



Carnegie Mellon University

11-785 Introduction to Deep Learning

Recitation - Representations and Autoencoders

Anurag, Shentong

Overview

- The importance of representations
- Representation Learning
- What is a representation
- Applications
- Summary

The importance of representations

- The choice of representation is key to solving a given problem correctly and efficiently
- A few examples to motivate this (from the [Deep Learning book](#)):

The importance of representations


- The choice of representation is key to solving a given problem correctly and efficiently
- A few examples to motivate this (from the [Deep Learning book](#)):
 1. Long division:
"Divide 210 by 6"

The importance of representations

- The choice of representation is key to solving a given problem correctly and efficiently
- A few examples to motivate this (from the [Deep Learning book](#)):
 1. Long division:
"Divide 210 by 6" vs. "Divide CCX by VI"

The importance of representations

- A few examples to motivate this (from the [Deep Learning book](#)):
 1. Long division:
"Divide 210 by 6" vs. "Divide CCX by VI"
 2. Binary Classification using a Linear Classifier



The importance of representations

- A few examples to motivate this (from the [Deep Learning book](#)):
 3. Insert a number into a sorted list
 - a. Linked List $\rightarrow O(n)$
 - b. Balanced Tree $\rightarrow O(\log n)$

Representation Learning


- The task of "*learning* representations of the data that make it easier to extract useful information when building classifiers or other predictors" and has become "a field in itself"*
- A shift from hand-crafting complex features from data to being able to learn these features (with a neural network)
- Learned representations often desired to satisfy certain properties:
 - Useful in downstream tasks through **transfer learning**,
 - Interpretability**

* Bengio, Yoshua et al. "Representation Learning: A Review and New Perspectives." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35 (2013): 1798-1828.

What is a representation?

We have seen several instances of "representations" in the course already:

- HW1P2: Phoneme classification



What is a representation?

- HW2P2: Face classification

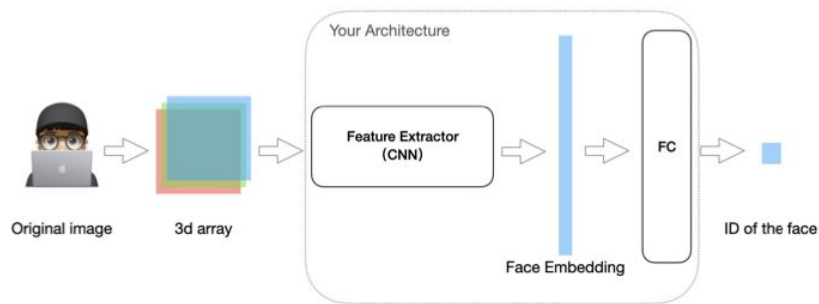



Figure 1: A typical face classification architecture


What is a representation?

- Recurrent Neural Networks (from Attention Recitation):



Applications


- Word Embeddings (word2vec*):
 - Make several downstream tasks easier - sentiment classification, question answering, etc.



*Mikolov, Tomas et al. "Distributed Representations of Words and Phrases and their Compositionality." *NIPS* (2013).

Applications

- Speech Utterance Embeddings (wav2vec*):
 - Useful for word/letter classification




*Schneider, Steffen et al. "wav2vec: Unsupervised Pre-training for Speech Recognition." *INTERSPEECH* (2019).

Applications

- A plethora of following work that:
 - **improves performance** on given task,
 - demonstrates usefulness of learned representations on **many other tasks** (ex. BERT, Mockingjay)
- A trend in learning representations for many problems - "anything2vec":
 - **speech2vec**: Chung, Yu-An and James R. Glass. "Speech2Vec: A Sequence-to-Sequence Framework for Learning Word Embeddings from Speech." *INTERSPEECH* (2018).
 - **node2vec**: Grover, Aditya and J. Leskovec. "node2vec: Scalable Feature Learning for Networks." *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2016): n. pag.
 - **(batter|pitcher)2vec**: Alcorn, Michael A.. "(batter|pitcher)2vec: Statistic-Free Talent Modeling With Neural Player Embeddings." (2018).

Applications

- Recent trends:
 - Interpretability of learned representations



Applications

- Recent trends:
 - Fairness of learned representations - recidivism, health insurance, etc
 - Important that the learned representations do not encode biases from demographic features,
 - Learning Fair Representations: **Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, Cynthia Dwork** ; Proceedings of the 30th International Conference on Machine Learning, PMLR 28(3):325-333, 2013.
 - Learning Adversarially Fair and Transferrable Representations: **David Madras, Elliot Creager, Toniann Pitassi, Richard Zemel** Proceedings of the 35th International Conference on Machine Learning, PMLR 80:3384-3393, 2018.

Summary

- Representations are learned:
 - end-to-end,
 - unsupervised manner (usually),
 - for some downstream task(s)
- Choosing representations:
 - Study the pre-training objective used to learn representations,
 - Consider model architectures used and suitability for your task

"Generally speaking a good representation is one that makes a subsequent learning task easier. The choice of representation will usually depend on the choice of the subsequent learning task"

- Chapter 15, Representation Learning, Deep Learning Book