



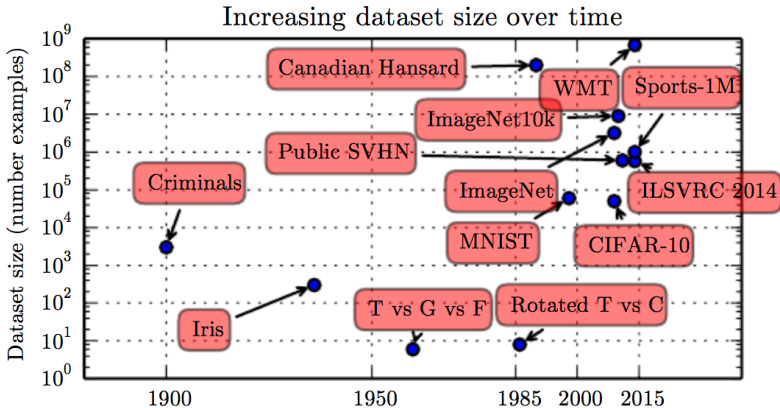
TRANSFER LEARNING

Vincent Barra
LIMOS, UMR 6158 CNRS, Université Clermont Auvergne

LACK OF UNLABELED DATA

Can we do Deep Learning with few labeled data?

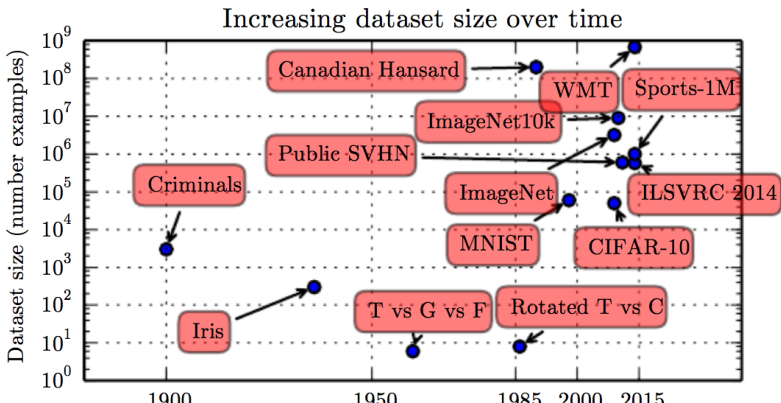
- ▶ Learn useful representation from unlabeled data
- ▶ Train on a nearby surrogate objective for which it is easier to generate labels



LACK OF UNLABELED DATA

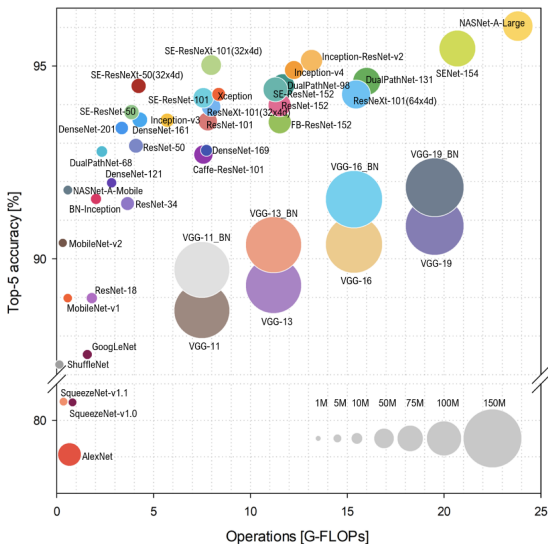
Can we do Deep Learning with few labeled data?

- ▶ Learn useful representation from unlabeled data
- ▶ Train on a nearby surrogate objective for which it is easier to generate labels
- ▶ **Transfer learned representation from a related task**



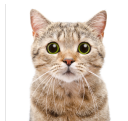
TRANSFER LEARNING

Training a Deep net from scratch on your dataset can be hard



TRANSFER LEARNING

Initial task/domain



Same domain



Same task



TRANSFER LEARNING

Concept

- ▶ Several networks have already been trained on a different domain for a different source task
- ▶ Adapt this network to the target class

Many variations

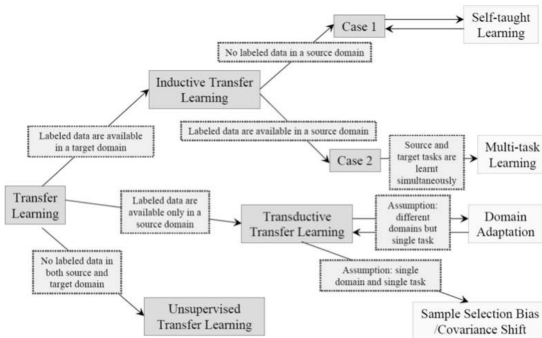
- ▶ Close domain, different task
- ▶ Different domain, same task
- ▶ Partial/full adaptation

See Lecture

“CNN Architectures”

TRANSFER LEARNING

A taxonomy of transfer Learning



Source: Pang & Yang, TKDE 2010)

We concentrate on domain adaptation/multi-task learning

DOMAIN ADAPTATION

Domain adaptation

- ▶ Domains are modeled as probability distributions over an instance space
- ▶ Task associated to a domain (classification, regression..)

The question

How can we learn a low-error classifier on a target data distribution, using labeled data from a source distribution ?

DOMAIN ADAPTATION

- ▶ $\mathcal{X}_S, \mathcal{X}_T$: source/target domain
- ▶ $\mathcal{Y}_S, \mathcal{Y}_T$: source/target label space

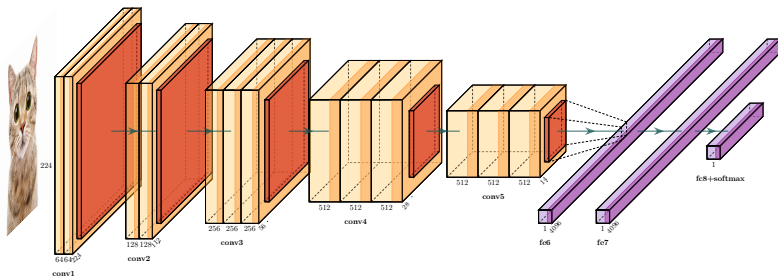
What can happen

- ▶ Data distribution change from \mathcal{X}_S to \mathcal{X}_T or $\mathbb{P}_S(x) \neq \mathbb{P}_T(x)$.
- ▶ Conditional probabilities may be different: $\mathcal{Y}_S \neq \mathcal{Y}_T$ or $\mathbb{P}_S(y|x) \neq \mathbb{P}_T(y|x)$



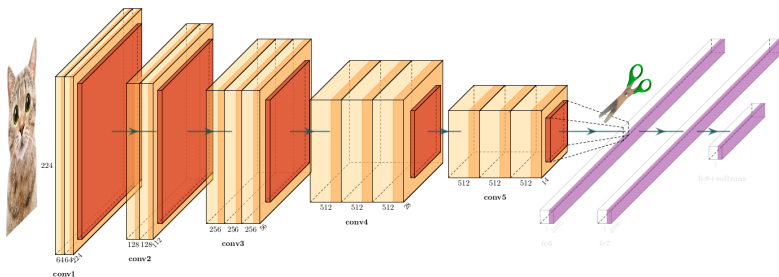
TRANSFER LEARNING

Take an already trained network



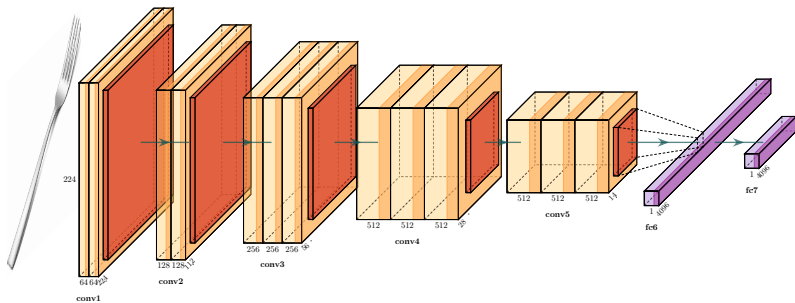
TRANSFER LEARNING

Use output of one or more layers as feature detector



TRANSFER LEARNING

Train a new classifier on these features



NOT ONLY IN COMPUTER VISION



Texts

??? **The end of the series.**
This book was written to provoke those who wanted Adams to continue the trilogy but I loved it. Author settled down on a bob fearing planet where he has acquired the prestigious...
[Read more](#)
Published on Mar 18 2002 by dan

??? **Mostly Harmless is Underrated**
I think most of the reviews for this book downplay it seriously. While the ending is kind of disappointing, the book overall is wonderful.
[Read more](#)
Published on Jan 22 2002 by A Big Adams Fan

??? **Please pretend this book was never written.**
I have long been a fan of the Hitchhikers series as they are comic genius. The book Mostly Harmless, however, should never have come about. It is frustration at its peak. [Read more](#)
Published on Jan 14 2002 by Paul Normd

??? **Kinda like horror movies...**
...in that the last one usually isn't all that appealing. I liked it fine, with some of Adams' wit, but it was a bit disappointing. [Read more](#)
Published on Nov 4 2001 by Kitesopher Vincent

??? **A Terrible End to A Great Series**
The ending for this books was so bad that I vowed never to read another Douglas Adams book. Adams was obviously sick and tired of the series and used this book to kill it off with...
[Read more](#)
Published on Oct 17 2001 by David A. Lesniou



Movies

-1 **An insult to Douglas Adams' memory**
I agree entirely with "darkgenius" comments. This movie is a travesty of the book and the TV series; a cutesy version totally lacking in the wit and satire of the original. [Read more](#)
Published 5 months ago by John W Beane

+1 **Don't Panic!**
If you haven't listened to the BBC radio-play, this isn't bad! Purists, no doubt, will dispute my verdict but the fact of the matter is THGTTG (see title) does have Douglas Adams'...
[Read more](#)
Published on Mar 13 2011 by Sid Matheson

+1 **On Blu-ray, even better**
I've seen this movie on TV and wanted to add it to my collection. I couldn't find it locally so when I saw it on amazon and on Blu-ray, I picked it up. [Read more](#)
Published on April 18 2009 by J. W. Little

-1 **An insult to Douglas Adams' memory**
The filmmaker's reverence for Adams' legacy? What kind of rubbish statement is that? As a loyal fan of Douglas Adams for more than a quarter of a century, I was appalled and...
[Read more](#)
Published on Aug 22 2006 by Daniel Jolley

Algorithm

Classifier

TRANSFER LEARNING

It works well...

... But can we do better than off-the-shelf features ?

TRANSFER LEARNING

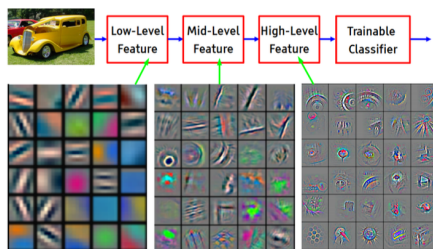
It works well...

... But can we do better than off-the-shelf features ?

Fine tuning

- ▶ Change the classification layer to match the problem
- ▶ Retrain some/all layers of the whole network

FINE TUNING



What layers to choose ?

- ▶ In computer vision
 - First layers detect simpler and more general patterns
 - Deeper layers capture more specific patterns related to data
- ⇒ Allow the last block(s) of convolution/pooling to be retrained.
- ▶ In sequential data: may keep the last few layers

FINE TUNING

Freeze / Finetune

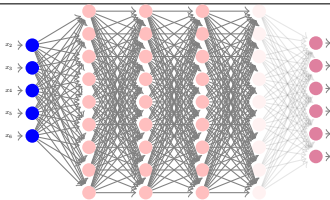
- ▶ Frozen layers: not updated during backpropagation
 - ▶ Finetuned layers: updated during backpropagation
 - ▶ Depends on the target task:
 - freeze if target task labels are scarce and no overfitting
 - finetune if more target labels.
- set learning rates to be different for each layer to find a tradeoff



Taille des données

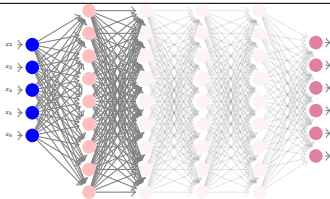
Réseau

Petite



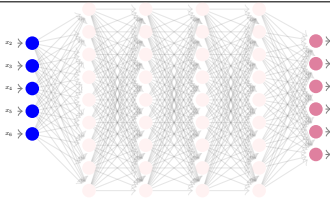
Geler toutes les couches, entraîner
le nouveau classifieur

Moyenne



Geler certaines couches

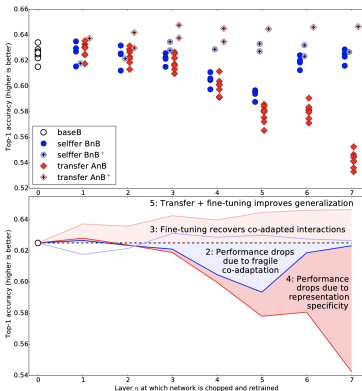
Grande



Geler toutes les couches.
Initialisation avec les poids pré-entraînés.

FINE TUNING

- ▶ If sufficient examples are available, fine tuning improves generalization
- ▶ Transfer learning and fine tuning can serve as an initialization process
- ▶ Very often better performance than training from scratch



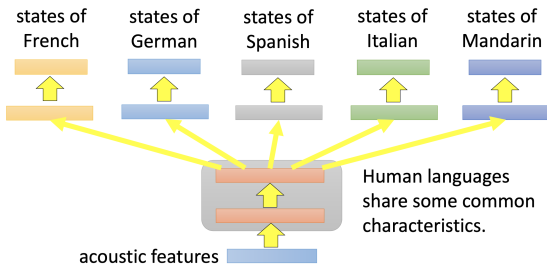
Yosinski et al, NIPS 2014

Source: 500 classes from ImageNet

Target, another 500 classes from ImageNet.

MULTITASK LEARNING

The multilayer architecture of Deep Neural Network makes them suitable for multitask learning.



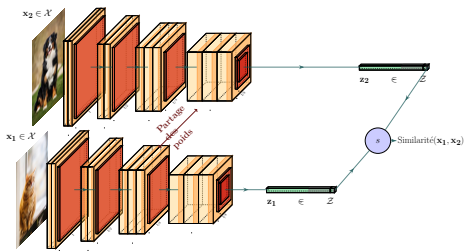
Huang & al, ICASSP 2013.

The training task can also depend on the number of training examples in the target domain.

- ▶ Sufficient number of examples : OK
- ▶ No examples: unsupervised domain adaptation
- ▶ Few number of examples (Few-shot learning)
 - Embedding learning
 - Data augmentation
 - Data generation
 - Semi supervised domain adaptation

EMBEDDING LEARNING

Embedding learning: siamese, triplet network...



See lecture

Matching networks

DATA AUGMENTATION

Given one example, generate n new ones using transformations (rotation, scaling, noise adding, nonlinear transformations, color processing...)



DATA GENERATION

Model the data distribution into the target space in order to be able to generate new and unseen samples.

- ▶ Generative Adversarial Networks (GANs)
- ▶ Variational Autoencoders
- ▶ ...

See lecture

Generative models

UN- OR SEMI-SUPERVISED DOMAIN ADAPTATION

- ▶ Matching source distributions
- ▶ Combination of fine tuning & unsupervised adaptation

Out of the scope of this lecture.

TRAINABLE/NON TRAINABLE

$$y = W^T x + b$$

```
dense = tf.keras.layers.Dense(7)
dense.build((None,5))
dense.trainable=False
print("weights:", len(dense.weights))
print("trainable_weights:", len(dense.trainable_weights))
print("non_trainable_weights:", len(dense.non_trainable_weights))
print(dense.weights)

weights: 2
trainable_weights: 0
non_trainable_weights: 2
<tf.Variable 'kernel:0' shape=(5, 7) dtype=float32, numpy=
array([[ 0.19174272,  0.1961686, -0.5426951,  0.27109373,  0.23231155,
         0.18705916, -0.15816867],
       [ 0.6983326,  0.702092,  0.67835873, -0.2961814, -0.49993056,
        -0.6466697,  0.3440327 ],
       [-0.1576004,  0.1708141,  0.01579124, -0.5328252, -0.27179474,
        -0.0705555, -0.7060735 ],
       [ 0.03460306,  0.03050655,  0.46395332, -0.29719362,  0.4411903,
        -0.17896783, -0.59996223],
       [-0.48393145,  0.38511735, -0.43845502, -0.485897,  0.6990538,
        -0.6045918,  0.4238811 ]], dtype=float32)>, <tf.Variable 'bias:0' shape=(7, ) dtype=float32, numpy=array([0., 0., 0., 0., 0., 0., 0.], dtype=float32)>
```

TRAINABLE/NON TRAINABLE

$$y = W^T x + b$$

```
dense = tf.keras.layers.Dense(7)
dense.build((None,5))
dense.trainable=False
print("weights:", len(dense.weights))
print("trainable_weights:", len(dense.trainable_weights))
print("non_trainable_weights:", len(dense.non_trainable_weights))
print(dense.weights)

weights: 2
trainable_weights: 0
non_trainable_weights: 2
<tf.Variable 'kernel:0' shape=(5, 7) dtype=float32, numpy=
array([[ 0.19174272,  0.1961686, -0.5426951,  0.27109373,  0.23231155,
         0.18705916, -0.15816867],
       [ 0.6983326,  0.702092,  0.67835873, -0.2961814, -0.49993056,
        -0.6466697,  0.3440327 ],
       [-0.1576004,  0.1708141,  0.01579124, -0.5328252, -0.27179474,
        -0.0705555, -0.7060735 ],
       [ 0.03460306,  0.03050655,  0.46395332, -0.29719362,  0.4411903,
        -0.17896783, -0.59996223],
       [-0.48393145,  0.38511735, -0.43845502, -0.485897,  0.6990538,
        -0.6045918,  0.4238811 ]], dtype=float32)>, <tf.Variable 'bias:0' shape=(7, 7) dtype=float32, numpy=array([0., 0., 0., 0., 0., 0., 0.], dtype=float32)>
```

IN KERAS

```
# Load a pretrained model
pretrained_model = tf.keras.applications.MobileNetV2(input_shape=[*image_size, 3],
                                                    include_top=False)

# Suppress the classification layer
pretrained_model.layers.pop()

# Freeze all or a subset of weights
for layer in pretrained_model.layers[:-stop_freeze]:
    layer.trainable = False

# Put our own classification layers
model = tf.keras.Sequential([pretrained_model, tf.keras.layers.Flatten(),
                             tf.keras.layers.Dense(5, activation='softmax')])
```


IN KERAS

Transfer Learning + Fine tuning

```
# Transfer Learning
pretrained_model = keras.applications.MobileNetV2(input_shape=(image_size,3),
                                                    include_top=False) # no classifier at the top.
pretrained_model.trainable = False

model = tf.keras.Sequential([pretrained_model, tf.keras.Flatten(),tf.keras.layers.Dense(5,activation='softmax')])
model.compile...
model.fit...

# Fine tuning
pretrained_model.trainable = True
model.compile... # Important to call it again if Trainable is changed
model.fit...
```