSE. FBeta, etc.

learn = cnn_learner(data, models.resnet34, metrics=error_rat

learn.fit_one_cycle(epocs)



ate)]	epoch	train_loss	valid_loss	error_rate	time
		0	1.427408	0.359479	0.117050	01:03
		1	0.615426	0.274810	0.086604	00:57
		2	0.390386	0.255367	0.085927	00:57
		3	0.282586	0.237271	0.073748	00:58

Step 5: Evaluation

Evaluate results. Improve. Rinse. Repeat.

Prediction/Actual/Loss/Probability

boxer/pug / 12.56 / 0.00



Bengal/Egyptian Mau / 5.61 / 0.00



Russian Blue/shiba inu / 5.27 / 0.01







miniature pinscher/chihuahua / 5.37 / 0.00



beagle/staffordshire bull terrier / 5.10 / 0.01



chihuahua/miniature pinscher / 6.59 / 0.00



Actual

leonberger/boxer / 5.33 / 0.00



great_pyrenees/samoyed / 5.05 / 0.01





0 0 0 0 0 Abyssinian 0 0 Bengal 0 0 0 Birman 0 0 0 0 Bombay 0 0 0 British Shorthair 0 Egyptian_Mau 0 0 0 0 0 0 0 Maine_Coon 0 0 0 Persian 0 0 Ragdoll 0 0 0 Russian_Blue 0 0 0 0 0 0 Siamese Sphynx american bulldog american pit bull terrier 0 0 0 0 0 basset hound 0 0 beagle 0 0 0 0 boxer 0 0 0 chihuahua 0 english cocker spaniel 0 0 0 1 0 0 0 english_setter german shorthaired 0 0 great_pyrenees 0 havanese japanese chin keeshond leonberger miniature_pinscher newfoundland 0 0 0 pomeranian 0 0 0 0 0 pug 30 0 0 saint bernard 0 39 0 0 samoyed scottish_terrier 0 0 00 0 0 1 0 shiba inu staffordshire bull terrier wheaten terrier yorkshire terrier miniature_pinscher newfoundland pomeranian Abyssinian Bengal Birman Bombay keeshond samoyed 0 0 Maine_Coon Persian Ragdoll Russian_Blue Siamese Sphynx american_bulldog can_pit_bull_terrier British_Shorthair Egyptian_Mau english_sette: havanese dhihuahui beagle german_shorthaire great_pyrenee 3 saint_bernar japanese_chi leonberge scottish_terrie basset_houn pox english_cocker_spani

Predicted

Confusion matrix

shiba_inu 0	-					ŕ.
0 0 0 0 0 <					0	
0 0 0 0 0 <					0.8	
0 0 0 0 0 <					261	
0 0 0 0 0 <					282	
0 0 0 0 0 <		0	0	0	0	
0 0 0 0 0 <		0	0	0	0	
0 0 0 0 0 <		0	0	0	0	
0 0 0 0 0 0 0 0 0 2 0 0 0 2 0 0 0 4 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 <		0	0	0	0	
0 0 0 0 0 2 0 0 0 2 0 0 0 4 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 <		0	0	0	0	
0 0 0 0 0 2 0 0 0 4 0 0 0 0 0 0 0 <		0	0	0	0	
0 2 0 0 0 4 0 0 0 0 0 0 0 <		0	0	0	0	
0 4 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 <td></td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td></td>		0	0	0	0	
riter - 6 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0	2	0	0	
riter - 6 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0	4	0	0	
Inu 0 0 Inu 0 0 0 Inu 0 0 0 Inu 0 0 0 Inu 0 0 0 0 Inu 0 0 0 0 Inu 0 0 0 0 0 Inu 0 0 0 0 0 0 Inu 0 0 0 0 0 0 0 Inu 0 0 0 0 0 0 0 0 Inu 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0		0	0	0	0	
Inu 0 0 Inu 0 0 0 0 Inu 0 0 0 0 0 Inu 0 0 0 0 0 0 Inu 0 0 0 0 0 0 0 Inu 0 0 0 0 0 0 0 0 Inu 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0	0	0	0	
riter - 6 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0	0	0	0	
Inu 0 0 0 Inu 0 0 8 Inu 0 0 8 Inu 0 0 8 Inu 0 0 8 Inu 0 0 0 0 Inu 0 0 0 0 0 Inu 0 0 0 0 0 0 Inu 0 0 0 0 0 0 0 Inu 0 0 0 0 0 0 0 0 Inu 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0		0	0	0	0	
rinu - 0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1		0	0	0	0	
Inu 0 0 0 mier 0 0 0 0 mier 0 0 0 0 0 0 0 0 0 0 0 0 mier 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 10 0 0 0 0 0 0 0 0 0 10 0		0	0	0	0	
Inu 0 0 0 mier 0 0 0 0 mier 0 0 0 0 0 0 0 0 0 0 0 0 mier 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 10 0 0 0 0 0 0 0 0 0 10 0		0	0	0	0	
Inu 0 0 0 mier 0 0 0 0 mier 0 0 0 0 0 0 0 0 0 0 0 0 mier 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 10 0 0 0 0 0 0 0 0 0 10 0		0	0	0	0	
Inu 0 0 0 mier 0 0 0 0 mier 0 0 0 0 0 0 0 0 0 0 0 0 mier 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 10 0 0 0 0 0 0 0 0 0 10 0		0	0		1	
inu - 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0		0	0	
inu - 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0			0	
inu - 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0	0	0		
inu - 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0	0	0	0	
inu - 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0	0	0	0	
inu - 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0	0	0	0	
inu - 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0		0		0	0	
		0	0	0	0	
		0	0	0	0	
		0	0	0	0	
		30	0		0	
	2	0	36	0	0	
		0		32	0	
	1	0	0	0	39	
		shiba_inu -	staffordshire_bull_terrier -	æ,	-	



Breeds can be accurately recognized!





course.fast.ai



Part 2 Industrial Applications



Even if algorithms have been known for more than three decades, today we're at the nexus of converging opportunities

Backpropagation with ReLU

Computing

Power



Huge dataset availability

Product auto-tagging and visual search

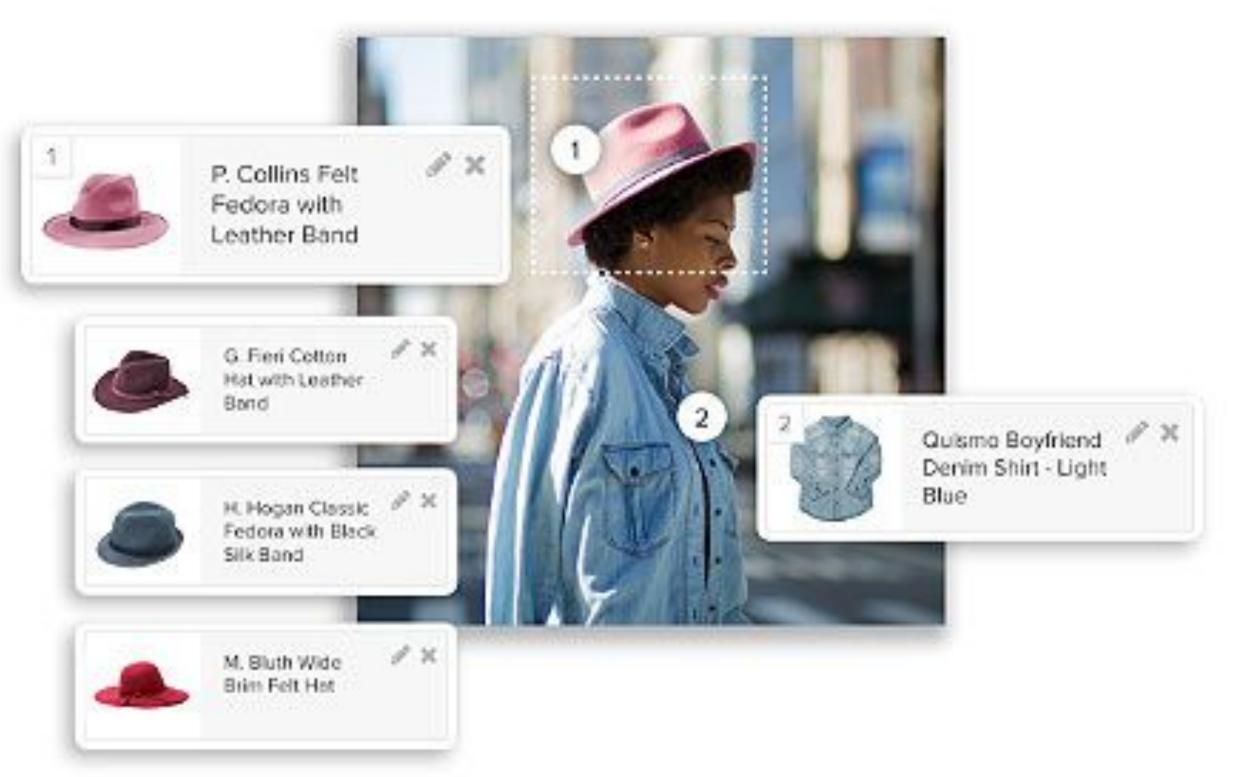
Use plain ResNet or VGG with transfer learnin catalogs or customer pictures.

- Automatically tag products
- Cut down on workload to categorize product;
- Show related products
- Find cheaper version of high end products
- Find complimentary products
- Find products usage on social media

https://www.kaggle.com/paramaggarwal/fashion-product-images-dataset



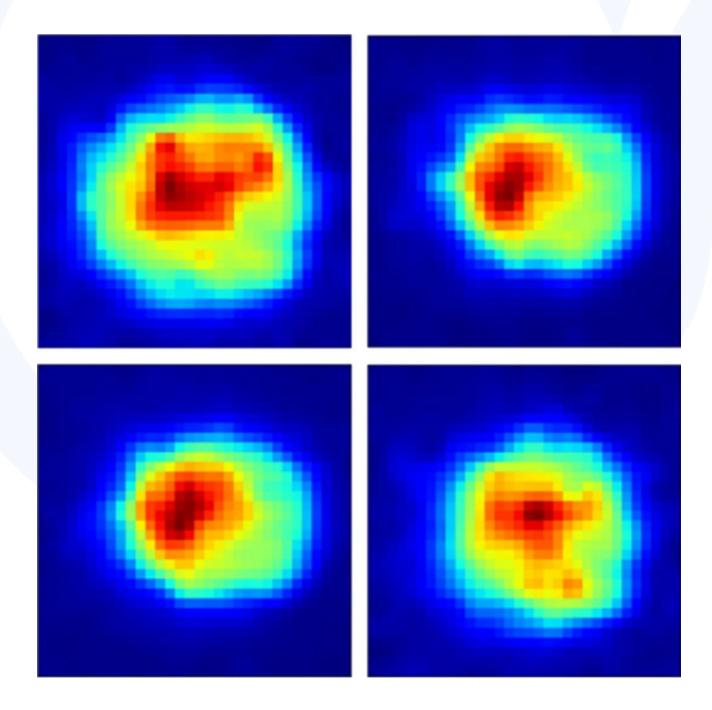
Use plain ResNet or VGG with transfer learning to find products within images coming from



Quality assurance

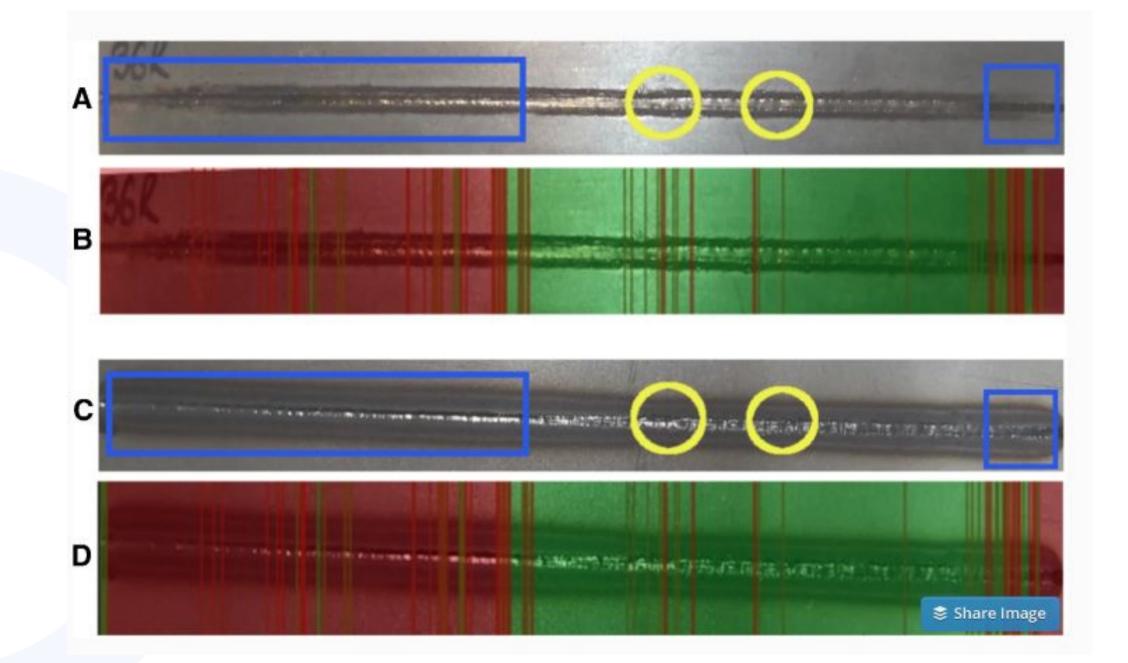
Detect items not compliant with accepted sizes/shapes/colors.

CNNs approaches are capable of analysing MWIR thermal images to extract parameters of laser processes and quality indicators.









Real-time defect detection on a laser weld bead, a and c show two side views of the weld bead where the blue rectangles mark a defective section in the first and final segments due to undercuts and the yellow ellipses mark a region where some points have excessive porosity



Self Driving cars

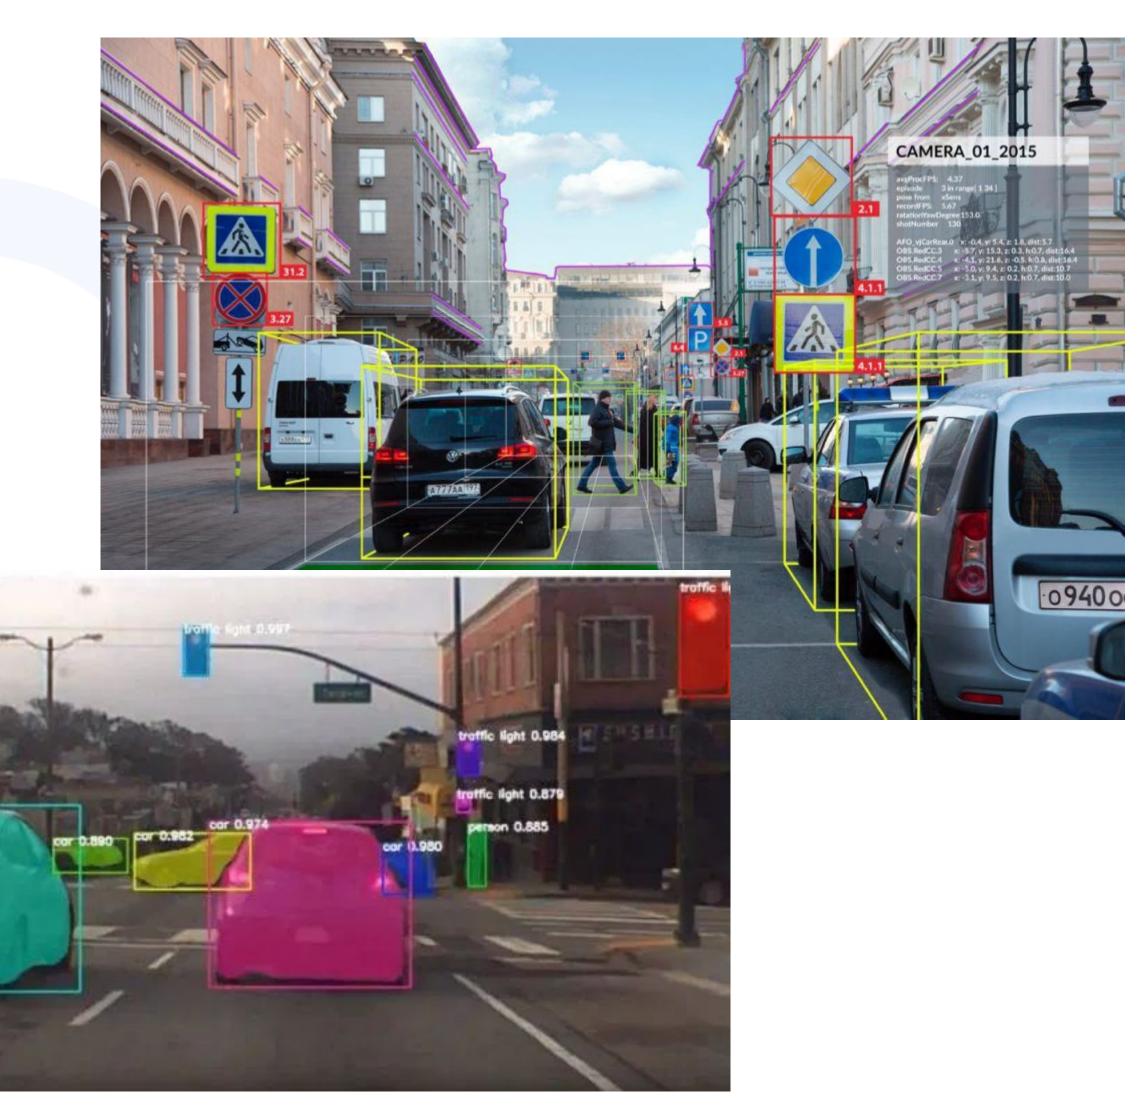
Uses a model ensemble to leverage segmentation properties of CNNs. CNNs to identify and segment, other ML models to track cars and respond to inputs

Lyft and Uber are experimenting self driving cars for public transportation in big cities such as Las Vegas.





Deep usage in security: detect accesses to restrict areas, detect people unhealthy behavior or



Payments using FaceID

Use customer face as key to unlock credit card informations in a third party store

Facebook Pay is experimenting payments with face recognition.

AliPay just updated its proprietary algorithm for face recognition to unlock payments in store and personalized advertising.

Libraries such as DLIB offer face embeddings extraction and recognition with an accuracy over 90%









DLIB a face recognition library

Multi-stage feature extraction and face rekognition. A CNN trained with triplet loss function

Sometimes we have to train a network not to recognize a given object, but to tell whether an image is or *is not* a given person of interest.

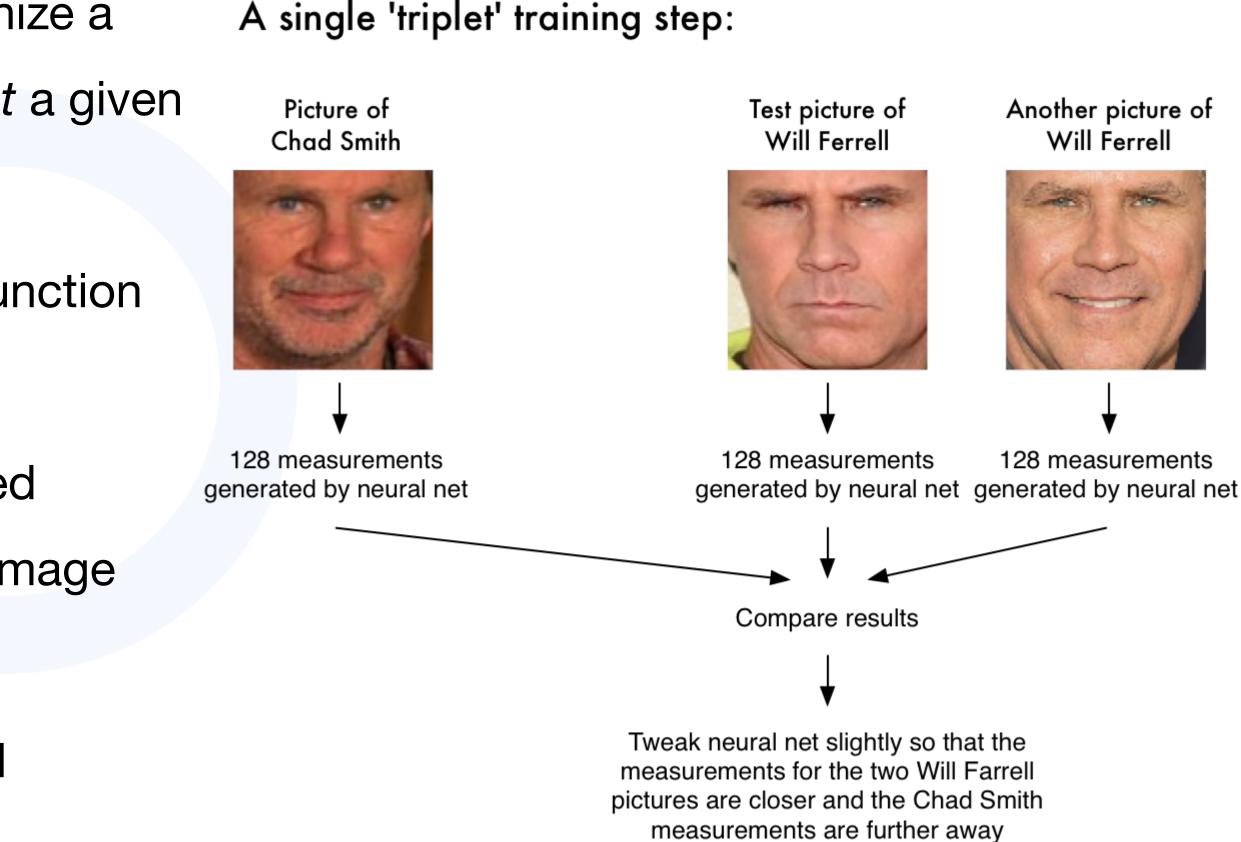
A common technique is to define a particular loss function named **Triplet Loss**.

DLIB network extracts landmarks from a face (named measurements), then trains a network wit a known image and two unknown different images.

This process makes the network able to understand differences between pictures of any face.

Neosperience









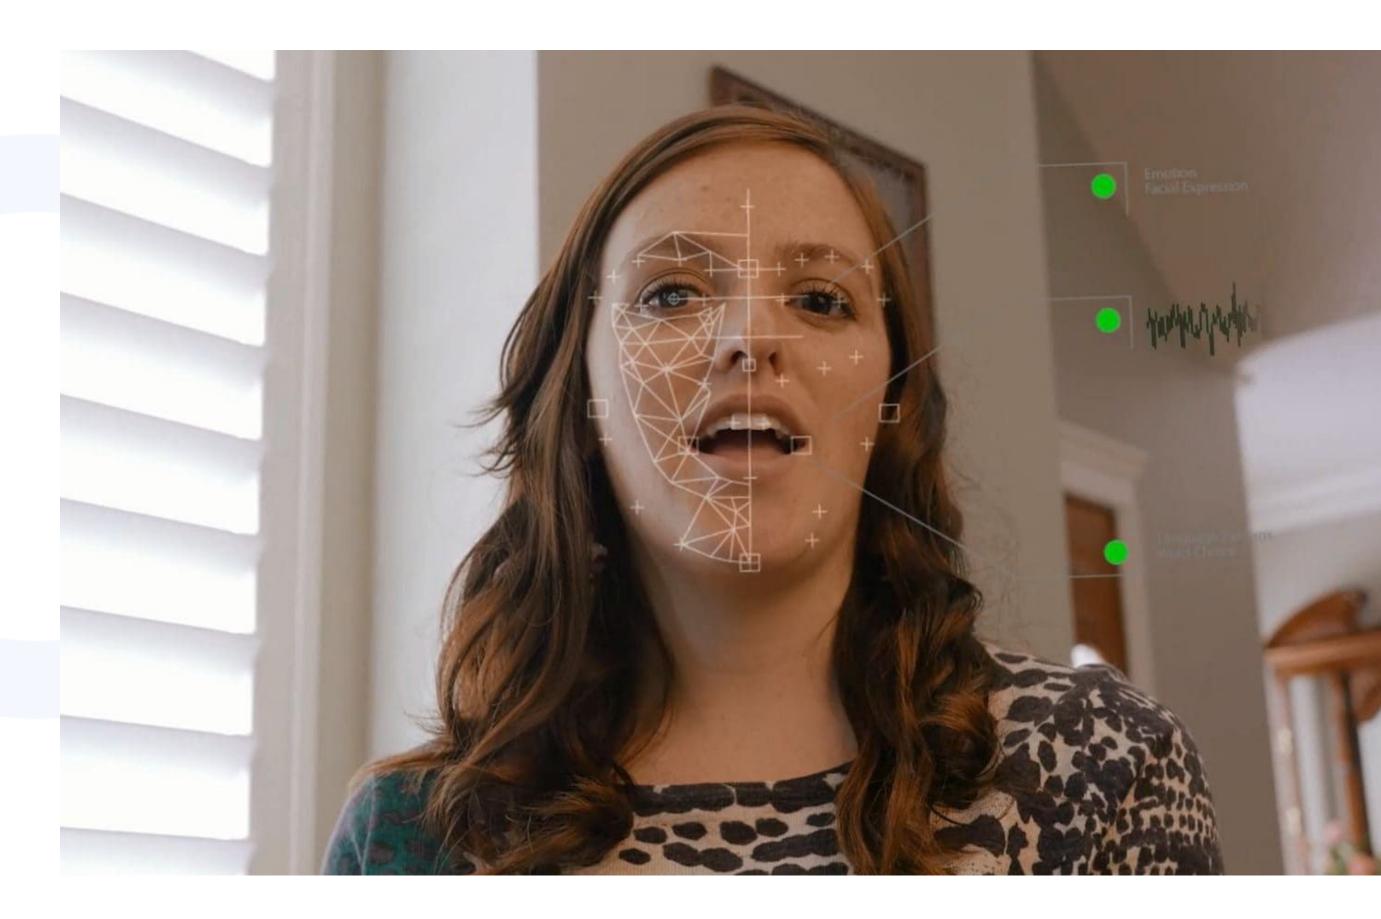
CNNs used in recruiting

Al used for first time in job interviews in UK to find best applicants

Unilever is among companies using AI technology to analyse the language, tone and facial expressions of candidates when they are asked a set of identical job questions which they film on their mobile phone or laptop.

The algorithms select the best applicants by assessing their performances.





CNNs in education

performance monitoring

CNNs are used by China schools to monitor students attention and posture, thus avoiding injuries or being too distracted



China is the current biggest investor on Computer Vision applications, with focus on schools and



https://youtu.be/JMLsHI8aV0g?t=52





Environmental Sound Classification

Use CNNs to classify different sounds in an open environment

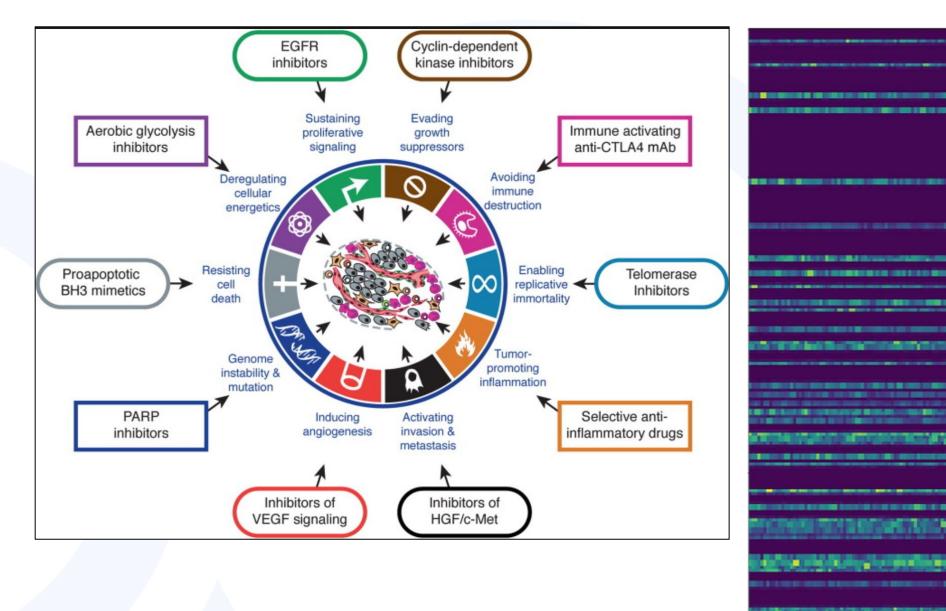
Represent sound frequencies as images,	0.90				I	
then classify different types of spectrum	^{28.0}				 + 	T
to better classify sounds in an	08.0 ACCU	Т		Т		
environment	Classification 0.75					
	O.65	T			*	
	0.60		*	*	i	
	0.00	SKM	PiczakCNN	SB-CNN	SKM (aug)	SB-CNN (aug





Neural Networks Applications in real life problems

Cancer Type Classification using CNN and Fast.Al



https://towardsdatascience.com/the-mystery-of-the-origin-ca cation-using-fast-ai-libray-212eaf8d3f4e

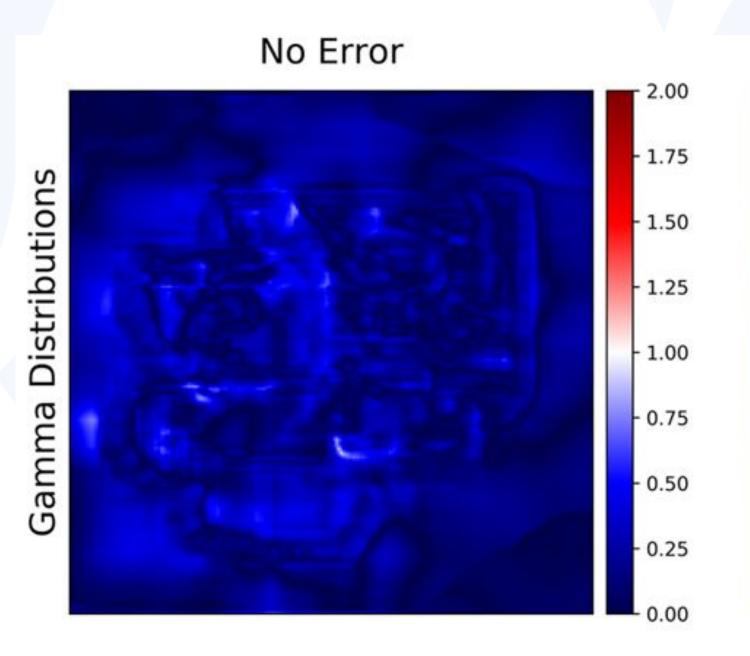


															1	Confus	sion mat	rix													
	ACC -	21	2	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0 0	0	0	0
	BLCA -	0	111	0	2	0	0	1	0 0	1	1	0	0	0	0	0	1	2	0 0	0	0	0	1	1	0	0	0	0 0	1	0	0
	BRCA -	1	0	318	2	0	0	0	0 0	0	0	1	0	0	0	0	0	2	0 1	0	0	0	0	3	0	0	0	0 0	0	1	0
	CESC -	0	0	1	83	0	0	0	0 0	2	0	0	0	0	0	0	0	1	0 0	0	0	0	0	0	0	0	0	0 0	5	0	0
	CHOL -	0	0	0	0	11	0	0	0 0	0	0	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0 0	0	0	0
	COAD -	0	0	1	0	0	66	0	0 0	0	0	0	0	0	0	0	0	1	0 0	0	0	0	15	0	2	0	0	0 0	1	0	0
	DLBC -	0	1	0	1	0	0	11	0 0	0	0	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	1	0	0 0	0	0	0
	ESCA -	0	0	0	0	0	0	0 4	46 0	0	0	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	9	0	0 0	0	0	0
	GBM -	0	0	0	0	0	0	0	0 45	0	0	0	0	0	1	0	0	1	0 0	0	0	0	0	0	0	0	0	0 0	2	0	0
	HNSC -	0	3	1	4	0	0	0	0 0	141	L 0	0	0	0	0	0	0	7	0 0	0	0	0	0	0	0	0	0	0 0	0	0	0
	KICH -	0	0	0	0	0	0	0	0 0	0	17	2	1	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0 0	0	0	0
	KIRC -	0	2	0	0	0	0	0	0 0	0	2	150	5	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0 0	0	0	0
	KIRP -	0	1	0	0	0	0	0	0 0	0	0	2	82	0	0	0	0	0	0 0	0	0	0	0	0	0	1	0	0 0	1	0	0
	LAML -	0	0	0	0	0	0	1	0 0	0	0	0	0	51	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0 0	0	0	0
	LGG -	0	0	0	0	0	0	0	0 0	0	0	0	0	0	156	0	0	0	0 0	0	0	0	0	0	0	0	0	0 0	1	0	0
	LIHC -	0	0	0	0	2	0	1	0 0	0	0	0	0	0	0	107	0	0	1 0	0	0	0	0	0	0	0	0	0 0	0	0	0
	- DAUL -	0	1	0	1	0	0	0	0 0	0	0	1	0	0	0	0	144	7	0 0	0	0	0	0	0	0	0	0	0 0	0	0	0
Street in the local part of th	م LUSC -	0	1	0	0	0	0	0	0 0	3	0	0	0	0	0	0	7 1	138	0 0	0	0	0	0	0	0	0	0	0 0	0	0	0
	MESO -	0	1	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0 1	25 0	0	0	0	0	0	0	0	0	0 0	0	0	0
	ov -	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0	0 126	0	0	0	0	1	0	0	0	0 0	1	0	0
Niesterheiten in die Arthürsterheiten vollternertistellen mit betreinen Arthurie met getrei	PAAD -	0	0	0	0	0	1	0	0 0	0	0	0	0	0	0	0	0	0	0 0	52	0	0	0	0	0	1	0	0 0	0	0	0
	PCPG -	2	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0	0 0	0	52	0	0	1	0	0	0	0 0	0	0	0
	PRAD -	0	1	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0	0 0	0	0	147	0	0	0	0	0	0 0	1	0	0
	READ -	0	0	0	0	0	11	0	0 0	0	0	1	0	0	0	1	1	0	0 0	0	0	0	14	0	0	0	0	0 0	0	0	0
	SARC -	0	0	0	0	0	1	0	0 0	1	1	0	0	0	0	0	0	0	0 0	0	0	0	0	74	1	0	0	0 0	0	1	0
	SKCM -	0	0	0	0	0	0	0	0 0	1	0	0	0	0	0	0	0	0	0 0	0	0	0	0	3	137	0	0	0 0	0	0	0
	STAD -	0	1	0	0	1	0	0 1	12 0	0	0	0	0	0	0	0	0	1	0 0	0	0	0	0	0	0	109	0	0 0	0	0	0
	TGCT -	0	1	0	0	0	1	0	0 0	0	0	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	39	0 0	0	0	0
	THCA -	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	2	0	0 0	0	0	0	0	0	1	0	0 1	49 0	2	0	0
<u>ancer-type-classifi</u>	тнум -	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0 36	0	0	0
$\frac{1}{1001} \frac{1}{1000} \frac{1}{1000$	UCEC -	0	0	1	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0	0 0	0	0	0	0	0	1	0	0	0 0	49	3	0
	UCS -	1	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0	0 0	0	0	0	0	1	0	0	0	0 0	4	11	0
	UVM -	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0	0 0	0	0	0	0	0	0	0	0	0 0	0	0	24
		ACC -	BLCA -	BRCA -	CESC -	CHOL -	COAD -	DLBC -	GBM -	HNSC -	KICH -	KIRC -	KIRP -	- TAML -	- 991	- HHC -	LUAD -	LUSC -	ME50 -	PAAD -	PCPG -	PRAD -	READ -	SARC -	- WDXS	STAD -	TGCT -	THYM -	UCEC -	ucs -	- WW
																Pr	edicted														

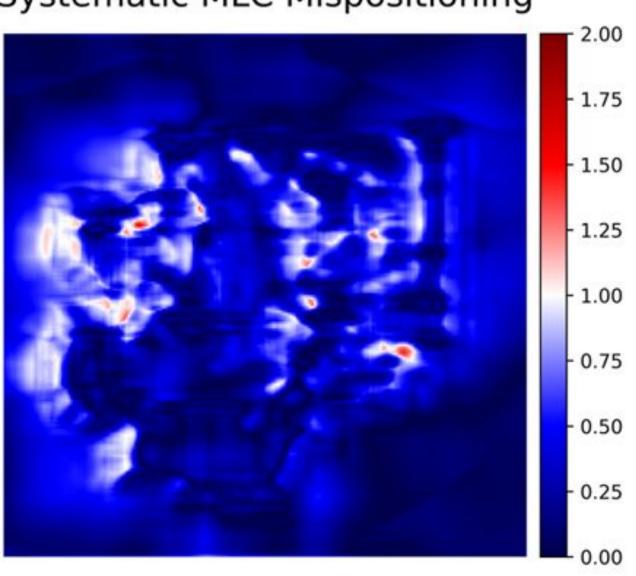
Quality assurance in radiotherapy

radiomic analysis of gamma images with convolutional neural networks

CNNs can be used to detect operational errors when exposing patients to radiotherapy and provide a better upfront correction of medical errors.



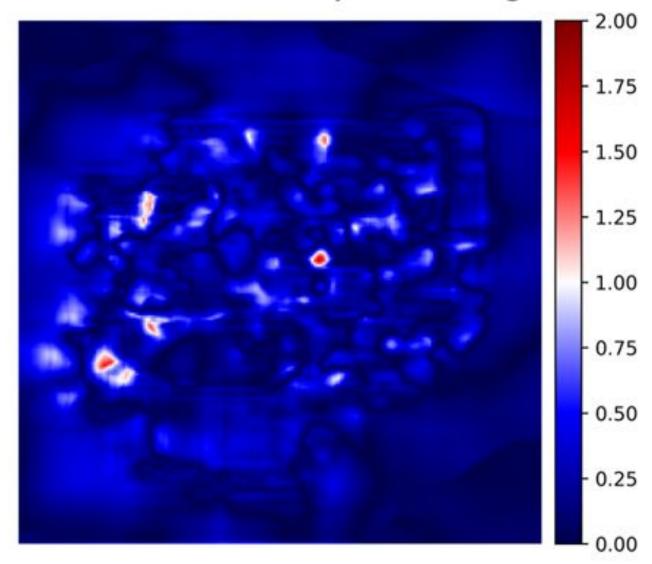
Systematic MLC Mispositioning





Deep learning for patient-specific quality assurance: Identifying errors in radiotherapy delivery by

Random MLC Mispositioning







Deep dive on **Computer Vision Applications** powered by Neosperience

Neosperience Image Memorability



Image Memorability — A business perspective

What is a memorability score?

Memorability is a measure of how much an image sticks into the memory of an average customer respect to average baseline images

A memorability score is a number representing memorability of an image, compared to the average capability of a human to remember an image which is 0.72

Images with a score higher than 0.72 have high memorability and are suitable for campaigns

Images with a score lower than 0.72 underperform and should be avoided because are not remembered

Neosperience



memorability score: 0.79

Image Memorability — A business perspective

A memorable image is a good image?

High memorability score is a good starting point, but using it to select an image could be too naive

More relevant than memorability itself is understanding which feature makes an image memorable

Assigning a score to each pixel of the image regarding its contribution to the resulting score

In this case memorability analysis outperforms humans because it is able not only to tell the score

but also to understand what makes this score

Neosperience





Image Memorability — A technical perspective

How to detect scores and heat maps?

Build an experiment to measure memorability (ground truth)

Deep Learning comes into help with CNNs

A CNN learns from experiment dataset how to estimate a memorability score

From a given inference, finding layer activations (through back propagation)

Convolutions and back propagation are compute intensive tasks that require GPUs even with inference

GPU inference is achieved through DeepLearning AMIs and on-premise instances

We needed an architecture to support inference through GPU in production in a scalable and cost effective way











Neosperience https://image.neosperience.com





Alisea Visual Clean

Alisea – Transfer learning example

PROBLEM: Classify images of air duct/pipes as 'dirty' or 'clean'

Step 1: Exploratory analysis

Dataset composed of hundreds of images of different air pipes, taken with different cameras, in different sizes.

Balanced dataset: 50% labelled 'dirty', 50% labelled 'clean'. RGB color channel.

Which images size to use? Which color channels?







Step 2: Data Cleaning

Choose which images are appropriate for your training dataset. Remove photos that would add 'noise'. In our case MANUALLY!

Considered image size:

- 128x128x3
- 256x256x3
- 320x320x3
- 480x480x3

Color channels:

• RGB, HSV







Not appropriate images for our dataset

Step 3 & 4: Data augmentation and training

Data augmentation to increase image size.

model architecture. Based on your need you can choose to keep the model as it is or:

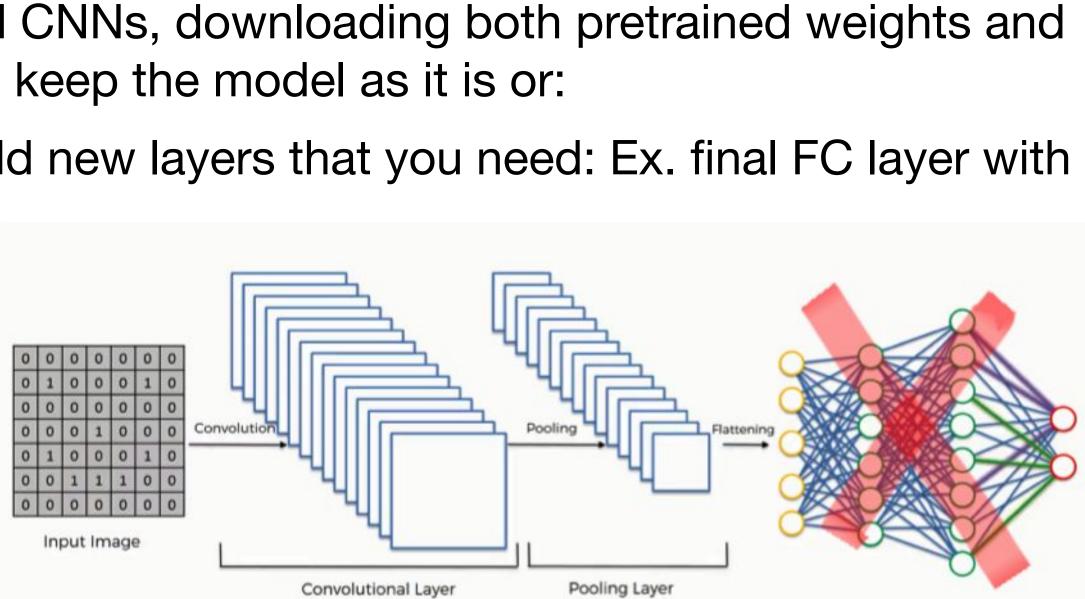
- remove the fully connected (FC) layers at the end and add new layers that you need: Ex. final FC layer with more output classes.
- Keep all the weights or train them all over again

Considered CNN architectures:

ResNet34, ResNet50, ResNeXt50

Trained several models using different image sizes to notice if there was a difference in our results. Best models in our case: ResNet50 and ResNeXt50 Best size: 256x256x3, bigger images need more computing power and longer training time Best color channel: RGB Final score: ~92% accuracy **Neosperience**

Keras and other libraries allow you to import already trained CNNs, downloading both pretrained weights and



What does the model see?

False negative images (truth: clean, predicted: dirty)

661_7_Dopo





810 131 Dopo

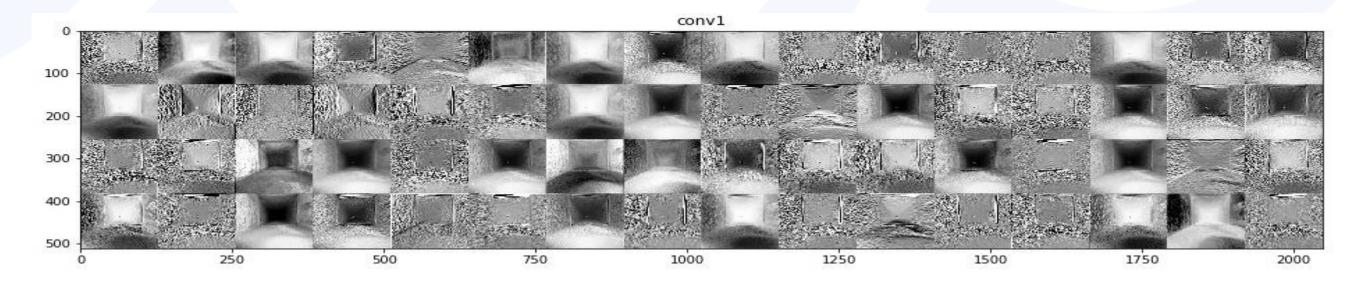


591_15_Dopo









Feature Map of first Conv Layer





Dirty: CAM Heatmap Layer Activation_144

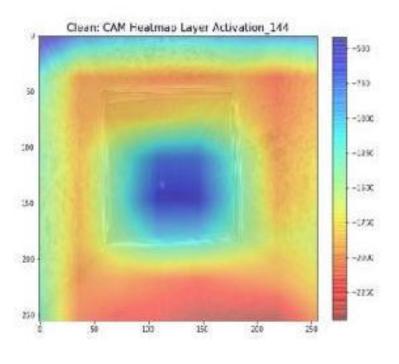


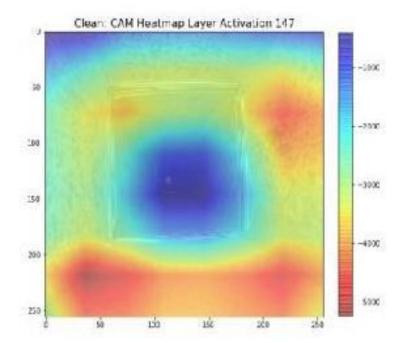
569_17_Dopo



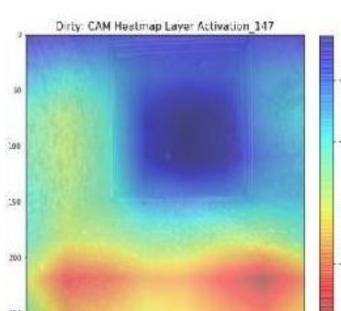
200







Attention Heatmap



130

150

200

30







Neosperience People Analytics

Why count people in store?

People number is a KPI used to estimate ROI

- Understanding the number of people is considered a good way to estimate the average return of a given store
- The daily income of a store divided by the overall number of people detected gives a ROI
- Understanding high traffic stores can led to strategic decisions
- Low traffic or lower ROI can be closed or moved





People counting is broken

Commonly used devices are flawed and lack accuracy

IR sensors are cheap but inaccurate

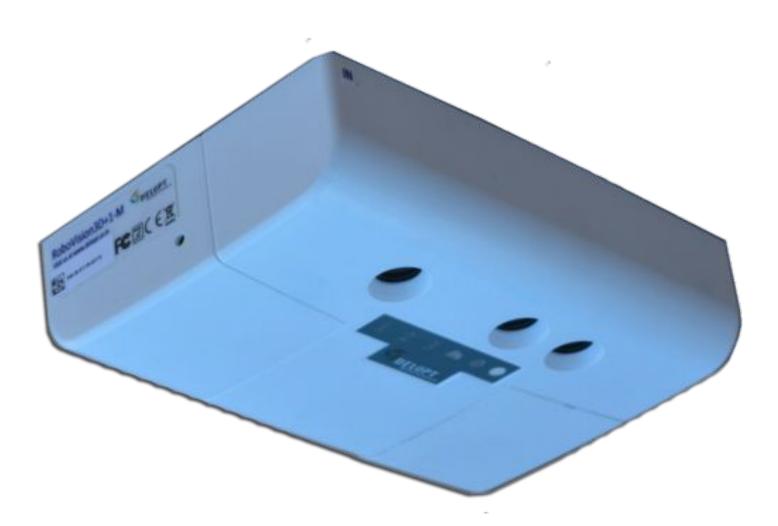
- no unique counting
- no information about what happens in store
- false positives

laser / thermal counters are expensive and do not provide relevant information except counting

- unique counting
- no information about what happens in store
- lesser false positives

Neosperience



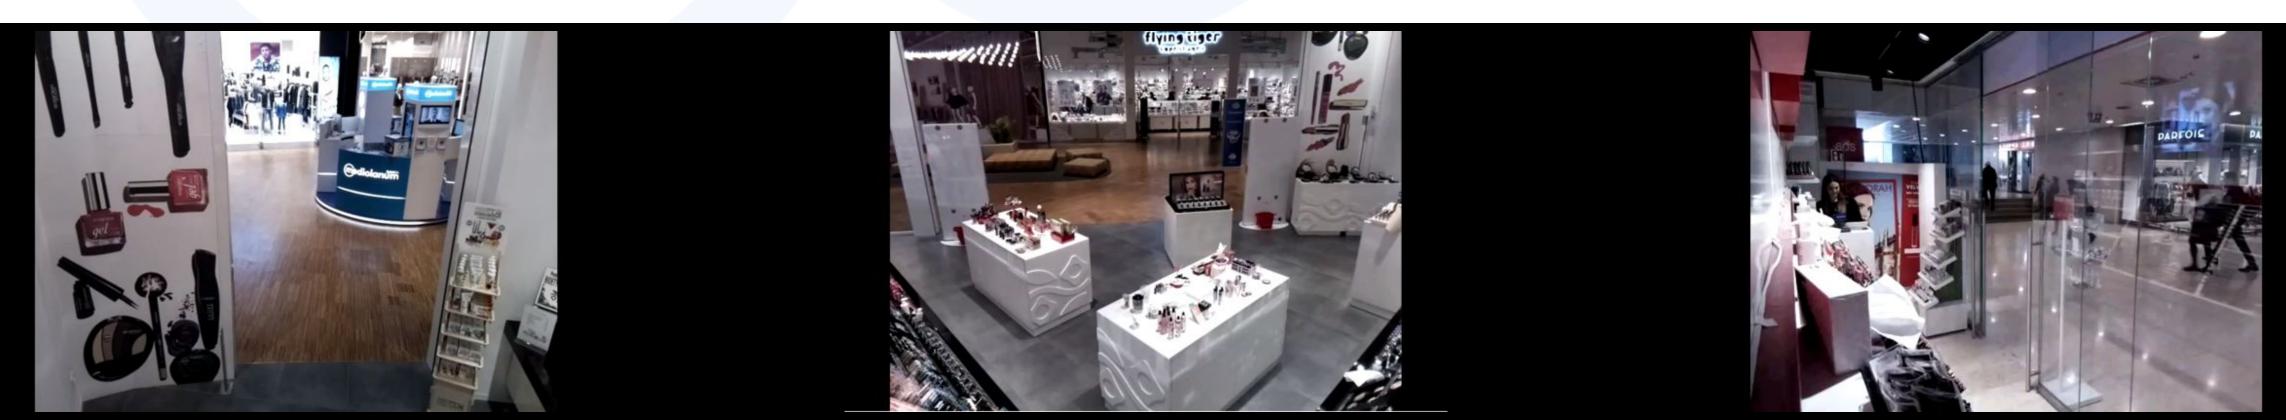


Introducing Neosperience Store Analytics

Detect relevant insights about your customers in stores using cameras

Neosperience Store Analytics is the SaaS solution to extract meaningful informations about people visiting stores in an accurate and reliable way

- Uses both standard cameras and dedicated hardware with a cost effective profile
- Dedicated Hardware is projected to optimise costs, heat management and reliability
- Stream acquisition is achieved in cloud
- Allows for multiple people counting, detects unique visits
- Enables advanced insights extraction



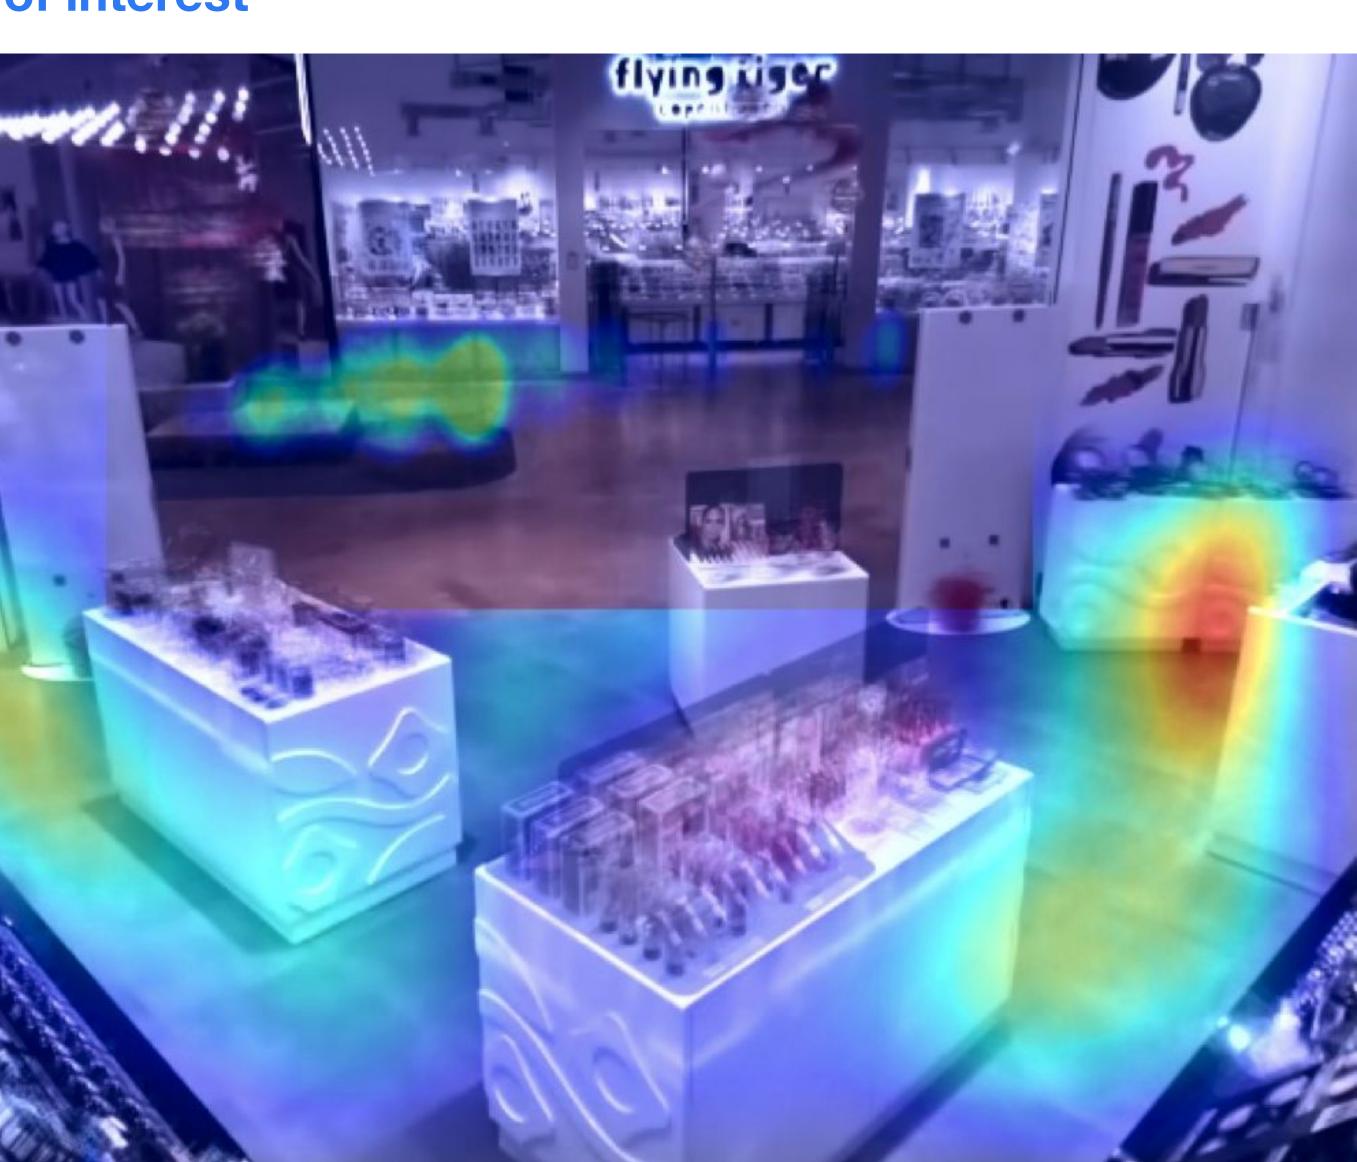
Results: people heatmaps, trajectories, insight

Mapping people presence within a given area of interest

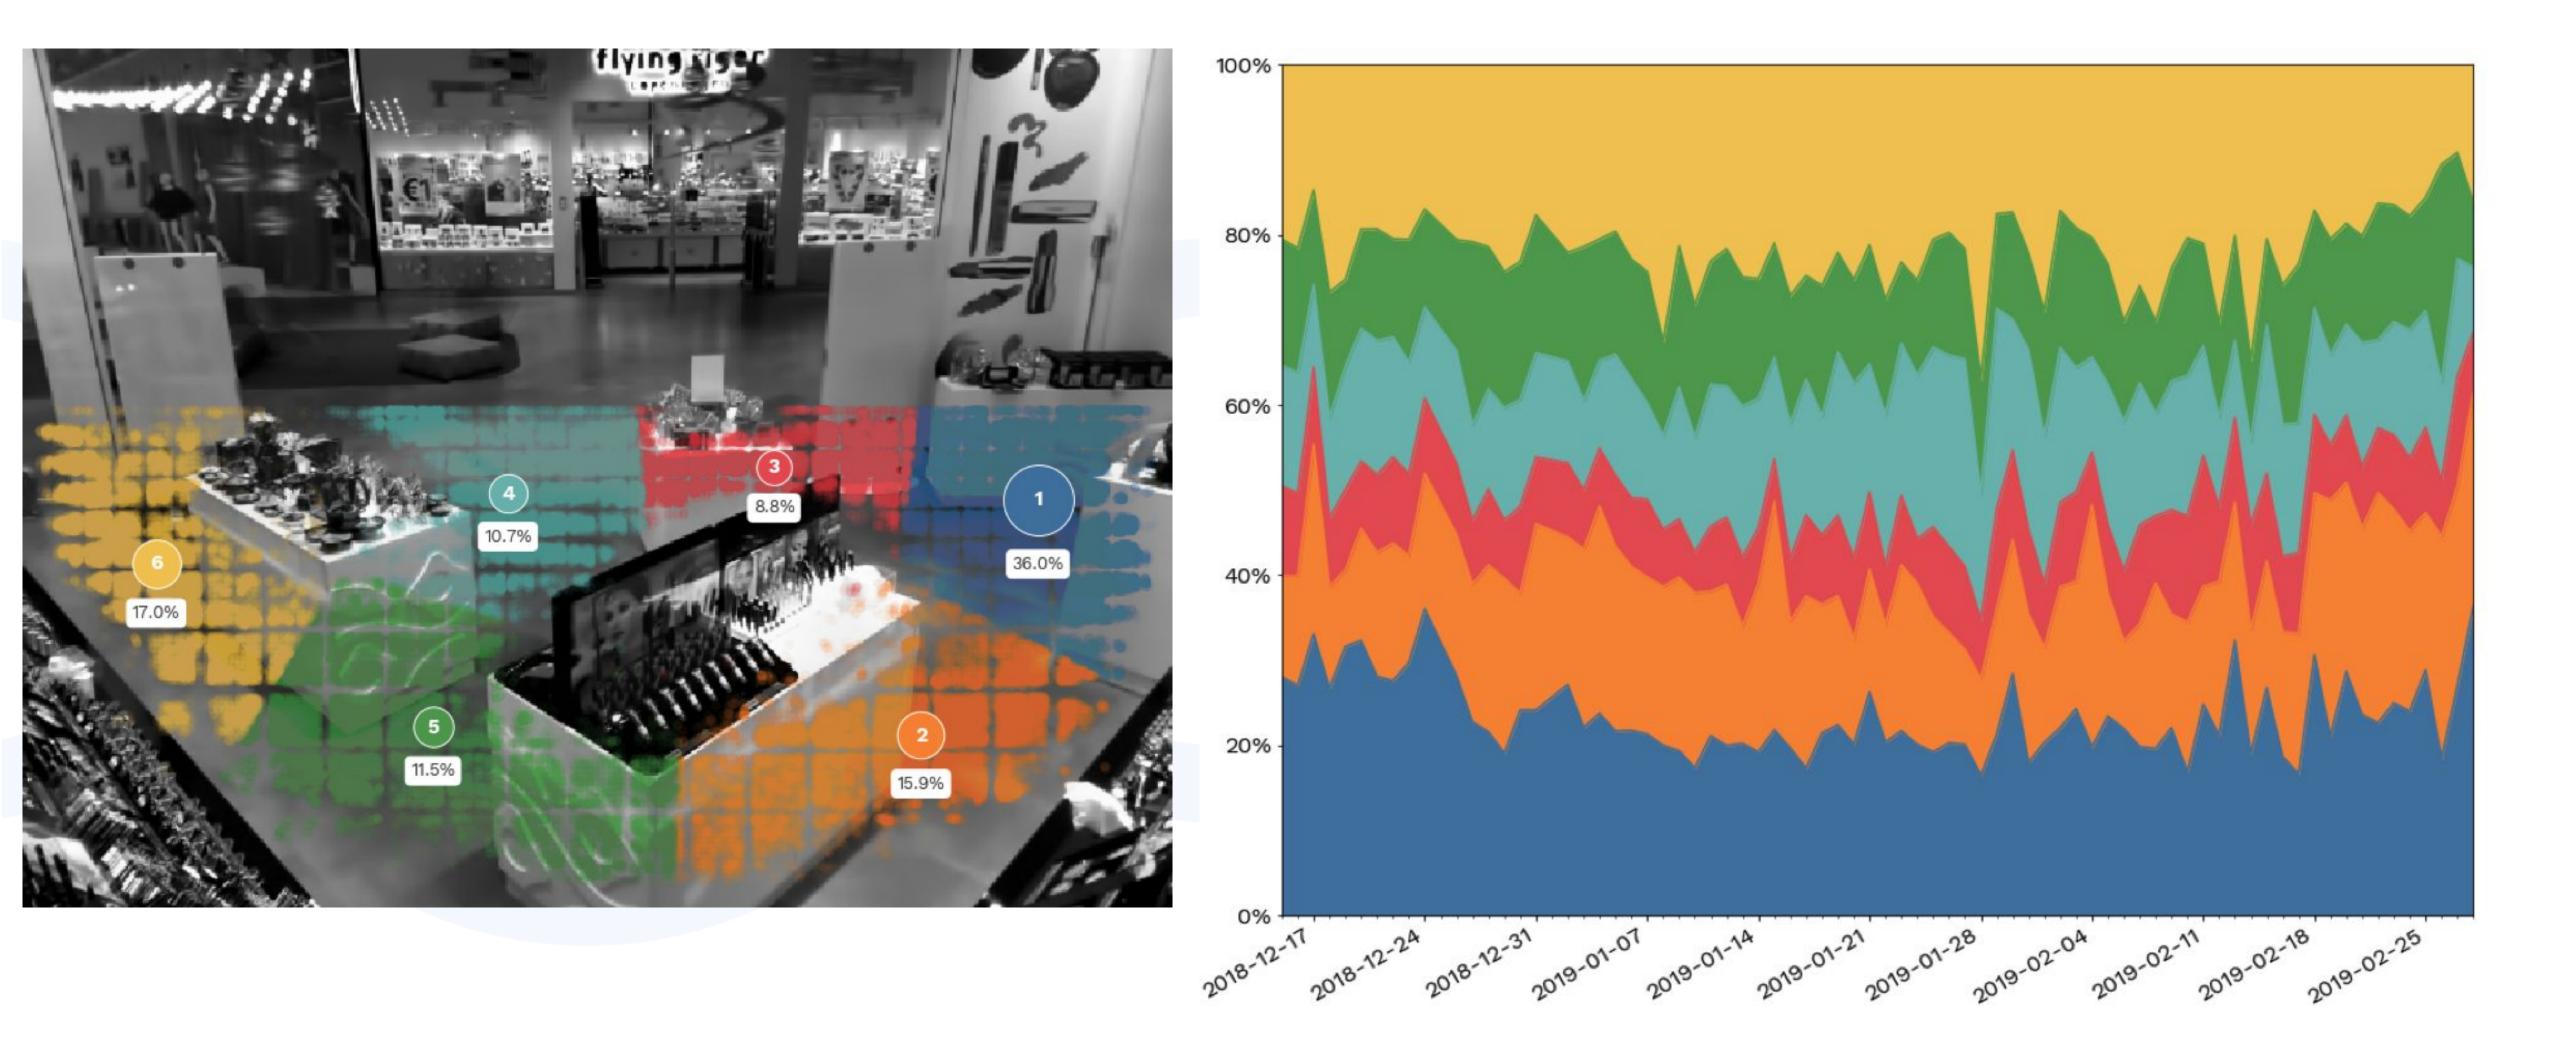
Being able to recognise people and track their movements in front of a camera leds to interesting results not only related to people counting

- Store managers can obtain a clear view of the preferred areas inside a store
- And event the overall amount of people that do not enter the store
- Store Analytics over delivered about store understanding, delivering a different but more meaningful metric

Neosperience

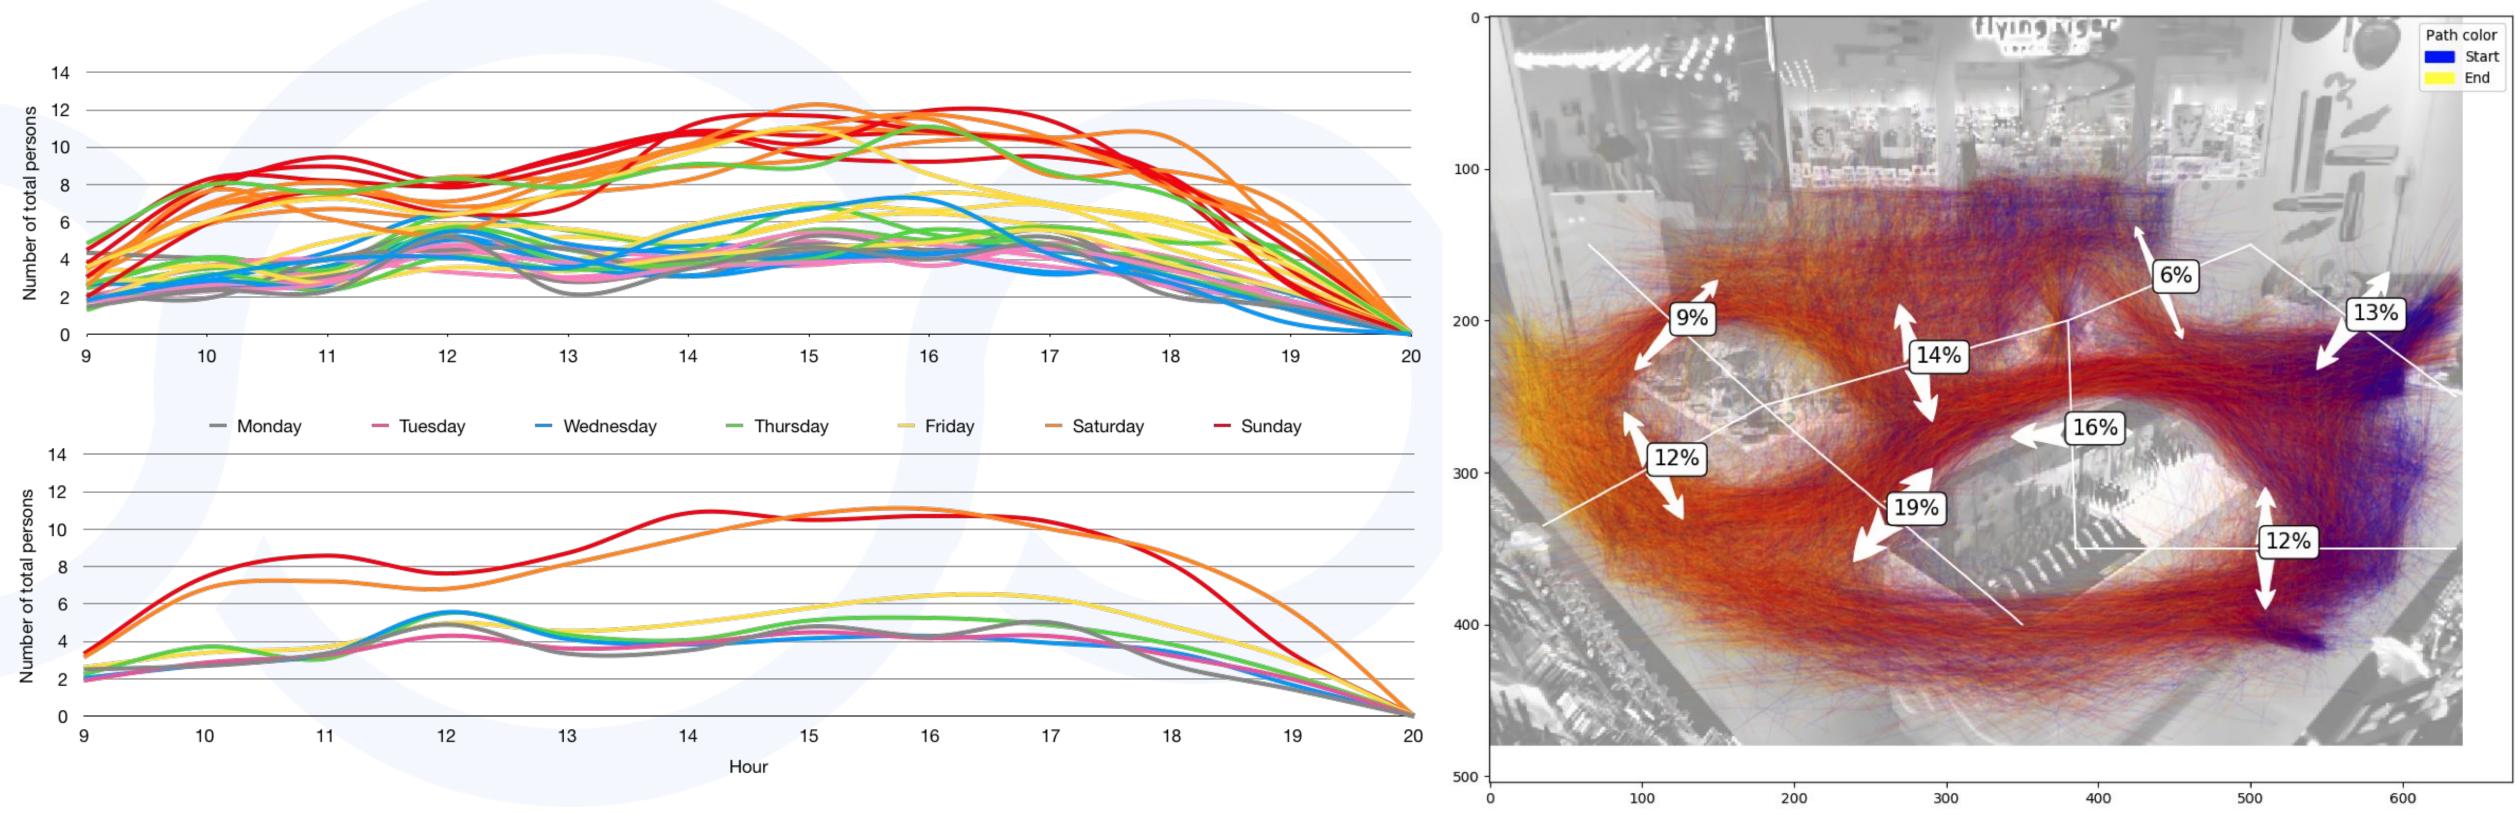














Neosperience Psychographics



Detecting Psychographic traits

Understand customer behavior based on their shared content on social networks

Locus of Control

Belief to have control over the outcome of events in our lives, as opposed to external forces beyond our control.

Sensation Seeking

Search for experiences and feelings that are varied, novel, and intense, and the readiness to take risks for the sake of such experiences.

Agreeableness

General concern for social harmony. Agreeable people are considerate, kind, generous, trusting and trustworthy, helpful, and willing to compromise their interests with others.

Need for Uniqueness

Pursuit of differentness relative to others that can be achieved through the acquisition, utilization, and disposition of consumer goods.

Need to Belong

Human emotional need to belong and be an important part of a social group greater than ourselves.

Need for Cognitive Closure

Motivation to find a firm answer to an ambiguous situation, enhanced by the perceived benefits of obtaining closure. Closed-mindedness, intolerance of ambiguity, preference for order, and predictability.

Openness to Experience

General appreciation for art, emotion, adventure, unusual ideas, imagination, curiosity, and variety. Tendency to be more creative, unconventional, and aware of our inner feelings.

Conscientiousness

Preference for planned rather than spontaneous behavior, the tendency to display self-discipline, act dutifully, and strive for achievement against measures or outside expectations.

Impulsiveness

Tendency to act on a whim, displaying behavior characterized by little or no forethought, reflection, or consideration of the consequences.

Need for Affect

Motivation to approach or avoid emotion-inducing situations.

Self-centredness, craving for admiration, and grandiose view of one's talents.

Belief in ours ability to succeed in specific situations or accomplish a task.

Emotional Stability

The tendency to be calm, eventempered, and less likely to feel tense or rattled.





Extraversion

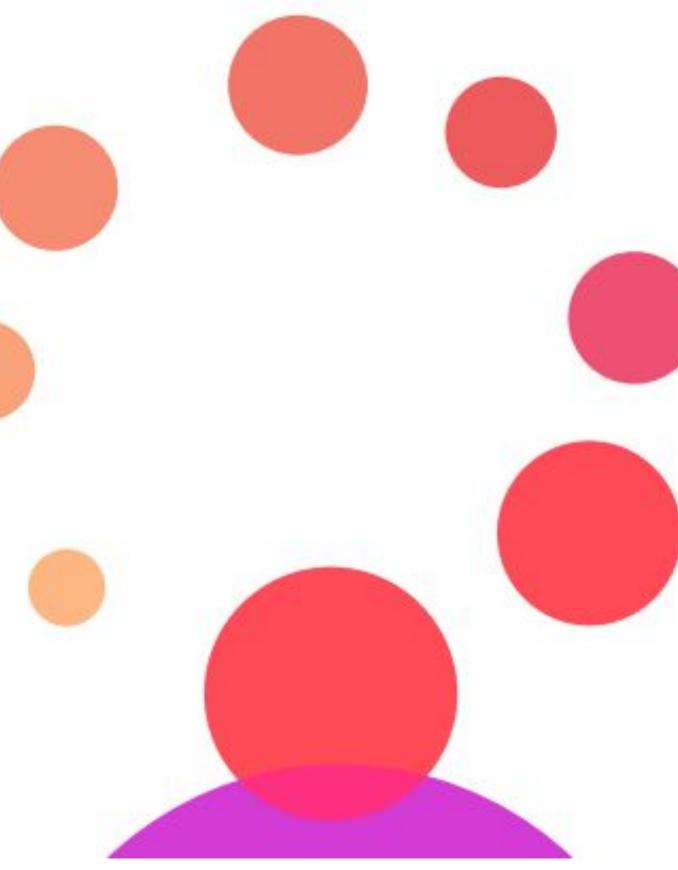
Pronounced engagement with the external world. Extroverts enjoy interacting with people and tend to be enthusiastic, they possess high group visibility, like to talk, and assert themselves.

Need for Cognition

Extent to which we are inclined towards effortful cognitive activities.

Narcissism

Self Efficacy

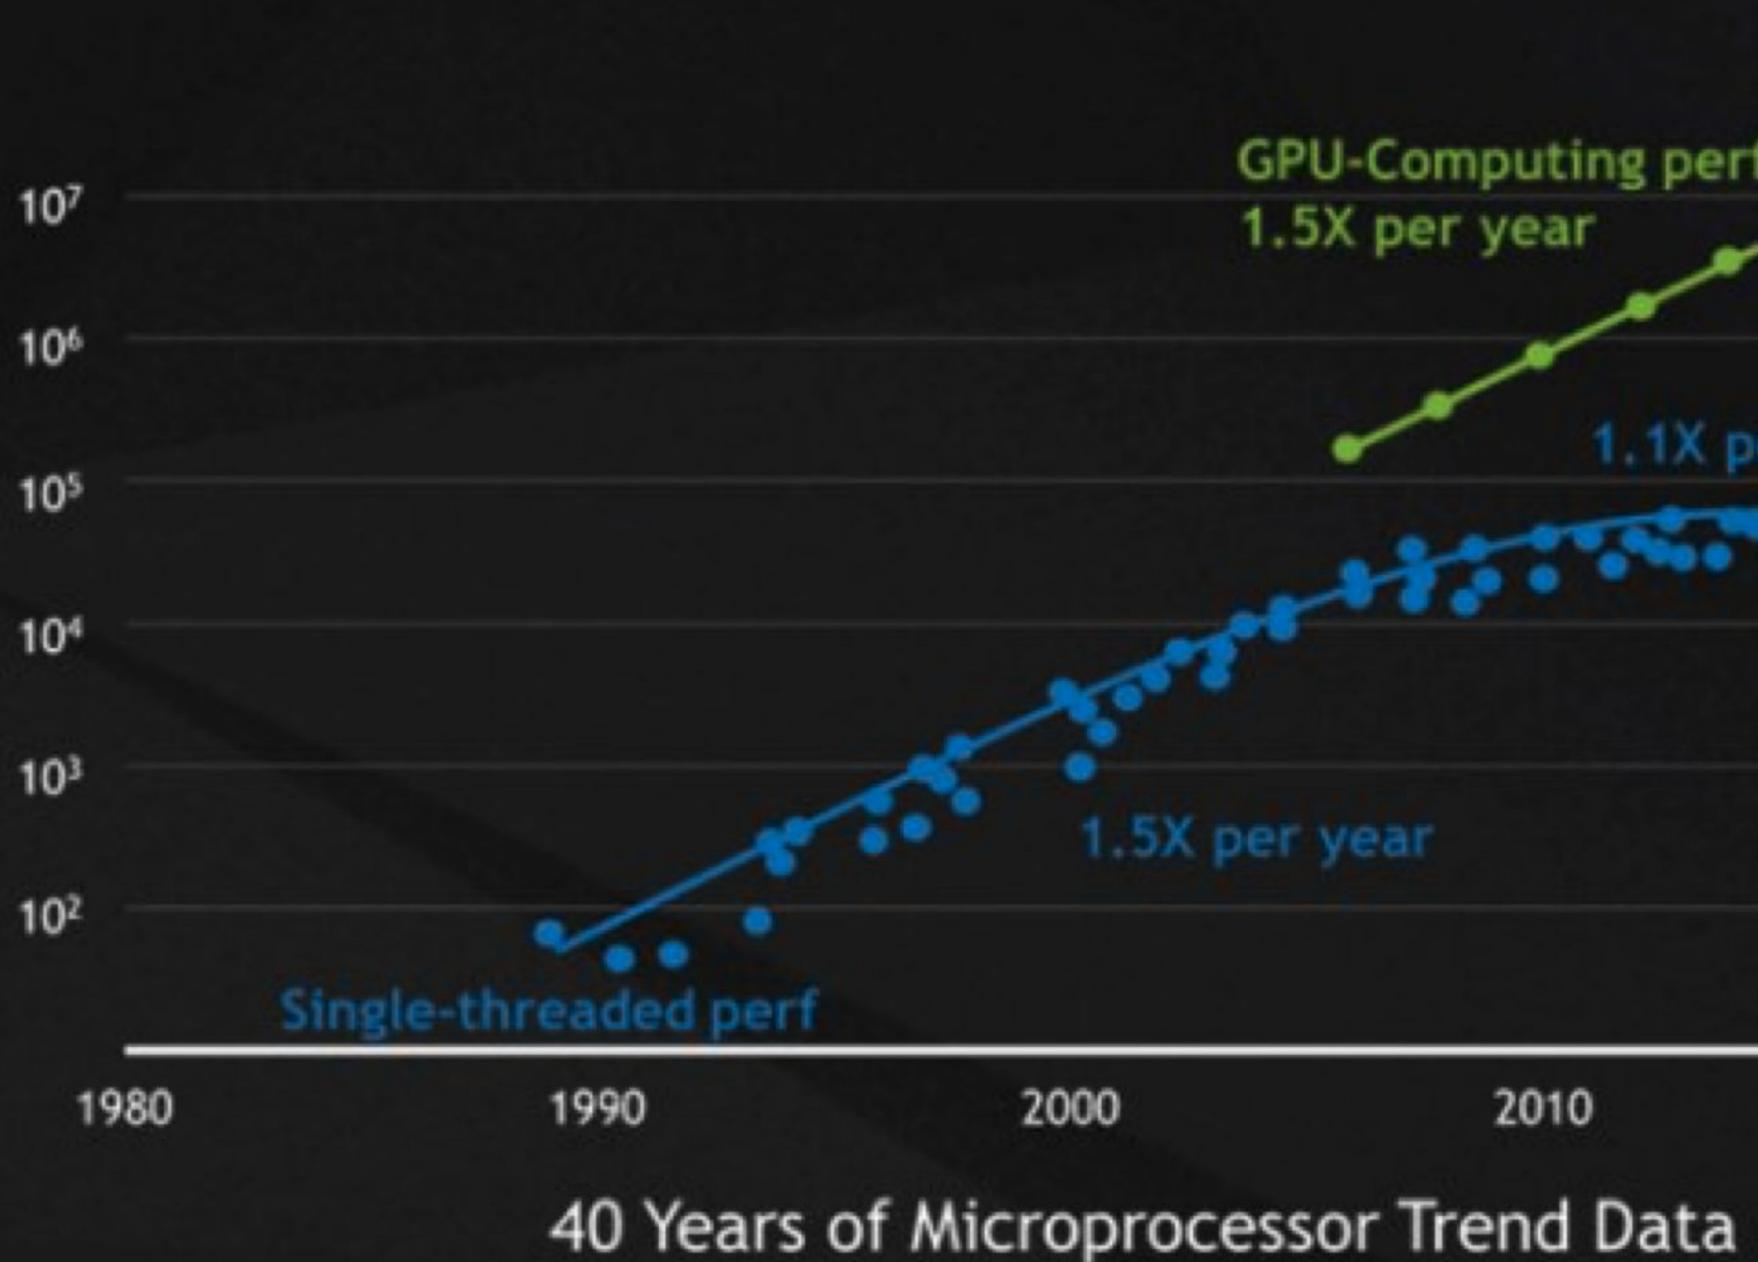






Hardware for Computer Vision





GPU-Computing perf 1.5X per year

1.1X per year

1.5X per year

2010

2020

by 2025

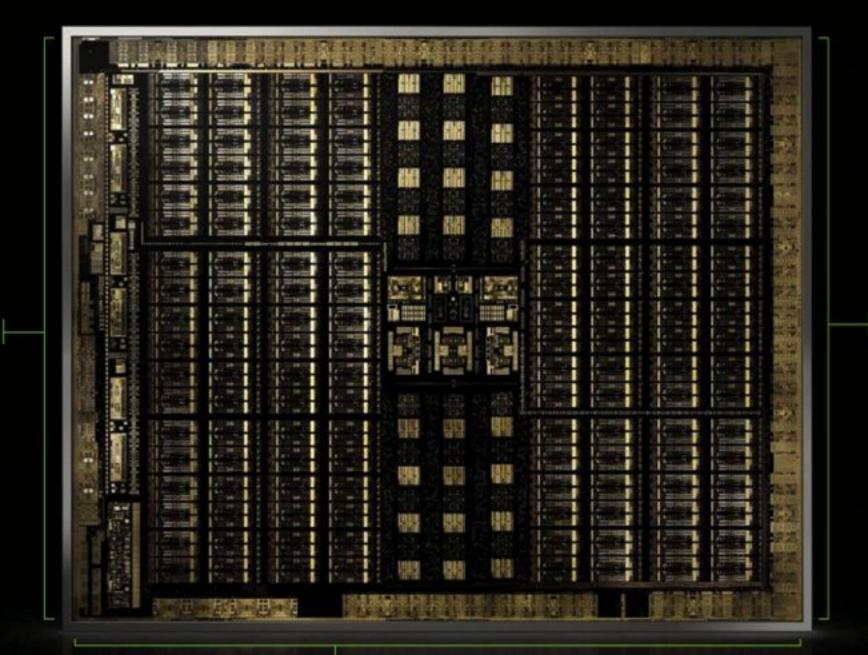




TURING BUILT FOR RTX Greatest Leap Since 2006 CUDA GPU

Turing SM

14 TFLOPS + 14 TIPS Concurrent FP & INT Execution Variable Rate Shading



Tensor Core 114 TFLOPS FP16

228 TOPS INT8

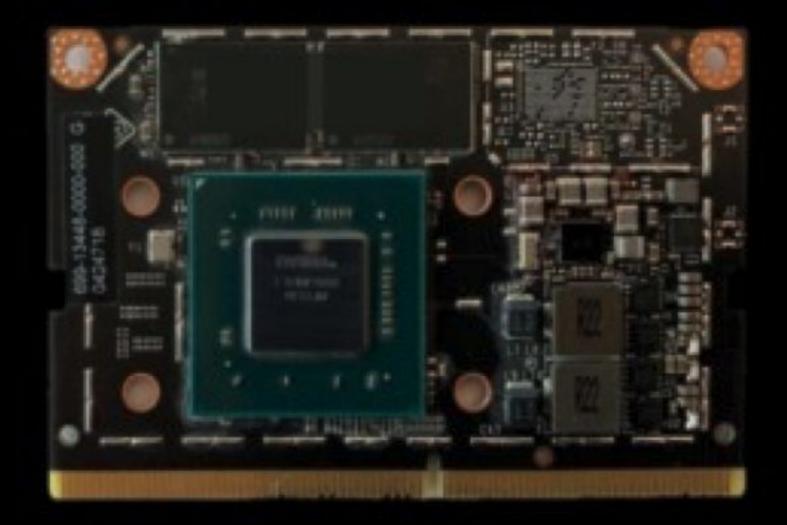
455 TOPS INT4

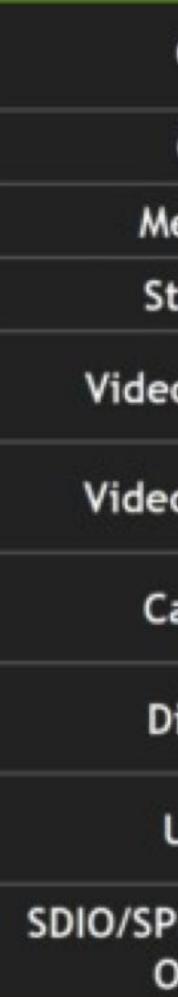
NGX

RT Core

10 Giga Rays/sec Ray Triangle Intersection BVH Traversal

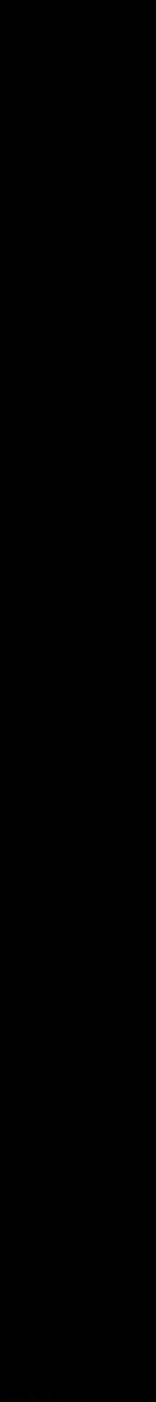
JETSON NANO SPECIFICATIONS



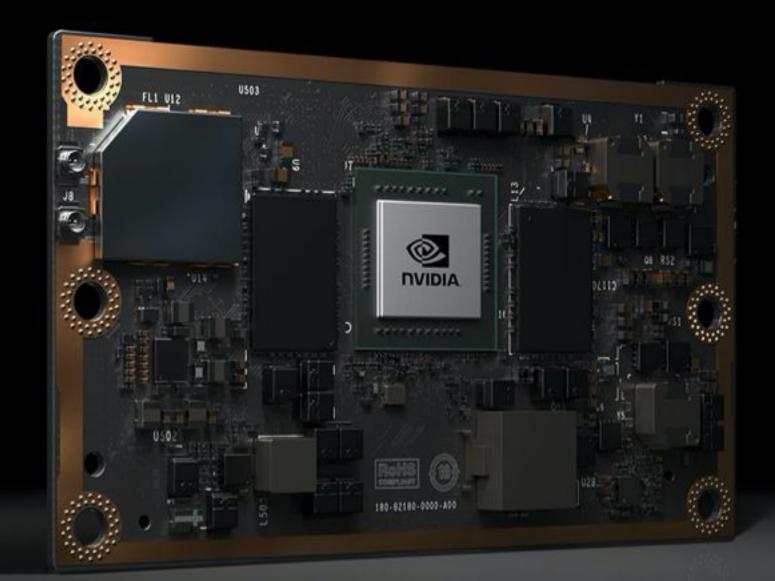


GPU	128 Core Maxwell 472 GFLOPs (FP16)
CPU	4 core ARM A57 @ 1.43 GHz
lemory	4 GB 64 bit LPDDR4 25.6 GB/s
torage	16 GB eMMC
eo Encode	4K @ 30 4x 1080p @ 30 8x 720p @ 30 (H.264/H.265)
eo Decode	4K @ 60 2x 4K @ 30 8x 1080p @ 30 16x 720p @ 30 (H.264/H.265)
Camera	12 (3x4 or 4x2) MIPI CSI-2 DPHY 1.1 lanes (1.5 Gbps)
Display	HDMI 2.0 or DP1.2 eDP 1.4 DSI (1 x2) 2 simultaneous
UPHY	1 x1/2/4 PCIE 1 USB 3.0
PI/SyslOs/GPI Os/I2C	1x SDIO / 2x SPI / 5x SysIO / 13x GPIOs / 6x I2C

DVIDIA



NVIDIA JETSON TX2

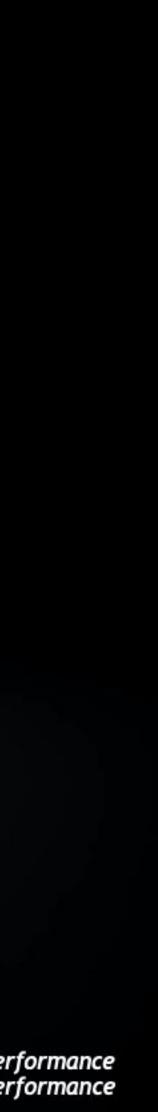




EMBEDDED AI SUPERCOMPUTER

Advanced AI at the edge JetPack SDK < 7.5 watts full module Up to 2X performance or 2X energy efficiency

Max-Q operating mode (< 7.5 watts) delivers up to 2x energy efficiency vs. Jetson TX1 maximum performance Max- P operating mode (< 15 watts) delivers up to 2x performance vs. Jetson TX1 maximum performance

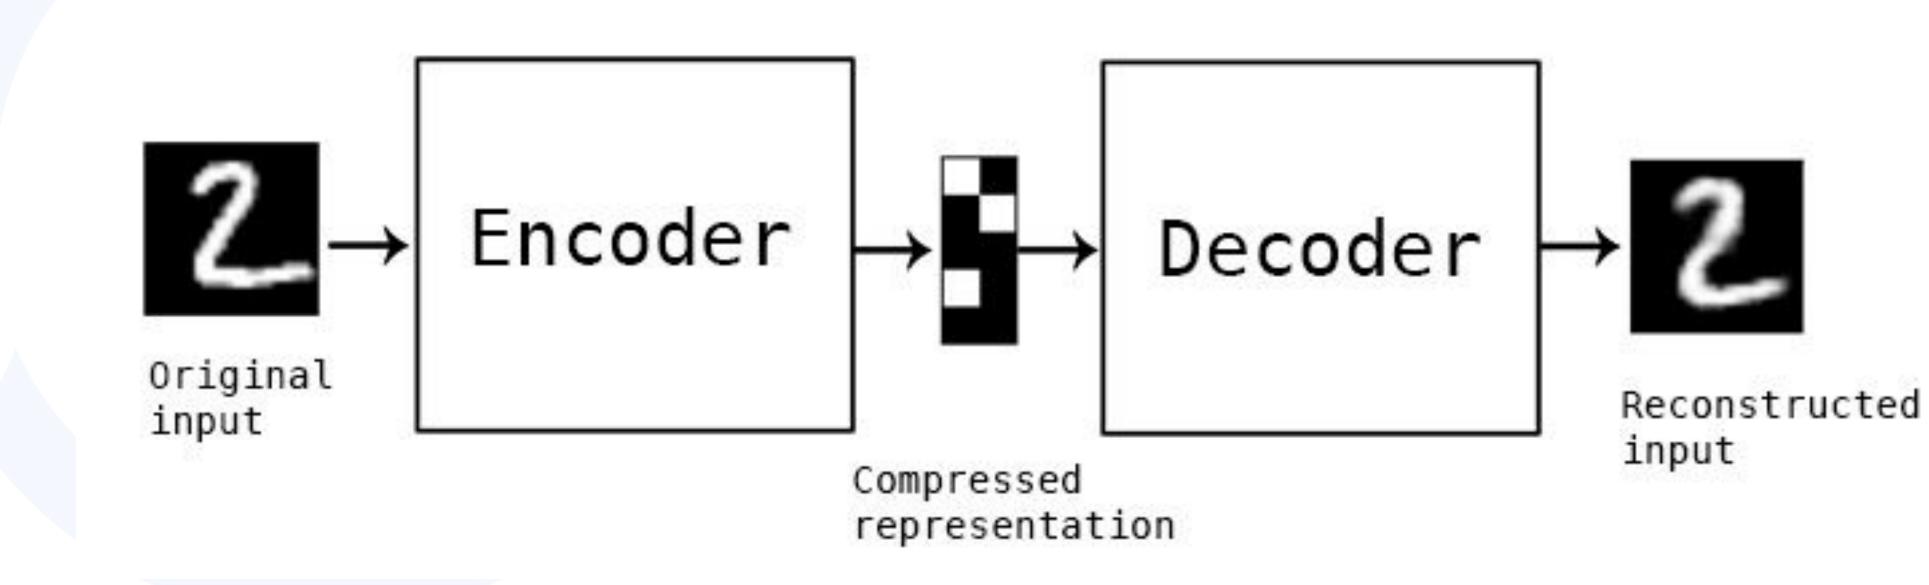


What's after CNNs? (an overview)





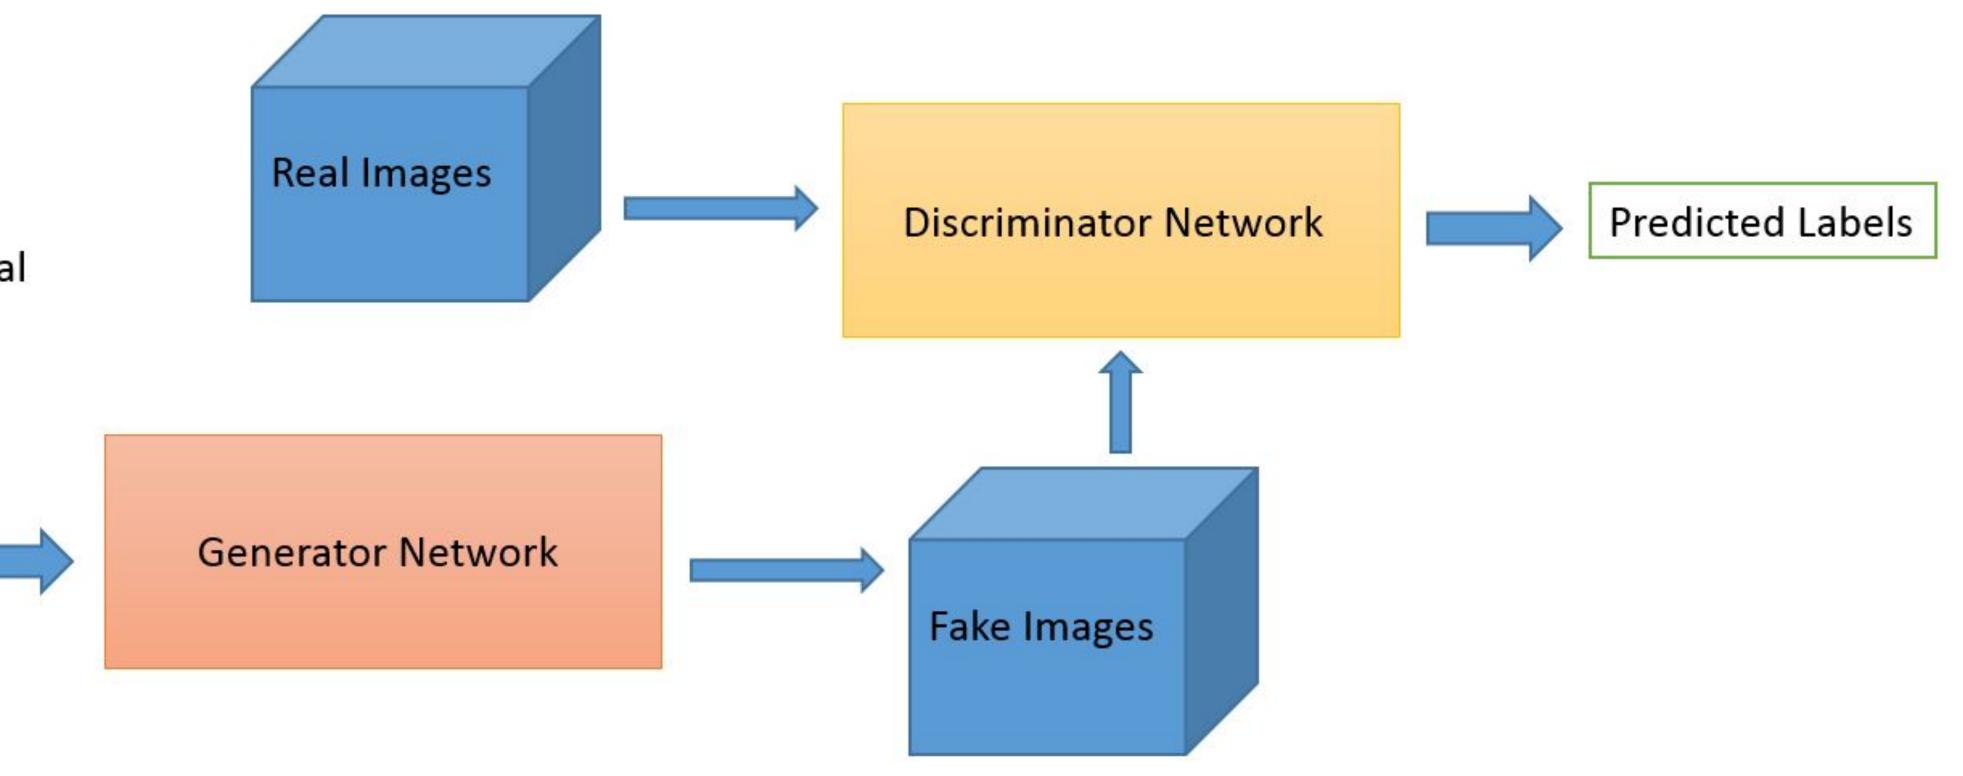
Used mainly for data compression, a network is trained to represent an image with less information than the original and reconstruct back the input minimizing the loss.



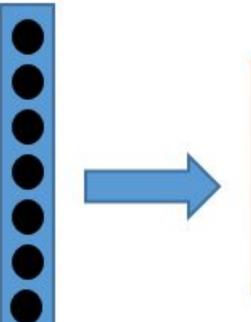


Generative Adversal Networks (GANs)

Generative networks train both a generator network to fake data from noise and detect whether an image is real or not



D-dimensional noise vector

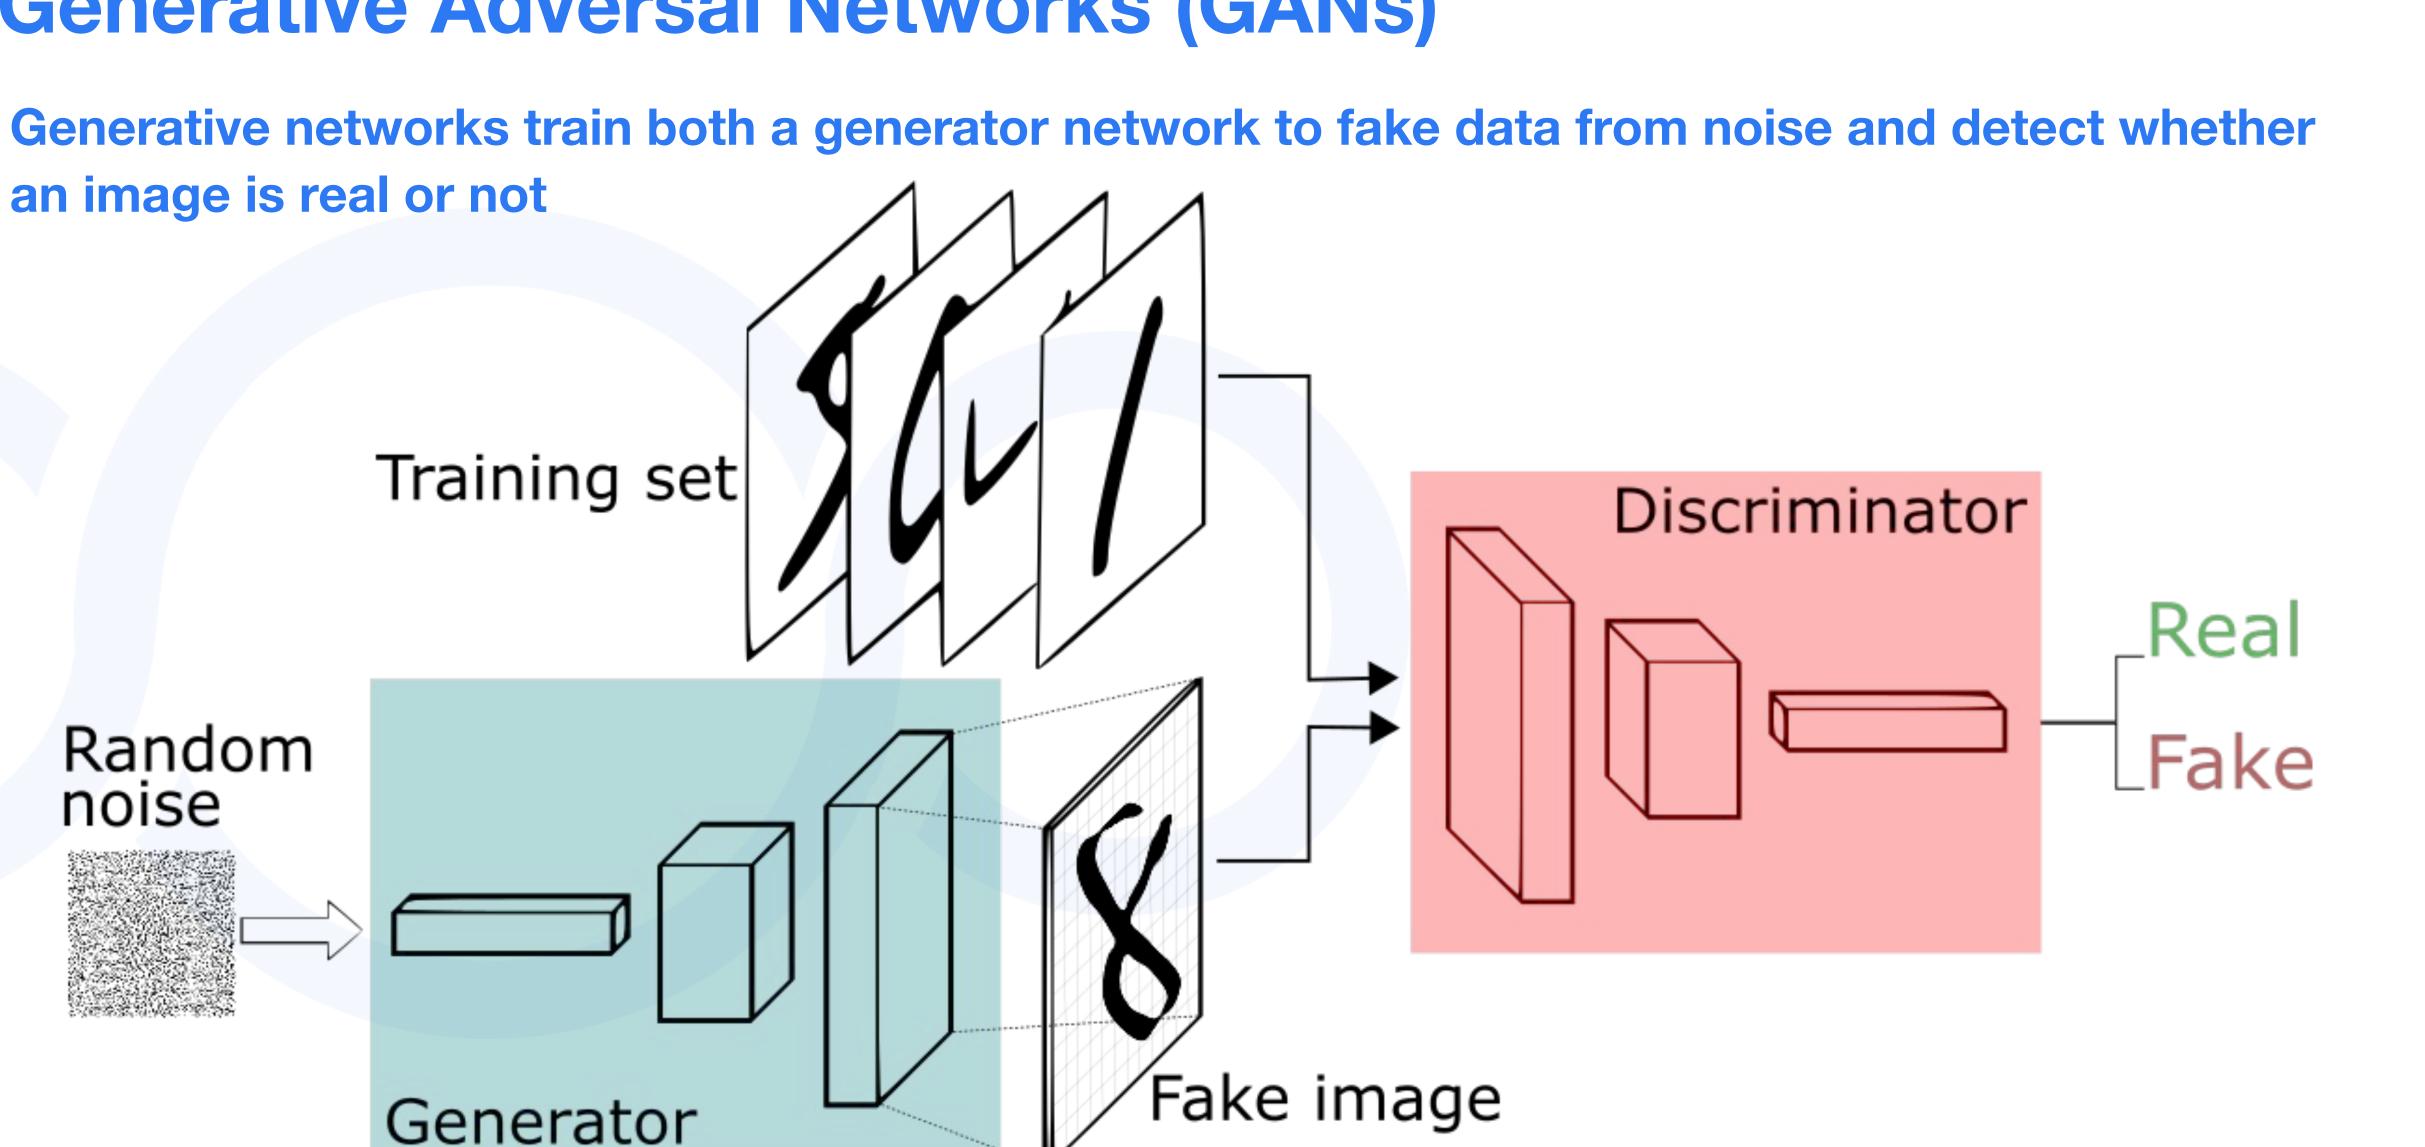






Generative Adversal Networks (GANs)

an image is real or not





GAN to generate artificial images

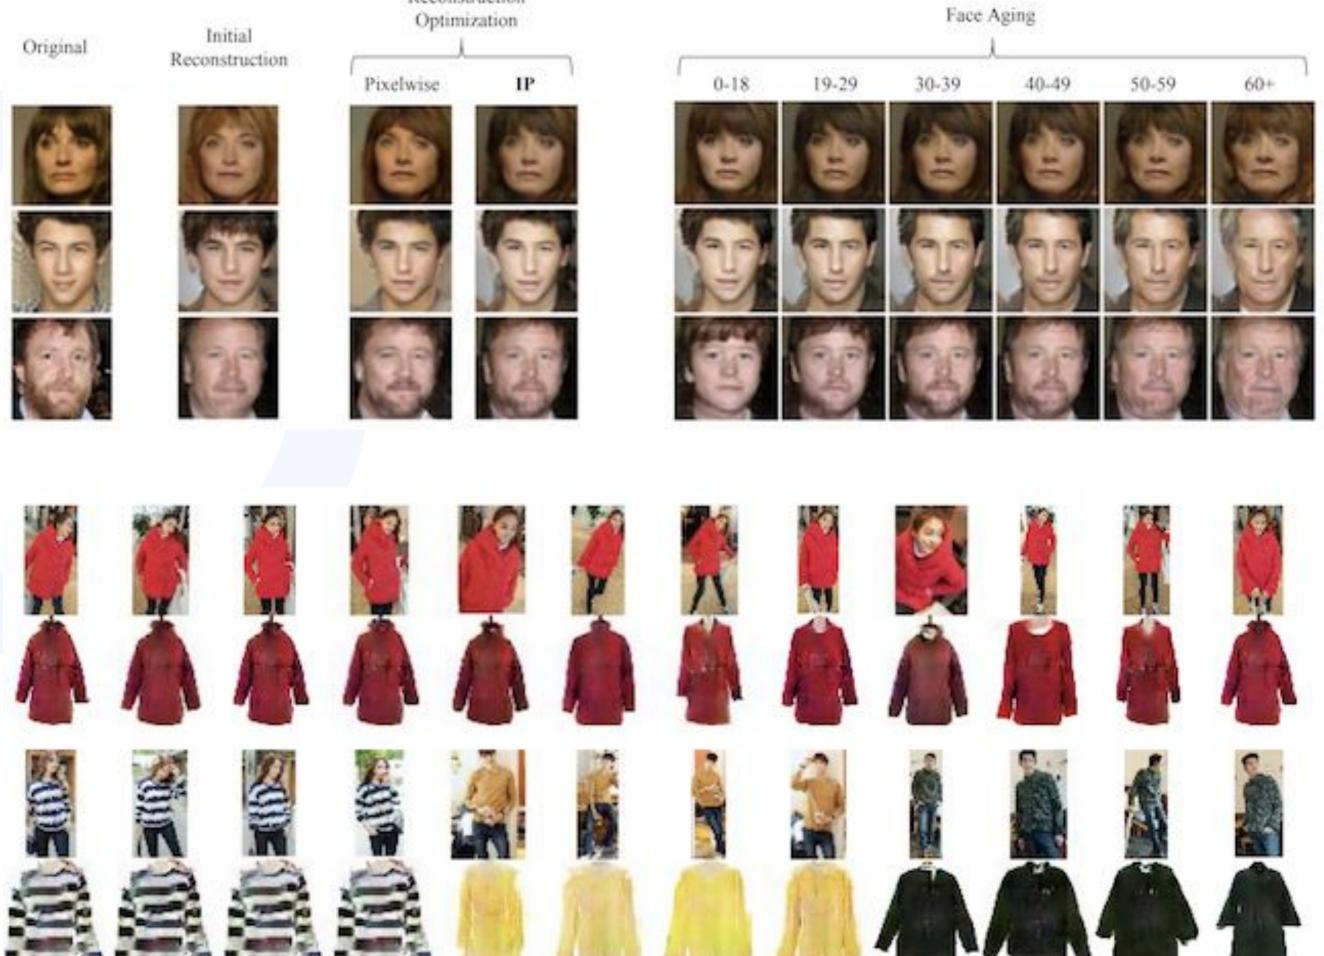
GAN can be used to simulate face aging of people in a natural and consistent way.

https://ieeexplore.ieee.org/document/8296650

GANs to generate photographs of clothing as may be seen in a catalog or online store, based on photographs of models wearing the clothing.

https://arxiv.org/abs/1603.07442 Neosperience





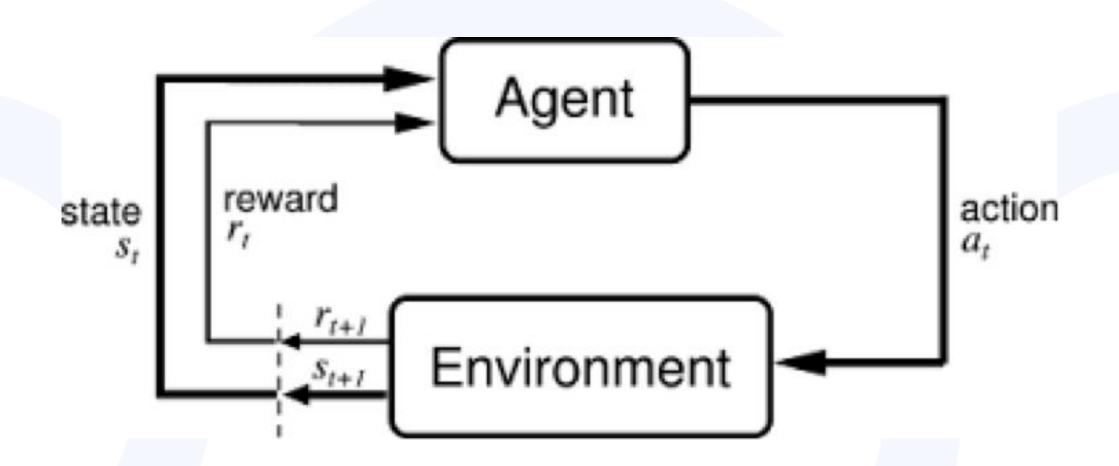






Reinforcement Learning

Used to train models in autonomous feedback autonomous driving agents.

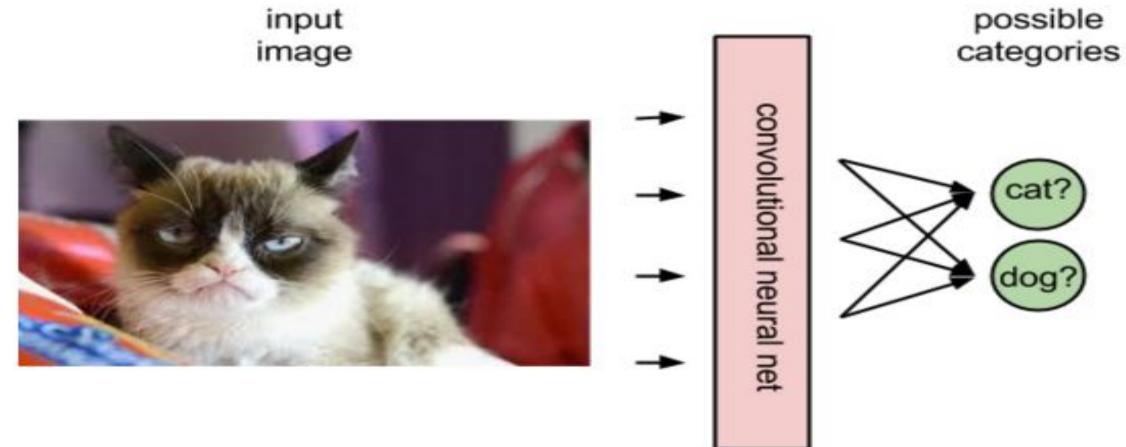


Reinforcement Learning has a wide range of applications from classification with a small dataset, to playing video games, firewall/system parameters tuning, personalizing reccomendations, automatic bidding.

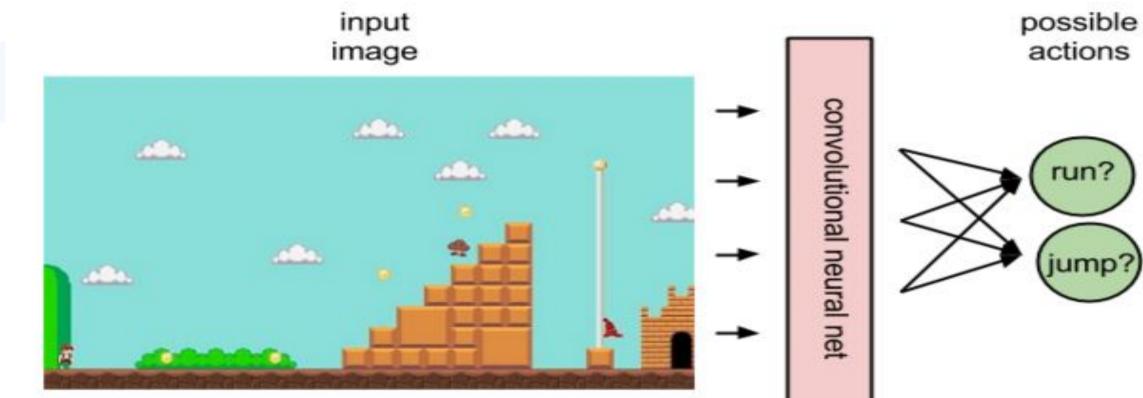
Neosperience

Used to train models in autonomous feedback-guided loops. It is used to implement variations of

Convolutional Classifier



Convolutional Agent





Amazon DeepRacer

Deploy a Deep Learning model on a Deep Racer, then win a race in a global competition

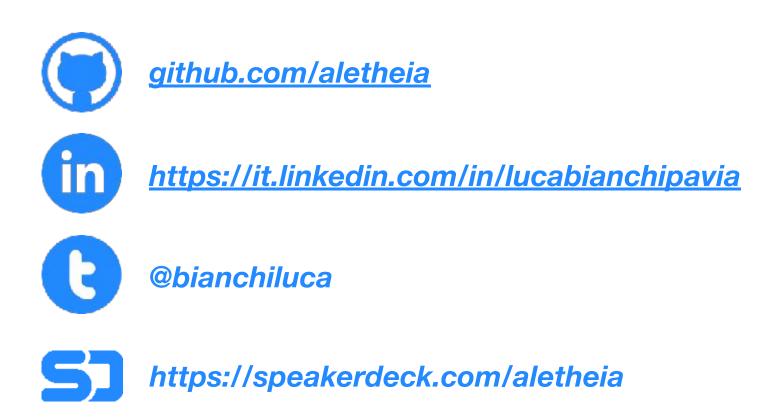
Developers of all skill levels can get hands on with machine learning through a cloud based 3D racing simulator, fully autonomous 1/18th scale race car driven by reinforcement learning, and global racing league.

AWS DeepRacer is an autonomous 1/18th scale race car designed to test RL models by racing on a physical track. Using cameras to view the track and a reinforcement model to control throttle and steering, the car shows how a model trained in a simulated environment can be transferred to the real-world.









Thank you. http://bit.ly/unipv-deep-cv









www.neosperience.com | blog.neosperience.com | info@neosperience.com

Neosperience