MACHINE LEARNING FOR HEALTH

Evolution and Foundations of AI in Health

From expert systems to human–machine collaboration

TECHNICAL HANDOUT · SESSION 01

Professor Flávio Luiz Seixas
Graduate Program in Computing

Academic teaching material based exclusively on the PDFs available in the project

Table of Contents

Introduction

Artificial intelligence in medicine has evolved from programs that explicitly represented knowledge and inference to systems that learn relationships from data. This shift broadened the technical repertoire but did not remove the need to understand clinical work, the quality of evidence, and the division of responsibilities between people and machines. [1]

This session reconstructs the historical cycles and compares knowledge-based systems with machine learning. The central question is less “which paradigm won?” and more “what information does each paradigm use, what output does it produce, and how does that output enter clinical cognition and workflow?” [1]

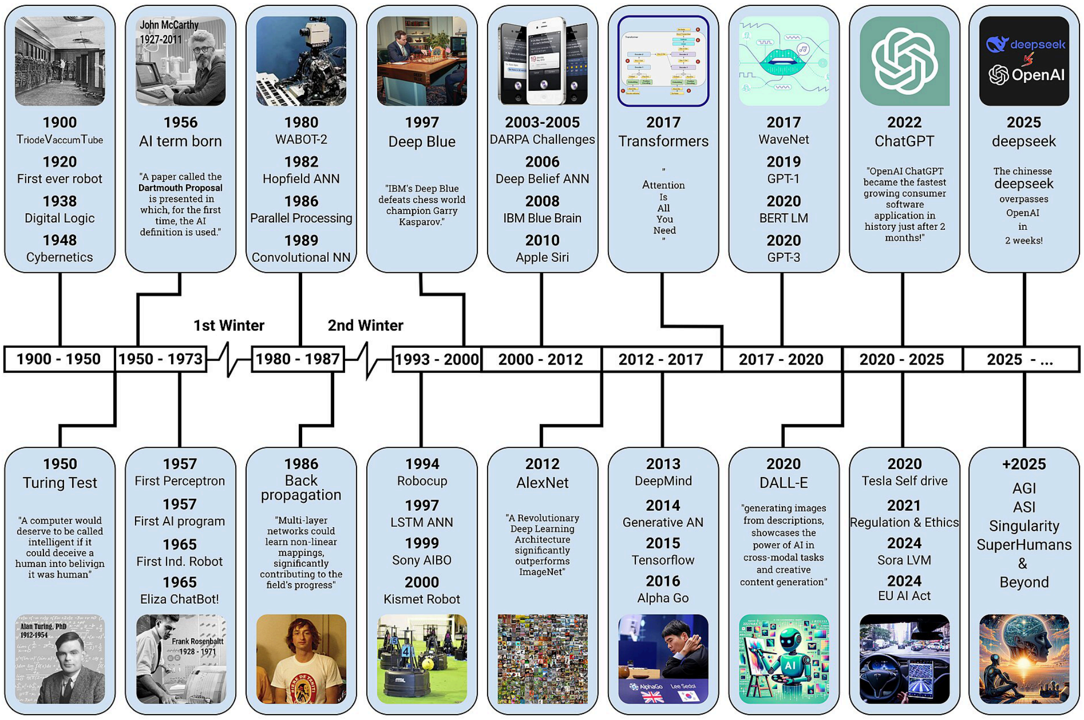

Figure 1 — Timeline used to organize cycles in AI. Visual source: credits recorded in the course material; technical content supported by [1] and [2].

1. From formalizing knowledge to learning

The term artificial intelligence became associated with the Dartmouth proposal and the hypothesis that aspects of learning and intelligence could be described precisely enough to be simulated by machines. In medicine, the field grew in close association with computational support for clinical reasoning [1].

Early expert systems sought to capture expert knowledge and represent it in structures that programs could manipulate. Rules and inference mechanisms made part of the reasoning inspectable, but acquiring and maintaining knowledge was labor-intensive; domain coverage and the handling of exceptions remained limiting factors [1].

KEY CONCEPT
The term artificial intelligence became associated with the Dartmouth proposal and the hypothesis that aspects of learning and intelligence could be described precisely enough to be simulated by machines.

2. Expert systems and learned models

MYCIN exemplifies a rule-based system designed to provide antimicrobial advice, whereas INTERNIST-1 was conceived to support diagnosis in internal medicine. The difference in task matters: selecting a therapy under specified conditions is not the same as organizing a broad differential diagnosis [2].

In machine learning, behavior is estimated from examples. The input is a representation of data; the output depends on the task, such as a class, risk, or value. The quality of this function depends on the problem definition, sample, and label. Predictive complexity does not, by itself, imply an understanding of the case [3].

3. Clinical cognition and human–machine collaboration

Cognitive informatics examines how people represent knowledge, solve problems, and interact with systems. The relevant unit of analysis may be the human–machine team because a model's output is interpreted in a context of goals, memory, attention, communication, and workflow [4].

A prediction estimates something from inputs; a decision requires combining that estimate with alternatives, consequences, and preferences. Systems can complement human limitations, but they can also redirect attention and introduce new errors. Interaction design is therefore part of system safety [4].

4. From performance to benefit

Evaluating a model only in the laboratory does not demonstrate clinical benefit. Validity, utility, integration, and effects on patients and professionals must also be examined. Explanations should likewise be evaluated according to their purpose: debugging, auditing, communication, and support for action may require different forms of transparency [5].

Session summary

From formalizing knowledge to learning: the term artificial intelligence became associated with the dartmouth proposal and the hypothesis that aspects of learning and intelligence could be described precisely enough to be simulated by machines.
Expert systems and learned models: mycin exemplifies a rule-based system designed to provide antimicrobial advice, whereas internist-1 was conceived to support diagnosis in internal medicine.
Clinical cognition and human–machine collaboration: cognitive informatics examines how people represent knowledge, solve problems, and interact with systems.
From performance to benefit: evaluating a model only in the laboratory does not demonstrate clinical benefit.

References

[1] Cohen, T. A.; Patel, V. L.; Shortliffe, E. H. Introducing AI in Medicine. In: Intelligent Systems in Medicine and Health. Springer, 2022. External link

Local file: pdfs/3-Intelligent Systems/I-Introducing AI in Medicine.pdf · PDF pages used: 1, 2, 5, 11.

[2] Shortliffe, E. H.; Buchanan, B. G.; Feigenbaum, E. A.; Shah, N. H. AI in Medicine: Some Pertinent History. In: Intelligent Systems in Medicine and Health. Springer, 2022. External link

Local file: pdfs/3-Intelligent Systems/II-AI in Medicine- Some Pertinent History.pdf · PDF pages used: 2, 15, 17, 20.

[3] Subramanian, D.; Cohen, T. A. Machine Learning Systems. In: Intelligent Systems in Medicine and Health. Springer, 2022. External link

Local file: pdfs/3-Intelligent Systems/VI-Machine Learning Systems.pdf · PDF pages used: 2, 4, 17, 67.

[4] Patel, V. L.; Cohen, T. A. Clinical Cognition and AI: From Emulation to Symbiosis. In: Intelligent Systems in Medicine and Health. Springer, 2022. External link

Local file: pdfs/3-Intelligent Systems/V-Clinical Cognition and AI- From Emulation to Symbiosis.pdf · PDF pages used: 1, 2, 18, 23.

[5] Shortliffe, E. H.; Sepúlveda, M.-J.; Patel, V. L. Framework for the Evaluation of Clinical AI Systems. In: Intelligent Systems in Medicine and Health. Springer, 2022. External link

Local file: pdfs/3-Intelligent Systems/XVII-Framework for the Evaluation of Clinical AI Systems.pdf · PDF pages used: 1, 2, 10, 20.