Certificate I: Understanding AI and Machine Learning in Africa

Course AIML01: Artificial Intelligence – Past, Present, and Future

Module 3: Example Applications

Lecture 1: AI Applications in Medicine

Carnegie Mellon University
Africa

Learning Objectives

1. Identify several applications of AI in the fields of medicine, biochemistry, and genomics
2. Highlight the impact of using artificial intelligence in these applications
3. Discuss the technical and ethical challenges posed by deploying artificial intelligence in these fields

Lecture Contents

1. Early expert systems
2. The impact of machine learning
3. Applications based on deep learning
4. Technological and ethical challenges
5. Lecture summary
6. Recommended reading & references

Early Expert Systems

DENDRAL

Work on DENDRAL began in 1965 at Stanford University by
Edward Feigenbaum, Bruce G. Buchanan, Joshua Lederberg, and Carl Djerassi

Written in the Lisp programming language

- Considered to be the first expert system
  - It was the first rule-based system to be applied to a real-world problem (Lindsay et al. 1993)

- The primary aim was to study the formation of hypotheses
  - Identify unknown organic molecules
  - Using the knowledge of expert organic chemists

- Many systems were derived from DENDRAL, including MYCIN

Early Expert Systems

MYCIN

Work on MYCIN began in 1972 at Stanford University by Edward Shortliffe

- A rule-based expert system for medical diagnosis
  - Designed to assist physicians
  - By recommending treatments for infectious diseases
    - Identify the bacteria that cause blood infections and other infections such as meningitis
    - Identify the correct antibiotics and the proper dosage
- Encoded the therapeutic decision rules that physicians found useful in their clinical practice

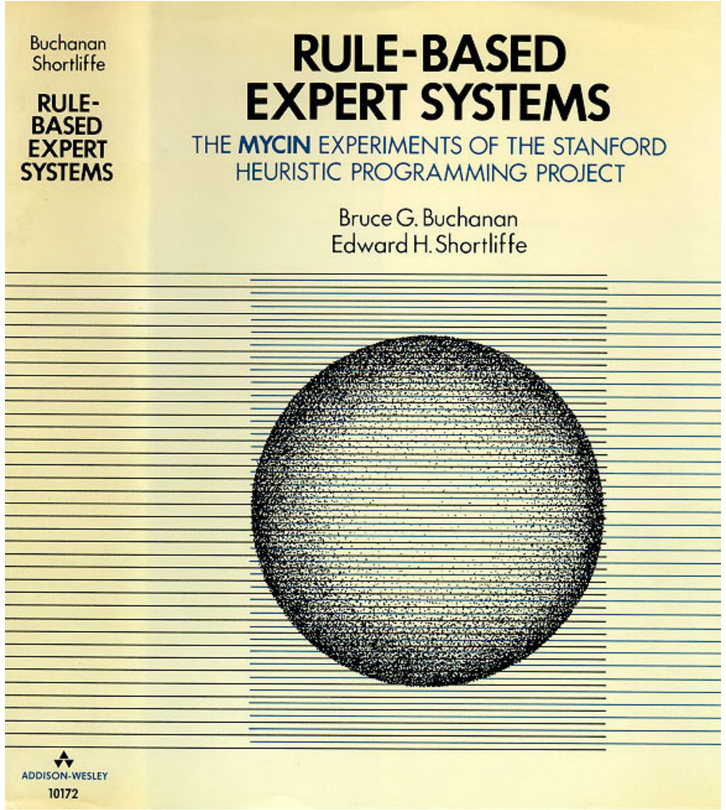

http://www.shortliffe.net/Buchanan-Shortliffe-1984/MYCIN%20Book.htm

Machine Learning in Medicine

More recently, deep machine learning has led to a resurgence of the development of medicine AI systems (Miotto et al., 2018)

- Diagnosis using clinical imaging
- Analysis of electronic health records
- Modelling genomics
- Analysis of medical sensor data

Machine Learning in Medicine

These machine learning systems use a variety of deep learning architectures

- CNN: convolutional neural network
- RNN: recurrent neural network
- LSTM: long short-term memory
- GRU: gated recurrent units
- RBM: restricted Boltzmann machine
- AE: autoencoder

Applications based on Deep Learning

Example applications include

- Prediction Alzheimer disease based on analysis of magnetic resonance imaging (MRI) scans
- Diagnosis of benign and malignant breast nodules from ultrasound images
- Classification of different types of skin cancer
- Prediction of congestive heart failure and chronic obstructive pulmonary disease from patient clinical health records
- Prediction of future clinical events from electronic health records
- Prediction of suicide risk
- Classification of cancer from gene expression profiles

Miotto et al. (2018) survey 32 different medical applications that use deep learning

Applications based on Deep Learning

Deep learning methods have been widely used for image-based cancer detection and diagnosis (Hu et al., 2018)

- Breast cancer (13)
- Lung cancer (11)
- Skin cancer (10)
- Prostate cancer (11)
- Brain cancer (9)
- Colorectal cancer (10)
- Other types of cancer, e.g., cervical, urinary tract, liver (12)

The bracketed numbers denote the number of papers in the survey by Hu et al. (2018)

Applications based on Deep Learning

These cancer detection and diagnosis systems mainly use convolutional neural networks (CNNs) as well as a variety of other architectures

- CNN (63)
- FCN (6) ← Fully convolutional network, a variant of a CNN which can perform pixel-wise classification
- SSAE (6) ← Stacked sparse autoencoder which can learn rich feature representations of the data
- DBN (4) ← Deep belief network, a probabilistic generative model which is constructed with a stack of restricted Boltzmann machines

Here, the bracketed numbers denote the instances of the model in the survey by Hu et al. (2018)

Applications based on Deep Learning

Example: skin cancer diagnosis

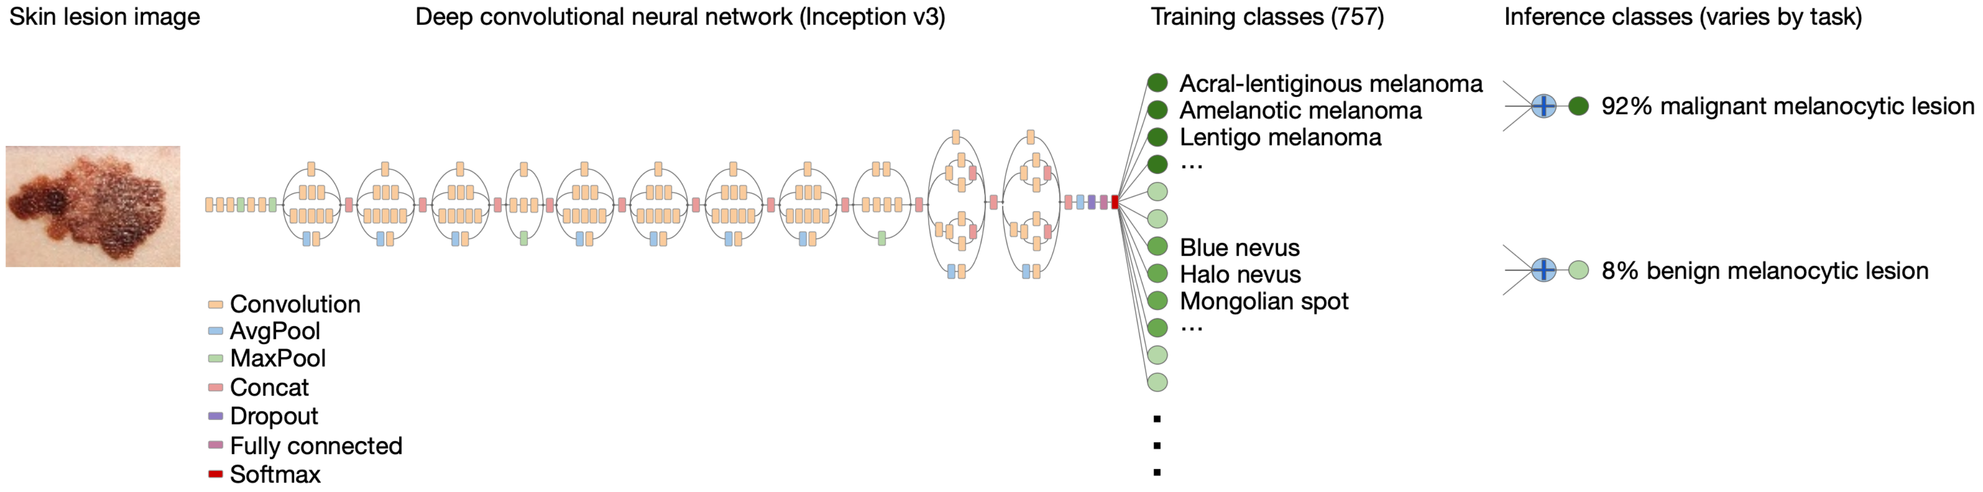

(Esteva et al., 2017)

Applications based on Deep Learning

- The performance in classifying biopsy-proven clinical images was similar to that of 21 board-certified dermatologists

- The CNN was trained using a dataset of 129,450 images with 757 disease classes

- For a more detailed explanation, see (Esteva et al., 2017)

Applications based on Deep Learning

Protein Folding

"Proteins are essential to life, supporting practically all its functions.

They are large complex molecules, made up of chains of amino acids, and what a protein does largely depends on its unique 3D structure.

Figuring out what shapes proteins fold into is known as the 'protein folding problem', and has stood as a grand challenge in biology for the past 50 years."

https://www.deepmind.com/blog/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology

Applications based on Deep Learning

Protein Folding

"Many of the world’s greatest challenges, like developing treatments for diseases or finding enzymes that break down industrial waste, are fundamentally tied to proteins and the role they play."

Applications based on Deep Learning

A recent landmark achievement in medicine and biochemistry is the AlphaFold AI model for protein folding, i.e., protein structure prediction (Senior et al. 2020)

https://www.deepmind.com

- Developed by Google DeepMind

- Winner of the biennial Critical Assessment of protein Structure Prediction (CASP13) competition in 2018

Applications based on Deep Learning

AlphaFold

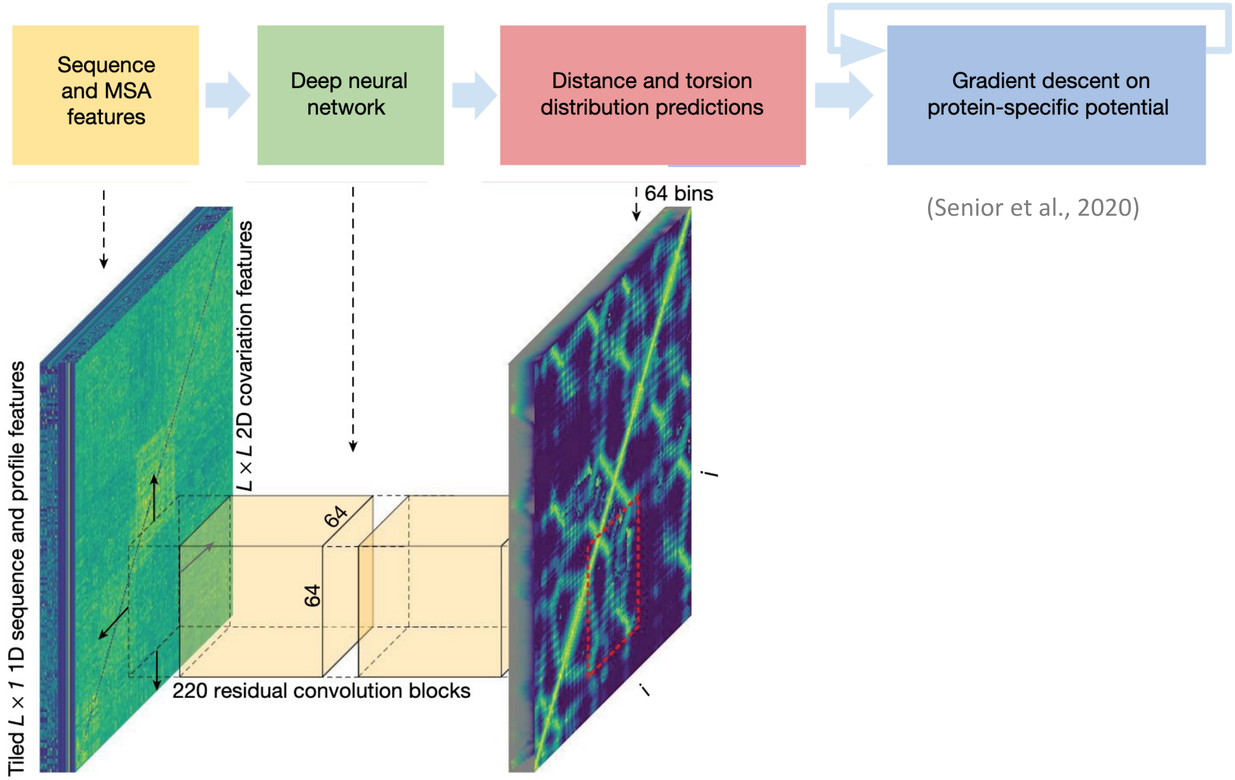

(Senior et al., 2020)

Applications based on Deep Learning

AlphaFold2 wins CASP14 in 2020 by a huge margin and is recognized as the solution to a 50-year-old grand challenge in biology

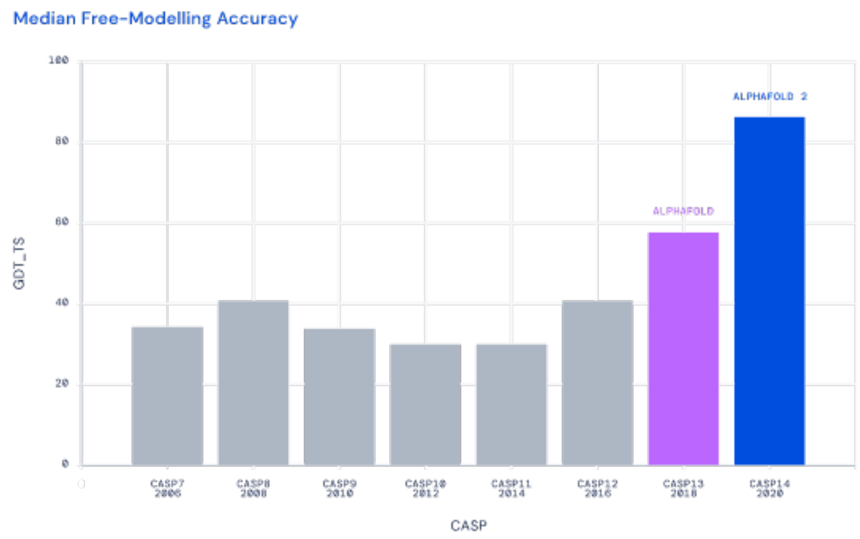

https://www.deepmind.com/blog/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology

Applications based on Deep Learning

AlphaFold2

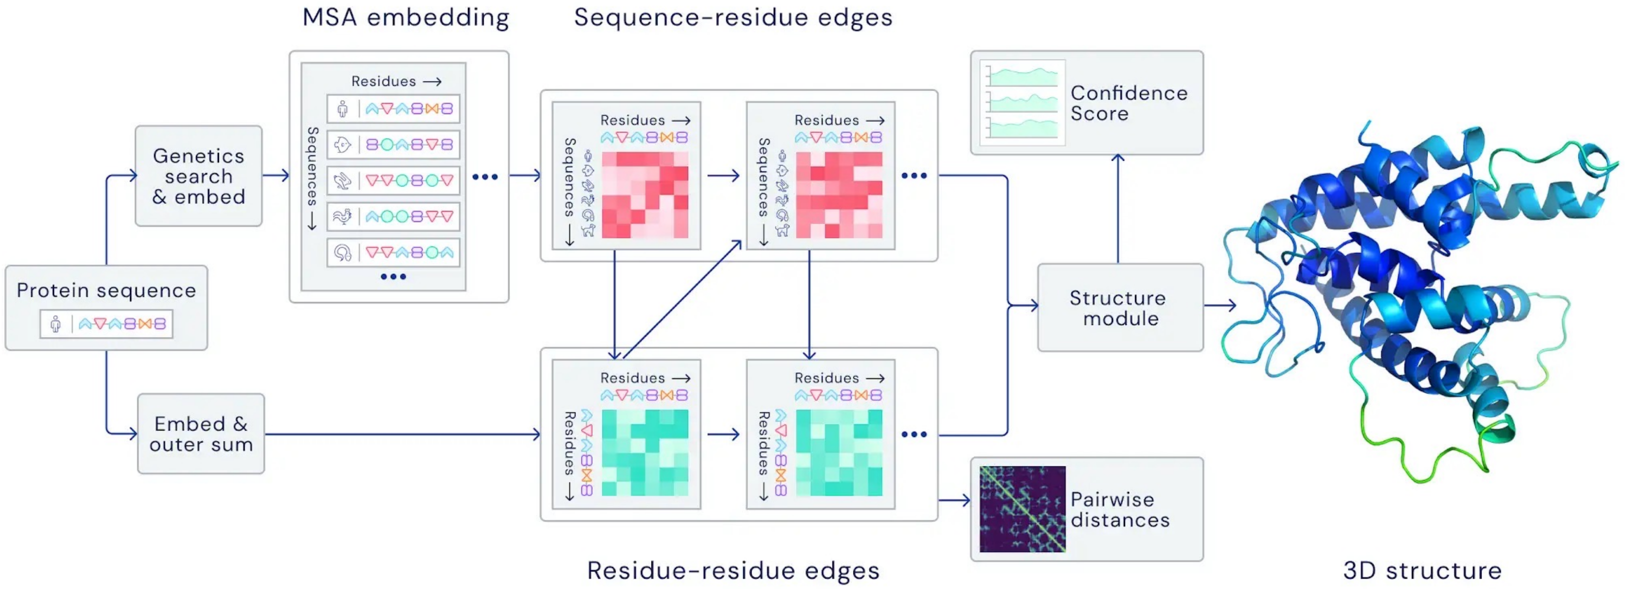

https://www.deepmind.com/blog/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology

For a timeline of the continuing development of AlphaFold and its impact, see
https://www.deepmind.com/research/highlighted-research/alphafold/timeline-of-a-breakthrough

Technological and Ethical Challenges

Medical AI applications present significant technological and ethical challenges

One key issue is the reliance on the quality and variety of the training data

Healthcare datasets typically are

- Sparse
- Noisy
- Heterogeneous
- Time-dependent

Technological and Ethical Challenges

New methods and tools are needed

- To enable interactive machine learning to interface with healthcare information workflows
- To keep the human in the loop

Technological and Ethical Challenges

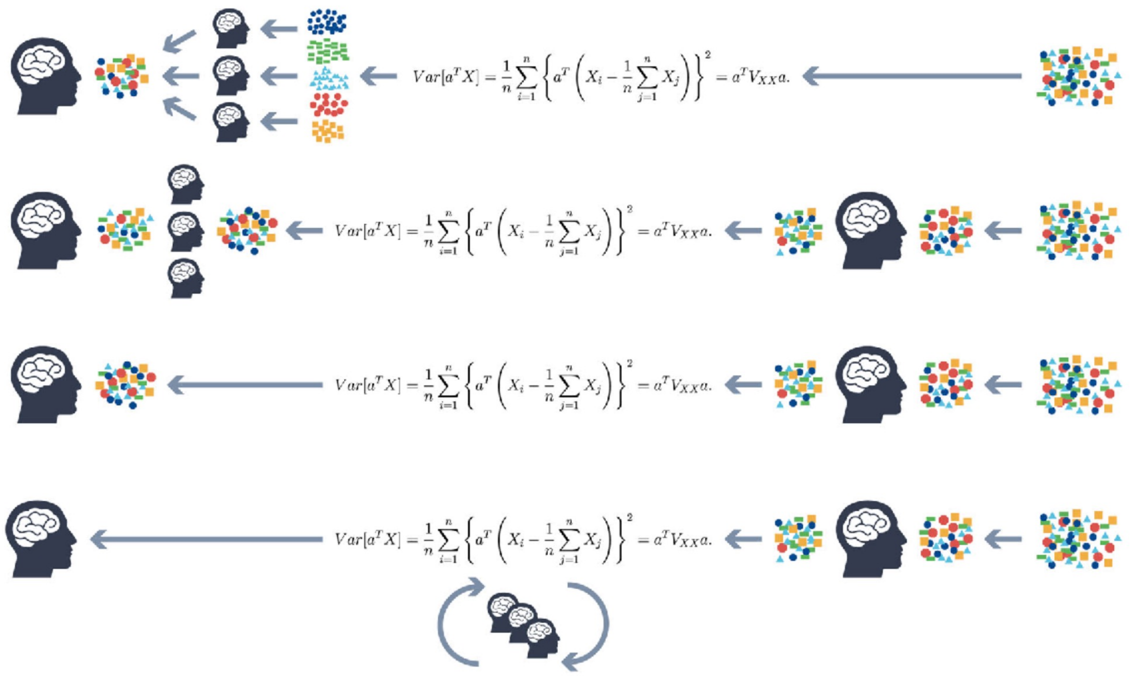

Unsupervised learning

Supervised learning:
humans provide labels for training data sets or selecting features

Semi-ssupervised learning

Human-in-the-loop approach:
humans are not only involved in pre-processing, by selecting data or features, but also during the learning phase, directly interacting with the algorithm

(Holzinger, 2016)

Technological and Ethical Challenges

There are also important ethical considerations

- The need for explainable systems
- So that clinicians (both novice and expert doctors) can access causal explanations of the AI's decision-making process

We return to this issue in a later course dedicated to trustworthy, explainable AI

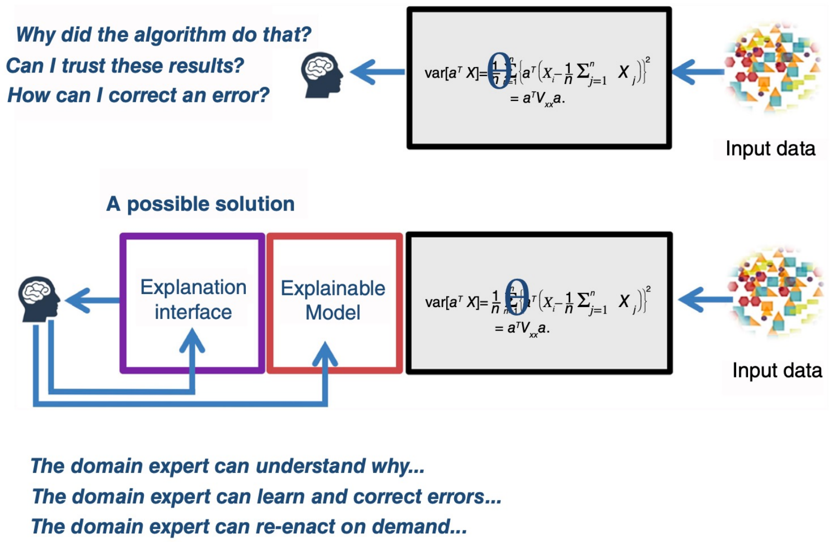

(Holzinger et al., 2019)

Lecture Summary

1. Early AI solutions, i.e. expert systems, had the goal from the beginning of being comprehensible, understandable, and thus explainable because they encoded the therapeutic decision rules that physicians found useful in their clinical practice

2. Unfortunately, they had limited success, leading to the AI Winter in the 1980s

3. The advent of machine learning, in particular statistical data-driven machine learning and deep neural networks heralded the present AI Summer and has yielded many breakthroughs and successful applications of AI in medicine

4. They also bring challenges, both technological and ethical, regarding their integration in human-centred processes and their ability to provide explainable solutions

Recommended Reading

Cangelosi, A. and Vernon, D. (2022). "Artificial Intelligence: Powering the Fourth Industrial Revolution", in EPS Grand Challenges: Physics for Society at the Horizon 2050, coordinated by the European Physical Society.
http://vernon.eu/publications/2022_Cangelosi_Vernon.pdf

Holzinger, A., Langs, G., Denk, H., Zatloukal, K., & Müller, H. (2019). Causability and explainability of artificial intelligence in medicine. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 9(4), e1312.
https://wires.onlinelibrary.wiley.com/doi/full/10.1002/widm.1312

Hu, Z., Tang, J., Wang, Z., Zhang, K., Zhang, L., Sun, Q. (2018). Deep learning for image-based cancer detection and diagnosis − A survey, Pattern Recognition, Volume 83, pp. 134-149.
https://doi.org/10.1016/j.patcog.2018.05.014

Miotto R., Wang F., Wang S., Jiang X., Dudley J. T. (2018). Deep learning for healthcare: review, opportunities and challenges. Brief Bioinform. Vol.19, No. 6, pp.1236-1246.
https://pubmed.ncbi.nlm.nih.gov/28481991/

References

Buchanan, B. G. and Shortliffe, E. H. (1984). Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project, Addison Wesley, Reading, MA.
http://www.shortliffe.net/Buchanan-Shortliffe-1984/MYCIN%20Book.htm

Esteva, A., Kuprel, B., Novoa, R. et al. (2017) Dermatologist-level classification of skin cancer with deep neural networks. Nature 542, pp. 115–118.
https://doi.org/10.1038/nature21056

Holzinger, A. (2016). Interactive machine learning for health informatics: when do we need the human-in-the-loop?. Brain Informatics, 3(2), 119-131.
https://link.springer.com/article/10.1007/s40708-016-0042-6

References

Lindsay, R. K., Buchanan, B. G., Feigenbaum, E. A., and Lederberg, J. (1993). DENDRAL: A case study of the first expert system for scientific hypothesis formation, Artificial Intelligence, Volume 61, Issue 2, pp. 209-261. https://doi.org/10.1016/0004-3702(93)90068-M

Senior, A. W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., Qin, C., Žídek, A., Nelson, A. W. R., Bridgland, A., Penedones, H., Petersen, S., Simonyan, K., Crossan, S., Kohli, P., Jones, D. T., Silver, D., Kavukcuoglu, K., Hassabis, D. (2020). Improved protein structure prediction using potentials from deep learning. Nature, 577(7792), pp. 706-710. https://www.nature.com/articles/s41586-019-1923-7.epdf