

Autonomous and Adaptive Systems

Introduction to Deep Learning and Neural Architectures I

Mirco Musolesi

mircomusolesi@acm.org

Overview and Historical Notes

- ▶ Initial AI systems are based on the definition of formal systems (logic, knowledge base, etc.).
 - ▶ Several artificial intelligence projects have sought to hard-code knowledge about the world in formal language.
 - ▶ Difficult to list rules for very large number of situations.
 - ▶ Some rules might also not be possible to be codified given the sheer complexity of the world.
- ▶ However note: recent developments in combining deep learning and symbolic AI.
 - ▶ Very open field of research at the moment.

CYC: A Large-Scale Investment in Knowledge Infrastructure

Douglas B. Lenat



Since 1984, a person-century of effort has gone into building CYC, a universal schema of roughly 10^5 general concepts spanning human reality. Most of the time has been spent codifying knowledge about these concepts; approximately 10^6 commonsense axioms have been handcrafted for and entered into CYC's knowledge base, and millions more have been inferred and cached by CYC. This article examines the fundamental assumptions of doing such a large-scale project, reviews the technical lessons learned by the developers, and surveys the range of applications that are or soon will be enabled by the technology.

One can think of CYC as an expert system with a domain that spans all everyday objects and actions. For example:

- You have to be awake to eat.
- You can usually see people's noses, but not their hearts.
- Given two professions, either one is a specialization of the other or else they are likely to be independent of one another.
- You cannot remember events that have not happened yet.
- If you cut a lump of peanut butter in half, each half is also a lump of peanut butter; but if you cut a table in half, neither half is a table.

By codifying reams of commonsense knowledge, CYC automates the white space in documents to help standardize—and make more efficient—information retrieval, integration, and consistency checking.





Abstract/Formal Tasks vs Intuition

- ▶ Abstract and formal tasks that are among the most difficult undertakings for a human being are among the easiest for a computer.
- ▶ Computers have long been able to defeat even the best chess player but only recently have begun matching some of the abilities of average human beings to recognise objects or speech.
- ▶ Much of human knowledge is about unstructured “inputs” (e.g., sensory data).
- ▶ Computers need to capture this same knowledge in order to behave in an intelligent way.

Extracting Patterns from Raw Data

- ▶ AI systems need the ability of acquiring their own knowledge by extracting patterns from raw data.
 - ▶ Usually, this capability is referred to as *machine learning*.
 - ▶ Machine learning allows computer systems to learn through data and experience.
 - ▶ Examples: naive Bayes, logistic regression, decision tree, random forest, etc.

Representation and Features

- ▶ The performance of these algorithms depends heavily on the *representation* of the data they are given.
- ▶ For example, in order to determine if a patient has a certain disease or not we can input certain physiological indicators (with thresholds, for example, temperature higher than 37.5C) and/or the presence of absence of certain symptoms.
- ▶ Each piece of information included in the representation of the patient is called known as a *feature*.
- ▶ A machine learning algorithm (let's say logistic regression) learns how each of these features of the patient are linked to a certain condition.

Mail thinks this message is Junk Mail.

Move to Inbox

☆ **MARK ZUCKERBERG**

Junk - Google August 24, 2018 at 10:48 AM



WINNING AMOUNT

Reply-To: MARK ZUCKERBERG

WINNING AMOUNT


My name is Mark Zuckerberg, A philanthropist the founder and CEO of the social-networking website Facebook, as well as one of the world's youngest billionaires and Chairman of the Mark Zuckerberg Charitable Foundation, One of the largest private foundations in the world. I believe strongly in 'giving while living' I had one idea that never changed in my mind - that you should use your wealth to help people and i have decided to secretly give {\$1,500,000.00} to randomly selected individuals worldwide. On receipt of this email, you should count yourself as the lucky individual. Your email address was chosen online while searching at random. Kindly get back to me at your earliest convenience, so I know your email address is valid. (mzuckerberg2444@gmail.com) Email me Visit the web page to know more about me: https://en.wikipedia.org/wiki/Mark_Zuckerberg/ or you can google me (Mark Zuckerberg)

Regards,
MARK ZUCKERBERG


Representation and Features

- ▶ In this case features might be occurrences or not of certain words, formatting, length of the message or other information related to the email protocols.
- ▶ We can build a vector of values (continuous and discrete) representing each email. Each element of the vector will be associated to one feature.
- ▶ Many artificial intelligence tasks can be solved by designing the right set of features.
- ▶ However, for many tasks it is difficult to know what features should be extracted.
 - ▶ Thinks about identification of a photo, emotion in a voice of a speaker, understanding images of a road, playing a complex videogame (e.g., Starcraft).
- ▶ Note: it is not only about the features themselves, but how the information is structured and “represented”:
 - ▶ Machine learning on Roman numbers is probably not a good idea.

Cartesian coordinates



Polar coordinates



Credit: Goodfellow et al. 2016

Representation Learning

- ▶ One solution to this problem is to use machine learning *to discover not only the mapping from representation to output but also the representation itself.*
- ▶ This approach is usually known as *representation learning.*
- ▶ Learned representation often results in much better performance than can be obtained with hand-made representations.

Factors of Variation


- ▶ When designing features or algorithms for learning features, our goal is to separate the *factors of variation* that explain the observed data.
- ▶ In this context, we use word *factors* to refer to separate sources of information that are useful for the machine learning task at hand.
- ▶ Such factors are often quantities that are not directly observed.
 - ▶ They are often unobserved (or latent) and they affect the observable ones.
 - ▶ Some of them might be linked to human constructs (e.g., the colour of an object) and other might not. In the latter case the factors might not easily interpreted by a human (see also the problem of AI interpretability).

Factors of Variation

- ▶ Some factors of variations affect all the piece of information we have (for example angle of view of a car).
- ▶ We need to disentangle the factors that allow us to successfully perform the machine learning task and extract representations that are not affected by factors of variation that are not “useful” for the task.
- ▶ For example, if we need to classify a car vs truck, the angle itself is not fundamental for the classification.
- ▶ We need a tool that learns to “ignore” that factor of variation.

Deep Learning


- ▶ Deep learning address this problem of representation learning by introducing representations that are expressed in terms of other, simpler representations.
- ▶ Classic example of deep learning model is the feedforward deep network (or multi-layer perceptron).
- ▶ Please note: a multilayer perceptron at the end is a (complex) mathematical function mapping input values to output values.
- ▶ This function is the result of combining several simpler functions in the intermediate nodes.




Credit: Zeiler and Fergus 2014

Deep Learning

- ▶ Until now, we have considered one of the possible interpretation of deep learning, i.e., that it allows to learn the right representation for the data.
- ▶ Another possible perspective on deep learning is that depth enables a computer to learn a multi-step computer program.
- ▶ Each layer of the network can be thought as the state of the computer's memory after executing a set of instructions in parallel.



Credit: Goodfellow et al. 2016



Credit: Goodfellow et al. 2016

References

- ▶ Chapter 1 of Ian Goodfellow, Yoshua Bengio and Aaron Courville. Deep Learning. MIT Press. 2016.