CS 7150 - Deep Learning

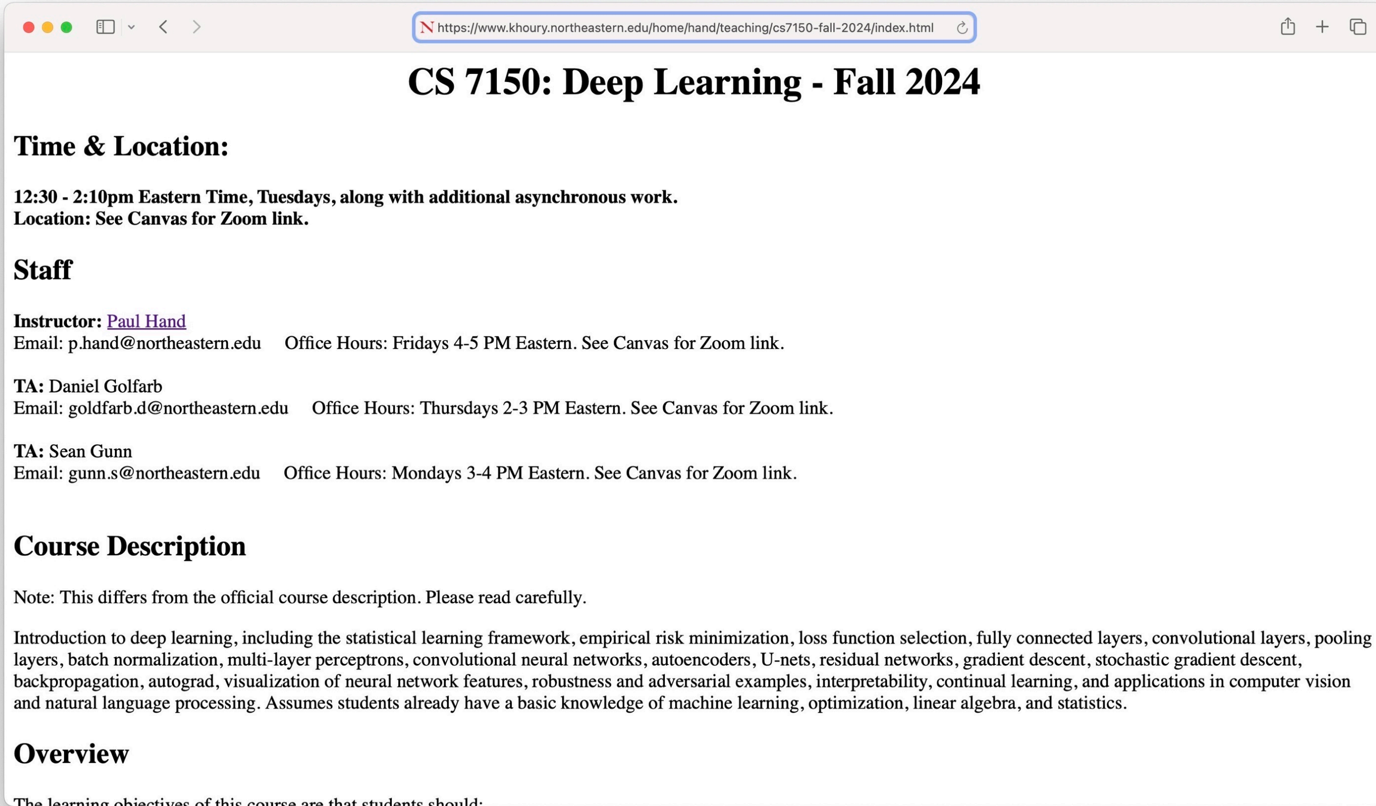

https://www.khoury.northeastern.edu/home/hand/teaching/cs7150-fall-2024/index.html

CS 7150: Deep Learning - Fall 2024

Time & Location:

12:30 - 2:10pm Eastern Time, Tuesdays, along with additional asynchronous work.
Location: See Canvas for Zoom link.

Staff

Instructor: Paul Hand
Email: p.hand@northeastern.edu   Office Hours: Fridays 4-5 PM Eastern. See Canvas for Zoom link.

TA: Daniel Golfarb
Email: goldfarb.d@northeastern.edu   Office Hours: Thursdays 2-3 PM Eastern. See Canvas for Zoom link.

TA: Sean Gunn
Email: gunn.s@northeastern.edu   Office Hours: Mondays 3-4 PM Eastern. See Canvas for Zoom link.

Course Description

Note: This differs from the official course description. Please read carefully.

Introduction to deep learning, including the statistical learning framework, empirical risk minimization, loss function selection, fully connected layers, convolutional layers, pooling layers, batch normalization, multi-layer perceptrons, convolutional neural networks, autoencoders, U-nets, residual networks, gradient descent, stochastic gradient descent, backpropagation, autograd, visualization of neural network features, robustness and adversarial examples, interpretability, continual learning, and applications in computer vision and natural language processing. Assumes students already have a basic knowledge of machine learning, optimization, linear algebra, and statistics.

Overview

The learning objectives of this course are that students should:

Deep Learning with PyTorch

33 reviews

Eli Stevens, Luca Antiga, and Thomas Viehmann
Foreword by Soumith Chintala
July 2020 · ISBN 9781617295263 · 520 pages
printed in black & white
Available translations: Complex Chinese, Japanese, Korean, Russian, Simplified Chinese
Data

eBook	print	online + audio	subscription
pdf, ePub, online	includes eBook	read and listen	from $19.99

Every other day we hear about new ways to put deep learning to good use: improved medical imaging, accurate credit card fraud detection, long range weather forecasting, and more. PyTorch puts these superpowers in your hands, providing a comfortable Python experience that gets you started quickly and then grows with you as you—and your deep learning skills—become more sophisticated. *Deep Learning with PyTorch* will make that journey engaging and fun.

Agenda:

- Data Center Energy Efficiency Example
- Q1- Discussion: Perceiving, learning, abstracting, reasoning
- Q2 - Hand-crafted knowledge and its limitations
- Q3 - Learning from examples
- Q4 - Benefit of context in ML systems

Examples of ML Problem: data center energy efficiency

Machine Learning Applications for Data Center Optimization
Jim Gao, Google

Mechanical Plant at a Google Data Center:

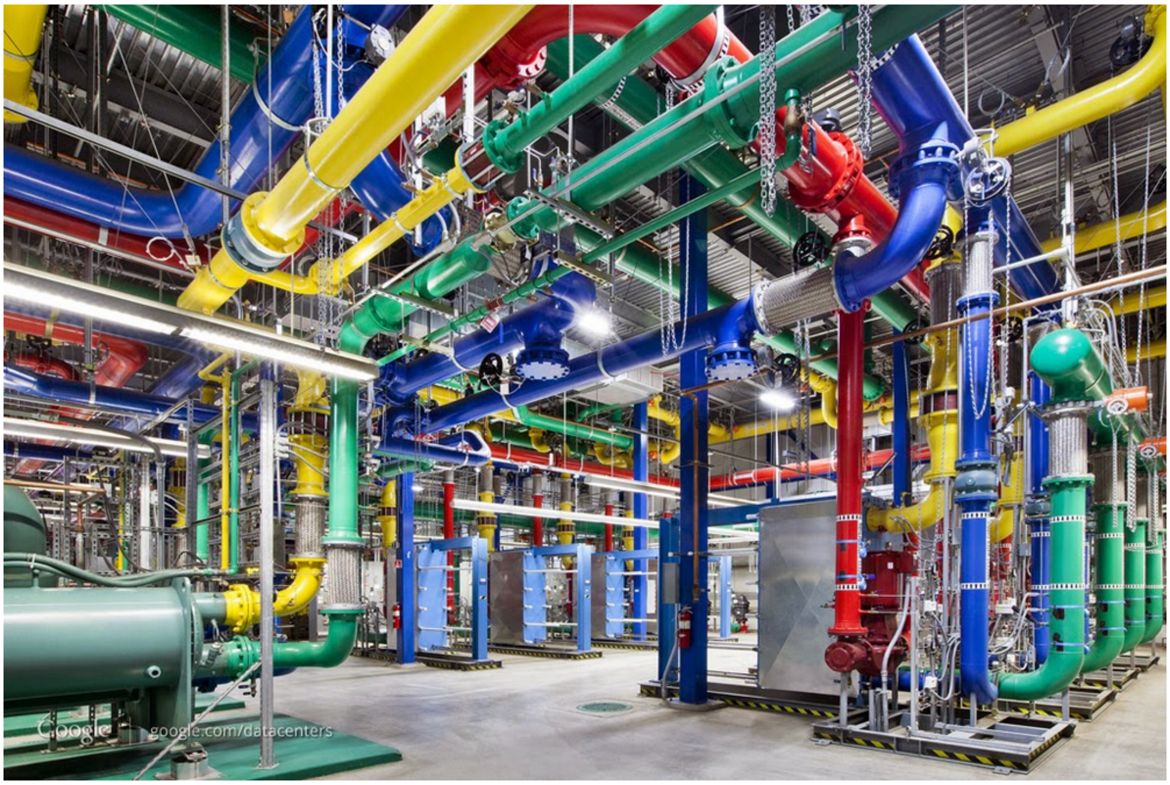

Energy Efficiency of the data center (Ideal value is 1)

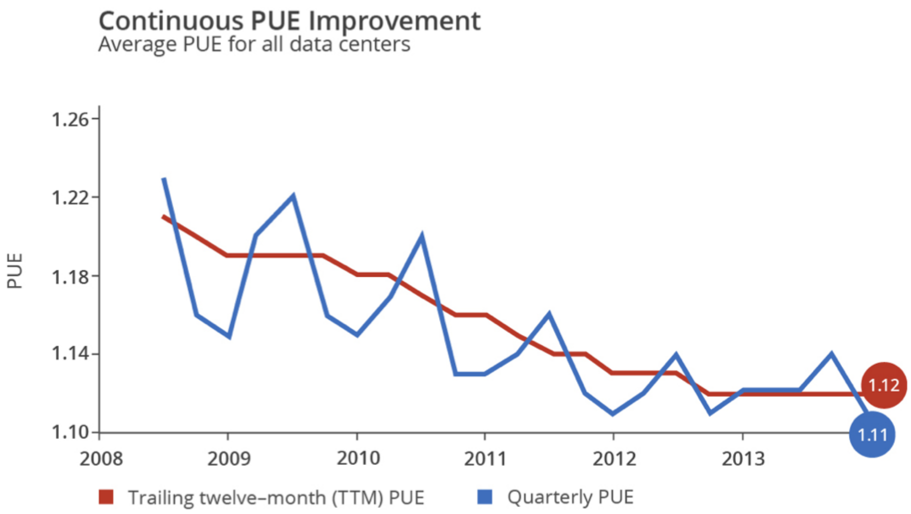

Fig 1. Historical PUE values at Google.

Features that affect the energy efficiency of the data center

1. Total server IT load [kW]
2. Total Campus Core Network Room (CCNR) IT load [kW]
3. Total number of process water pumps (PWP) running
4. Mean PWP variable frequency drive (VFD) speed [%]
5. Total number of condenser water pumps (CWP) running
6. Mean CWP variable frequency drive (VFD) speed [%]
7. Total number of cooling towers running
8. Mean cooling tower leaving water temperature (LWT) setpoint [F]
9. Total number of chillers running
10. Total number of drycoolers running
11. Total number of chilled water injection pumps running
12. Mean chilled water injection pump setpoint temperature [F]
13. Mean heat exchanger approach temperature [F]
14. Outside air wet bulb (WB) temperature [F]
15. Outside air dry bulb (DB) temperature [F]
16. Outside air enthalpy [kJ/kg]
17. Outside air relative humidity (RH) [%]
18. Outdoor wind speed [mph]
19. Outdoor wind direction [deg]

Question:
Some of these features can be controlled by the data center. Some can not. How would you use measurements of all of these features in order to make the data center as energy efficient as possible?

Results from a machine learning approach using neural networks

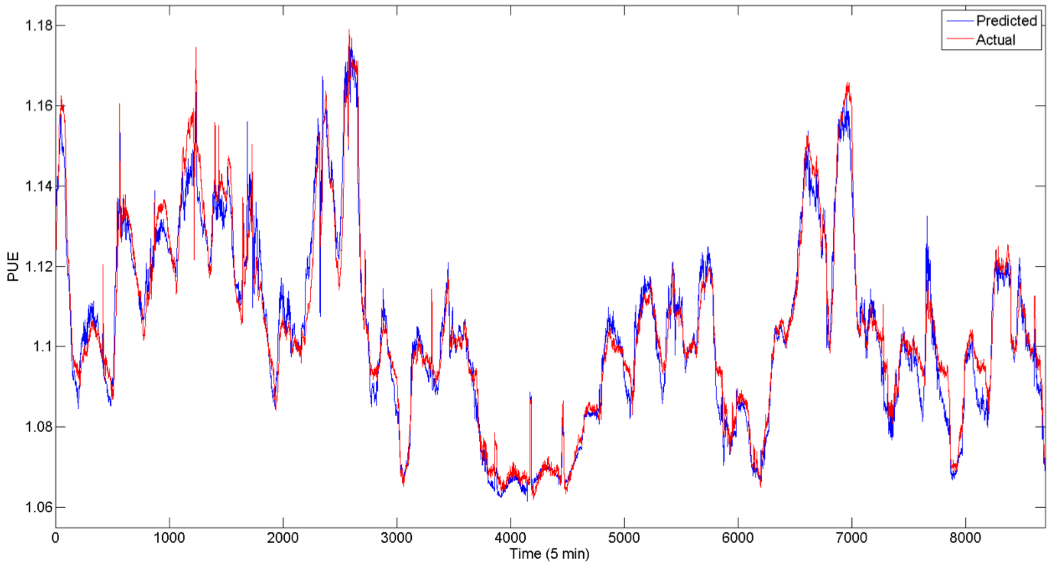

Fig. 3 Predicted vs actual PUE values at a major DC.

Provide a meaning and example of the following terms:

perceiving

Meaning: ___ identify / discern objects in the world. identifying what is in a signal
Examples: ___ identify cat or face in, image classification, object detection, voice recognition

learning

Meaning: ___ forming an association between input and output. forming a rule for predicting output from input. Improving performance on a specific task by experience
Examples: ___ learning what a cat is, recommender systems
Non-examples: ___
Is it possible for neural networks to be useful even when they aren't used for learning (that is, even if no data from the real world is provided)? Yes, it is possible. Look up "Deep Image Prior" if interested.

abstracting

Meaning: ___ identifying a concept that unifies observations in a natural way. Taking knowledge from one domain and applying it to another
Examples: ___ There is a canonical shape of a letter or number, independent of the specific way it was drawn (typed, handwritten, drawn in sand). Determining the laws of physics
Non-examples: ___ build a NN that predicts the evolution of a flame

reasoning

Meaning: ___ making chained deductions based on logic
Examples: ___
Non-examples: ___

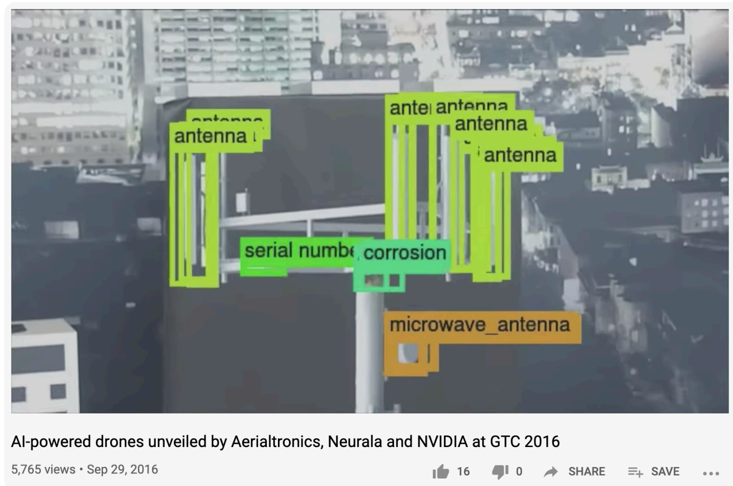

Example of perceiving

This deluge of data calls for automated methods of data analysis, which is what **machine learning** provides. In particular, we define machine learning as a set of methods that can automatically detect patterns in data, and then use the uncovered patterns to predict future data, or to perform other kinds of decision making under uncertainty (such as planning how to collect more data!).

What is machine learning and how does it related to AI and to Deep Learning?

The first definition of *machine learning* dates back to 1959, from American AI pioneer Arthur Samuel:

> *Machine learning is the field of study that gives computers the ability to learn without being explicitly programmed.*

— Kevin Murphy

— Mauro and Valigi

Relationship of ML to AI and Deep Learning

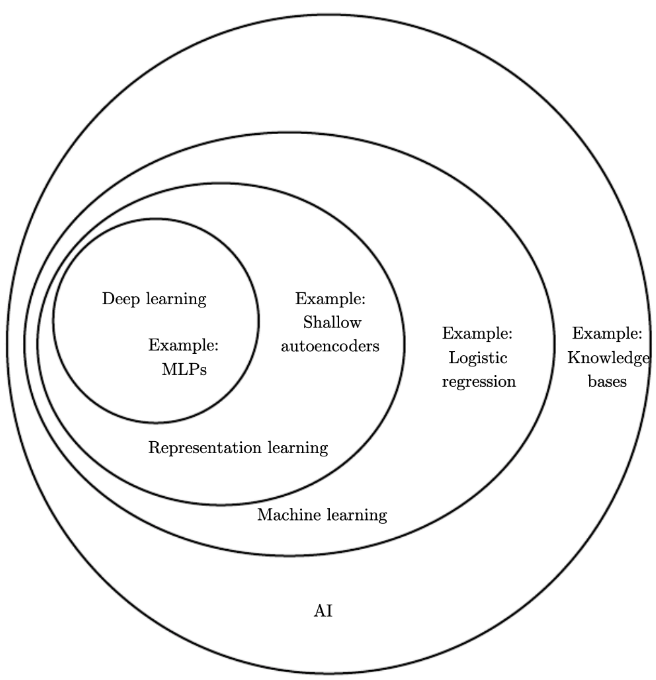

Figure 1.4: A Venn diagram showing how deep learning is a kind of representation learning, which is in turn a kind of machine learning, which is used for many but not all approaches to AI. Each section of the Venn diagram includes an example of an AI technology.

Question 2: Provide an example of what handcrafted knowledge could be in the context of understanding audio or image data

____

________________Audio (humans have emotions, identifying frequency bands of a voice, there are different accents)

____

Comparison of deep learning with conventional machine learning and the limitations of hand-crafted knowledge

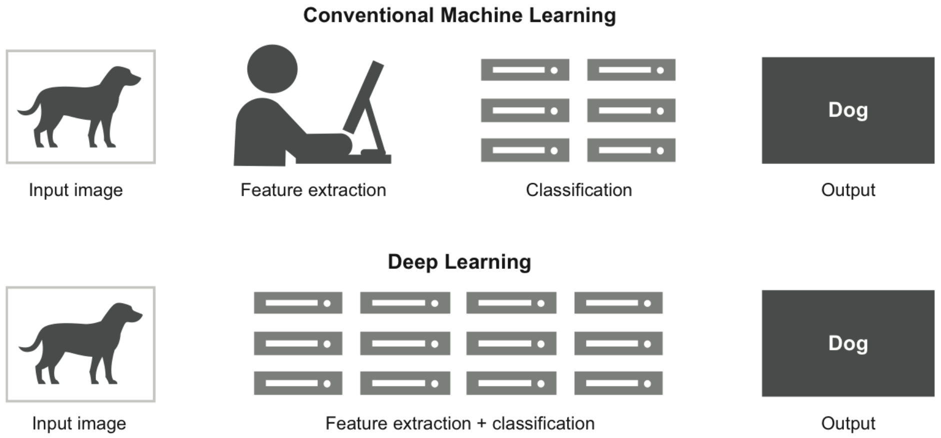

Figure 4.5 In traditional ML, engineers have to develop algorithms to extract features that can be fed to the model. A DL model doesn't need this complex preliminary step.

Example of feature engineering: HOG (Histogram of Gradients) Features

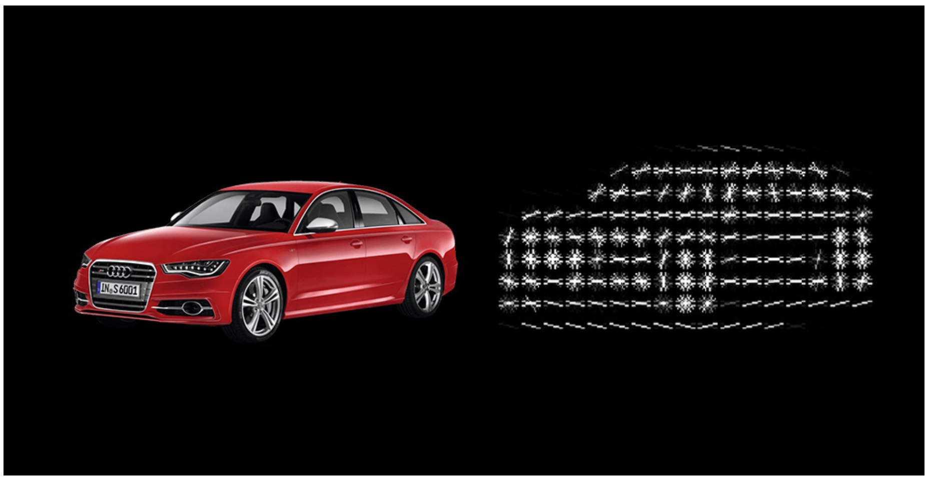

Example of feature engineering: SIFT (Scale-invariant feature transform) Features

Example of hand-crafted knowledge in natural language processing: Sentiment analysis

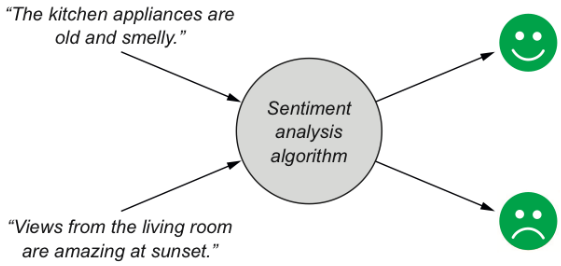

Figure 5.6 The basic concept of sentiment analysis

Let's suppose that we don't know machine learning, and we need to explain to a computer how to rate these sentences by developing an algorithm. A good rule of thumb would be to look at certain telltale words, such as "terrible" and "amazing." In fact, this is exactly how the earliest algorithms for sentiment analysis worked: researchers painstakingly built a dictionary of important words, and labeled each as positive, negative, or neutral. For example, such a word-sentiment dictionary might look like this:

- *Delighted*—Positive
- *Killed*—Negative
- *Shout*—Negative
- *Desk*—Neutral

Once you have an *emotional glossary* like that, you can classify each sentence by counting the number of positive and negative words and get a final score for the sentence.

This simplistic approach has a bunch of problems. Language is extremely complex, and we use it in very different ways from person to person; the same word can be used in different ways to communicate completely opposite messages. Let's say you listed "nice" as a positive-sentiment word, and one of the reviewers writes this:

It would be nice if they completely removed that tapestry.

Even if the sentence has an overall negative opinion of the house, our naive system would consider it positive, because of the positive connotation of "nice." Maybe you could try improving this system by adding more-complex rules, something like this:

'It would be [POSITIVE WORD] if ..' => negative

Although this rule would work on the preceding snippet, it's still easy to fool. For example, the following sentence is actually extremely positive, but would be ranked as a negative opinion:

It would be nice if I could move into this place right away!

Should we keep adding hardcoded rules? The game is already becoming complicated (and boring), yet it's still easy to fool our system. Notice also that we're still playing with just a few words; we haven't even started mapping the vast ocean of the English vocabulary. Even if we did get to the bottom of this, we would almost need to start all over again for other languages.

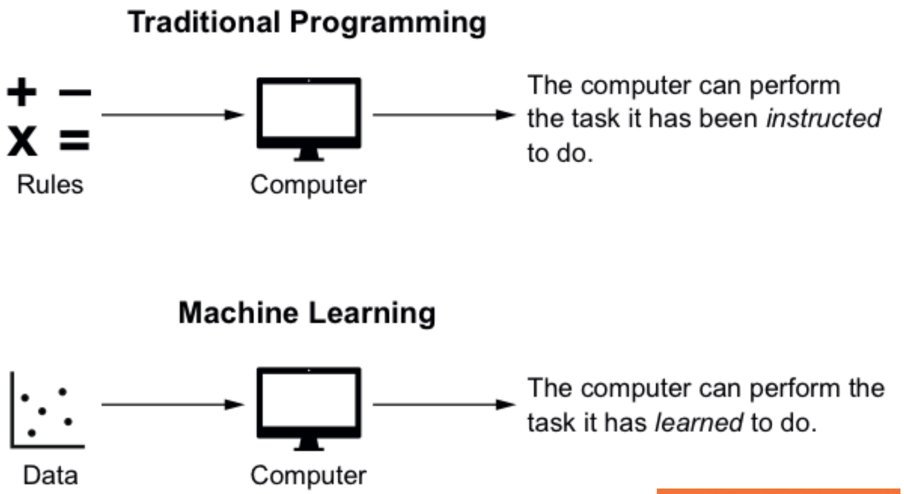

Figure 1.1 The difference between the traditional programming approach and machine learning: the first relies on precise rules and instructions, the latter on data and learning.

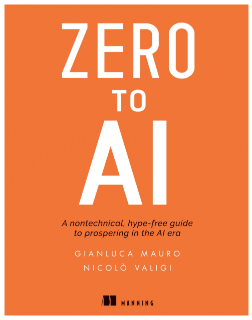

Other challenges of traditional programming approaches:

May not be adaptable
Large amount of computing (e.g. deep search tree)
Limited to expertise of programmers
The application be too complex for humans to understand/prescribe

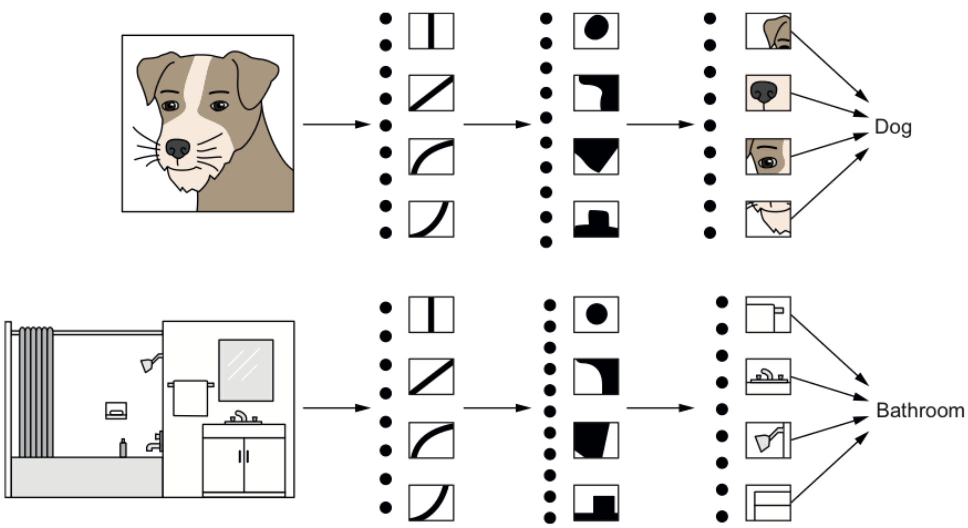

Figure 4.10 Similarities between DL-based classifiers for rooms and animals. The first layers learn to identify the same general geometric patterns.

Examples of ML Problems: Playing games like chess and go

Non-ML method beat Gary Kasparov at Chess in the 1990's

ML method beat Lee Sedol at Go in 2010's.

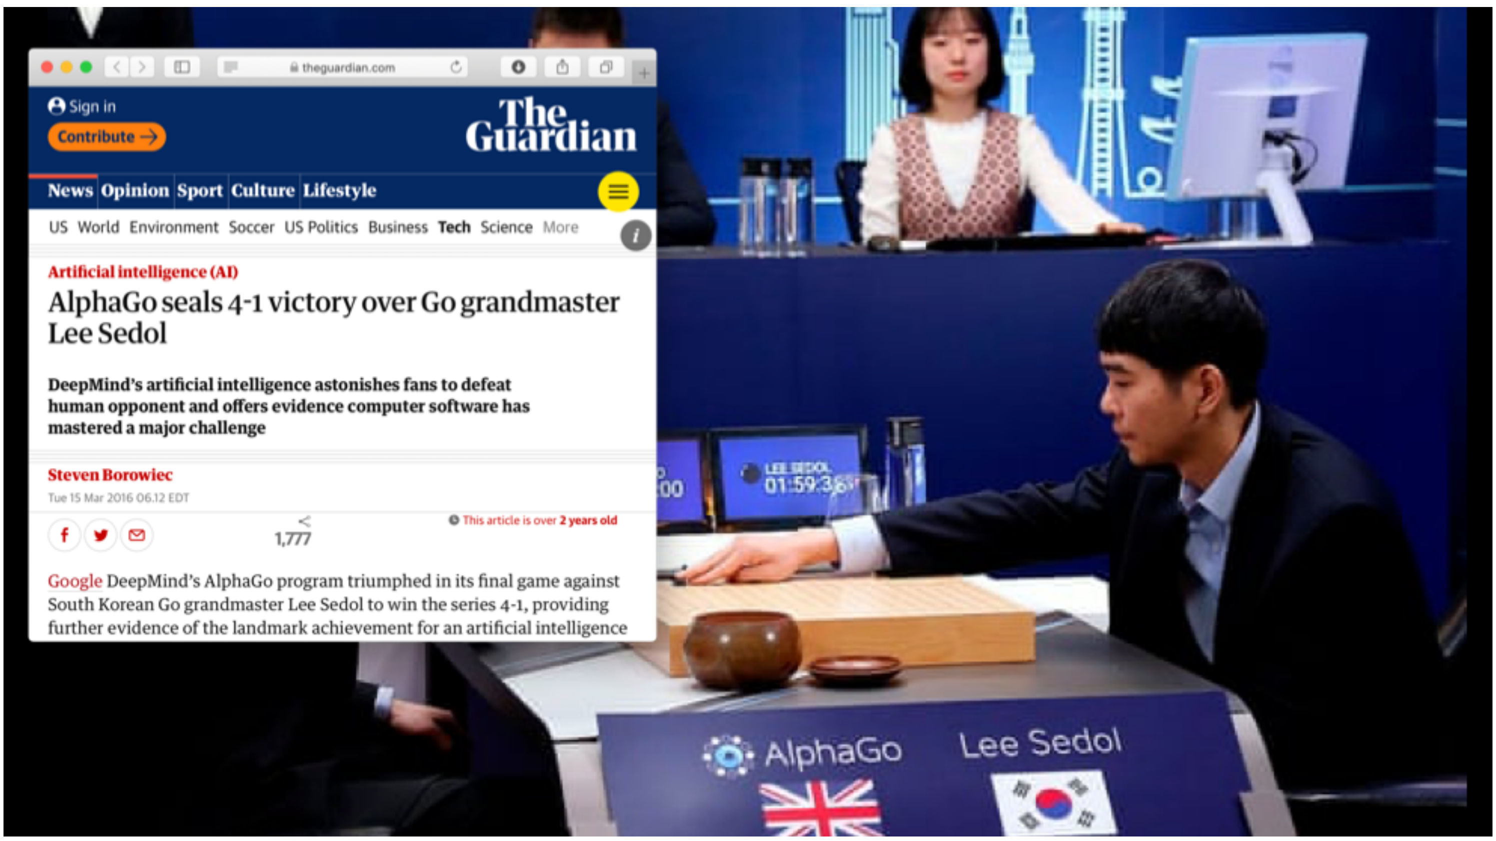

Examples of ML Problems:
Prediction of 3d protein shape given its composition

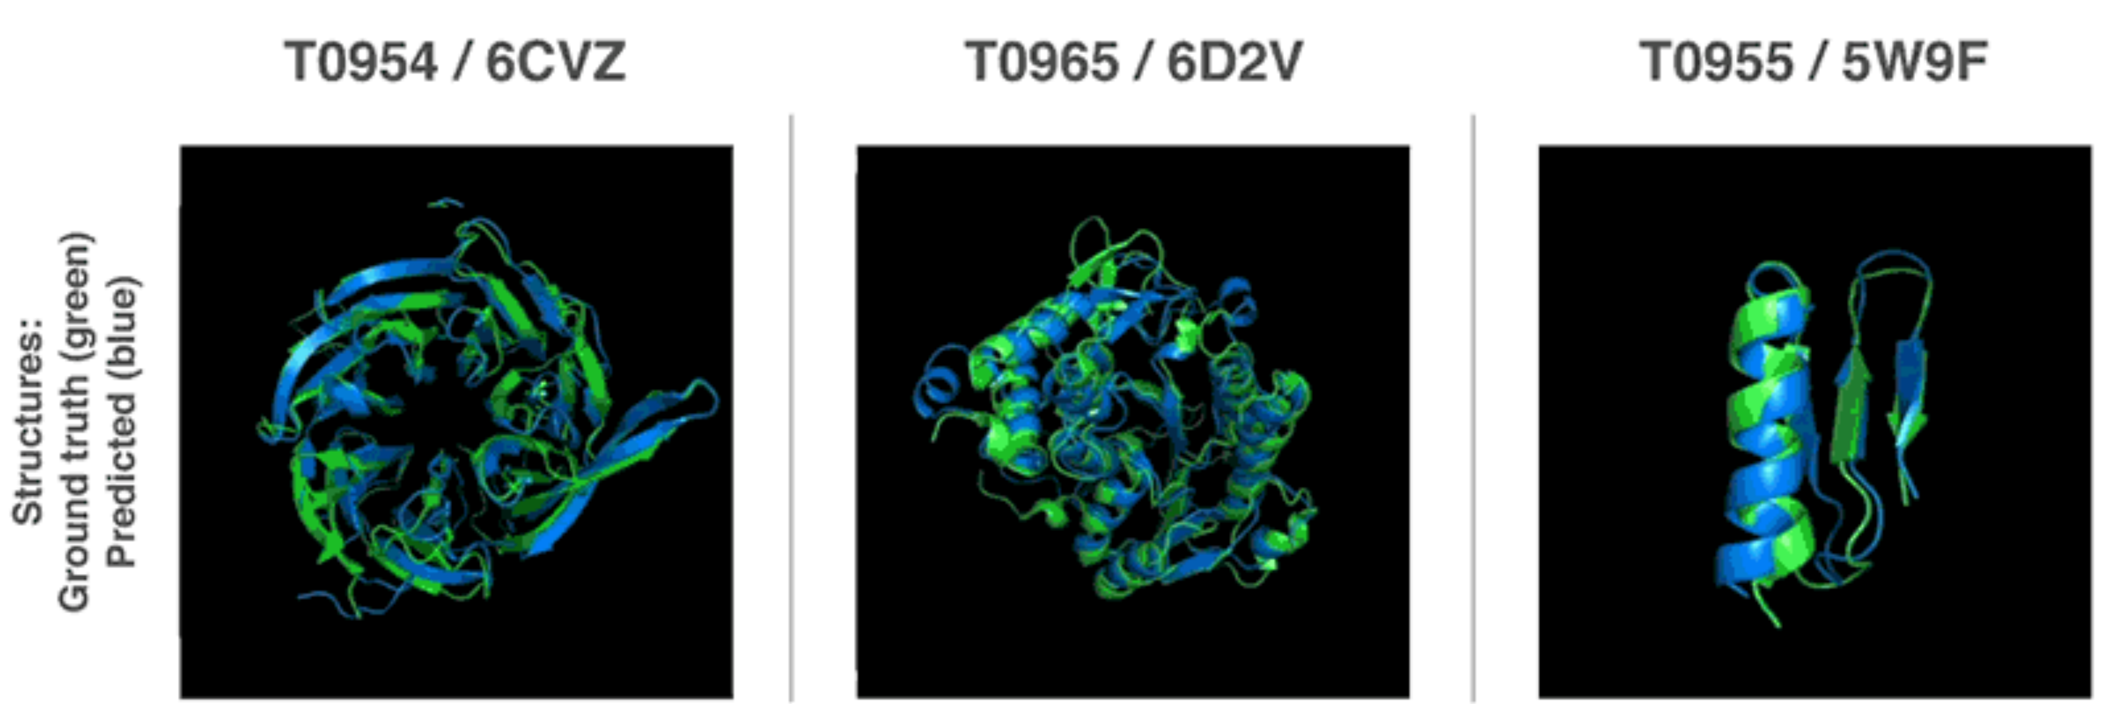

Protein Folding by AlphaFold

Example: Image generation from a text description - DALL E

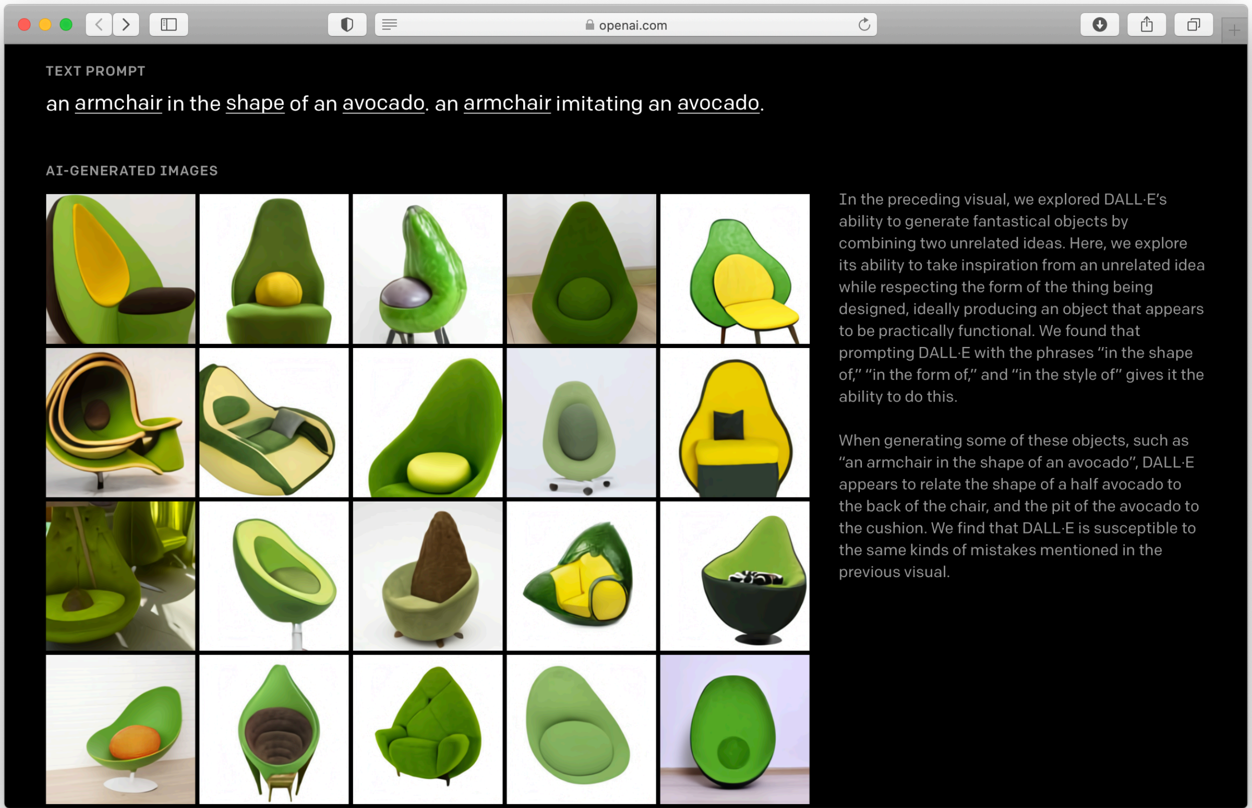

Q3 The speaker makes the point that humans can learn to identify objects when shown only one or two examples, in contrast to handwriting recognition algorithms which may need 50k-100k examples. What are your thoughts about this comparison?

____

____

____

____

Q4 The speaker suggests that a handwriting recognition system that understands context via a handwriting model (sequential strokes of a pen) could improve classification decisions. Provide and walk through an example of a specific image where a misclassification may be made without such a model and the correct classification may be made with the model.

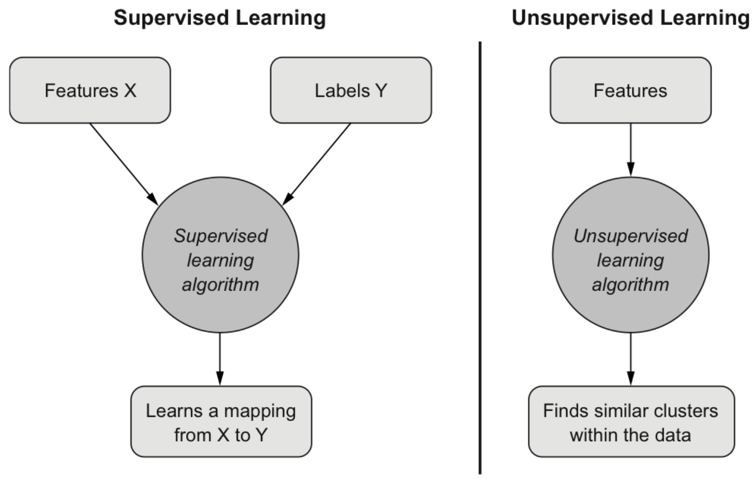

Figure 3.6 The differences in input and output of supervised and unsupervised algorithms

training data

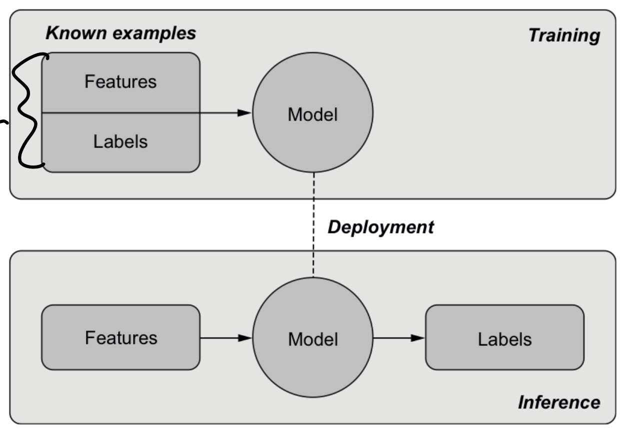

Figure 2.3 The two phases of machine learning: training and inference

Supervised Learning Pipeline: