Machine Learning Graz

Transfer Learning

31-Jul-2018

Adrian Spataru
Data Scientist at Know-Center
adrian@spataru.at
https://www.fb.me/adrian.spataru.5

Outline

- Quick Neural Network Refresher
- What is Transfer Learning?
  - Bottlenecking & Fine-Tuning
  - Multitask Learning
  - Domain-adversarial Training
  - Zero-Shot Learning
- Resources for Pretrained models

Perceptron

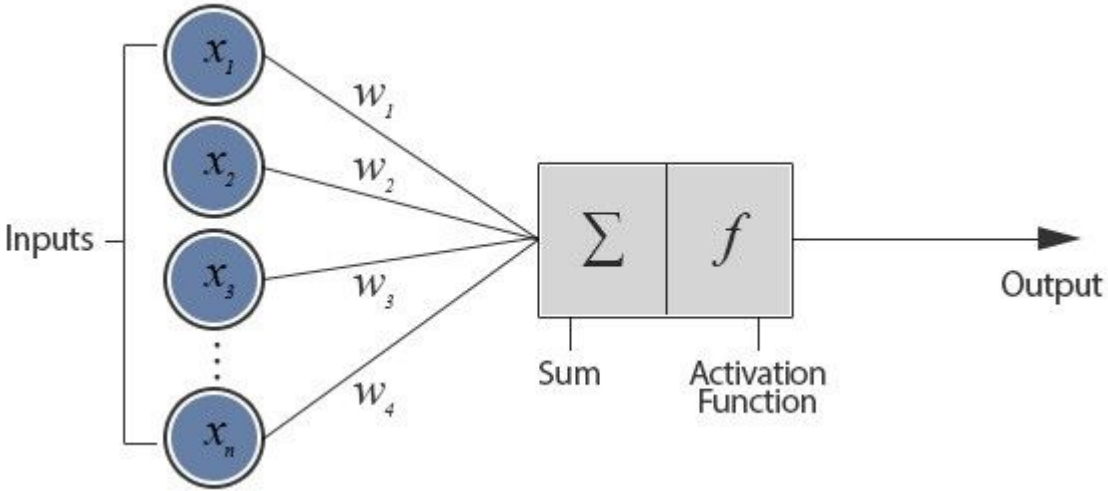
Inputs
x_1
x_2
x_3
x_n
w_1
w_2
w_3
w_4
Σ
f
Sum
Activation Function
Output

Multilayer Perceptron

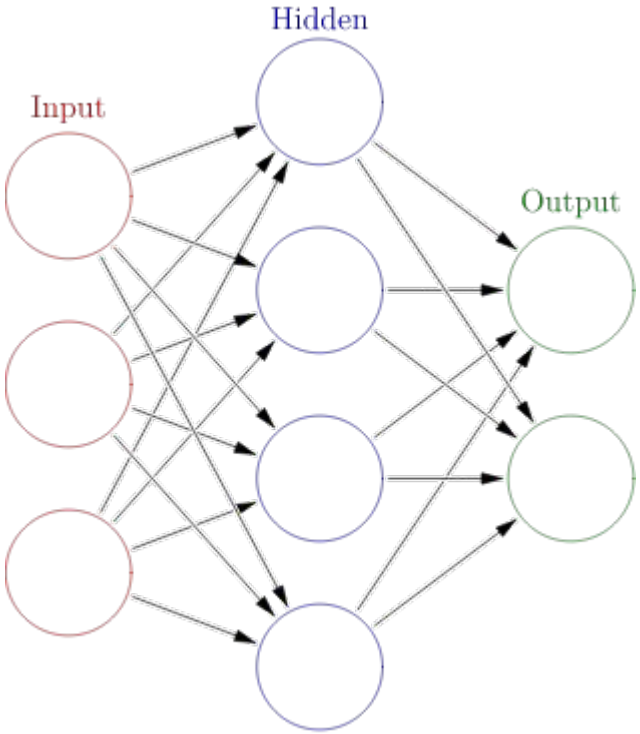

Convolutional Neural Network

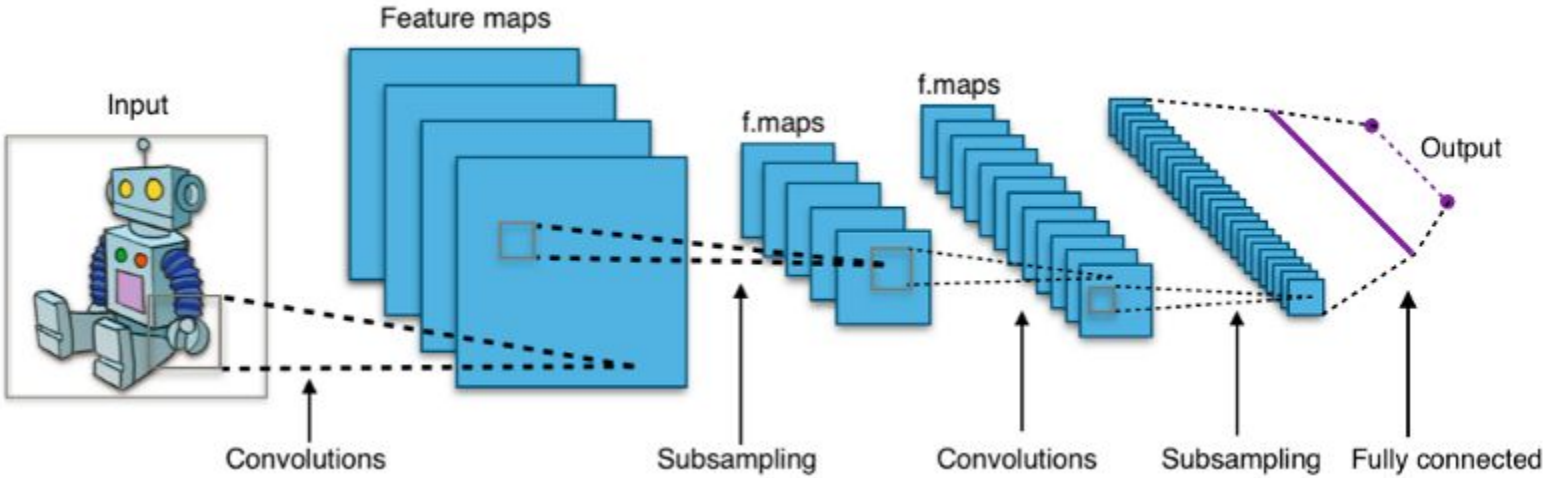

Traditional ML

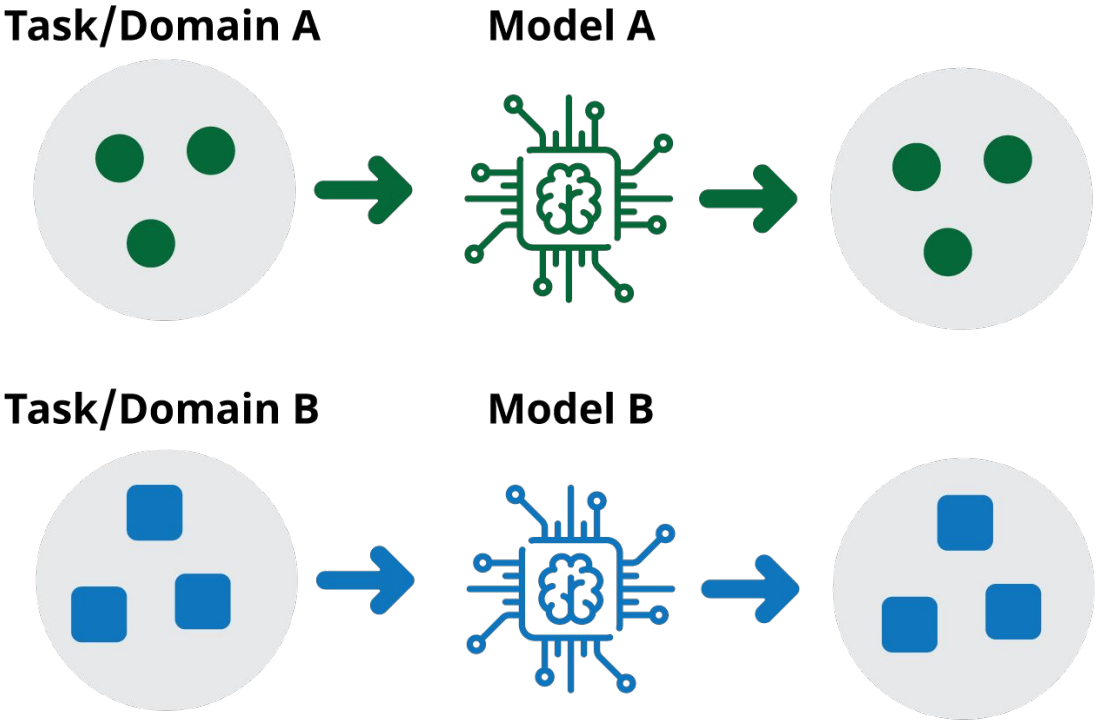

Transfer Learning

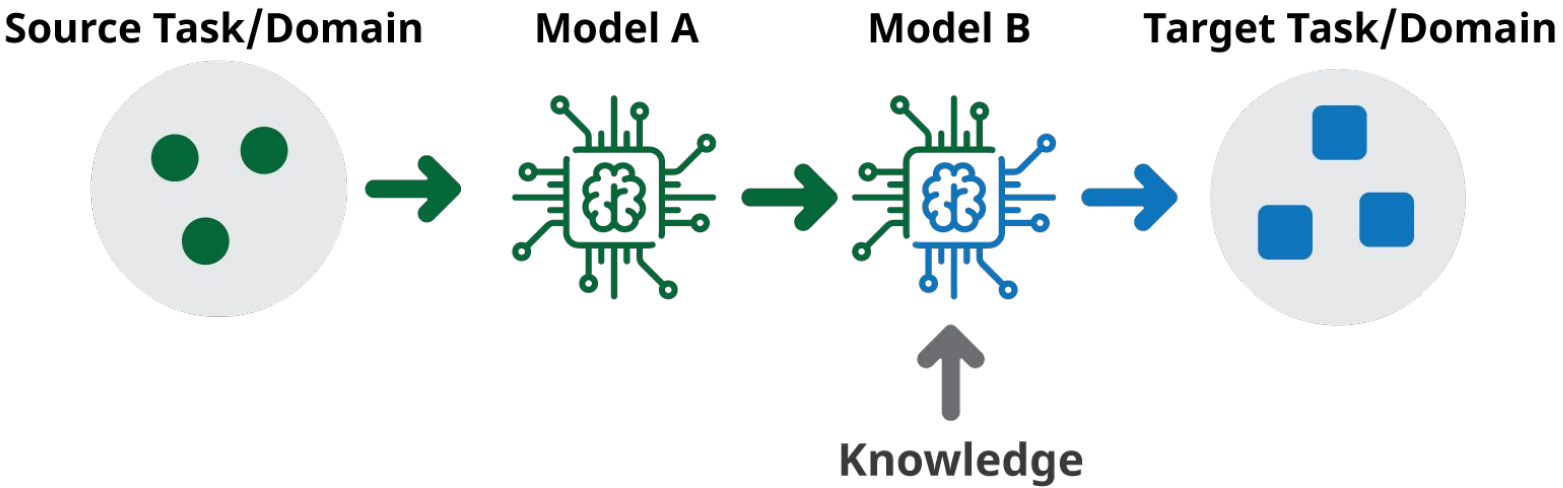

Transfer Learning Landscape

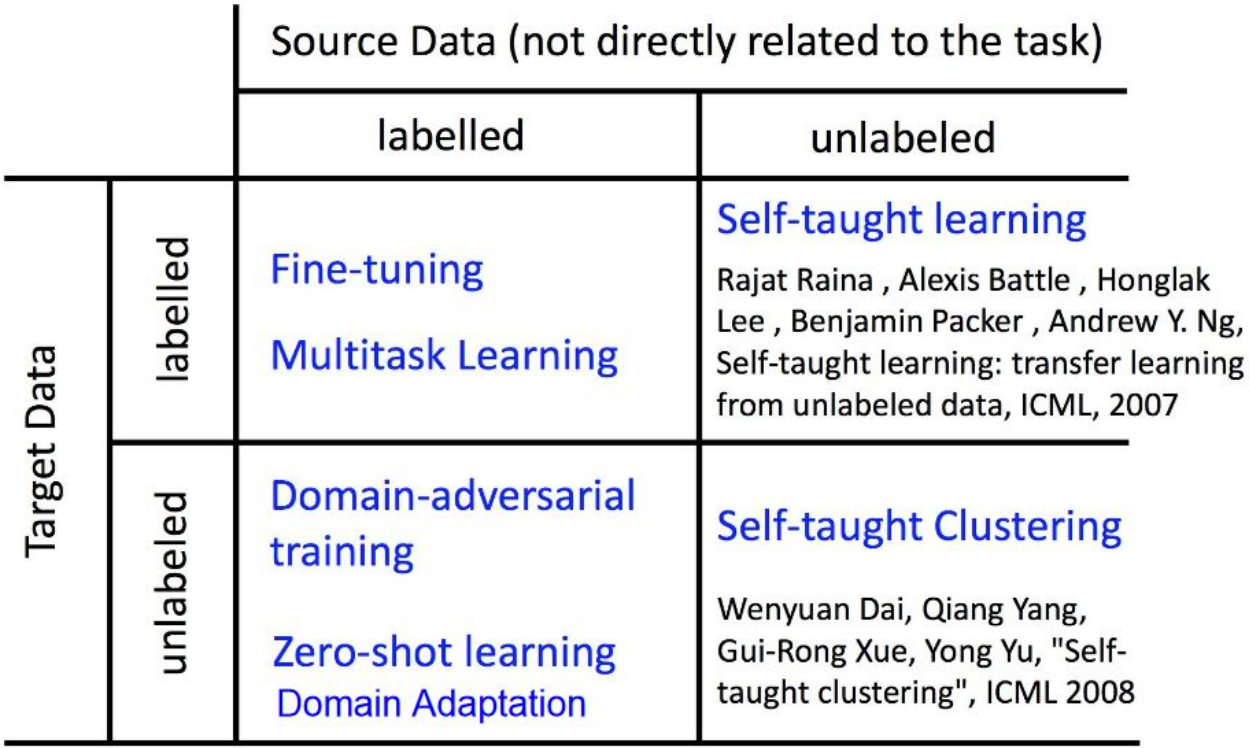

Transfer Learning Tutorial (Hung-yi Lee)

Fine-Tuning

- Scenario
  - Lot of labeled source data
  - limited labeled target data
- Idea: train a model by source data, then fine-tune the model with the target data.
- Why?
  - Training on target data only, will likely overfit.
  - May reduce training time with pretrained models

Fine-Tuning Example

- Kaggle Imaterialist Challenge
- Labeling of Household items (glass,chair etc)
- Multiclass Classification Problem
- 128 Classes
- Training 190k Images

Typical Image Classifier

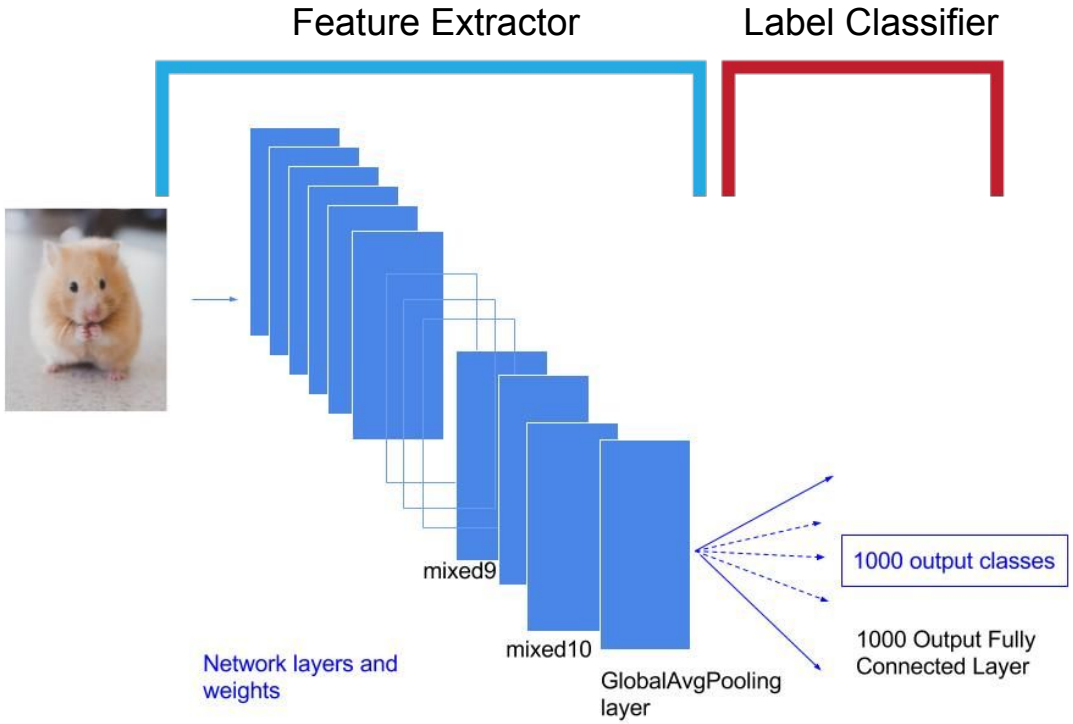

Fine-Tuning

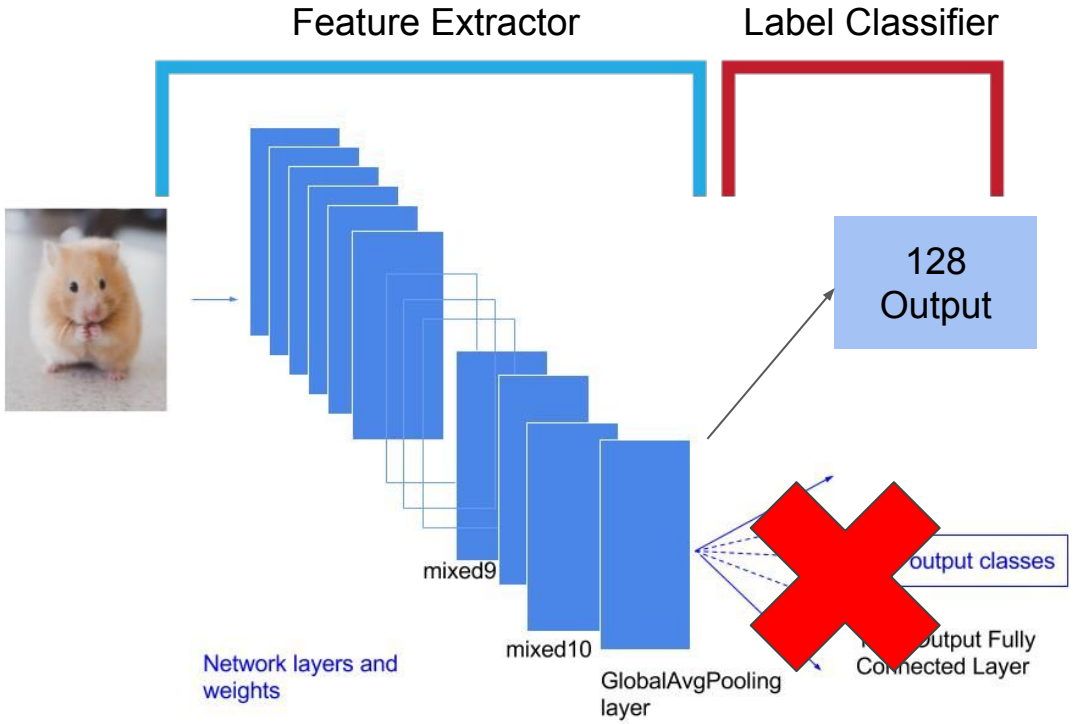
Feature Extractor
Label Classifier
128
Output
mixed9
mixed10
GlobalAvgPooling
layer
output classes
Network layers and
weights

Results

- Using a pretrained ResNet50 model on ImageNet and finetuning
  - 84% Accuracy - Top 20%
- Avg Ensemble of 11 different pretrained models and finetuning
  - 89% Accuracy - First Place
- ResNet50 takes 4 days to train on a GTX1060

Bottlenecking - Fine-Tuning of the poor.

- Idea: train a model by source data -> use the model to extract features for the target data -> train a new model with the extracted features

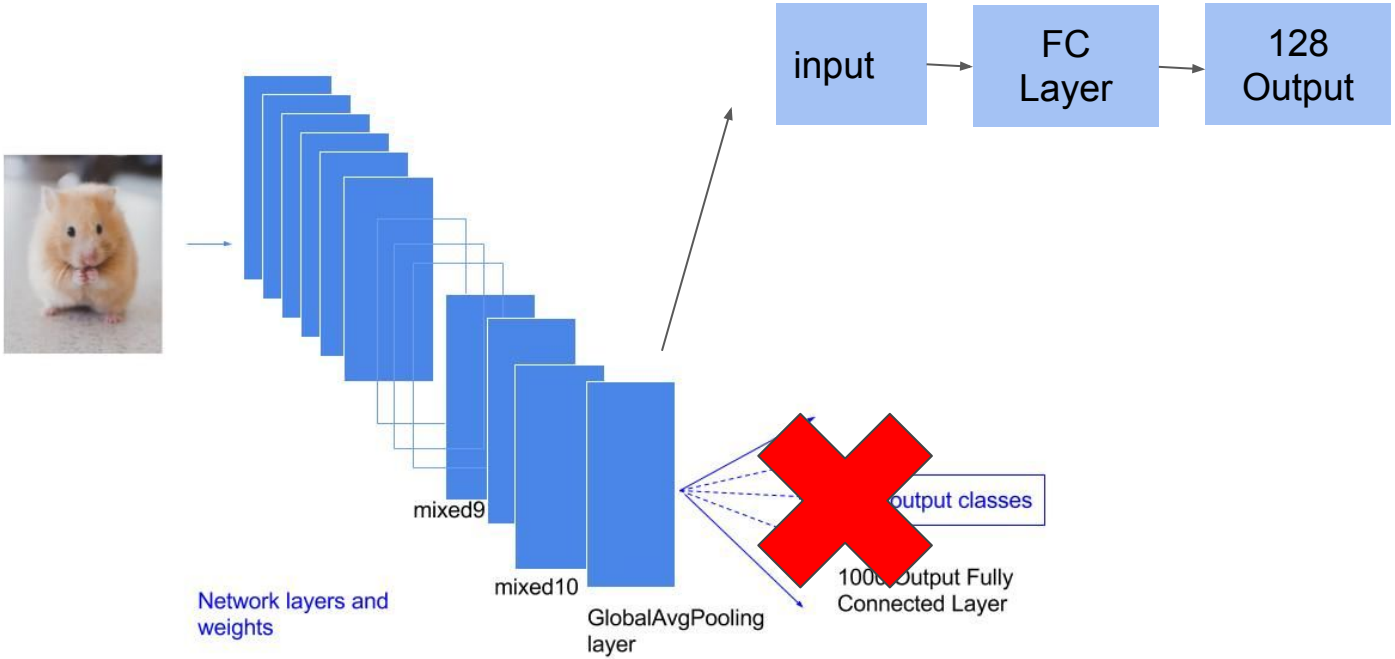

Results

- Initial feature extraction takes 6 hours on a GTX 1060
- Training a models takes 20min
- Reduces the data from 30gb to 2-3gb
- 82% Accuracy - Avg Ensemble of 5 bottleneck models
- Top 38% Place - 159/436

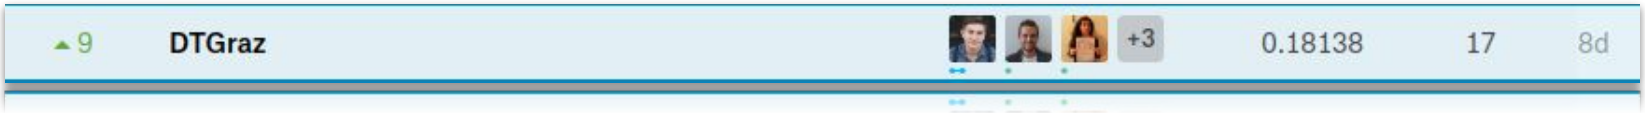

CODE DEMO


AND NOW FOR SOMETHING
COMPLETELY DIFFERENT

Multitask Learning

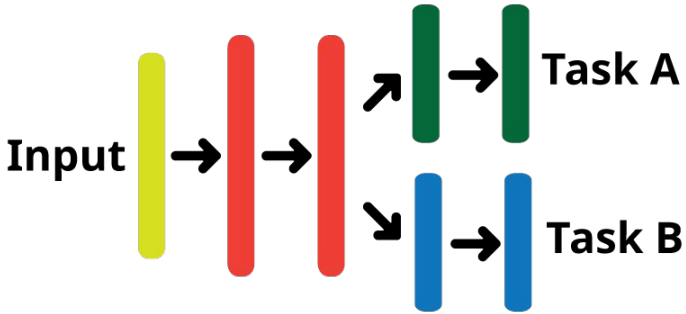

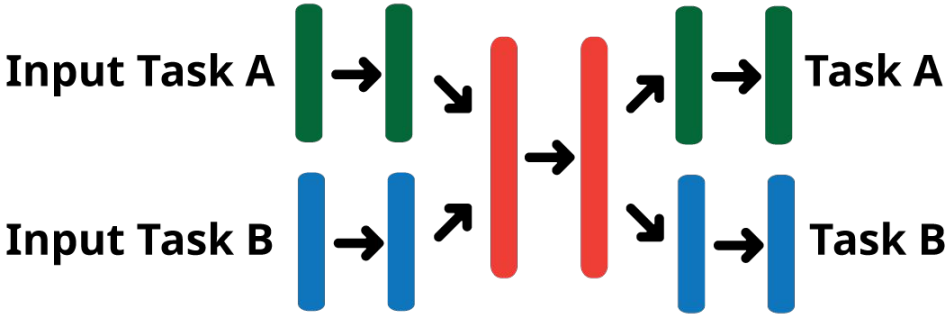

Multitask Learning (2)

- Why consider it?
  - Attention focusing
  - Eavesdropping
  - Regularization

- Lot of Architectures
  - Cross-stitch Networks
  - Fully-Adaptive Feature Sharing
  - etc

Multiple Language Translation

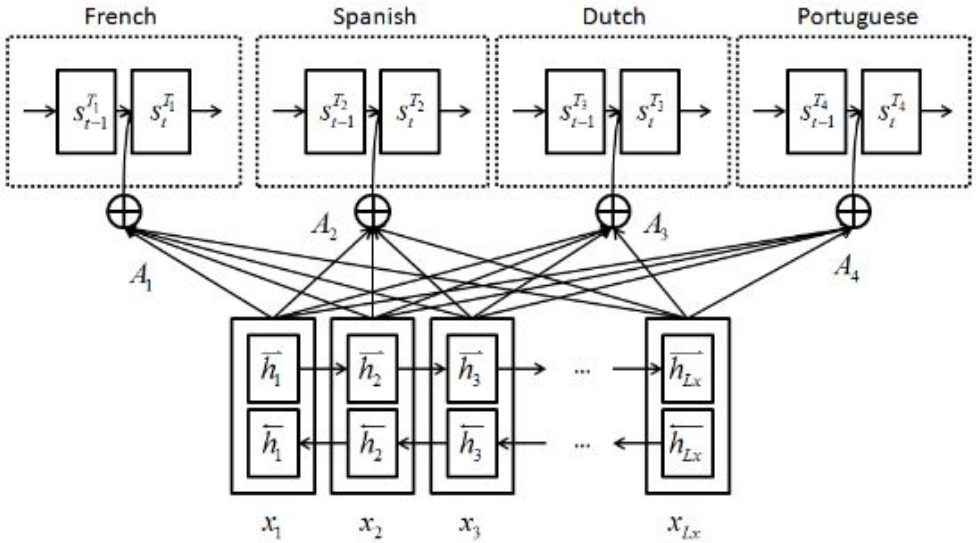

Figure 2: Multi-task learning framework for multiple-target language translation

Lang-Pair	En-Es	En-Fr	En-Nl	En-Pt
Single NMT	26.65	21.22	28.75	20.27
Multi Task	28.03	22.47	29.88	20.75
Delta	**+1.38**	**+1.25**	**+1.13**	**+0.48**

Table 3: Multi-task neural translation v.s. single model given large-scale corpus in all language pairs

Domain Adversarial Training

- Scenario
  - lot of labeled source data
  - lots of unlabel target data
- Goal: Train a model which performs well on unlabeled data.

Domain Adversarial Training

- Scenario
  - lot of labeled source data
  - lots of unlabel target data
- Goal: Train a model which performs well on unlabeled data.
- Goal 2.0: The distribution of the features extracted are similar

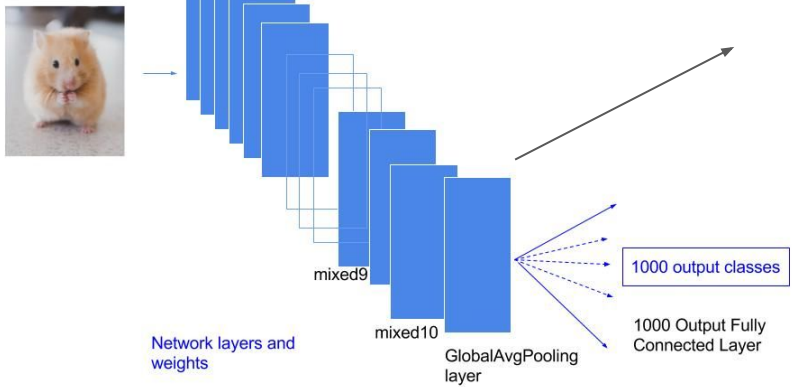

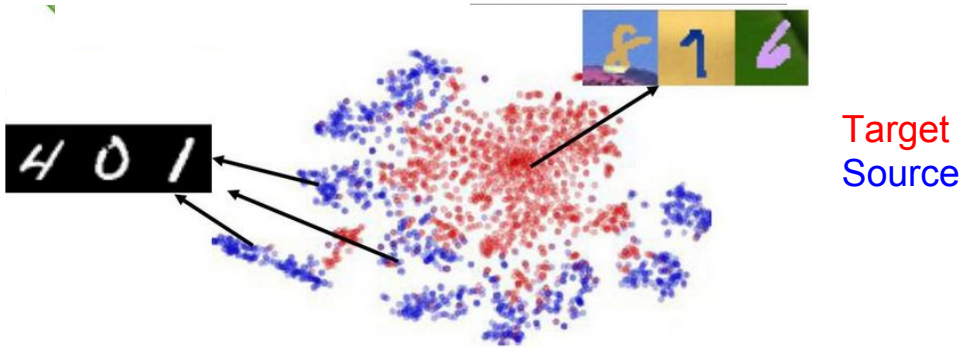

Target
Source

Domain Adversarial Training

- Scenario
  - lot of labeled source data
  - lots of unlabel target data
- Goal: Train a model which performs well on unlabeled data.
- Goal 2.0: The distribution of the features extracted are similar

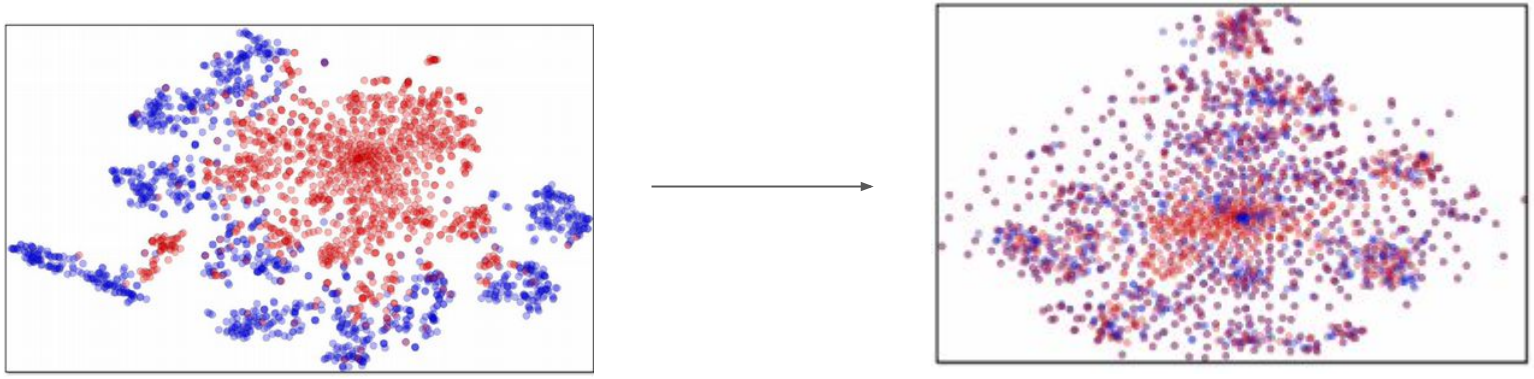

Target
Source

Domain Adversarial Training

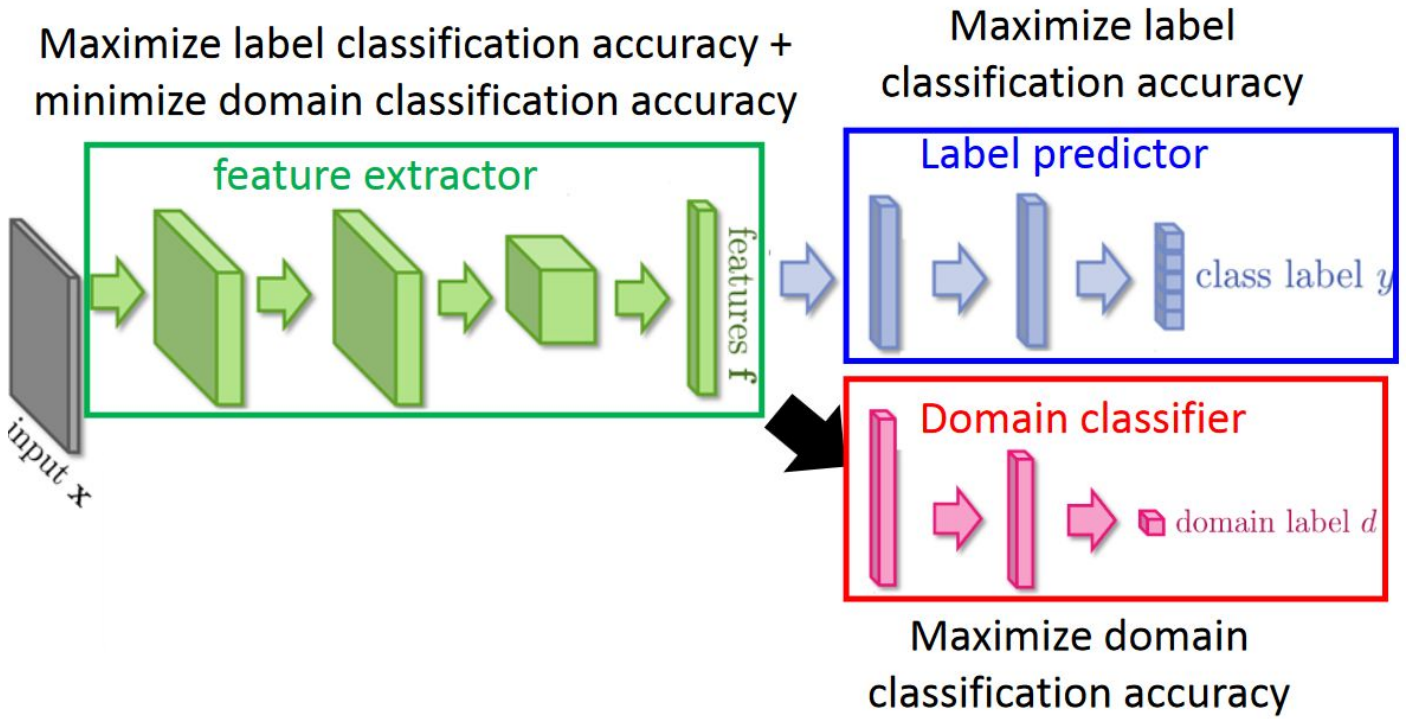

Transfer Learning (Hung-yi Lee)

Domain Adversarial Training Example

No Adapt - 87% Acc
With Adapt - 91% Acc

Zero Shot Learning

- Scenario
  - lot of labeled source data
  - unlabeled target data
- Goal: Train a model which performs well on target data.
- How? Inference through attributes, metadata etc

Zero Shot Learning Attributes

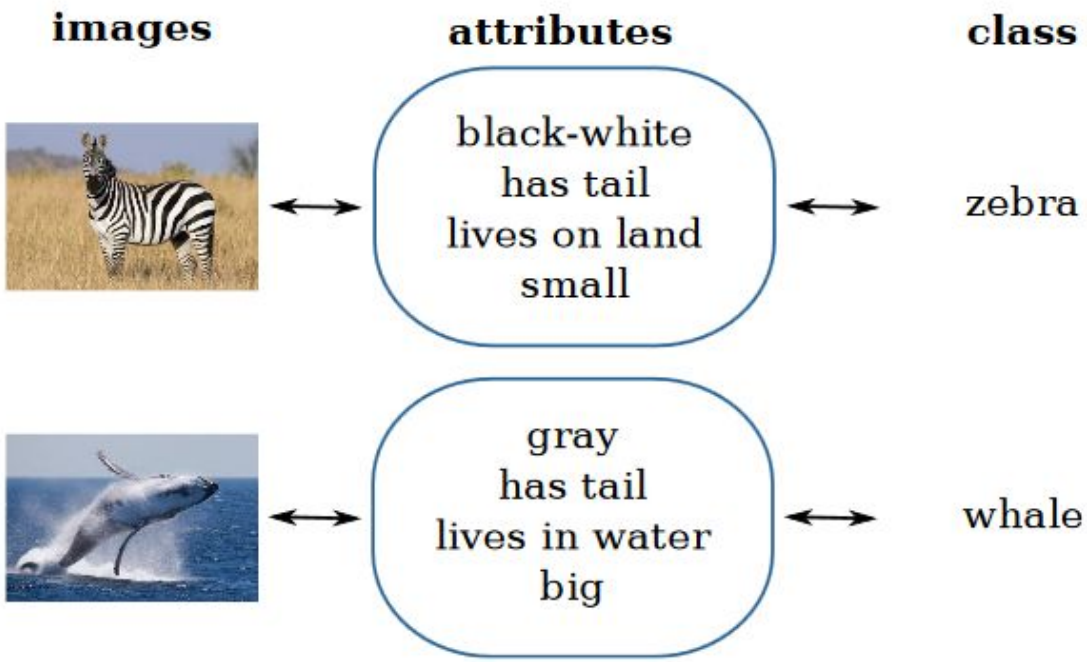

Zero Shot Learning Attributes

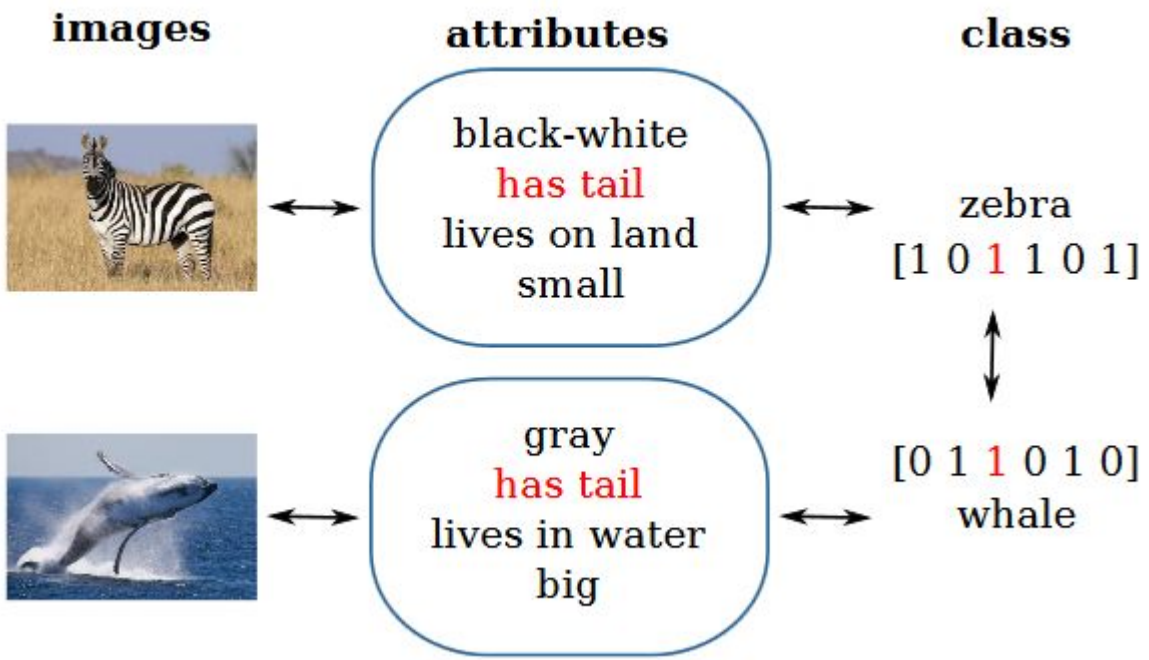

Multimodal Embeddings

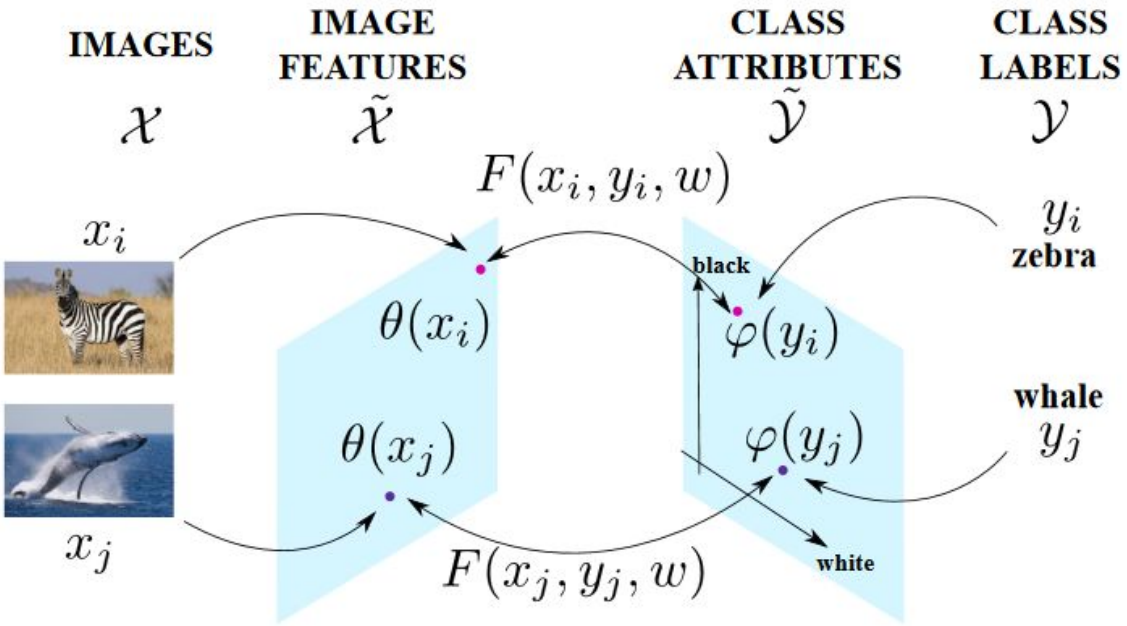

Zero Shot Learning Attributes

- Can Classify unseen classes.
- If 1:1 Mapping can be ensured
- Ex: Weimaraner Dog
  - gray, has tail, is on land, small
  - [0,1,1,1]

Zero Shot Learning Attributes Wikipedia

- Use Wikipedia and Word2Vec/GloVe articles as object description

Whale

From Wikipedia, the free encyclopedia

This article is about a marine mammal. For other uses, see Whale (disambiguation).
For further information, see Cetacea.

Whales are a widely distributed and diverse group of fully aquatic placental marine mammals. They are an informal grouping within the infraorder Cetacea, usually excluding dolphins and porpoises. Whales, dolphins and porpoises belong to the order Cetartiodactyla with even-toed ungulates and their closest living relatives are the hippopotamuses, having diverged about 40 million years ago. The two parvorders of whales, baleen whales (Mysticeti) and toothed whales (Odontoceti), are thought to have split apart around 34 million years ago. The whales comprise eight extant families: Balaenopteridae (the rorquals), Balaenidae (right whales), Cetotheriidae (the pygmy right whale), Eschrichtiidae (the grey whale), Monodontidae (belugas and narwhals), Physeteridae (the sperm whale), Kogiidae (the dwarf and pygmy sperm whale), and Ziphiidae (the beaked whales).

Whales are creatures of the open ocean; they feed, mate, give birth, suckle and raise their young at sea. So extreme is their adaptation to life underwater that they are unable to survive on land. Whales range in size from the 2.6 metres (8.5 ft) and 135 kilograms (298 lb) dwarf sperm whale to the 29.9 metres (98 ft) and 190 metric tons (210 short tons) blue whale, which is the largest creature that has ever lived. The sperm whale is the largest toothed predator on earth. Several species exhibit sexual dimorphism, in that the females are larger than males. Baleen whales have no teeth; instead they have plates of baleen, a fringe-like structure used to expel water while retaining the krill and plankton which they feed on. They use their throat pleats to expand the mouth to take in huge gulps of water. Balaenids have heads that can make up 40% of their body mass to take in water. Toothed whales, on the other hand, have conical teeth adapted to catching fish or squid. Baleen whales have a well developed sense of "smell", whereas toothed whales have well-developed hearing – their hearing, that is adapted for both air and water, is so well developed that some can survive even if they are blind. Some species, such as sperm whales, are well adapted for diving to great depths to catch squid and other favoured prey.

Whales have evolved from land-living mammals. As such whales must breathe air regularly, although they can remain submerged under water for long periods of time. Some species such as the sperm whale are able to stay submerged for as much as 90 minutes.[1] They have blowholes (modified nostrils) located on top of their heads, through which air is taken in and expelled. They are warm-blooded, and have a layer of fat, or blubber, under the skin. With streamlined fusiform bodies and two limbs that are modified into flippers, whales can travel at up to 20 knots, though they are not as flexible or agile as seals. Whales produce a great variety of vocalizations, notably the extended songs of the humpback whale. Although whales are widespread, most species prefer the colder waters of the Northern and Southern Hemispheres, and migrate to the equator to give birth. Species such as humpbacks and blue whales are capable of travelling thousands of miles without feeding. Males typically mate with multiple females every year, but females only mate every two to three years. Calves are typically born in the spring and summer months and females bear all the responsibility for raising them. Mothers of some species fast and nurse their young for one to two years.

Once relentlessly hunted for their products, whales are now protected by international law. The North Atlantic right whales nearly became extinct in the twentieth century, with a population low of 450, and the North Pacific grey whale population is ranked Critically Endangered by the IUCN. Besides whaling, they also face threats from bycatch and marine pollution. The meat, blubber and baleen of whales have traditionally been used by indigenous peoples of the Arctic. Whales have been depicted in various cultures worldwide, notably by the Inuit and the coastal peoples of Vietnam and Ghana, who sometimes hold whale funerals. Whales occasionally feature in literature and film, as in the great white whale of Herman Melville's *Moby Dick*. Small whales, such as belugas, are sometimes kept in captivity and trained to perform tricks, but breeding success has been poor and the animals often die within a few months of capture. Whale watching has become a form of tourism around the world.

Contents [hide]

Whales

Whales are not a taxon, they are an informal grouping of the infraorder Cetacea

Southern right whale

Information

Classification of Cetacea	• Kingdom: Animalia • Phylum: Chordata

Image example of zero shot learning

Animals with Attributes (AWA) [Lampert et.al. CVPR'09]	50 cls	85 att	
Caltech UCSD-Birds (CUB) [Wah et.al.'11]	200 cls	312 att	

Input Embeddings $\theta(x)$: 1K-dim GoogLeNet features

Output Embeddings $\varphi(y)$: `att, w2v, glo,`

Image example of zero shot learning

	AWA	CUB
w2v	51.2	28.4
glo	58.8	24.2
att-	**60.1**	**29.9**
att+	**73.9**	**51.7**

- Attributes & Wikipedia & WordNet are complementary

Backup

Google Translate - Zero shot Learning

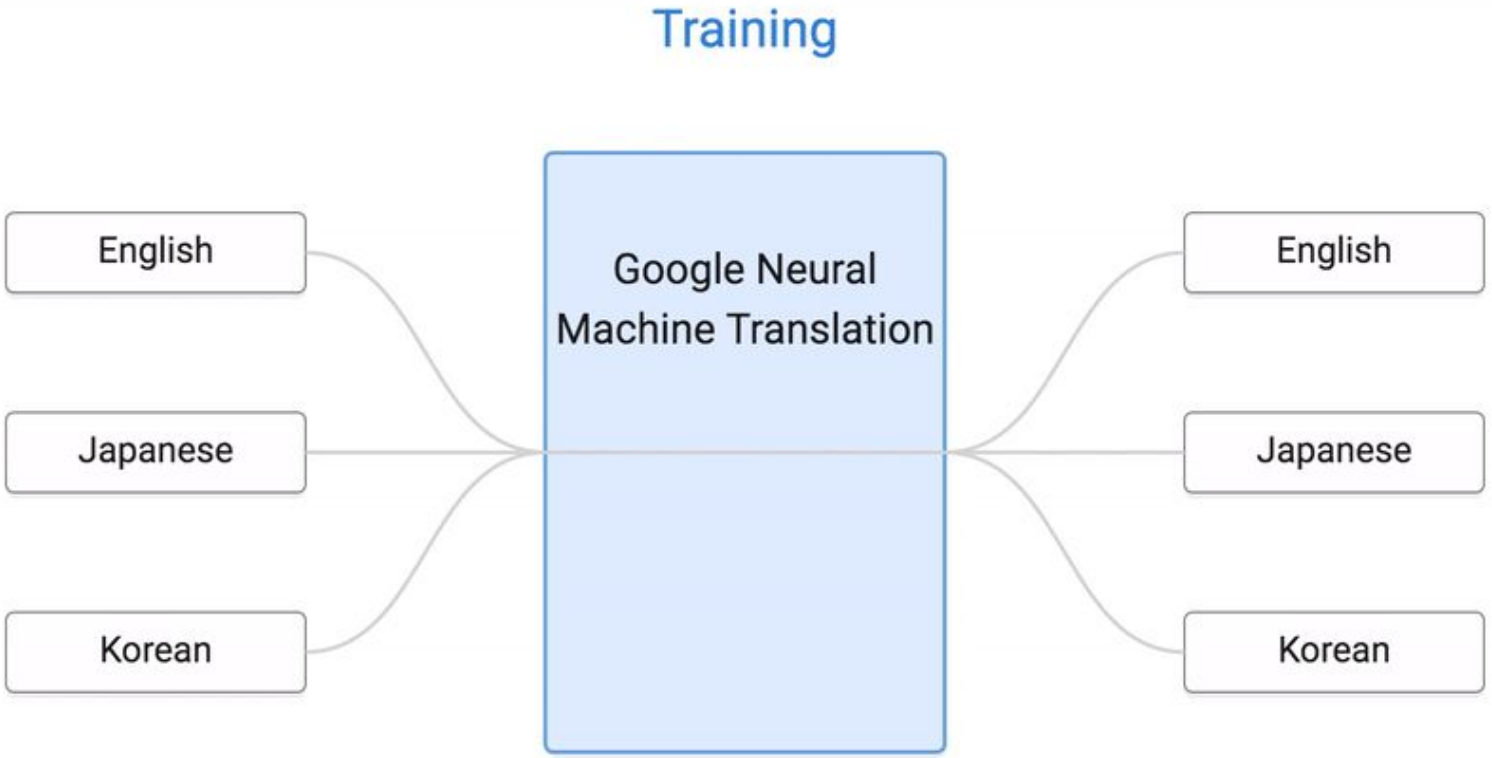

Google Translate - Zero shot Learning

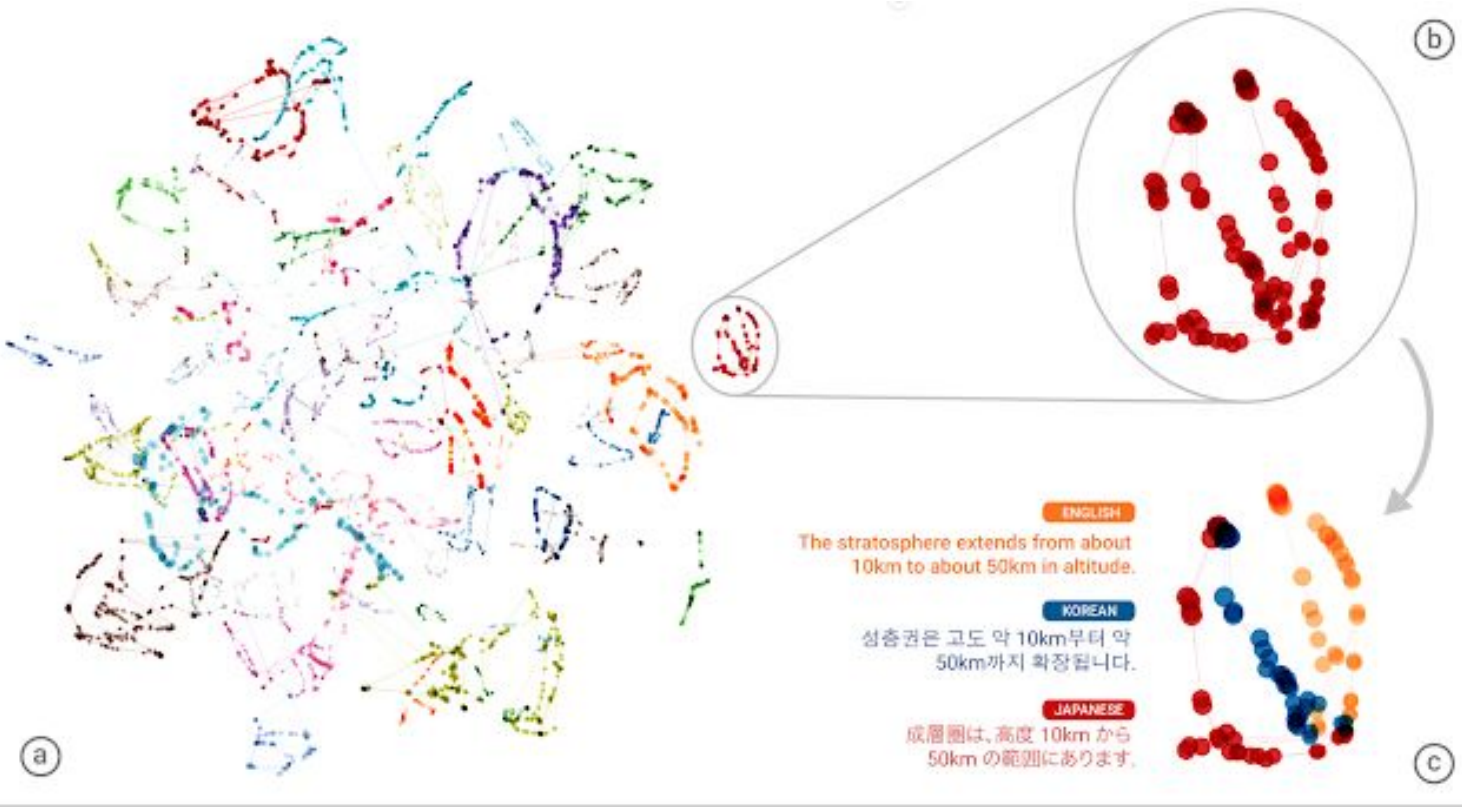

Pretrained models

https://modeldepot.io/

http://pretrained.ml/

https://keras.io/

https://github.com/Cadene/pretrained-models.pytorch

https://nlp.stanford.edu/projects/glove/

https://github.com/pumpikano/tf-dann

References

Weiss, Karl, Taghi M. Khoshgoftaar, and DingDing Wang. "A survey of transfer learning." *Journal of Big Data* 3.1 (2016): 9.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." *Advances in neural information processing systems*. 2012.

Dong, Daxiang, et al. "Multi-task learning for multiple language translation." *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Vol. 1. 2015.

Ganin, Yaroslav, et al. "Domain-adversarial training of neural networks." *The Journal of Machine Learning Research* 17.1 (2016): 2096-2030.

Johnson, Melvin, et al. "Google's multilingual neural machine translation system: enabling zero-shot translation." *arXiv preprint arXiv:1611.04558* (2016).

Zhang, Ziming, and Venkatesh Saligrama. "Zero-shot learning via semantic similarity embedding." *Proceedings of the IEEE international conference on computer vision*. 2015.

Thank you for your Attention!

Adrian Spataru
Data Scientist at Know-Center GmbH
adrian@spataru.at
https://www.fb.me/adrian.spataru.5