

50° CONGRESSO NAZIONALE ANMDO

Intelligenza umana ed Intelligenza artificiale: quale futuro?

Rita Cucchiara Università degli studi di Modena e Reggio Emilia

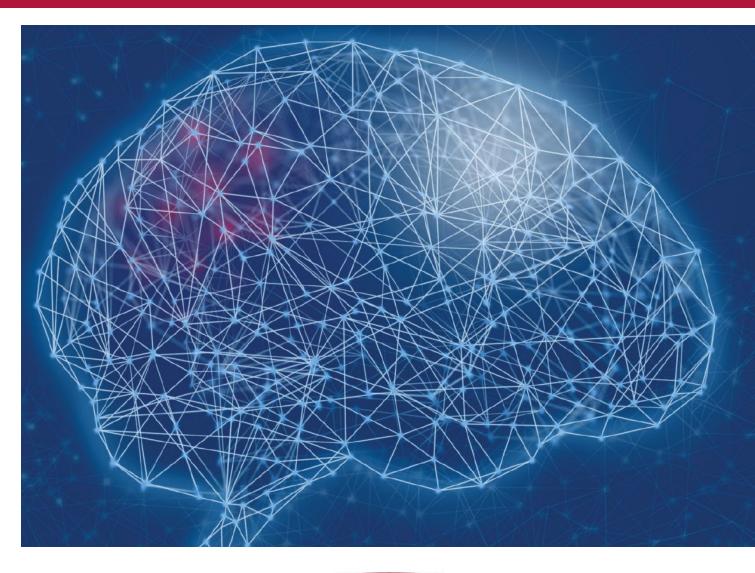




Prologue: Human vs Artificial Intelligence

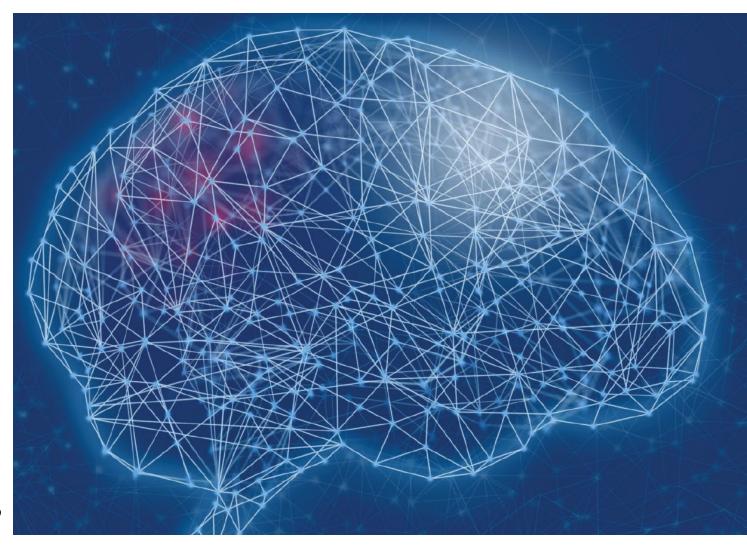
Intelligence is broadly linked to our ability to **think**, and consequently **understand**, **reason**, **and act effectively within the world**.

What does the machine think?





Nothing.



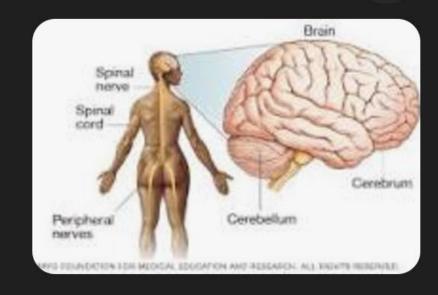
Are we sure?

Prologue

Is AI so different to human thinking?

How does the thinking brain work?

Neurotransmitters pass through the gap between two nerve cells, known as the synapse. They then attach to receptors on the receiving cell. This process repeats from neuron to neuron as the impulse travels to its destination. This web of communication that allows you to move, think, feel and communicate.







What Machines Do Instead

1.Process Information:

Machines perform computations, analyze data, and identify patterns using algorithms and mathematical models.

- **2.Follow Rules**: They operate within the constraints of their programming or training data, executing pre-defined tasks.(also backpropagation in training is a set of rules).
- 3. **Create representations:** With Neural Networks Machines receives external stimula, process them by parameters learned by data, transmit information among neurons, create an internal representation and produce output when required...

Why It Seems Like Thinking:

- Pattern recognition. Computer Vision,
- Audio understanding, Generating audio, video
- •Statistical probabilities, Logical processing.
- •Reasoning on temporal series...
- Plays with words
- Plan, reason and act

Is it really SO different?





Human vs Al: a quantitative comparison

Aspect	Human Brain	GPT-4 (Foundation Model)	
Memory Capacity	≈ 2.5 billion GB (long-term)	≈ 1,000 billion GB (training + context)	
Processing Units	86 billion neurons, 100 trillion synapses	1.8 trillion parameters (≈ synaptic weights)	
Processing Speed	~20 Hz neural firing ($\approx 10^{15}$ ops/sec parallel)	~10 ¹⁵ ops/sec on GPU clusters (petaflop range)	
Learning Speed	~20 years to full cognitive maturity	~3–6 months training on trillions of tokens	
Energy Consumption	20 W (≈ 7.3 MWh over 20 years)	≈ 1,750 MWh (~160 homes/year)	
Learning Mode	Continuous, sensory, emotionally grounded	Batch training; statistical pattern learning	
Adaptability	Learns from few examples, real-time updates	Needs re-training; limited continuous adaptation	
Error Correction	Self-adjusting via feedback and emotion	Probabilistic fine-tuning via reinforcement	





Intelligent human and machine together

NOW, Al-based systems can act as humans in many tasks, in different fields and applications and improve day-by-day...

what about Healthcare?



Al in Healthcare now

Credits: a fruiful conversation With ER President july 2025

Many application directions but the same technical models.

Al in Healthcare: Three Functions



Al for Clinical/Scientific Analysis & Individual Problem Solving



- Data analysis
- Diagnostic assistance
- Personalized treatment

Al for Citizen-Hospital Interface



- Customer care chatbot
- Pre-diagnosis
- Telemedicine

Al for Healthcare System Management & Optimization



- Resource allocation
- Data analysis ad prediction
- Healthcare System





Al for scientific/clinical data & individual problem solving

.. 1. Al can understand a lot



write please in a single long sentence what an AI can understand in medicine when "see an MRI imageand what thinks about



An AI, when presented with an MRI image, can analyze the intricate patterns of tissue density, signal intensity, and anatomical structures, identifying potential abnormalities like tumors, lesions, or degenerative changes, comparing these findings to vast datasets of normal and pathological images to assess the likelihood of specific conditions, and suggesting potential diagnoses or treatment strategies while considering factors like patient history, symptoms, and relevant medical guidelines.

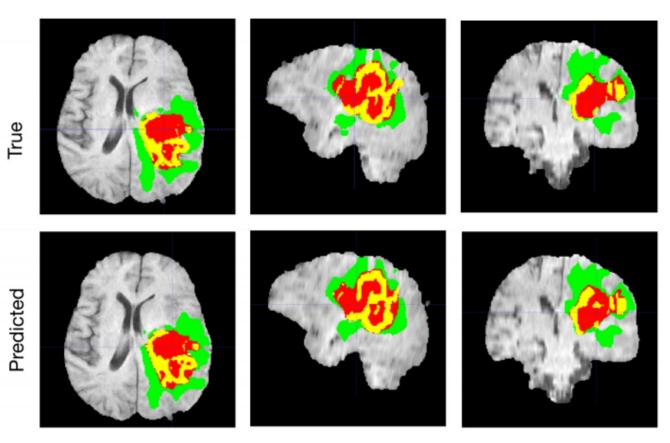
→ A ready but still not a mature technology in generality, explaination, convergence with human protocols, Learning rare patterns, defining specific activities..



An example

NVIDIA Autoencoder network (2024) winner of the BRATS19 Multimodal Brain Tumor Segmentation Challenge

256x20x24x16



https://developer.nvidia.com/blog/automatically-segmenting-brain-tumors-with-ai/



Examples

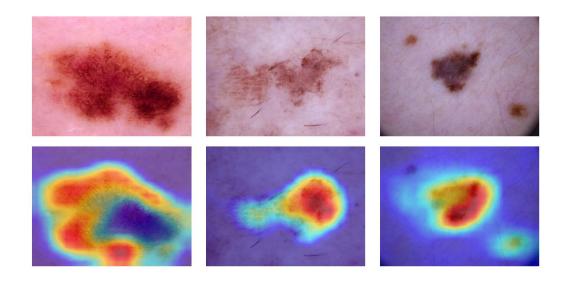
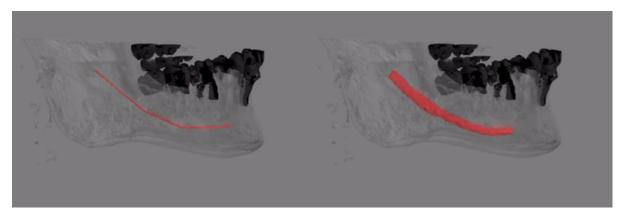


Fig. 7: Samples from the 2019 ISIC dataset and attention heatmaps obtained by means of the Grad-CAM method. Besides always

Grana et al . Almagelab 2019-2024







Bolelli et al. CVPR 2023 CBCT (Cone Beam Computed Tomography) 3D dataset





Al for citizenhospital interface

.. 2. Al can chat a lot

write in a short sencitence how an AI based system can help in the interface between citizen and hospitals or healthcare systems

An **AI-based system** can act as an intelligent bridge between citizens and healthcare institutions, making the entire interaction smoother and more responsive. It can manage **appointment scheduling**, remind patients of medications or follow-up visits, and provide **personalized health information** based on individual medical histories.

Through natural language interfaces like chatbots or voice assistants, it can help citizens **navigate hospital services**, locate specialists, and even perform **preliminary symptom assessments** before visits. It can also analyze population-level data to identify trends and **predict community health needs**, helping hospitals plan resources more effectively.

Overall, it turns the hospital-citizen relationship into a **continuous**, **data-informed dialogue**, improving accessibility, satisfaction, and preventive care.



1. Virtual Health Assistants & Chatbots

-handle communication, triage, and navigation for citizens. (Babylon Health UK, Ada Health D, Buoy Health US)

Impact: Reduces unnecessary hospital visits and improves triage efficiency.



2. Al Scheduling & Appointment Systems

-automate booking and help optimize hospital resource use. (ZocDoc US, Mayo Clinic's Al Appointments)

Impact: Shorter waiting times and smoother coordination between departments.



3. Al for Telemedicine & Continuous Monitoring

-connect patients to care remotely, often integrating data from wearables. (Teladoc Health, Sensely...)

Impact: Expands access to healthcare, especially for rural or mobility-limited patients.



Al citizen interfaces

AI LLMs

Pre-diagnosis chat, follow-up recommendations, FAQ automation, discharge guidance.

Multimodal LLs

Image-based triage, speech-based symptom capture, visual teleconsultation support.

Conversational AI tools

Smart triage assistants, 24/7 hospital chatbots, empathetic virtual nurses, voice-based hotline bots.

Warning: the costs are very high...

Citizen Hospital / Healthcare System

Al Large Language Models

- Analyze symptoms and provide information
- Answer questions about hospital procedures
- Generate instructions for patients
- Translate medical text into simplified language
- · Summarize clinical notes

Multimodal LLMs

- Interpret images and documents
- Analyze patient speech or text
- · Process sensor data
- Integrate multimodal data
- · Support visual telemedicine

Conversational Al Systems

- Assist with appointment scheduling
- Provide navigation of services
- Offer conversational support
- Escalate issues to human staff
- 24/7 virtual assistants





3. Al can do more...

write in few rows how an agentic AI system can work for optimization in hospitals, in helthcare systems (by considering all the pipeline, of planning reasoning acting reflecting..)

Al for hospital optimization an management

An agentic AI system in hospitals can autonomously plan, reason, act, and reflect to optimize care delivery and operations. It gathers and interprets real-time data from patients, staff, and infrastructure to reason about bottlenecks or risks, then plans optimal actions—like adjusting schedules, reallocating resources, or predicting patient surges. It can act directly through automation or support human decisions, and later reflect on outcomes, learning from feedback to improve future performance. Through this continuous cycle, the hospital evolves into a self-improving, data-driven ecosystem that enhances efficiency, safety, and patient care quality.

You Agentic AI doesn't just report data — **it interprets, plans, and acts.**It turns hospitals into **adaptive, self-optimizing ecosystems**, capable of learning from their own operation



Al agents in Hospital organization

1. Data Understanding and Integration

- -Collect, clean, and link data from multiple hospital sources EHRs, lab systems, IoT devices, and administrative platforms.
- Automatically detects inconsistencies or missing data.
- -Merges structured and unstructured information (e.g., clinical notes, images, lab results).
- Builds semantic knowledge graphs to represent patients, treatments, and outcomes as interconnected entities.
- Example: An agent identifies delays in lab result uploads and suggests process fixes to sync diagnostics with treatment timelines.

2. Predictive and Prescriptive Analytics

Unlike static models, agentic AI systems learn continuously from real-time hospital data.

- Forecast bed occupancy, emergency influx, or ICU demand.
- •Predict patient readmission risks, medication adherence, or complications.
- •Simulate "what-if" scenarios for policy or budget decisions.
- •Move from prediction to **prescription** by suggesting actionable steps (e.g., reassigning staff or adjusting supply orders).
- Example: The AI forecasts a 20% rise in cardiac admissions and recommends temporary nurse reallocation and resource pre-stocking.



Al agents in Hospital organization



3. Operational Optimization

Agentic AI can reason over hospital workflows and propose or even autonomously execute improvements.

- •Optimize **staff scheduling**, surgery block allocation, and outpatient flow.
- •Balance resources between departments based on dynamic priorities, detect process bottlenecks
- Example: The system reconfigures appointment slots after identifying underused diagnostic equipment during afternoon.



4. Strategic Decision Support

The system can assist management in **long-term planning and scenario modeling**.

- •Analyze population health trends to guide preventive care investments.
- •Generate impact projections of new clinical pathways or technologies.
- Example: It identifies correlations between local demographics and hospital admission patterns,

5. Coordination and Governance

Agentic AI can serve as a **collaborative management layer**, connecting analytics with action.

- •Acts as an intelligent mediator between departments, ensuring data-driven decisions are consistent across systems.
- •Monitors compliance with safety or quality standards. Summarizes insights in natural language.
- Example: When infection rates rise in one ward, the system cross-analyzes cleaning schedules, staff rotations, and patient transfers, then alerts administrators with a prioritized action list



A short journey in Al discovery

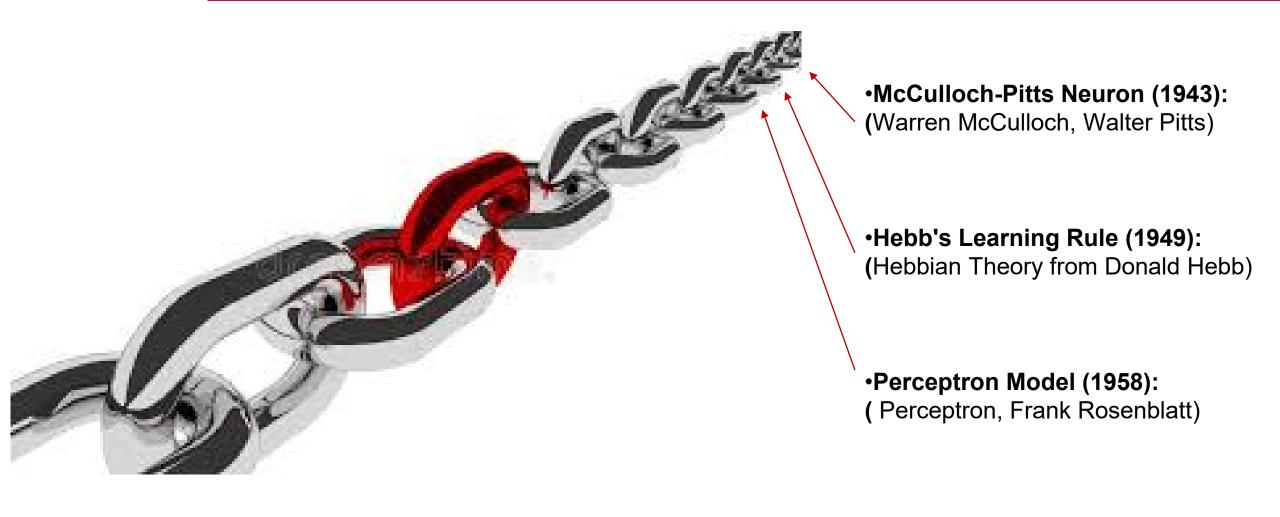
A short journey in AI discovery

- From a simple Neural-like computation model
- To supervised /unsupervised learning
- To Generative AI Models
- To Foundation models
- To Agentic AI models





The origin: the simple model





The 80s -90s- 20s



•Backpropagation Algorithm (1986): (David E. Rumelhart, Geoffrey E. Hinton, Ronald J. Williams)

•Autoencoder (1986):(Geoffrey Hinton, David Rumelhart)

•LeNet-5 (1998): (LeNet-5 Yann LeCun)

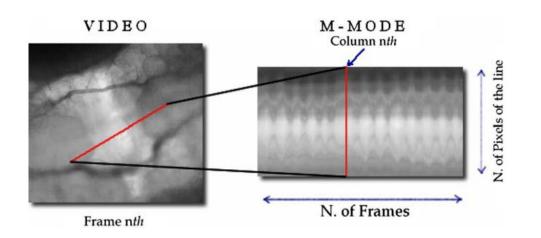
The statistical Machine Learning
ad Pattern Recognition time
(Bayesian classifiers, SVMs...)

2019 Turing Award 2024 Nobel



A n old example

 <2010 very rare use of Neural Networks for Medicine, discriminative Machine Learning



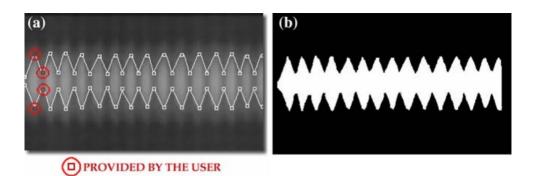


Table 3 Comparison of the performance of the neural network classifier (NNC), the Bayesian classifier (BAC), the hidden Markov

chain classifier (HMC), and the proposed multiclassifier (MCS)

Parameters NNC BAC HMC MCS

Parameters	NNC	BAC	HMC	MCS
$N_{D_{ m mean}}$ $N_{S_{ m mean}}$ ξ Points added Points changed	3.45(0.81) 2.07(0.92) 0.04(0.012) 0(0) 3.57(0.74)	1.82(0.49) 2.38(0.81) 0.041(0.013) 0.19(0.11) 3.45(0.53)	1.64(0.36) 2.19(1.31) 0.036(0.018) 0.067(0.067) 2.79(0.55)	0.64(0.16) 0.78(0.44) 0.015(0.0070) 0(0) 1.19(0.35)
Time (s) T_A	13	6.333	8.533	16.067

Figures in the first three main rows are the average error. In brackets the standard deviations

Luca Bertelli, Rita Cucchiara, Giovanni PaternostromAndrea PratiA semi-automatic system for segmentation of cardiac M-mode imagesPattern Analysis and Applications 2006-01-0



The deep learning expoit



- •Imagenet Dataset (2010)
- •(Fei Fei Li)
- •AlexNet (2012):
- •(Alex Krizhevsky, Geoffrey Hinton, Ilya Sutskever)
- •VGG (2014):

(VGGNet Karen Simonyan, Andrew Zisserman)

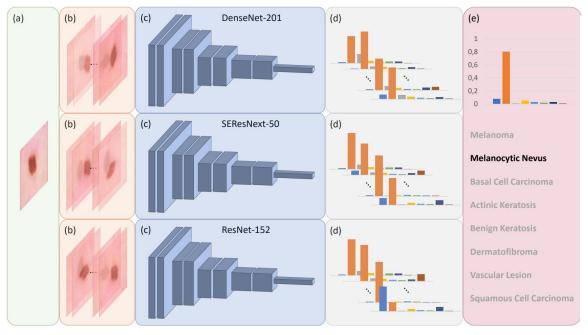
- •U-Net (2015):
- •(U-Net Olaf Ronneberger, Philipp Fischer, Thomas Brox) → for medical imaging
- •ResNet (2015):
- •(Residual Neural Network: Kaiming He, Xiangyu Zhang, Shaoqing Ren, Jian Sun)



New classifiers based on CNNs

Supervised networks

mostly for discriminative tasks



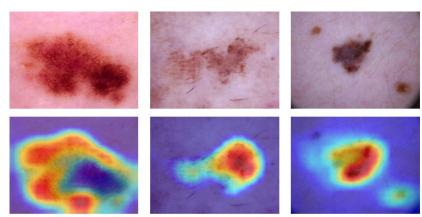


Fig. 7: Samples from the 2019 ISIC dataset and attention heatmaps obtained by means of the Grad-CAM method. Besides always locating the skin lesion, the network can aim its attention at specific sections such as lesion borders (leftmost image) or darker patches (rightmost image).

 Ensemble of CNN for secure Melanoma diagnosis and use of Grad-CAM for attention heatmaps in explainability

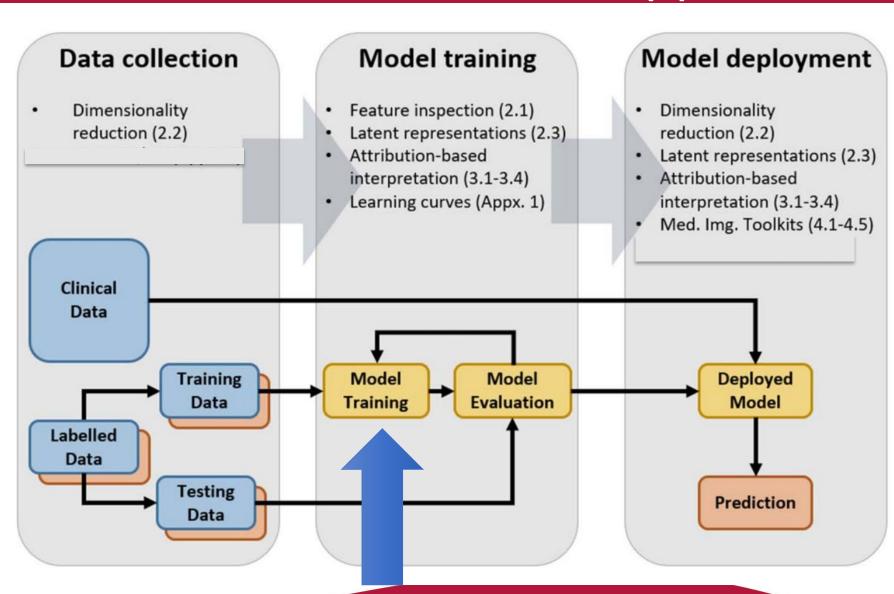
A Deep Analysis on High Resolution Dermoscopic Image Classification Pollastri... Grana IET 2021 (https://iris.unimore.it/retrieve/handle/11380/91227178/312990/2021_IETCV_A_Deep_Analysis_on_High_Resolution_Dermoscopic_Images_Classification.pdf



Common cues for all approaches

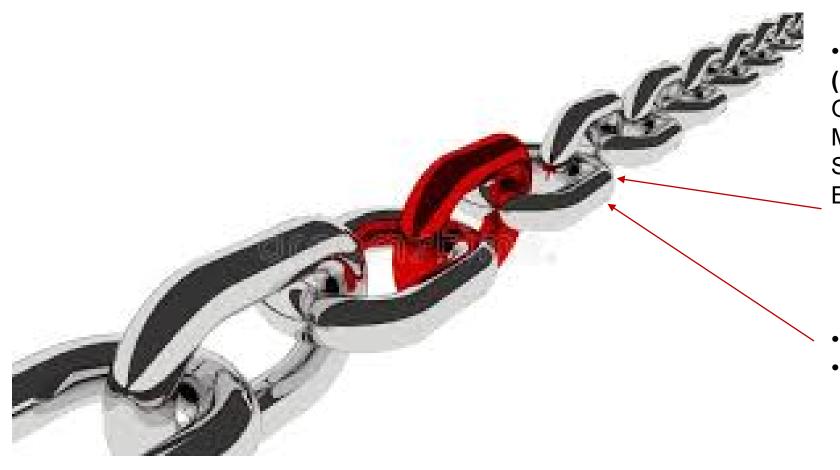
"Classic" DL- based supervised approaches

Discriminative vs. Generative Machine Learning





Generative Al for images (and video)



•GANs (2014):

(Generative Adversarial Networks Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, Yoshua Bengio.

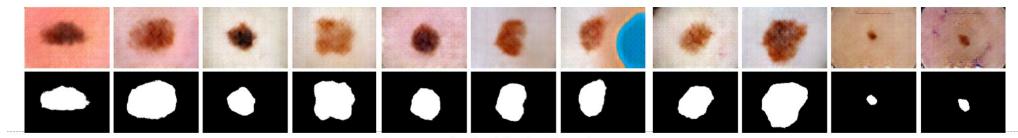
•Diffusion Models (2020-2021):

•(Jonathan Ho, A Jain, P Abbeel)

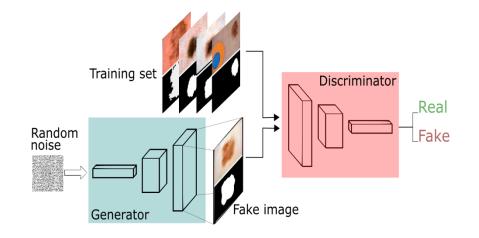


Generation for Data augmentation

• 2017-2019 SYNTHETIC DATA AUGMENTATION to solve the data challenge e .g use of GAN



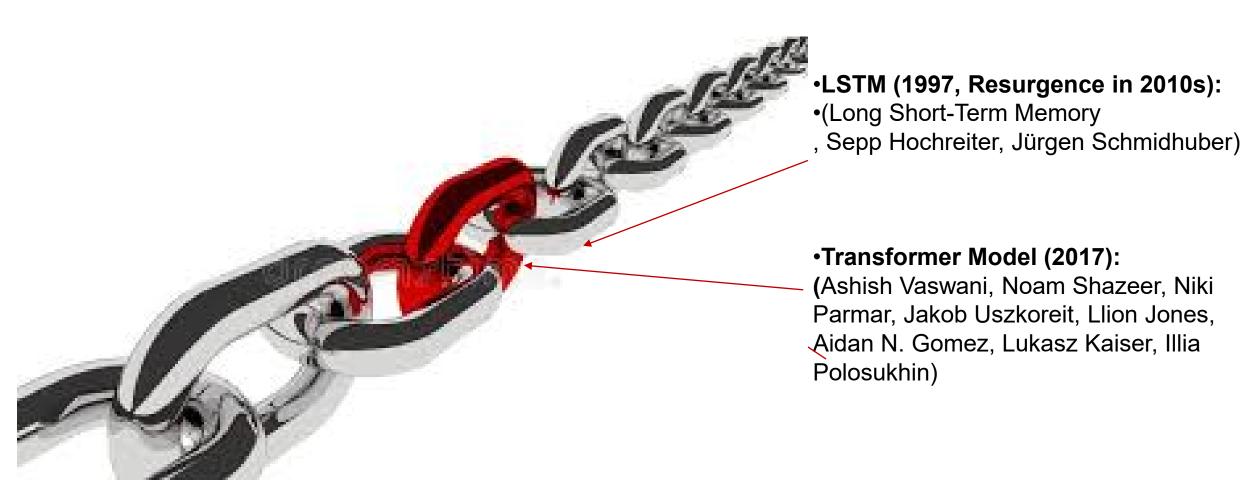
Generated skin lesion images



Thanks Grana et al 2019 EU DECIDER project UNIMORE



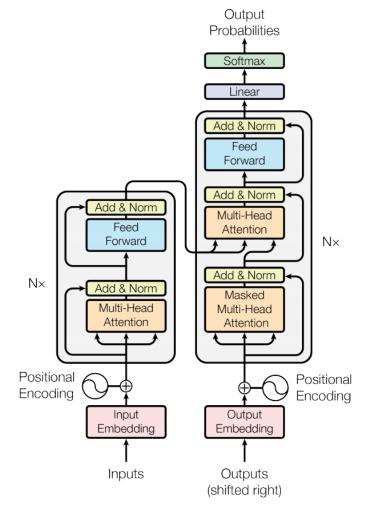
Networks for sequence generation



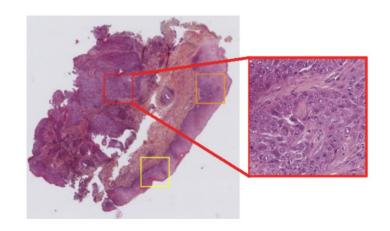


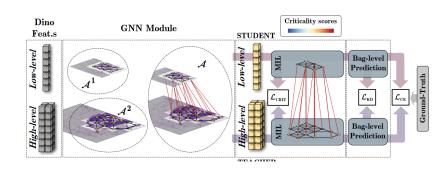
Attention: the Transformer

- 2020 Definition of new problems also in Medicine e.g. question
 answering, summarization, captioning, Focus on human centric AI: the
 role of explanation
- >2022 The role of self-attentive and self-supervising architectures (from Transformer to...), language translation and then
- Interactive, continual, Generative, cooperative learning.... Foundational ,LLM, VLM Models









Histological Whole-slide Images (DECIDER project)

- Very large images 100K x 100K. new MULTI-INSTANCE LEARNING approaches
- Correlation among spatial cues: use Graph Attention networks
- Extraction the best Features: exploit Facebook DINO ViT for selfsupervised attention



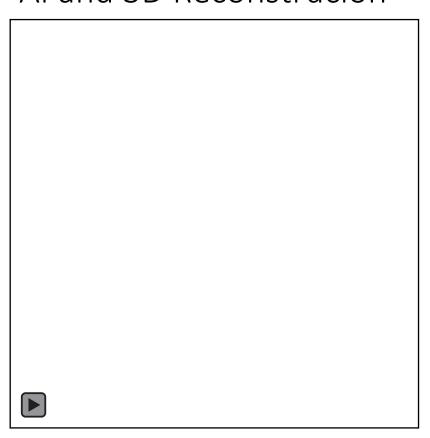
Generative for 3D reconstruction

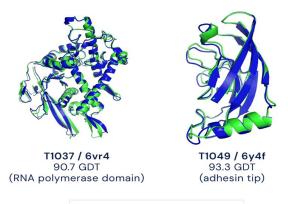


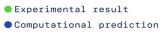


The role of Unsupervised Learning

- Al and Drug Discovery (Alphafold 1,2 3)
- Al and 3D Reconstrucion









31



Foundation Models and (M)LLMs



- Unsupervised Learning
- •BERT (2018):
- •(Jacob Devlin et al.)
- •GPT-3 (2020), Dall-e (2021) and GPT-4 (2023), GPT-5° (2024)
- •ChatGPT (2022+):
- OpenAl Team

CLIP (2021), OpenCLIP...

- OpenAl Team etc
- •AlphaFold (2021):
- •(AlphaFold 2 DeepMind Team)

LLaMA. LLaMA3 (META Team)

DINO V3 (2025) Meta



Foundation models

ΣĒ Language Text Specific Based Model Domain Image Specific Based Model Multi Audio and Modal Speech **(**)

A foundation model is a large-scale AI model trained on a massive and diverse dataset (text, images, audio, or multimodal data) that can be adapted to many downstream tasks with minimal additional training.

It's called "foundation" because it serves as a **base layer of intelligence** — once trained, it can be fine-tuned or prompted to perform specialized functions like medical diagnosis, language understanding, image analysis, or data summarization.

These models learn **general representations of knowledge and patterns**, rather than solving a single task, allowing them to reason, generalize, and transfer learning across different domains.

In healthcare: it can be

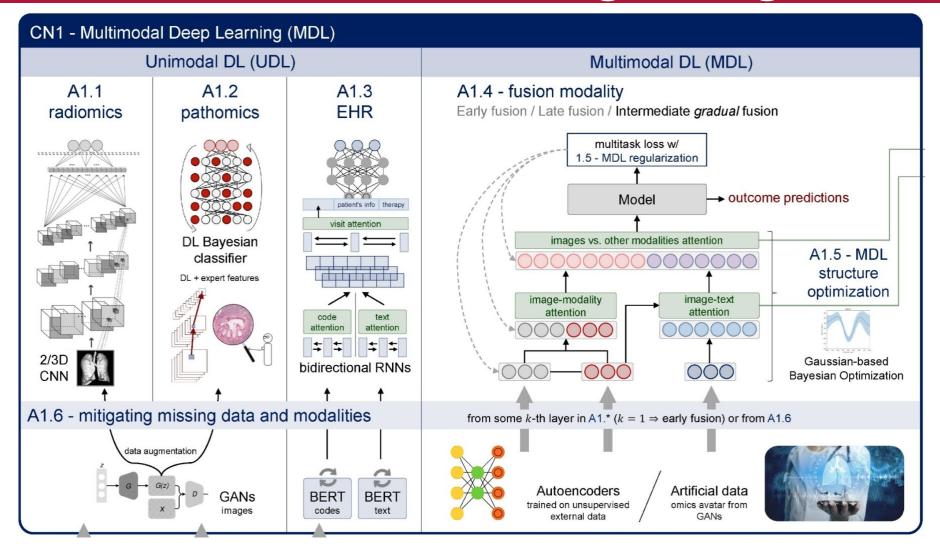
- trained on clinical notes, imaging data, and lab results,
- adapted to predict disease risk, assist in diagnosis, or optimize hospital workflows
- Il from the same core intelligence.



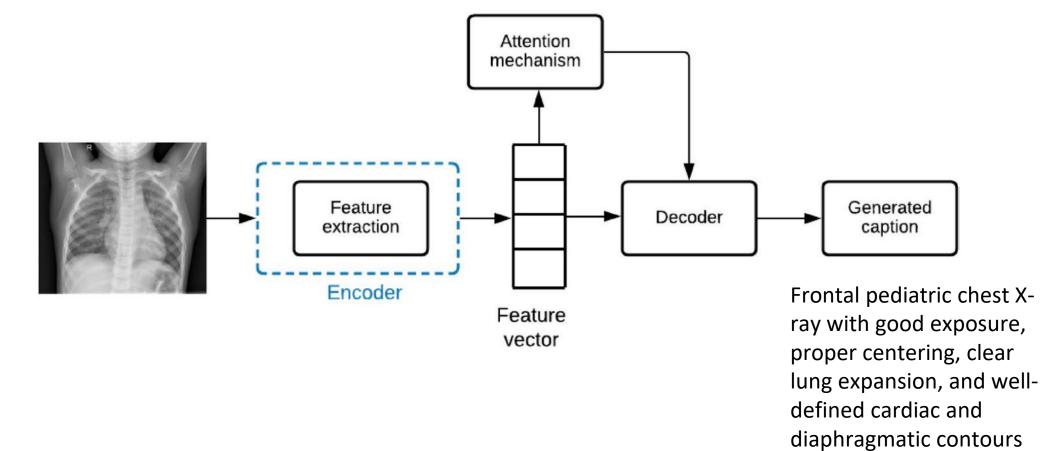
Putting all together

Integrated platforms for Multi-modalities

In understanding, Diagnosis support prediction



Captioning also in med images



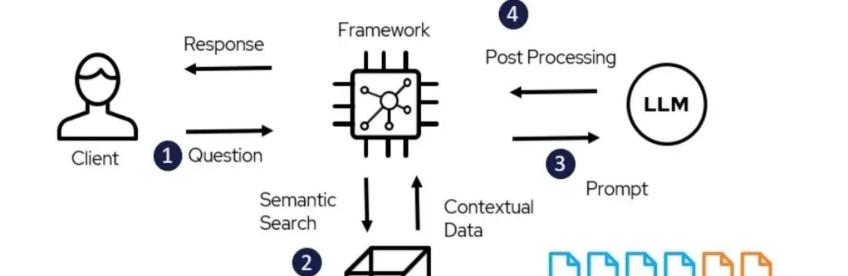
Automatic captioning for medical imaging (MIC): a rapid review of literature Baddlar et al Springer Nature 2022

Chatting with GPTs, Gemini, Claude,...

• LLM and RAG (Retrieval augmented Generation)

RAG Architecture Model

- The ways to specialize an LLM for a given scope:
- Fine-tuning
- RAG
- RAG and RAV



Vector Database

New Content

Original Content



Agentic Al



2023 — The Age of Adaptability
GPT-4 (OpenAI) and Gemini (Google
DeepMind), foundation models for
reasoning, coding, summarizing, and
problem solving — not just predicting words.
integrate tool use, memory, and planning,
showing early agentic behavior (acting
purposefully rather than only responding)

24-2025 Agentic AI Emerges

- models that can reason, plan, and take initiative based on LLMs and multimodal models as their cognitive core.
- Tools and APIs for action (like accessing databases or scheduling).
- Memory and goals for persistence over time.



Agentic Al

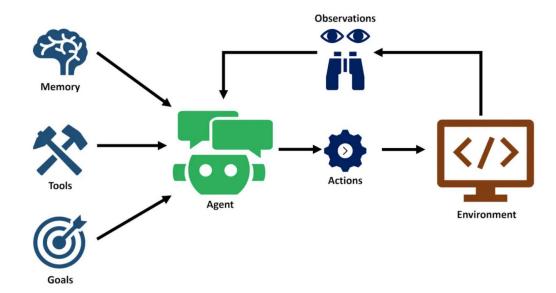
Agentic AI is « nothing more than»

High performance AI tools è memory+

Coordination systems for high level goals

Often human-in-the -loop

What is Agentic Al?





Al agents for healthcare

Table 1. Comparison of human clinicians, traditional medical AI, and medical AI agents			
Feature	Human clinician	Traditional medical Al	Medical Al agent
Core nature	biological intelligence with consciousness & empathy	a computational tool or algorithm	an autonomous or semi-autonomous system
Primary function	holistic patient care, complex reasoning, empathy	performs a specific, pre-defined task (e.g., classify, regression)	manages complex, multi-step tasks and workflows
Decision-making	experience-based, intuitive, and evidence-informed	rule-based or pattern recognition on a given dataset	data-driven, goal-oriented, and adaptive reasoning
Autonomy	full autonomy in clinical judgment	none; acts as a passive tool for human use	high; can operate independently to achieve goals
Learning & adaptation	lifelong learning from experience and education	static; requires retraining on new datasets to update	continuous learning from new data and interactions
Context awareness	high; understands patient history and context	low to none; typically context-agnostic	high; maintains context across interactions via memory
Interaction	direct interaction with patients and systems	human-invoked; typically a one-way information output	interactive; can query systems and collaborate with humans
Role in healthcare	the central decision-maker and caregiver	a supplementary tool to augment a specific task	a collaborative partner or assistant in the care process

Liu et al "A foundational architecture for Al agents in healthcare" September 2025 CELL Report in Medicine



So many agents...



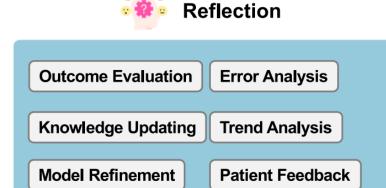


The agent activities









Feedback Loop

Comparative Analysis

Short Memory

Real-Time Data

Dynamic Updates

Immediate Access

Long Memory

Historical Data

Knowledge Base

Patient Trends



The Al Agents for patient management

1. Early Warning & Prevention

- •Collects behavioral, social, and wearable data.
- •Detects health risks and suggests lifestyle changes or medical intervention.

2. Patient Admission

- •Integrates data from EHRs, lab tests, imaging, genomics, and sensors.
- •Prepares unified patient profiles for diagnosis.

3. Diagnostic Al Agent

- •Performs multimodal data analysis and risk stratification.
- •Supports differential diagnosis and report generation.

4. Human-in-the-Loop Decision-Making

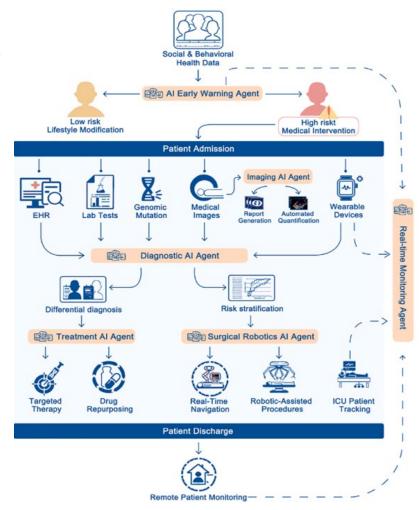
•Clinicians validate AI insights, ensuring safety and accountability.

5. Intervention & Treatment

- •Guides robotic or assisted procedures.
- •Recommends targeted therapies or drug repurposing.

6. Monitoring & Continuous Learning

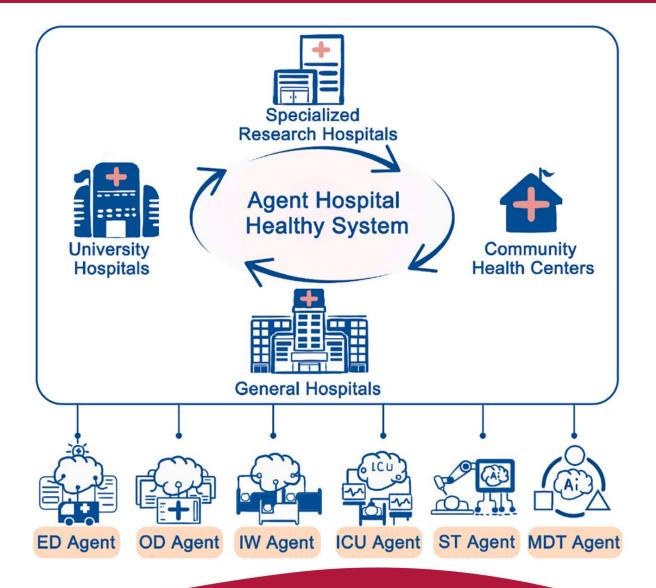
- •Tracks patients via real-time data streams.
- •Adapts protocols and improves future decision models.





The next step

ED, emergency department; OD, outpatient department; IW, inpatient ward; ICU, intensive unit care; ST, surgical theater; MDT, multidisciplinary team.





The life cycle











Defining the Problem and Scope

- Early Diagnosis
- Outcome Prediction
- Disease Staging
- Chronic Disease Management
- Imaging Data Analysis
- Customized Therapeutics
- Health Surveillance
- Robot-assisted Surgery

Select Agentic Framework

- LangChain
- Autogen
- CrewAl
- Swarm
- OpenAGI
- Flowise
- MetaGPT
- Camel
- Transformers2.0

Developing the Al Model

- Pre-Trained Models
- Fine-Tuning for Medical Context
- RAG
- Tool Developing
- Reinforcement Learning
- User Interface Design

Test and Validation

- In-Silico Simulations
- Real-World Clinical Trials
- Accountability for Outcomes
- Performance Metrics
- External Validation
- Longitudinal Validation

Deployment & Continuous Learning

- Clinical Environment Integration
- Adaptation to Evolving Data
- Feedback from Healthcare Professionals
- Real-Time Performance Optimization
- Data Security and Privacy

- •In Italy for instance:
- •Quibim, ReportAld, Brain Computing
- •U-Care Medical, SmairtHero, PatchAl
- •Almawave, GPI SpA, Dedalus, Memori.ai

From National Natural Science Foundation of China, the Macau Science and Technology Development Fund, Guangzhou National Laboratory et al 2025



But are there already experimented?

YES. In USA (supported by many companies, AWS, Google systems, Lena Health etc), in UK, in China, few examples in Germany.

Mitigating errors

All algorithms have been known to produce hallucinated outputs, resulting in errors.

To prevent or minimize these errors, engineers have developed various guardrails. And quality control officers are still needed

Barriers to integration

In many large hospital groups, the IT department does not control all systems used in daily hospital operations. Many hospitals have opted to partner with third-party providers. The cannot operate together..

Europe vs US:

Problems of workflow standardization in US (not in Europe)

"It is not feasible to simply implement an off-the-shelf Agentic AI architecture in a German hospital if it was previously trained and customized for a Czech hospital".

The system must be adaptable to different hospital infrastructures and varying healthcare regulations



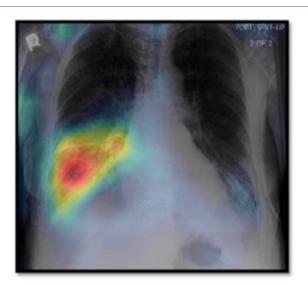
At the end...Al today

Now AI works well enough for ...

- UNDERSTANDING (data, correlations, causalities...)
- PREDICTING (in time, in space, in missing data... GenerativeAI)
- INTERPRETING and GENERATING (Visual, language... other signals)
- DISCOVERING (creative thinking, imagining)
- REASONING and ACTING (planning and moving)
- INTERACTING (by language mainly)

But together with problems of

- Bias
- Ethical issues
- Errors
- Hallucinations
- Danger
- Excessive Power

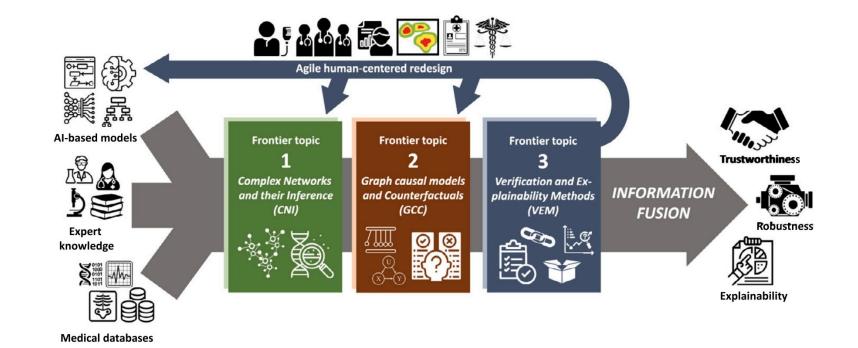




Trustworthy Medical Artificial intelligence*

A common discussion at European Level*

Al models and agents & Human experts Trusting each other



[*Andreas Holzinger, Matthias Dehmer, Frank Emmert-Streib, Rita Cucchiara, Isabelle Augenstein, Javier Del Ser, Wojciech Samek, Igor Jurisica, Natalia Díaz-Rodríguez "Information fusion as an integrative cross-cutting enabler to achieve robust, explainable, and trustworthy medical artificial intelligence" Information Fusion, Elsevier March 2022

And the next?

And After a possible integration...

- 1. Improvement of each specific agent
- 2. Continual and federated learning
- 3. Hyper personalized agent on single patient
- 4. Human digital twin

...

Now

we need tools, people and money; it is not a problem of data anymore

... good luck ...

"Progress is strongest when human wisdom and technological intelligence move forward together." GPT5 oct.2025



























thanks

direttore.AIRI@unimore.it
Aimagelab@unimore.it
[Costantino.grana, Elisa.Ficarra, Federico.bolelli @unimore.it]





