Life-long learning for vision and beyond

Timothy Hospedales

Edinburgh 24th Nov 2016

Lifelong Learning for Vision and Beyond

- Making (visual) machine learning more usefully humanlike by consolidating knowledge learned across multiple tasks and domains.
- Focus Directions:
  - Life-long / Transfer Learning
    - Cross-task. Cross-domain. Cross-modality.
  - Weakly + semi-supervised learning
  - Active + curriculum learning
- Applications:
  - Object/action/person recognition
  - Language + vision
  - (Robot Control, Medical Diagnosis, Finance, etc, etc)
- Techniques:
  - Probabilistic models.
  - Deep Learning.
  - Reinforcement Learning.

Why Transfer Learning?

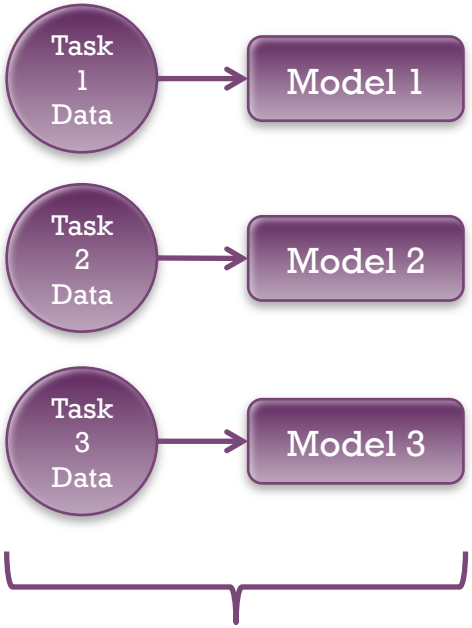

IID
Tasks or Domains

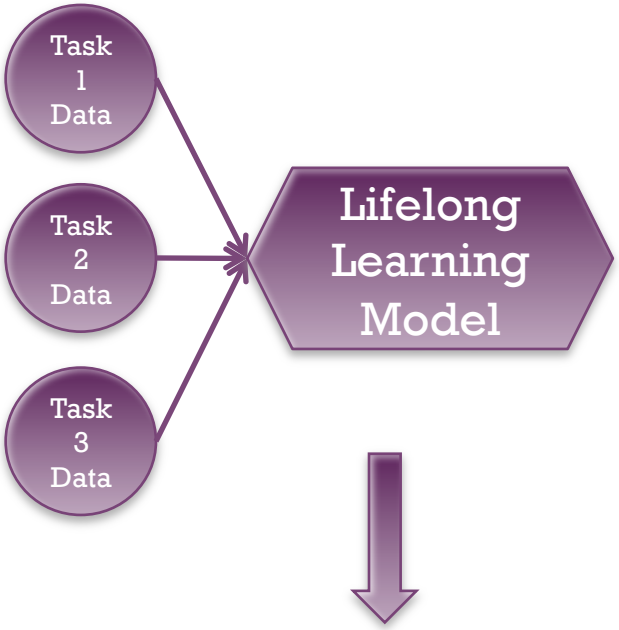

But…. Humans seem to exploit cross-task generalisation:

E.g., Crawl => Walk => Run => Scooter => Bike => Motorbike => Driving.

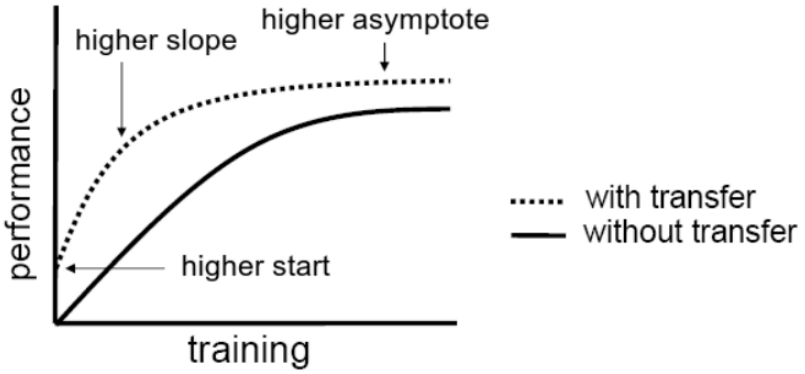

+

Cross Image Modalities

+ Deep Learning

Sketch-Based Image Retrieval with Deep Learning

- How to match free-hand sketches and photos?
  - Cross-domain, and cross-abstraction level matching problem.
- Three-branch CNN:
  - Data: Pairs of photos and sketches.
  - Annotation: Triplet rankings: photo p+ is more similar to sketch s than photo p-.
  - Train with triplet ranking loss: $L_\theta(t) = \max(0, \Delta + D(f_\theta(s), f_\theta(p^+)) - D(f_\theta(s), f_\theta(p^-)))$
    - => Domain-independent representation reflecting the ranking.
- Result:
  - [Video/Demo].
  - [Under Commercialization]

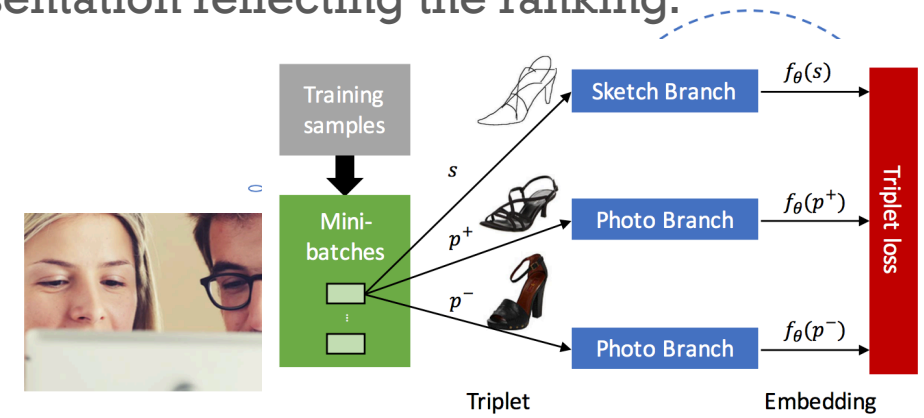

[Yu et al, CVPR, 2016]

Queen Mary
University of London

Upload Reset Submit

Queen Mary
University of London

Upload Reset Submit

Forensic Facial Sketch Matching

- Police sketch eyewitness description and match to mugshot database.
  - A sketch-photo cross-domain matching problem?
- Turns out the weak link is eyewitness memory.
  - Collect a database of forensic sketches with different amounts of forgetting.
- Train a regression model to correct memory errors.
  - Input: A (bad) sketch drawn a day after seeing a face.
  - Target: (Good) sketch drawn while looking at a face.
- Result:
  - 200%+ improvement of state of the art matching rate on 10,000 mugshot benchmark. [Under Commercialization]

[Ouyang et al, CVPR 2016]

Forensic Facial Sketch Matching

Cross Domain: Vision & Language

+ Weakly Supervised Learning
+ Deep Learning
+ Probabilistic Graphical Models

Zero-Shot Learning: Live Demonstration!

- Audience Task: Recognise the Wampimuk.
  - Impossible?
- Solution: Semantic Transfer
  - Domain Ontology:
    - Wampimuk := small, horns, furry, cute
  - Wikipedia Page:

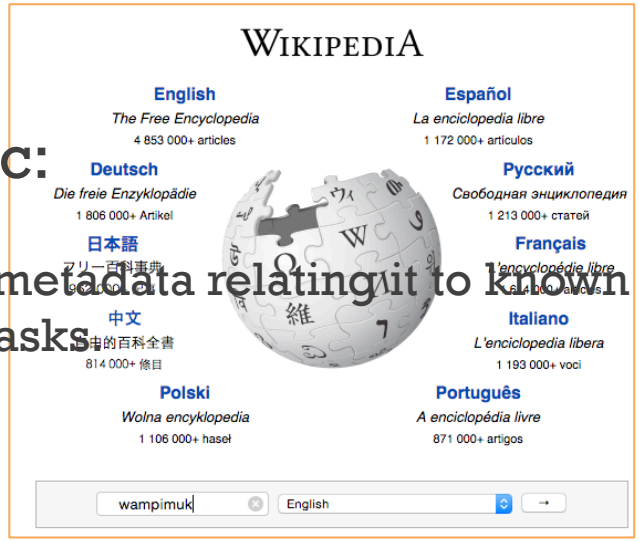

Zero-Shot problem spec:

1. Solve a novel task
2. …Via some (un)structured metadata relating it to known tasks

Crucial for scalability re: # of tasks

Weakly supervised object & attribute learning

Train

person, horse, cloth, furry

Test

Person Cloth
Horse Furry

person: cloth
horse : furry

(Detect, associate, Segment)

Indian Buffet Process:

Infinite Sparse Binary Matrix Factorization

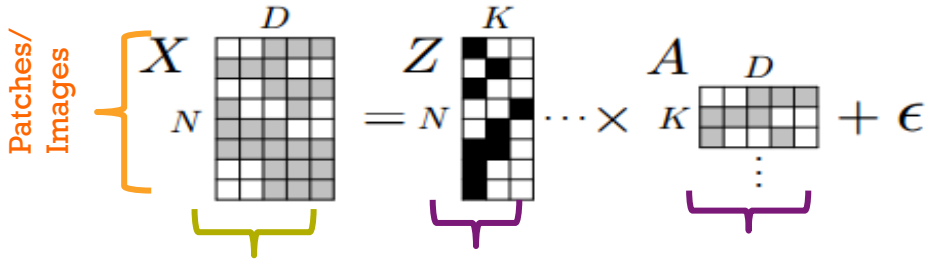

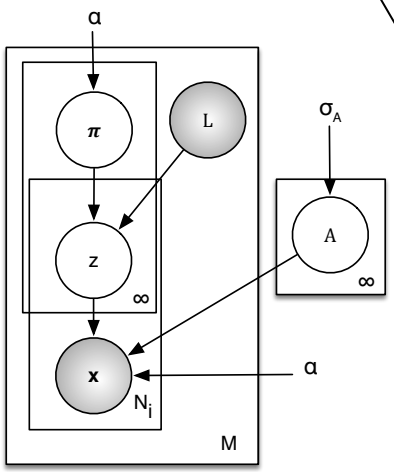

[Shi, ECCV' 14]

Captioning, QA, Visual Turing Test

What color are her eyes?
What is the mustache made of?

How many slices of pizza are there?
Is this a vegetarian pizza?

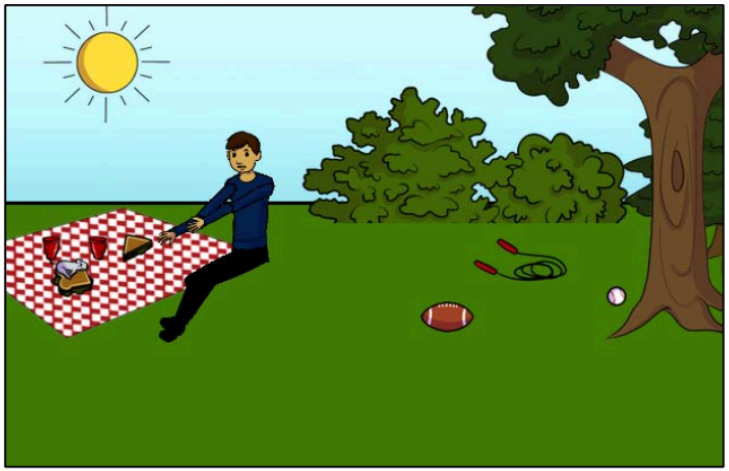

Is this person expecting company?
What is just under the tree?

Does it appear to be rainy?
Does this person have 20/20 vision?

+

Cross-Task Transfer

+ Deep Learning

Deep Transfer Learning

- Analysis of tensor-factorisation :: Deep Neural network architecture equivalence
  - => Application of transfer learning theory from ML to deep learning.
  - Share latent layers across different NNs for different tasks
- Parameterized Neural Networks
  - Synthesize context-dependent DNN weights on the fly.

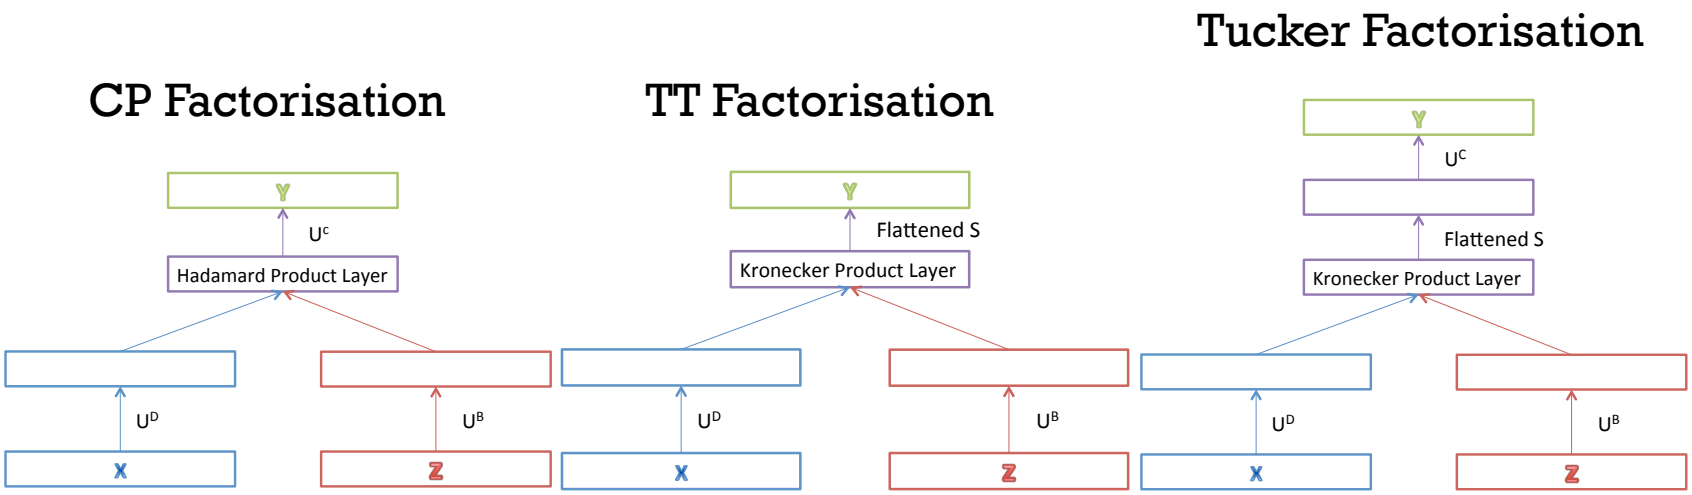

Deep Transfer Reinforcement Learning

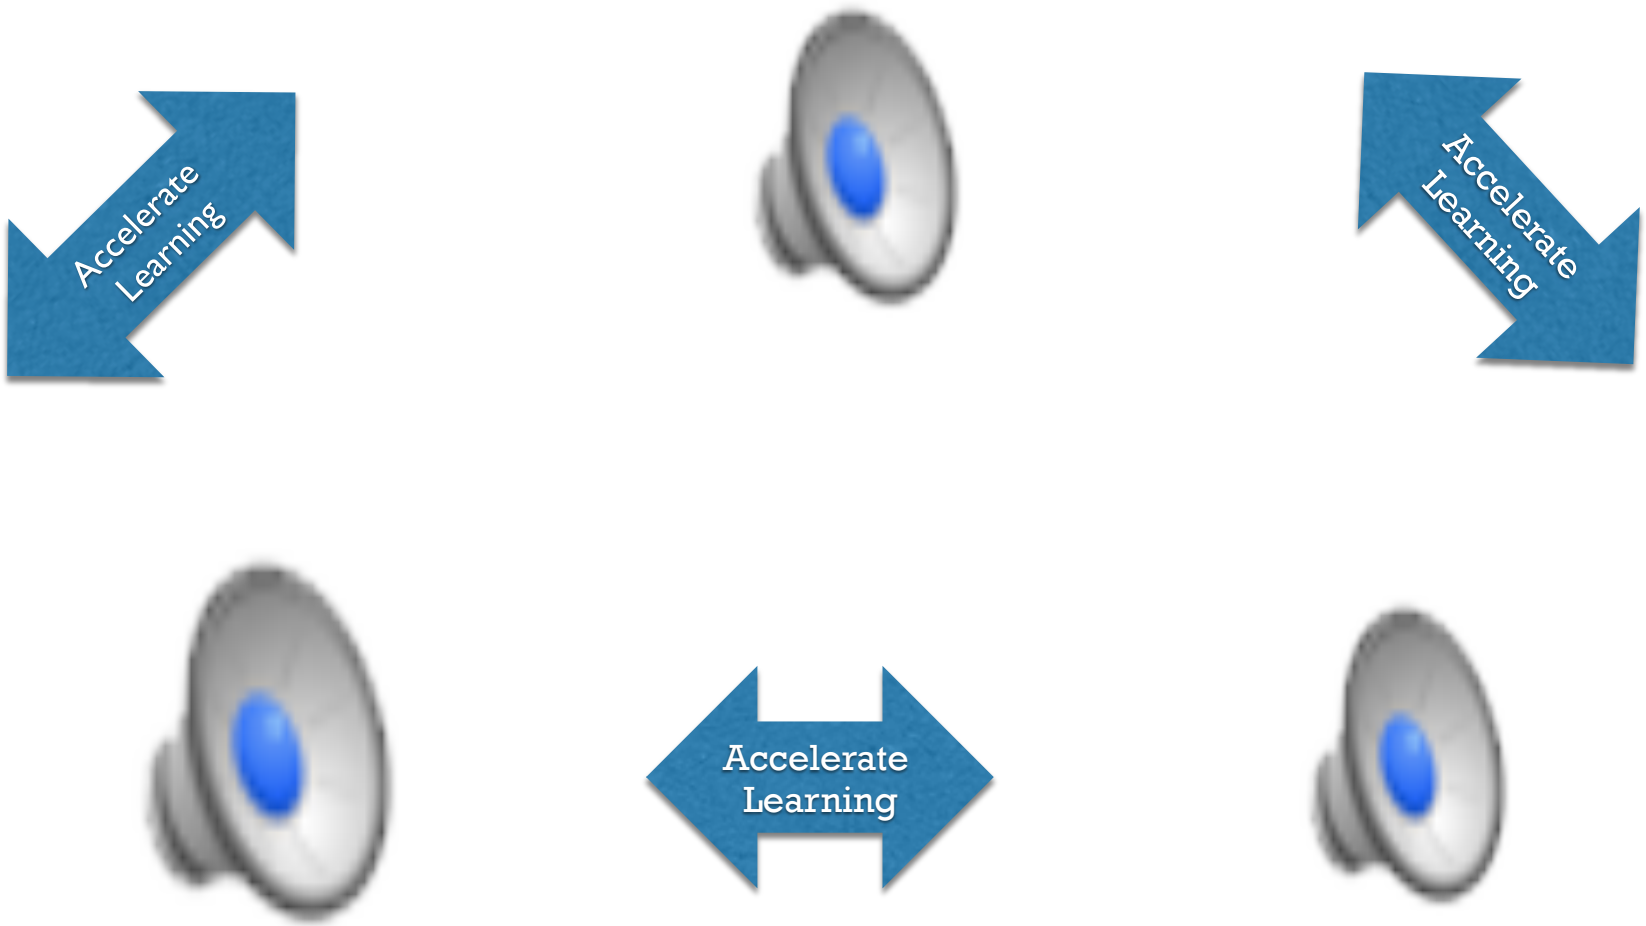

Deep Learning with Dynamic Weights: Option Pricing

- A company's stock is worth $100 at T=0. Option C gives the right buy the stock in T=5 days for $110. If the price is $120 in 5 days, I can execute C: buying and then immediately selling. Profit $120-110=10.
  - How much is option contract C worth at T=0?
- Sanity guarantees:
  - Price must go up with T.
  - Price = 0 at T=infinity, etc.
- Internal modeling of asset price PDF.

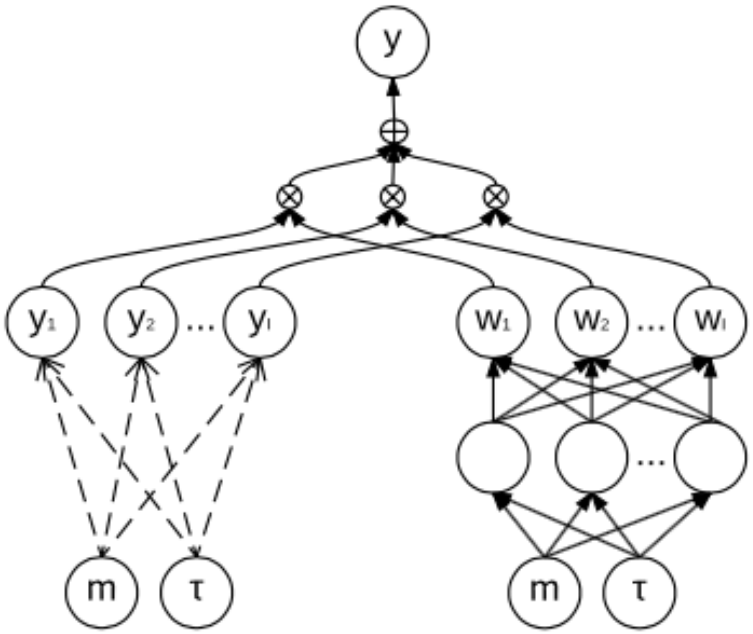

Summary

- http://homepages.inf.ed.ac.uk/~thospedales/
- https://www.inf.ed.ac.uk/teaching/courses/rtds/projects/2016/areas/hospedales.html
- Projects in
  - Vision
  - Vision & Language
  - (+Life Long) Deep Learning
  - (+Life Long) Deep Reinforcement Learning