

Carnegie Mellon University Africa
Certificate Program in AI and Machine Learning in Africa

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

Module 3: Example Applications
Lecture 1: AI Applications in Medicine

Welcome to Lecture 1 of Module 3, the first of four lectures that look at the applications of artificial intelligence. This lecture focusses on applications in medicine. In the three lectures that follow, we will look at applications in robotics, for the web and social media, and in sports. In a different micro course – AIML02 – we will look at applications of AI and machine learning specifically in Africa.

In this lecture, we look briefly at the early GOFAI expert systems, such as the MYCIN for infectious diseases and DENDRAL on discovery of chemical compounds. We will then highlight the impact of machine learning, and its focus on learning from data, which has led to a resurgence of the development of medicine AI systems. We look in particular at the impact of deep learning on areas such as image-based cancer detection and diagnosis and protein folding. We then identify the significant technological and ethical challenges posed by the application of AI in medicine, such as the need for systems that can explain their results. We will finish up by summarizing what we have covered and identifying the articles that you should read to consolidate what you have learned.

We have three learning objectives, so that, after studying the material covered in this lecture, you should be able to do the following.

1. Identify several applications of AI in the fields of medicine, biochemistry, and genomics.
2. Identify the dominant AI and machine learning techniques used in these applications.
3. Discuss the technical and ethical challenges posed by deploying artificial intelligence in these fields.

Slide 1 Welcome to Lecture 1 of Module 3, the first of four lectures that look at the applications of artificial intelligence. This lecture focusses on applications in medicine.

In the three lectures that follow, we will look at applications in robotics, for the web and social media, and in sports.

In a different micro course – AIML02 – we will look at applications of AI and machine learning specifically in Africa.

Slide 2 DENDRAL is considered to be the first expert system to be applied to a real-world problem.

Recall from Lecture AIML01-02-01 that an expert system is a program that represents the knowledge of the human expert in a specific domain,

using a set of IF-THEN production rules to encode knowledge and draw inferences or conclusions,

and to offer advice to non-experts or to provide solutions to experts.

The primary aim of DENDRAL was to study the formation of hypotheses, and specifically to identify unknown organic molecules using the knowledge of expert organic chemists.

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

DENDRAL was written in the Lisp programming language.

Many systems were derived from DENDRAL, including MYCIN.

Slide 3 MYCIN was the first rule-based expert system for medical diagnosis.

It was designed to assist physicians by recommending treatments for infectious diseases.

It could identify the bacteria that cause blood infections and other serious infections such as meningitis and then identify the correct antibiotics for that infection and the proper dosage.

It encoded the therapeutic decision rules that physicians found useful in their clinical practice, sometimes referred to as heuristics.

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

Slide 4 More recently, the advent of machine learning, with its focus on learning from data, has led to a resurgence of the development of medical AI systems.

These systems range through

the use of deep learning for diagnosis using clinical images,
the analysis of electronic health records and medical sensor data,
modelling genomics, and
understanding protein structure.

Slide 5 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

We met many of these earlier in the course in Lecture AIML01-02-02.

Slide 6 Example applications include

- Prediction of Alzheimer's disease
- Diagnosis of breast nodules
- Classification of skin cancer
- Prediction of heart failure
- Prediction of pulmonary disease
- Prediction of future clinical events
- Prediction of suicide risk
- Classification of cancer

Miotto et al. (2018) survey 32 different medical applications that use deep learning. You can find the reference at in the Readings at the end of the lecture.

Slide 7 Deep learning methods have been widely used for image-based cancer detection and diagnosis

- Breast cancer
- Lung cancer
- Skin cancer
- Prostate cancer
- Brain cancer
- Colorectal cancer
- Other types of cancer, e.g., cervical, urinary tract, liver cancer.

Hu et al. survey 76 applications. Again, you can find the reference at in the Readings at the end of the lecture.

Slide 8 These cancer detection and diagnosis systems mainly use convolutional neural networks (CNNs) as well as a variety of other architectures

CNN (63), FCN (6), SSAE (6), DBN (4)

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

A FCN is a fully convolutional network, a variant of a CNN which can perform pixel-wise classification

A SSAE is a stacked sparse autoencoder which can learn rich feature representations of the data

And a DBN is a deep belief network, a probabilistic generative model which is constructed with a stack of restricted Boltzmann machines.

- Slide 9 For example, a GoogleNet Inception v3 CNN architecture was used for skin cancer diagnosis by Esteva and colleagues.
- Slide 10 The performance in classifying biopsy-proven clinical images is similar to that of 21 board-certified dermatologists.
- The CNN was pre-trained on approximately 1.28 million images, with 1,000 object categories, from the 2014 ImageNet Large Scale Visual Recognition Challenge,
- and then trained using transfer learning on the skin cancer dataset of 129,450 images with 757 disease classes.
- For a more detailed explanation, see the paper by Esteva et al. (2017).
- Slide 11 The next application we look at is 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"
- Slide 12 "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.
- Slide 13 A recent landmark achievement in medicine and biochemistry is the **AlphaFold AI model for protein folding**.
- It was developed by Google DeepMind and was the winner in 2020 of the biennial Critical Assessment of protein Structure Prediction (CASP) competition.
- It is extremely difficult to predict protein structure but AlphaFold achieved a performance similar to the results from experimental methods.

- Slide 14 Protein sequences derived from DNA sequencing are aligned to generate a multiple sequence alignment (MSA).
- A deep neural network is trained to make accurate predictions about the structure of the protein given its sequence, constructing a learned, protein-specific potential.
- The neural network predictions include backbone torsion angles and pairwise distances between amino acid residues.
- The structure itself accurately predicted by minimizing the potential by gradient descent.
- Slide 15 AlphaFold2 wins CASP14 in 2020 by a huge margin and is recognized as the solution to a 50-year-old grand challenge in biology.
- Slide 16 AlphaFold2 uses a different approach to AlphaFold1
- Using attention-based deep neural networks, referred to as transformers, to interpret the structure of the spatial graphs used to represent the proteins
- Slide 17 We turn now to consider the technological and ethical challenges of using AI in medicine.
- Medical AI applications present significant technological and ethical challenges.
- One key issue is the reliance on the quality and variety of the training data.
- This poses a challenge because healthcare datasets typically are sparse, noisy, heterogeneous, and time-dependent.
- Slide 18 New methods and tools are needed to enable interactive machine learning to interface with healthcare information workflows, keeping the human in the loop.
- Slide 19 Humans are typically involved in supervised and semi-supervised learning, providing labels for training data sets or selecting features
- However, in the 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.

Slide 20 There are also important ethical considerations.

One is for example the need for explainable systems so that clinicians (both novice and expert doctors) can access causal explanations of the AI system's decision-making process.

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

.

To summarize:

1. Early AI solutions, called 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.

We will study these issues in more detail in later courses.

Recommended Reading

Here is some recommended reading.

Bassereau, P., Cangelosi, A., Cejkova, J., Gershenson, C., Goldstein, R., Martins, Z., Matthews, S., Ritort, F., Seager, S., Van Tiggelen, B., Vernon, D., and Westall, F., (2024) Chapter 4 Physics for Understanding Life, **Section 4.3 Artificial Intelligence: Powering the Fourth Industrial Revolution**, EPS Grand Challenges: Physics for Society at the Horizon 2050, IOP Publishing.

http://vernon.eu/publications/2024_Bassereau_et_al.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

Here are the references cited to support the main points in what we covered today.

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>

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](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>