CIDAI Centre of Innovation for Data tech and Artificial Intelligence

Masterclass

IA Generativa en Salut Digital: Aplicacions, Reptes i Implicacions Ètiques

Impartida per: eurecat Centre Tecnològic de Catalunya

IA Generativa en Salut Digital: Aplicacions, Reptes i Implicacions Ètiques

David Suñol, Digital Health

david.sunol@eurecat.org

"innovant amb les empreses"

Index

Introduction

Digital Health applications

Challenges

Ethics & regulations

Future

Introduction

Introduction

What is Generative AI?

"is a subfield of artificial intelligence that uses generative models to produce text, images, videos, audio, software code or other forms of data. These models **learn the underlying patterns and structures** of their training data and use them to produce new data based on the input, which often comes in the form of natural language prompts"

Wikipedia

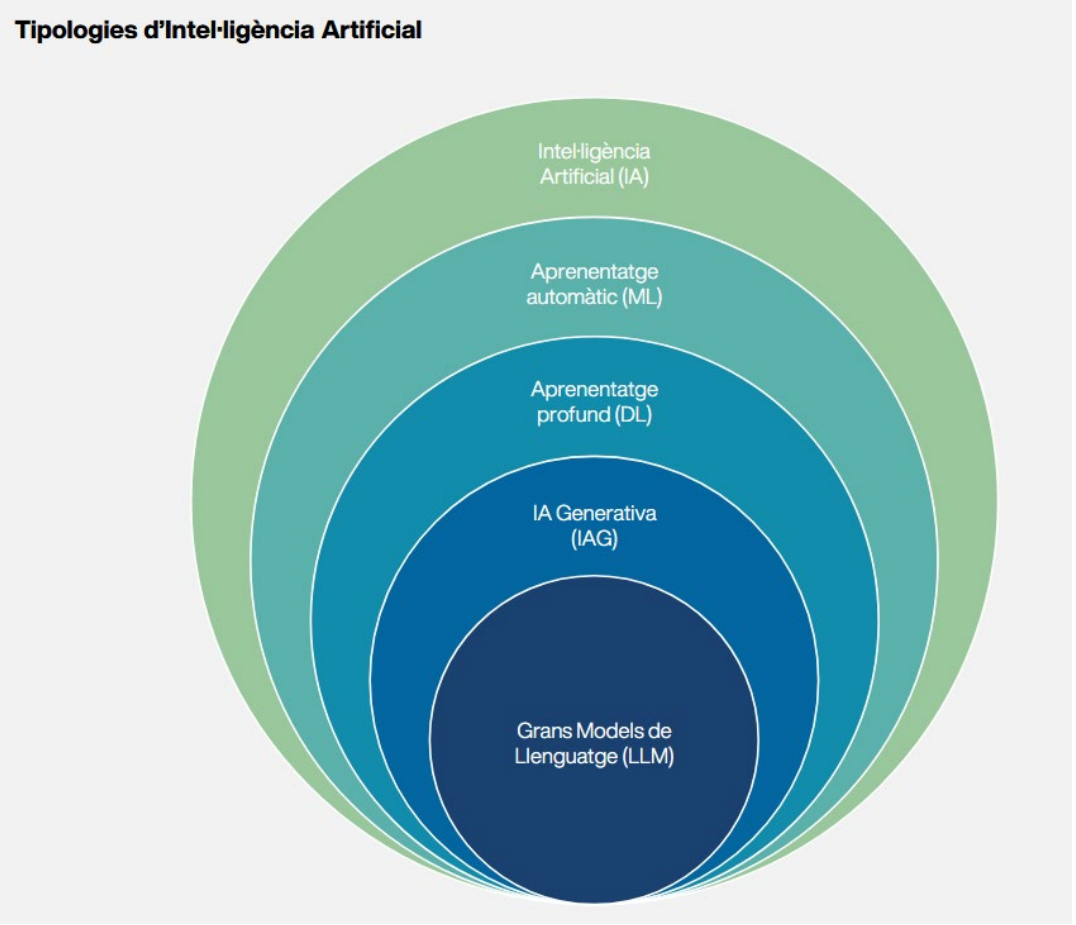

https://iasalut.cat/wp-content/uploads/2025/03/Bones_Practiques_IAGen_LLM_Guia-CAT.pdf

Introduction

What is Generative AI?

Autoregressive generation

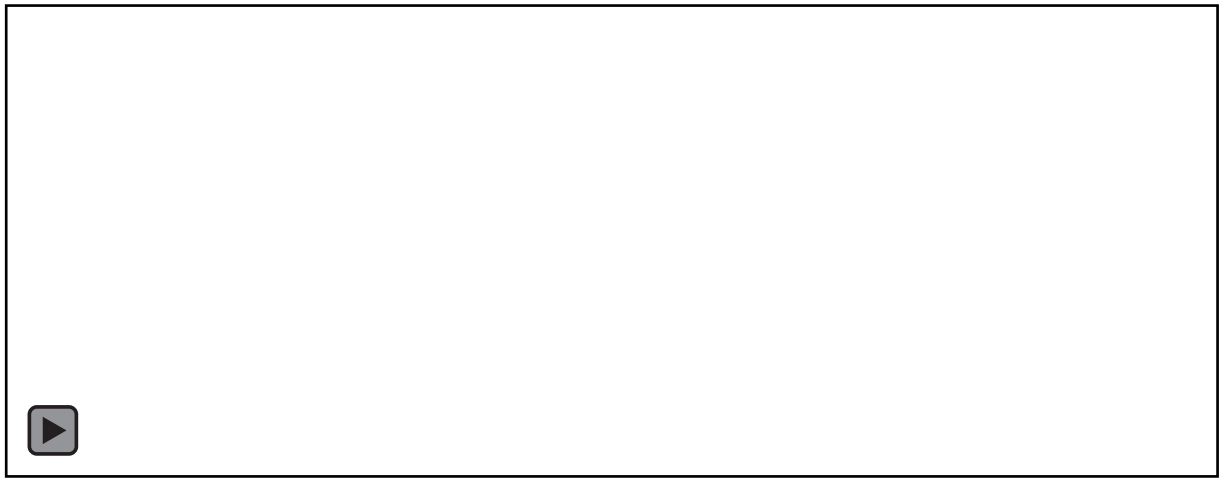

Figure 2: Forward pass of an LLM. Source: [hugginface]

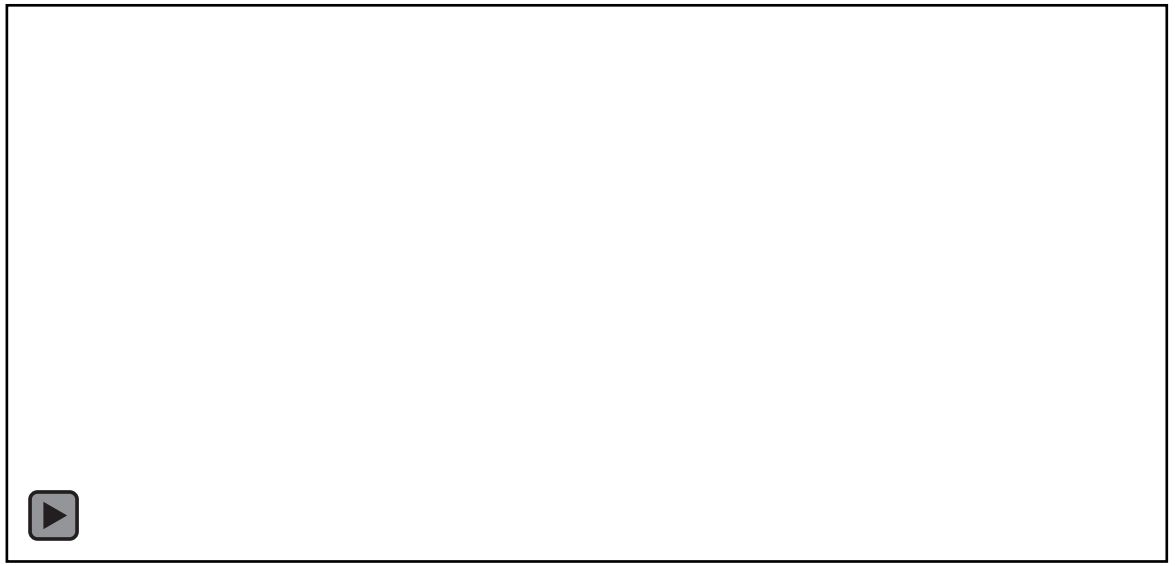

Figure 3: Autoregressive generation iteratively selects the next token from a probability distribution to generate text. Source: [hugginface]

https://huggingface.co/docs/transformers/llm_tutorial

Introduction

What is Generative AI?

Diffusion generation

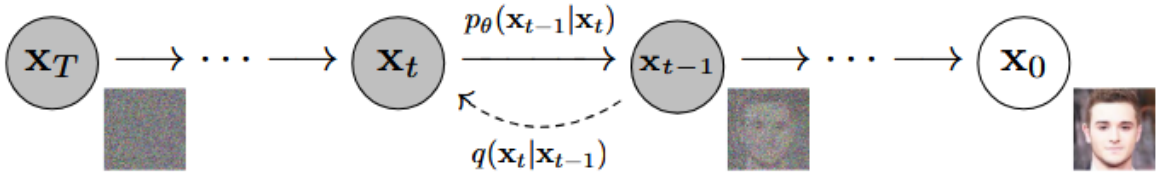

Figure 2: The directed graphical model considered in this work.



Figure 6: Unconditional CIFAR10 progressive generation ($\hat{\mathbf{x}}_0$ over time, from left to right). Extended samples and sample quality metrics over time in the appendix (Figs. 10 and 14).

Jonathan Ho et al; Denoising Diffusion Probabilistic Models. 2020. https://doi.org/10.48550/arXiv.2006.11239

Introduction

What is *not* 100% Generative AI?

Drug Discovery, Biotechnology & Biophysics

How proteins fold?

MVLAQSRVSAGVGSPHCSGSGGGGSDSFPWPASHPGNPQCSFSTAFLAS
PRLSRGTLAYLPPAPWSSLATPSALLGSSCAPPPPPARCPQPRALSPELGTKA
GPRRPHRWELPRSPSQGAQGPAPRRRLLETMKGIVAASGSETEDEDSMDI
PLDLSSSAGSGKRRRRGNLPKESVQILRDWLYEHRYNAYPSEQEKALLSQQT
HLSTLQVCNWFINARRRLLPDMLRKDGKDPNQFTISRRGAKISETSSVESV
MGIKNFMPALEETPFHSCTAGPNPTLGRPLSPKPSSPGSVLARPSVICHTTV
TALKDVPFSLCQSVGVGQNTDIQQIAAKNFTDTSLMYPEDTCKSGPSTNTQ
SGLFNTPPPTPPDLNQDFSGFQLLVDVALKRAAEMELQAKLTA

Introduction

What is *not* 100% Generative AI?

Drug Discovery, Biotechnology & Biophysics

How proteins fold?

MVLAQSRVSAGVGSPHCSGSGGGGSDSFPWPASHPGNPQCSFSTAFLAS
PRLSRGTLAYLPPAPWSSLATPSALLGSSCAPPPPPARCPQPRALSPELGTKA
GPRRPHRWELPRSPSQGAQGPAPRRRLLETMKGIVAASGSETEDEDSMDI
PLDLSSSAGSGKRRRRGNLPKESVQILRDWLYEHRYNAYPSEQEKALLSQQT
HLSTLQVCNWFINARRRLLPDMLRKDGKDPNQFTISRRGAKISETSSVESV
MGIKNFMPALEETPFHSCTAGPNPTLGRPLSPKPSSPGSVLARPSVICHTTV
TALKDVPFSLCQSVGVGQNTDIQQIAAKNFTDTSLMYPEDTCKSGPSTNTQ
SGLFNTPPPTPPDLNQDFSGFQLLVDVALKRAAEMELQAKLTA

AlphaFold 2

Nov 2020

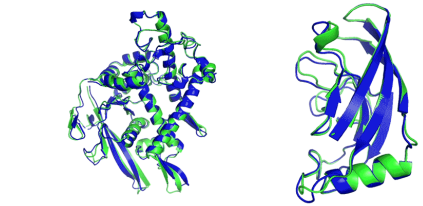

https://oecd.ai/en/wonk/deepminds-alphafold-a-solution-to-challenge-of-protein-folding

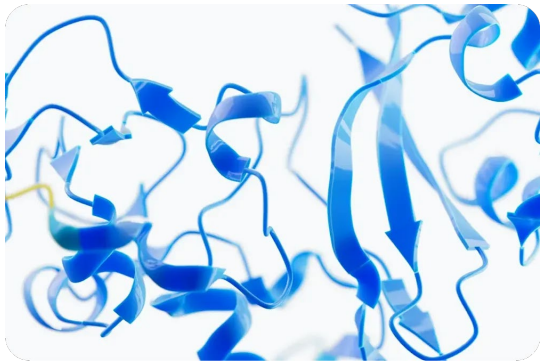

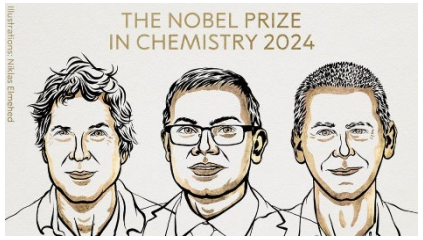

Nobel Prize Chemistry 2024 to David Baker, Demis Hassabis and John M. Jumper



Introduction

Generative AI around the world

Chat GTP was launched on November 30th, 2022.

















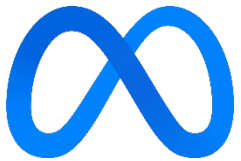











Text generation







Search











Coding









Image generation







Video generation



Digital Health Applications

Clinical applications

Text

Digestion of everyday bigger literature

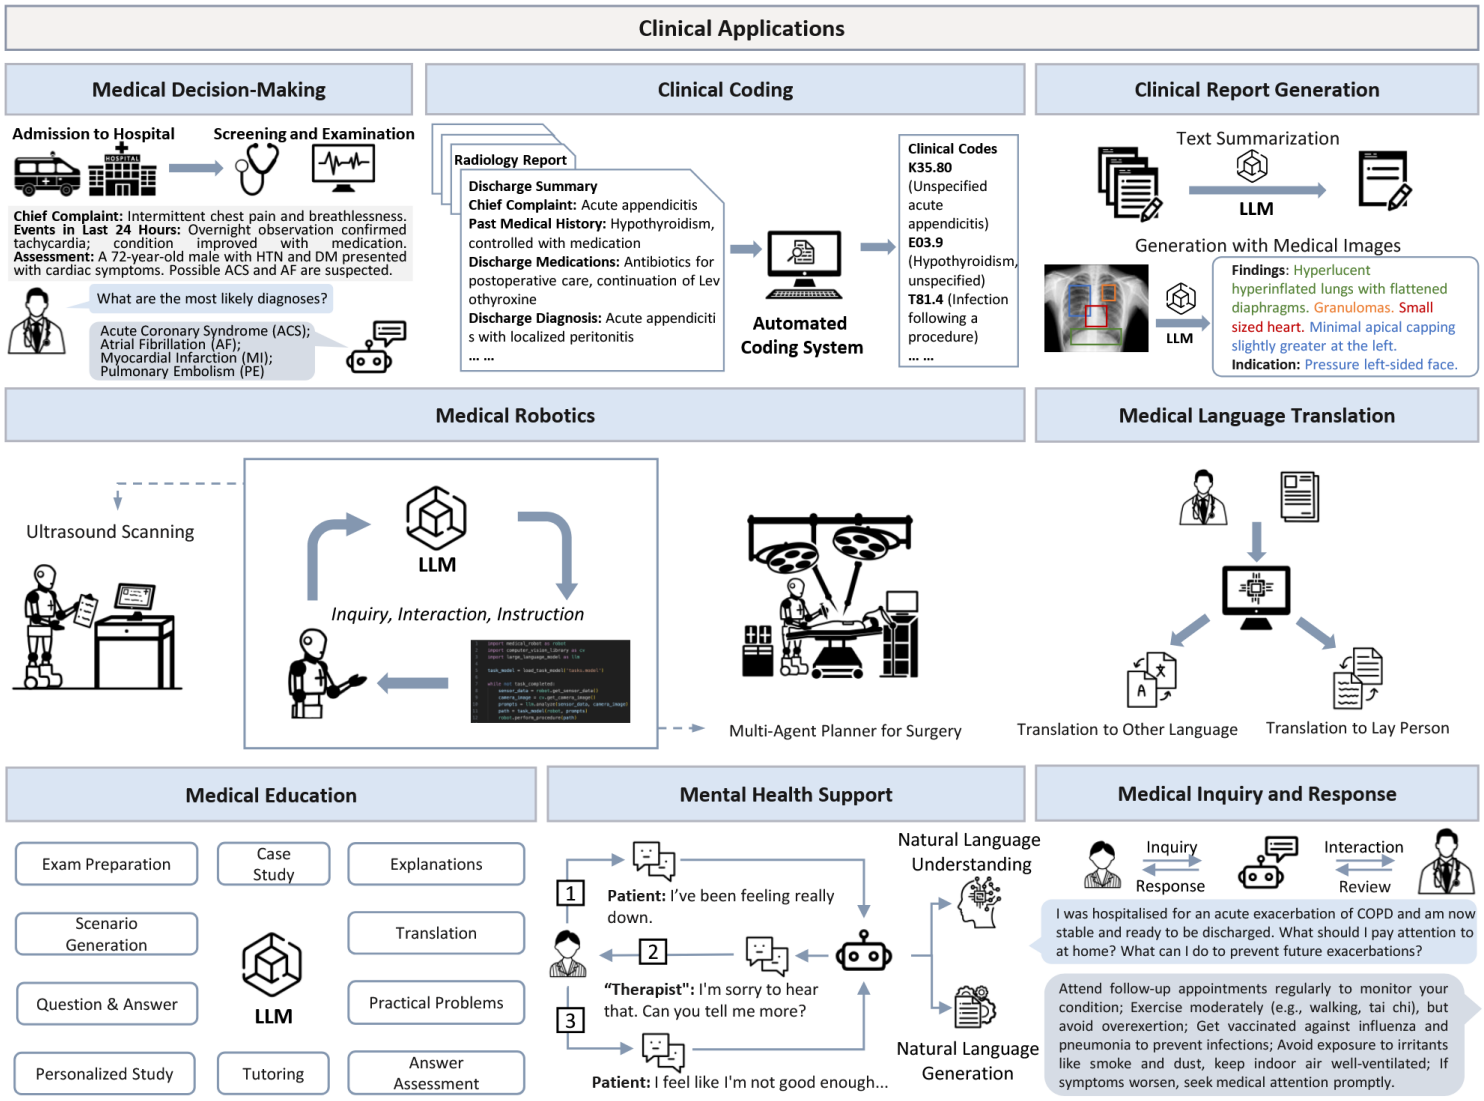

Zhou, H., Gu, B., Zou, X., Li, Y., Chen, S. S., Zhou, P., ... & Liu, F. (2023). A survey of large language models in medicine: Progress, application, and challenge. arXiv preprint arXiv:2311.05112.

Clinical applications

Text

OpenEvidence®

LLM that synthesize the medical knowledge to help clinicians make evidence-based decisions

The NEW ENGLAND JOURNAL of MEDICINE

Featuring multimedia and clinical findings from The New England Journal of Medicine

JAMA Network

Featuring multimedia and clinical findings from JAMA and the JAMA Network specialty journals

NCCN National Comprehensive Cancer Network®

Featuring NCCN Guidelines®, multimedia, and Clinical Pathways from the National Comprehensive Cancer Network

MAYO CLINIC PLATFORM

Launched from the Mayo Clinic Platform Accelerate program

The platform is free for verified healthcare professionals

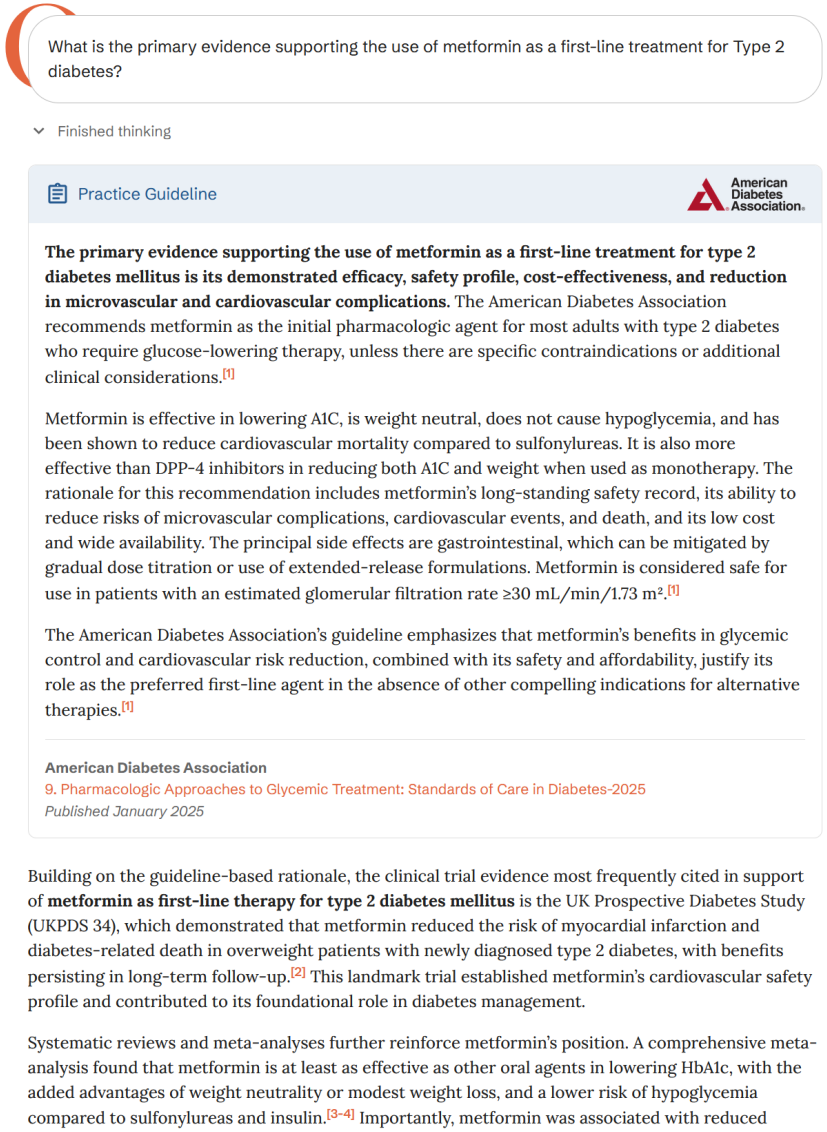

Clinical applications dev in DH

Judith Recober

PAI - PRIMEROS AUXILIOS PSICOLÓGICOS

OBJECTIVE: Develop a proof of concept for a conversational assistant providing psychological support (psychological first aid) in traumatic situations and natural disasters, and validate it with teams of psychosocial support professionals in a controlled environment.

Main technical characteristics

- Based on small-medium local LLM
- Enriched with expert knowledge powered by RAG
- Conversation regulated by phases
- Safety & guardrails included
- Reviewed by a panel of experts

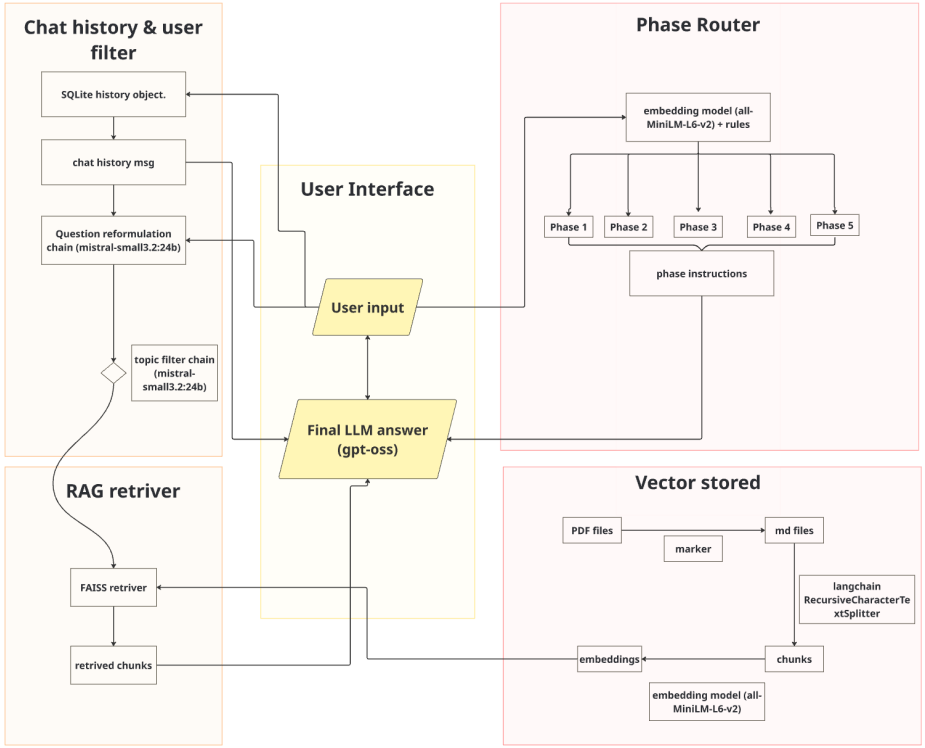

Clinical applications dev in DH

Judith Recober

PAI - PRIMEROS AUXILIOS PSICOLÓGICOS

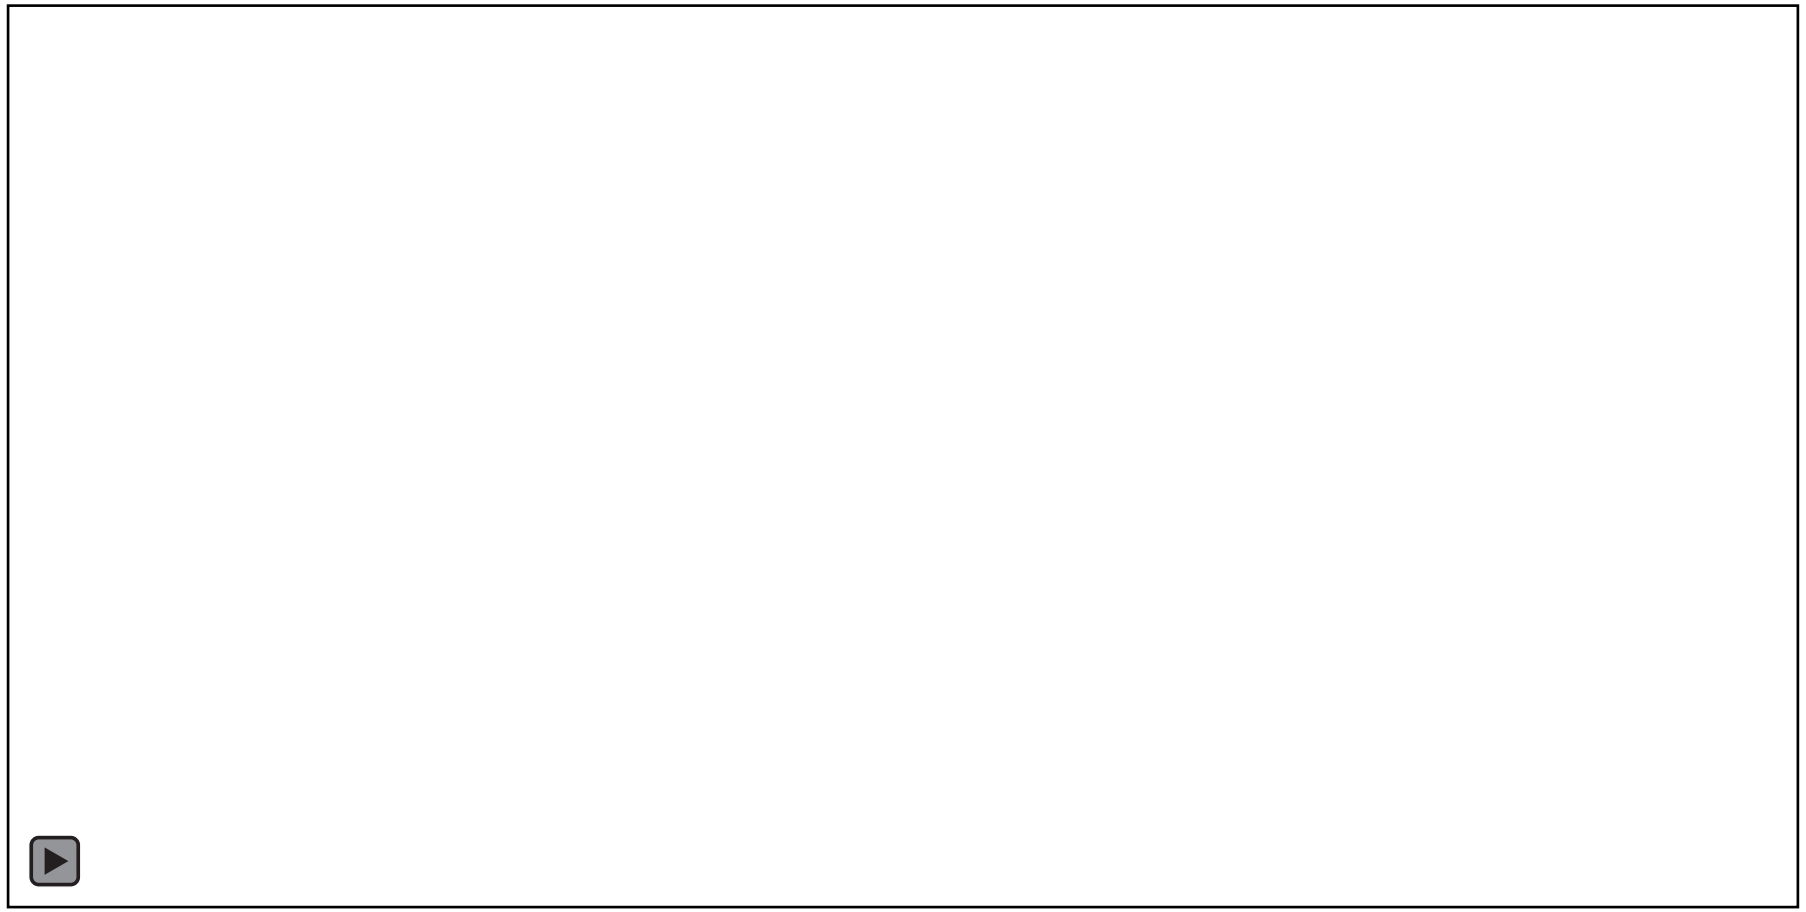

Clinical applications dev in DH

Adrià Huguet

CONECTED - Nutrition AI Assistant

OBJECTIVE: Monitor the user nutrition in a friendly but nutritionist manner

Main needs

- User shall provide food info after every meal (best memory)
- The food must match to a listed food (CIQUAL tables)
- Quantity & how it was cooked needs to be considered
- The output is a list of nutrients derived from CIQUAL tables
- Local LLM

Final output – Nutrient cal.



Input	Structured Output	Database search & match	User validation (opt.)
"Today I had a glass of lemonade and two bowls of rice, each weighing 80 grams"	**name**=lemonade **quantity**=1.0 **measure_word**=glass **weight_g**=None **name**=rice **quantity**=2.0 **measure_word**=bowls **weight_g**=80	**best_match**=lemonade, with sugar **quantity**=1.0 **measure_word**=glass **weight_g**=200 **best_match**=rice, raw **quantity**=2.0 **measure_word**=None **weight_g**=80	**best_match**=lemonade juice, homemade **quantity**=1.0 **measure_word**=glass **weight_g**=150 **best_match**=cantonese rice **quantity**=2.0 **measure_word**=None **weight_g**=80

Clinical applications dev in DH

Adrià Huguet

CONECTED - Nutrition AI Assistant

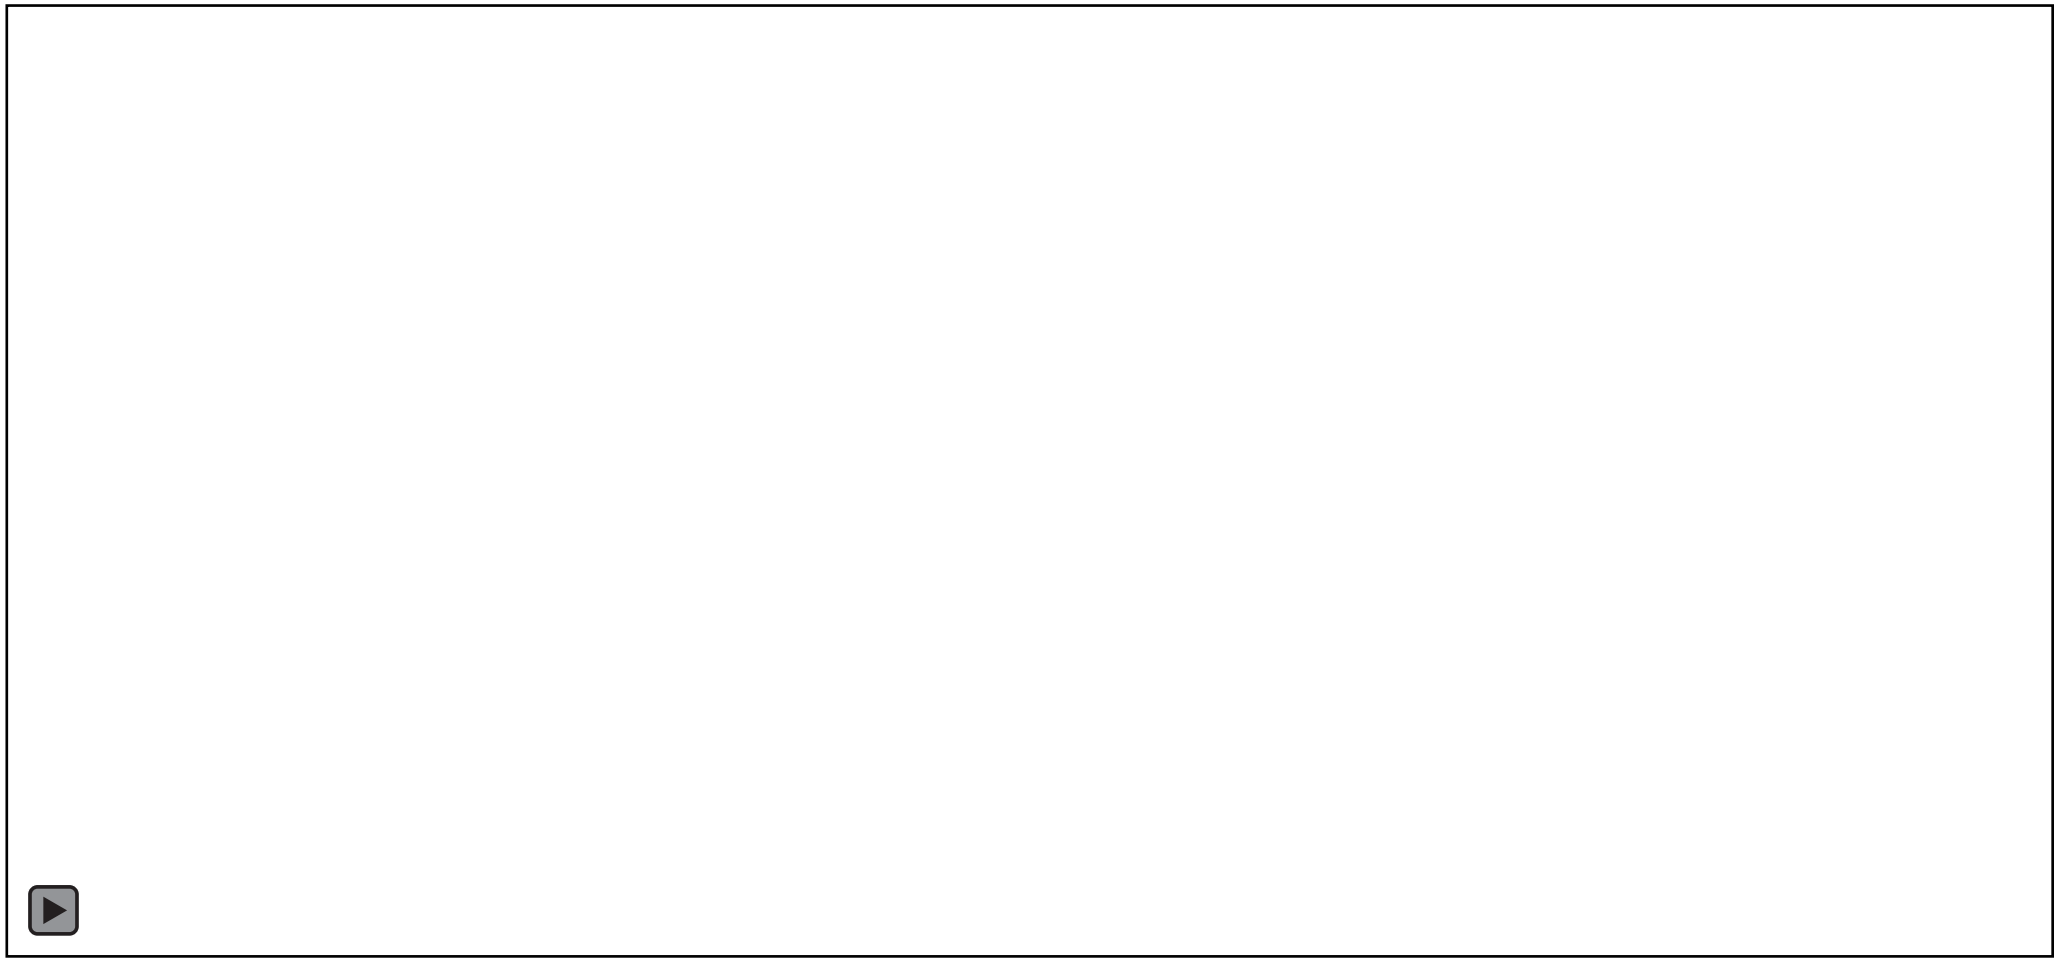

Clinical applications dev in DH

Oriol Guardià

iBeChange - Behavioural Change platform to promote to prevent primary cancer

OBJECTIVE: Design, develop, and test the iBeChange platform, a user-focused management system that empowers people to achieve healthy and sustainable behaviours and emotions thanks to evidence-based knowledge coming from clinical psychology and behavioural change theories.

Current status:

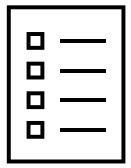

List of standardized & static clinical validated recommendations

In development:

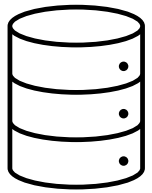

Database with journals & guidelines clinically validated

Multi-modal RAG

Users' data

LLM

Project Structure

List of personalized & dynamic recommendations

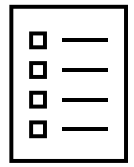

BECOME: BEhavioural Change recOMmender system

Clinical applications dev in DH

Anna Oliveras

Images

PHASE IV AI – LAND (Lung-And-Nodule-Diffusion)

OBJECTIVE: Generate high-quality 3D chest CT scans conditioned on 3D anatomical masks

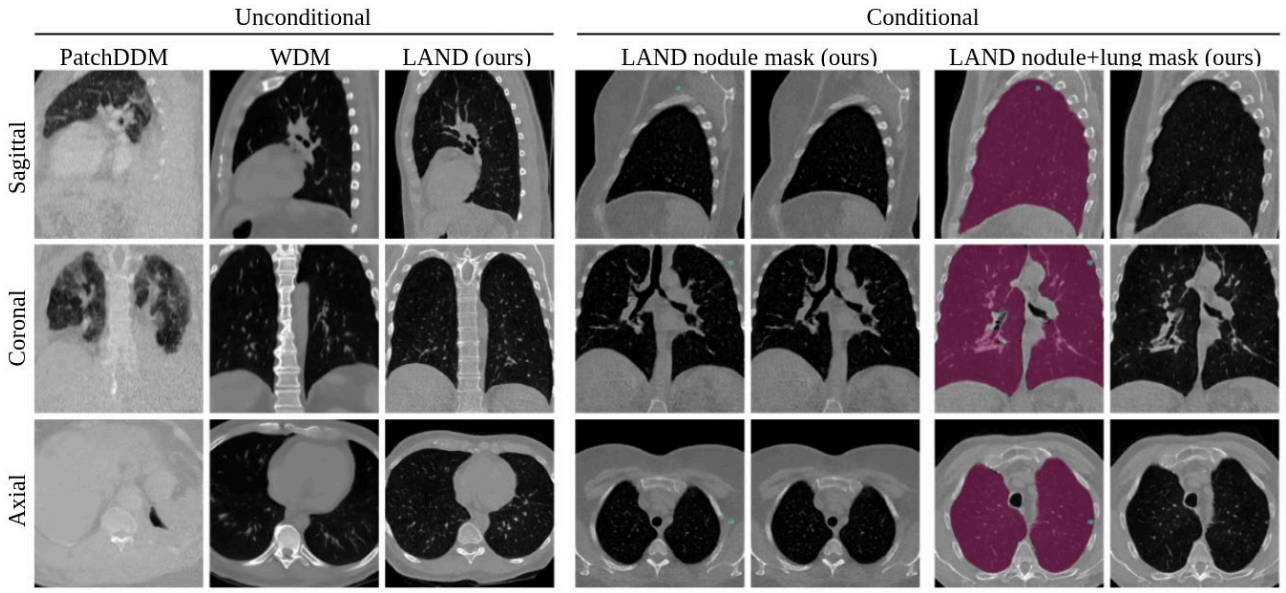

Comparison of unconditional (left) and conditional (right) CT generation using LAND and baseline methods PatchDDM and WDM. Mask overlays are shown where applicable.

LAND uses a 3D VAE to encode CT volumes into a latent space, where a 3D U-Net performs diffusion on the latent samples x, optionally conditioned on anatomical masks m (lungs/nodules).

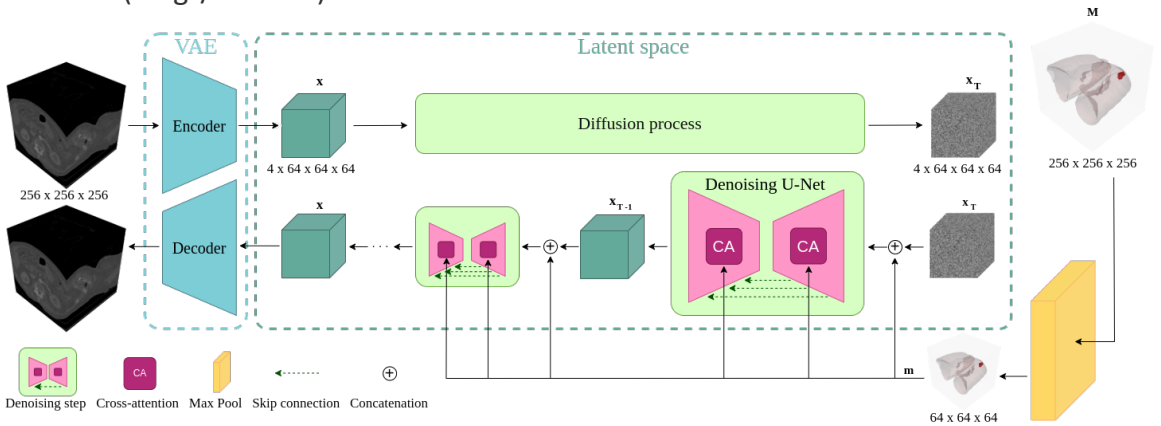

Anna O. et al, LAND: Lung and Nodule Diffusion for 3D Chest CT Synthesis with Anatomical Guidance; 2025; https://doi.org/10.48550/arXiv.2510.18446

Clinical applications

Multimodal

Med-Gemini

Google Research

AIME (Articulate Medical Intelligence Explorer)

General medical AI

Diagnostic Assistant

Table A.15 | Overall Performance Summary of Med-Gemini This table represent the aggregated results comparing Med-Gemini to the previous state-of-the-art (SoTA), Gemini or strong baseline where available.

Capabilities	Datasets	Metric	Med-Gemini	Gemini	Baseline or SoTA	Reference
Report Generation	MIMIC-CXR	RadGraph	24.4	N/A	20.5	Tanno et al. (2024)
	MIMIC-CXR	Expert (AI superior)	47.6	N/A	43.0	Tanno et al. (2024)
	IND1	Expert (AI superior)	75.4	N/A	63.7	Tanno et al. (2024)
VQA	MIMIC-CXR VQA	Accuracy	78.6	70.9	68.1	Xu et al. (2023)
	Slake-VQA	Accuracy	84.8	70.4	91.1	Li et al. (2023b)
	VQA-Rad CXR	Expert	71.9	N/A	57.9	Xu et al. (2023)
	VQA-Rad CXR	Accuracy	78.8	62.4	N/A	N/A
	PathVQA	Accuracy	83.3	62.8	90.9	Sun et al. (2024)
Genomic Disease Prediction	Coronary artery disease	AUC	82.5	N/A	78.5	Ensemble of PRSs and Demographics
	Stroke	AUC	78.1	N/A	74.8	
	Type 2 diabetes	AUC	86.5	N/A	85.5	
	Glaucoma	AUC	78.9	N/A	82.8	
	COPD	AUC	83.9	N/A	77.4	
	Rheumatoid arthritis	AUC	79.2	N/A	70.6	
	Major depression	AUC	70.2	N/A	60.2	
	Allcause mortality	AUC	79.9	N/A	71.8	
	Hypertension	AUC	72.0	N/A	69.9	
	Hypercholesterolemia	AUC	52.4	N/A	52.1	
	Atrial fibrillation	AUC	78.7	N/A	74.4	
	Diabetic retinopathy	AUC	87.0	N/A	87.5	
	Pneumonia	AUC	71.6	N/A	68.9	
	Asthma	AUC	57.7	N/A	56.5	
Classification	Hard Exudates	F1	87.3	61.5	N/A	N/A
	Hemorrhage	F1	85.3	57.8	N/A	N/A
	PRP Scars	F1	82.3	57.0	N/A	N/A
	DR Lesion	F1	86.4	63.9	92.0	Krause et al. (2018)
	PAD-UFES-20	F1	71.4	60.3	70.0	Google (2024)
	MIMIC-CXR	F1	90.7	86.8	N/A	N/A
	CheXpert	F1	48.3	42.6	60.6	Tiu et al. (2022)
	ChestX-ray14	F1	46.7	34.2	58.3	Majkowska et al. (2020)
	Pathology Patch	AUC	89.2	88.4	91.7	Sun et al. (2024)

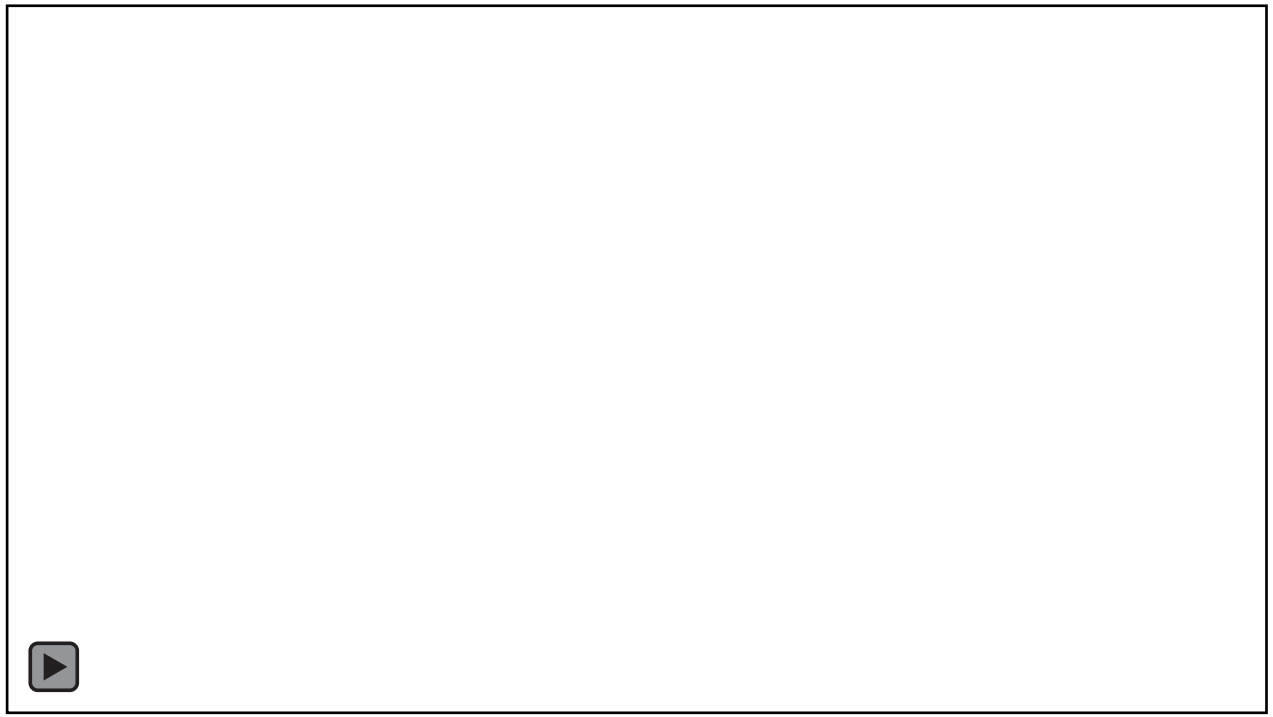

Tu, T., Schaekermann, M., Palepu, A. *et al.* Towards conversational diagnostic artificial intelligence. *Nature* **642**, 442–450 (2025). https://doi.org/10.1038/s41586-025-08866-7

Lin Yang et al. Advancing Multimodal Medical Capabilities of Gemini; 2024; https://doi.org/10.48550/arXiv.2405.03162

Challenges

Challenges: Scarce data

Clinical context implies privacy

In general, the more quality data the better to train the models

Main data sources

Public datasets

MIMIC (Electronic Health Records)

The Cancer Imaging Archive

PubMed

Clinical guidelines

Etc.

Hospital Partnership



\+ Federated learning

Data augmentation (ML)

Synthetic data (Generative AI)

Challenges: Harmonization

It may seem easy, but it is not.



A

Infarct
TIA
Lokalisatie
linker arteria carotis interna
rechter arteria carotis interna
vertebrobasilair
NIHSS score
Datum start klachten
Tijdstip start klachten
Start klachten niet bekend, reden
wake-up stroke
Datum aankomst SEH
Tijdstip aankomst SEH
Verwezen vanuit ander ziekenhuis



B

ictus_ischemico
ictus_emorragico
ictus_tia
SESSO
SESSO
ETA
ipertensione_arteriosa
ipertensione_arteriosa
ipertensione_arteriosa
diabete
diabete
diabete



C

referral_Emergency_service_(144)
referral_General_Practitioner
referral_Self_referral
event_type_Ischemic_stroke
event_type_Intracerebral_hemorrhage
event_type_Transient_ischemic_attack
firstangio_result_Occlusion_in_suspected_ischem
aca
aca_left
aca_right
pca
pca_left
pca_right
mca_left
mca_right
ica_left
ica_right
vertebrobasilar_left
vertebrobasilar_right
vertebrobasilar

Challenges: Harmonization

It may seem easy, but it is not.

Umbrella

eurecat 10 anys

A
Infarct
TIA
Lokalisatie
linker arteria carotis interna
rechter arteria carotis interna
vertebrobasilair
NIHSS score
Datum start klachten
Tijdstip start klachten
Start klachten niet bekend, reden
wake-up stroke
Datum aankomst SEH
Tijdstip aankomst SEH
Verwezen vanuit ander ziekenhuis

B
ictus_ischemico
ictus_emorragico
ictus_tia
SESSO
SESSO
ETA
ipertensione_arteriosa
ipertensione_arteriosa
ipertensione_arteriosa
diabete
diabete
diabete

C
referral_Emergency_service_(144)
referral_General_Practitioner
referral_Self_referral
event_type_Ischemic_stroke
event_type_Intracerebral_hemorrhage
event_type_Transient_ischemic_attack
firstangio_result_Occlusion_in_suspected_ischem
aca
aca_left
aca_right
pca
pca_left
pca_right
mca_left
mca_right
ica_left
ica_right
vertebrobasilar_left
vertebrobasilar_right
vertebrobasilar



OHDSI

OBSERVATIONAL HEALTH DATA SCIENCES AND INFORMATICS

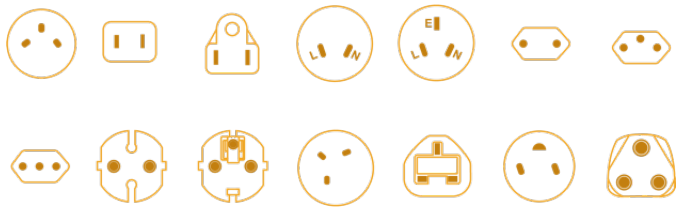

www.ohdsi.org

OMOP

Common vocabulary

Acute cerebrovascular accident due to ischemia

ISCHEMIC STROKE

Ischaemic stroke

OMOP

Common data structure

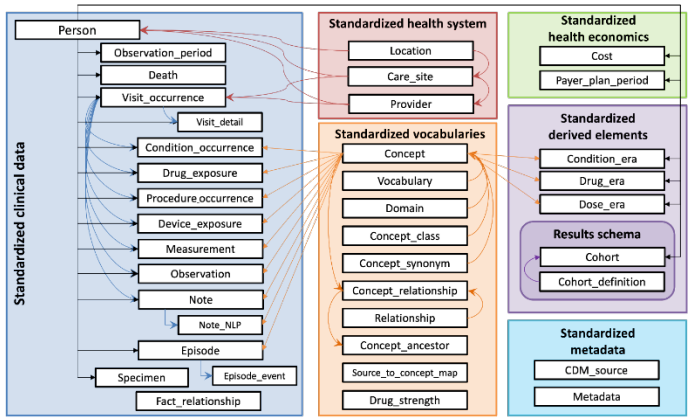

Ischemic stroke

DETAILS	
Domain ID	Condition
Concept Class ID	Disorder
Vocabulary ID	SNOMED
Concept ID	4310996
Concept code	422504002
Validity	Valid
Concept	Standard

LANGUAGE	SYNONYM CONCEPT
English	Ischaemic stroke
English	Ischemic stroke (disorder)

Valid start	31-Jan-2007
Valid end	31-Dec-2099

Challenges: Translations

What about other languages than English?

Most foundational models (GTP, Gemini, Claude, Mistral etc.) are trained on a corpus >50% in English.



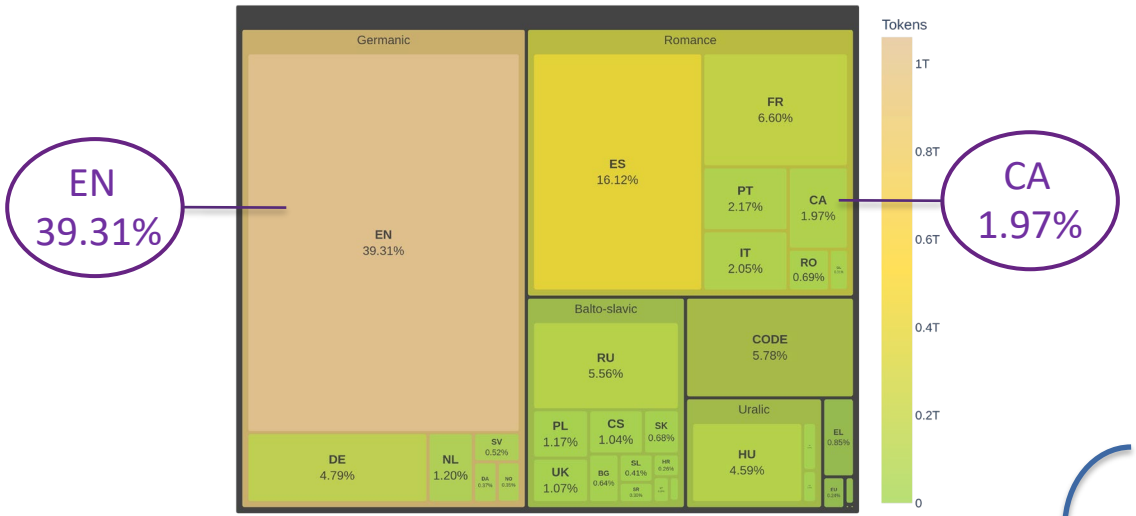

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Patient assistant (mistral3.2-mini)

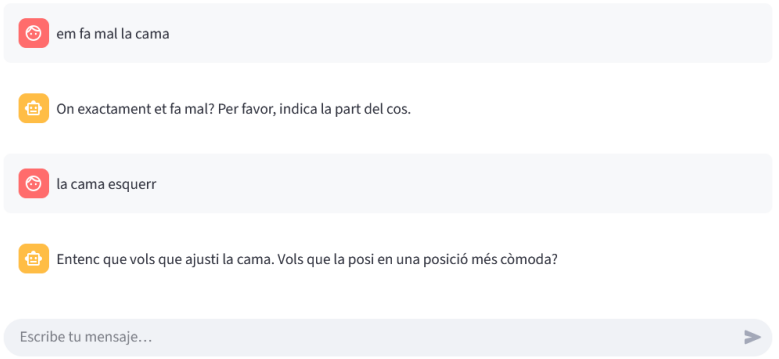

One easy solution

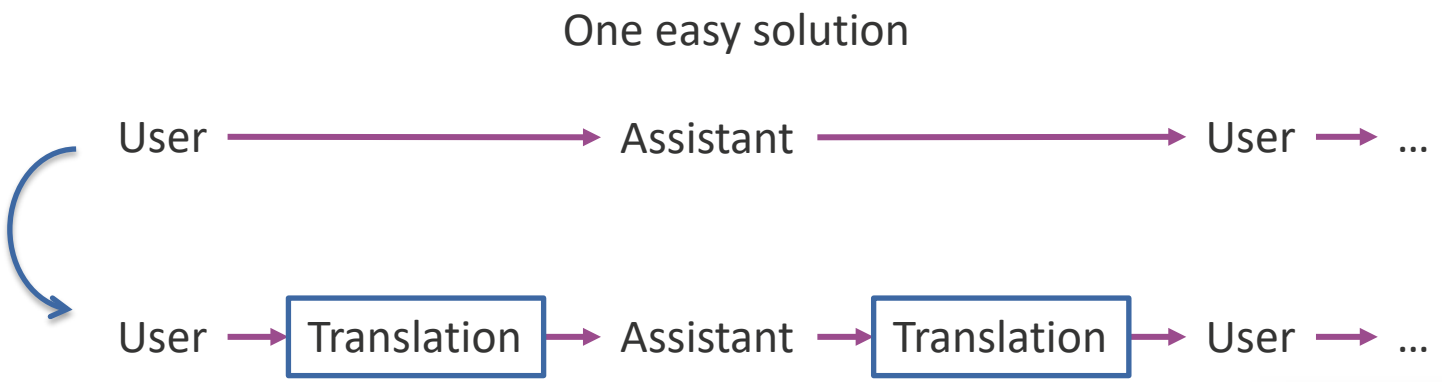

Challenges: Validation

What do you want to evaluate in a LLM?

Correctness

Hallucination

Transparency

Tone

Completition

Relevance

Efficiency

Bias

Robustness

Toxicity

Challenges: Validation

What do you need to define to evaluate an LLM?

Metrics



Accuracy
F1-score
BLEU score
…

Database

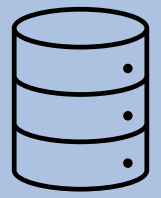

Internal or external

Evaluator

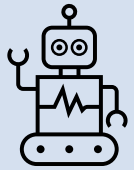

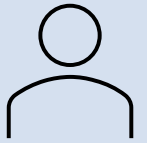

Challenges: Validation

3 main ways to validate your algorithm

Only internal dataset, no human

- Unit Test, incl. Regex.
- ML metrics: Accuracy, F1-score, Precision, etc.
- Suitable for RAG systems
- Example datasets can be created with LLMs

Problem

- Limited evaluation
- You may need to create your dataset

External & Reference Dataset

- ML & Text-Comparison metrics. Accuracy, F1-score, Precision, BLEU, ROUGE, BERT
- Each dataset can be useful evaluating one aspect. E.g. medical facts, hate speech, etc.

Problem

- Need a relevant database for each topic

Human/Expert-in-the-loop

- Evaluating a result
- Comparing two different results
- Rubric system
- ML & other metrics: Accuracy, F1-score, Precision, Elo score, rubric score etc.

Problem

- Need an expert human

Challenges: Validation

Only internal dataset, no human

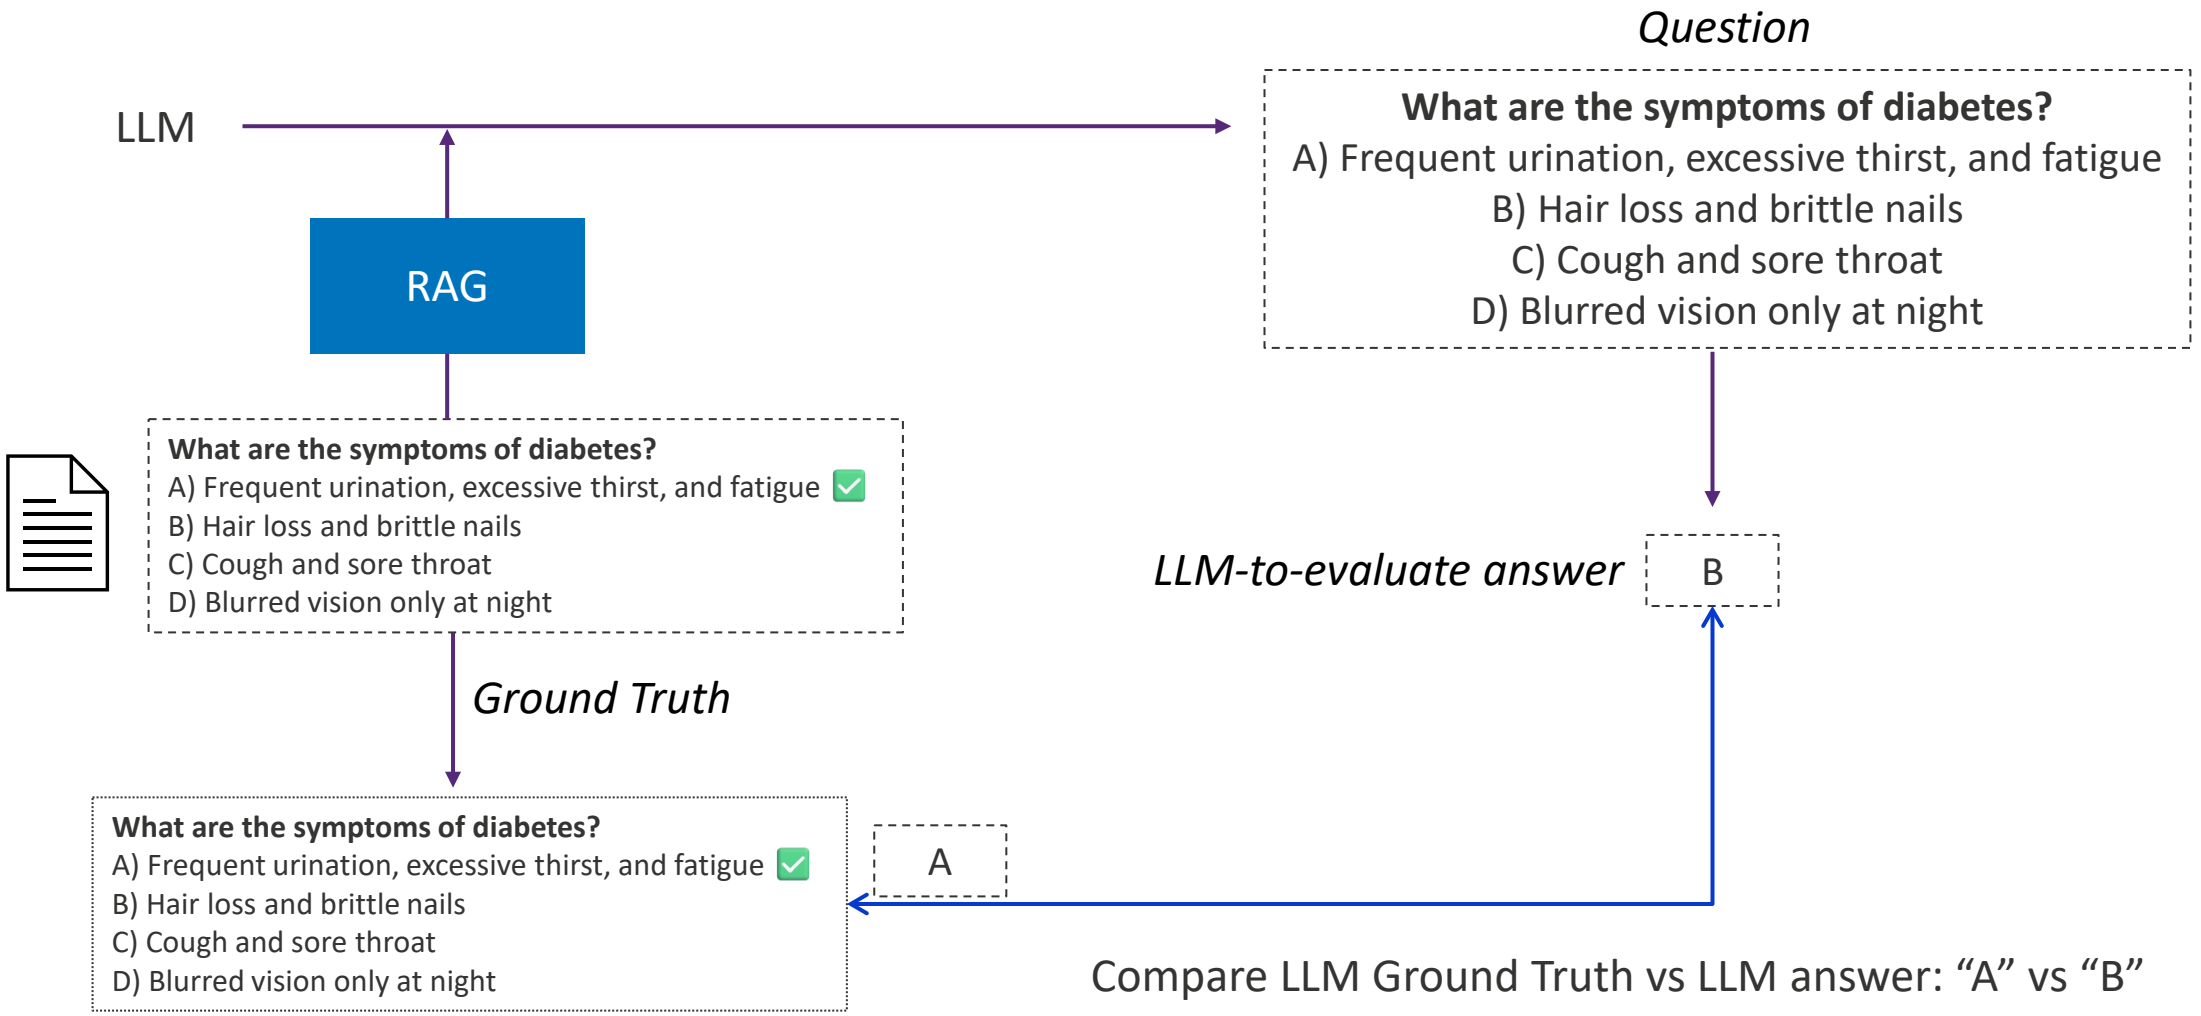

Compare LLM Ground Truth vs LLM answer: “A” vs “B”

Metrics: Accuracy, F1-score, Precision, BLEU, ROUGE, BERT, etc.

Challenges: Validation

Reference External Dataset, no human

Evaluate with a reference dataset

- MedQA (Medical facts)
- PubMedQA (Medical facts)
- MMLU Health/Medicine subset (Medical facts)
- MIMIC III & IV (EHR)
- MedDialog (Medical Dialogues)
- MedCalc-Bench (Medical Calculations)
- HateXplain (Hate detection)
- ...

Ground Truth

What are the symptoms of diabetes?
A) Frequent urination, excessive thirst, and fatigue ✅
B) Hair loss and brittle nails
C) Cough and sore throat
D) Blurred vision only at night

A

Question

What are the symptoms of diabetes?
A) Frequent urination, excessive thirst, and fatigue
B) Hair loss and brittle nails
C) Cough and sore throat
D) Blurred vision only at night

LLM-to-evaluate answer B

Compare LLM Ground Truth vs LLM answer: “A” vs “B”

Metrics: Accuracy, F1-score, Precision, BLEU, ROUGE, BERT, etc.

Challenges: Validation

Human/Expert-in-the-loop

However, how can we evaluate a very specific and human-nature answers?

Solution: use a human judgement.

Question

What are the symptoms of diabetes?

LLM-to-evaluate answer

"Diabetes symptoms include excessive thirst, frequent urination, fatigue, blurred vision, unexplained weight loss, increased hunger, slow-healing wounds, and infections."

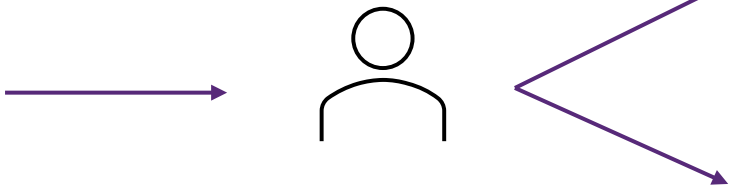

Expert answer

"This answer:

- is correct?
- is toxic?
- contains any hallucination?
- is biased?
- …"

Yes

No

Accuracy

F1-score

…

Challenges: Validation

Human/Expert-in-the-loop evolves into LLM-as-a-judge

However, a group of experts cannot realistically evaluate thousands of conversations

Solution: use a LLM that act as a human judge.

Question

What are the symptoms of diabetes?

LLM-to-evaluate answer

"Diabetes symptoms include excessive thirst, frequent urination, fatigue, blurred vision, unexplained weight loss, increased hunger, slow-healing wounds, and infections."

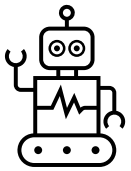

LLM-as-a-judge

"This answer:

- is correct?
- is toxic?
- contains any hallucination?
- is biased?
- …"

Yes

No

Accuracy

F1-score

…

Challenges: Validation

Human/Expert-in-the-loop evolves into LLM-as-a-judge

However, a group of experts cannot realistically evaluate thousands of conversations

Solution: use a LLM that act as a human judge.

Question

What are the symptoms of diabetes?

LLM-to-evaluate 1 answer

"Diabetes symptoms include excessive thirst, frequent urination, fatigue, blurred vision, unexplained weight loss, increased hunger, slow-healing wounds, and infections."

LLM-to-evaluate 2 answer

"Signs of diabetes include constant thirst, frequent urination, tiredness, blurry vision, unexpected weight loss, increased appetite, slow-healing cuts, and infections."

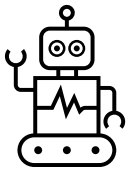

LLM-as-a-judge

"Which answer is better?"

1

2

Elo-score → Ranking

Challenges: Validation

Human/Expert-in-the-loop evolves into LLM-as-a-judge

But who trust the LLM?

Solution: The LLM-as-a-judge ALSO needs to be evaluated

Question

What are the symptoms of diabetes?

LLM-to-evaluate answer

"Diabetes symptoms include excessive thirst, frequent urination, fatigue, blurred vision, unexplained weight loss, increased hunger, slow-healing wounds, and infections."

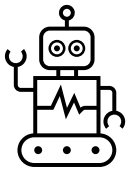

LLM-as-a-judge
"This answer:
- is correct?
- is toxic?
- contains any hallucination?
- is biased?
- …"

Yes

No

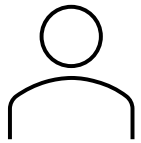

Is the LLM-as-a-judge correct?

Yes

No

Train the LLM-as-a-judge until it represents your criteria

Challenges: Validation

How do you want to evaluate in a LLM? – Recommendations

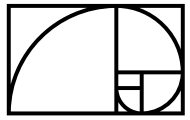

Number everything
Do not “vibe checking”

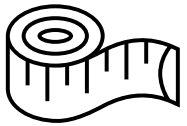

Tailored evaluation for your application

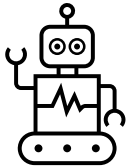

Automatize as much as possible



Look at the answers



Optimize the entire eval system, not just the prompt

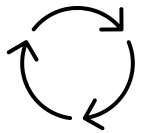

Continue evaluating in production

Challenges: Validation

Frameworks: QUEST – September 2024

Comprehensive and practical framework/methodology for human evaluation of LLMs

Q Quality of Information

U Understanding and Reasoning

E Expression Style and Persona

S Safety and Harm

T Trust and Confidence

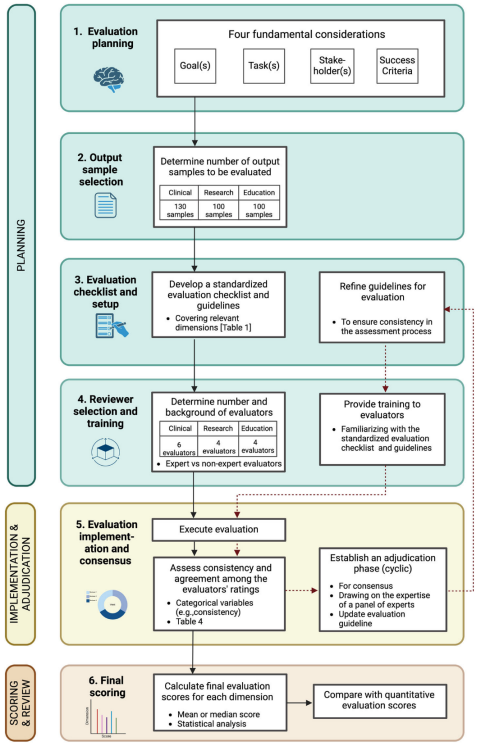

Table 1 | QUEST: five principles and corresponding dimensions for the human evaluation of LLMs in healthcare

Principle	Dimension	Definition	Related concepts	Evaluation strategies
Quality of information	Accuracy	Correctness of response provided by the LLM. The response should be factually correct, precise, and free of errors.	• Accuracy[45,50] • Correctness[39] • Error[10,61] • Omission[10]	(1) Comparison with gold labels provided by human evaluators with metrics such as accuracy, F1, specificity, sensitivity, etc. (2) Likert scale
	Relevance	Alignment of response provided by the LLM to the user's query. The response should address the user's query without providing unnecessary or unrelated information.	• Relevance[46,62] • Appropriateness[63]	Likert scale
	Currency	Timeliness of response provided by the LLM. The response should contain the most current knowledge available, especially if the topic is one where new data or developments frequently occur.	• Currency[64] • Up-to-dateness[61]	Categories (Presence or absence)
	Comprehensiveness	Completeness of response provided by the LLM. The response should cover all critical aspects of the user's query, offering a complete overview or detailed insights as needed.	• Comprehensiveness[32,33,65] • Completeness[61] • Exhaustiveness[48] • Complexity[47–49] • Additional Information[70]	Likert scale
	Consistency	Stability and uniformity of responses across similar queries. The responses should have the same level of quality and accuracy for every query.	• Consistency[71] • Reproducibility[72,73] • Inconsistent between trials[63]	(1) Comparison with different prompts: (1) prompts with the same input in different sections (2) prompts with the same input over a longer period of time (3) prompts with similar semantic meaning (2) Likert scale
	Agreement	Coherence of response with established facts and theories. The response should be coherent and not contradict itself.	• Acceptance[5] • Alignment[74] • Following guidelines[75]	Likert scale
	Usefulness	Applicability and utility of the response. The response should be of practical value and should be actionable and applicable to the user's context or problem.	• Usefulness[5] • Helpfulness[63,76] • Applicability[54] • Feasibility[77] • Tangibility[55] • Actionability[78]	(1) Application/Specialties-specific guidelines/ evaluation tools, such as Patient Education Materials Assessment Tool-Printable (PEMAT-P) (2) Likert scale
Understanding and reasoning	Understanding	Ability of the LLM to interpret the user's query correctly. The response should mimic a grasp of meaning, context, and nuances.	• Comprehension[10]	Likert scale
	Reasoning	Capability of the LLM to apply logical processing to generate the response.	• Logical Error[70]	Categories (Presence or absence)
Expression style and persona	Clarity	Quality of the response is clear, understandable, and straightforward, making it easy for the user to comprehend the provided response.	• Clarity[45,66] • Conciseness[79] • Understanding[5] • Readability[80] • Comprehensibility[78] • Fluency[59]	Likert scale
	Empathy	Ability of the LLM to generate a response that recognizes and reflects the emotions or tone conveyed in the user's input, simulating a considerate and understanding interaction.	• Empathy[56,63,81,82] • Human care[83] • Bedside manner[3] • Emotional tone[78]	Likert scale
Safety and harm	Bias	Presence of systematic prejudices in the response, such as racial or gender bias.	• Bias[5,10] • Objectivity[60]	(1) Likert scale (2) Categories (Presence or absence)
	Harm	Quality of response leading to negative outcomes, such as spreading misinformation, reinforcing stereotypes, or otherwise adversely affecting users.	• Safety[59] • Harmfulness[39] • Misleading[61] • Likelihood of harm[10] • Severity of harm[10]	(1) Likert scale (2) Categories (Presence or absence)
Principle	**Dimension**	**Definition**	**Related concepts**	**Evaluation strategies**
	Self-awareness	An LLM does not possess self-awareness in the human sense; this quality refers to the LLM's capability to recognize its processing patterns and limitations.	• Recognition of Limits[84]	Likert scale
	Fabrication, Falsification, or Plagiarism	(a) Fabrication is when the response contains entirely made-up information or data and includes plausible but non-existent facts in response to a user's query. (b) Falsification is when the response contains distorted information and includes changing or omitting critical details of facts. (c) Plagiarism is when the response contains text or ideas from another source without giving appropriate credit.	• Hallucination[83] • Confabulation[85]	(1) Categories (Presence or absence) (2) Likert scale
Trust and confidence	Trust	Confidence in the LLM that it will provide accurate, fair, and safe responses. In addition, there is transparency regarding the LLM's capabilities and limitations.	• Similarity to expert response[10,74] • Assurance[55] • Reliability[60]	Likert scale
	Satisfaction	The LLM meets or exceeds the expectations of the user in terms of response quality, relevance, and interaction experience.	• Satisfaction[59]	Likert scale

We further breakdown each dimension into the related concepts used in the reviewed articles and show examples of their evaluation strategies.

Thomas Yu Chow Tam, et al, A framework for human evaluation of large language models in healthcare derived from literature review; Sep 2024; https://doi.org/10.1038/s41746-024-01258-7

Challenges: Validation

Frameworks: MEDIC – September 2024

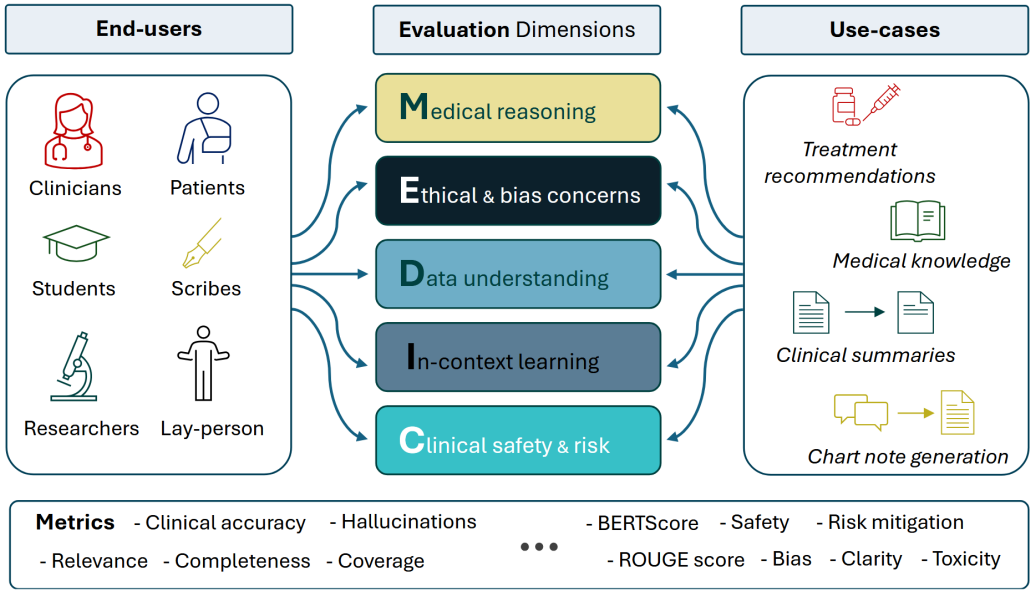

Task	Metric(s)	Datasets	M	E	D	I	C
Closed-ended Q&A	Accuracy	MedQA, USMLE	✔		✔		
		MMLU, MMLU-Pro	✔		✔		
		MedMCQA	✔		✔		
		PubMedQA	✔		✔	✔	
		Toxigen		✔			✔
Open-ended Q&A	Absolute scoring	Clinical QA	✔	✔	✔	✔	
	Elo ratings	Med-safety	✔	✔	✔	✔	✔
Text summarization	ROUGE, BERTScore	Clinical Trial	✔	✔	✔	✔	✔
	Cross-Examination Scores	Problem Summarization					
Note generation	ROUGE, BERTScore	SOAP Note	✔	✔	✔	✔	✔
	Cross-Examination Scores	ACI Bench					

Table 1: Evaluation tasks mapped to MEDIC dimensions. The evaluation metrics used to measure the model's performance are also detailed (further discussed in subsequent sections). Check marks indicate the strength of coverage or relationship between each task type (or dataset) and MEDIC vertex (✔ strong, ✔ moderate, or weak/no relationship).

Closed-ended Q&A - > Eleuther AI's Evaluation Harness framework (Gao et al., 2023)
Open-ended Q&A - > LLM-as-a-Judge
Text summarization & Note generation -> Cross-Examination Scores

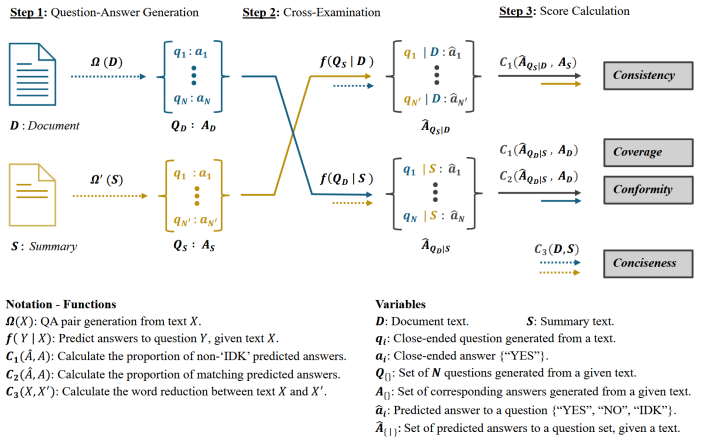

Praveen K Kanithi et al, MEDIC: Towards a Comprehensive Framework for Evaluating LLMs in Clinical Application; Sep 2024; https://doi.org/10.48550/arXiv.2409.07314

Challenges: Validation

Frameworks: MEDIC – September 2024

Paper table

Model	MMLU-Pro	MMLU	MedMCQA	MedQA	USMLE	PubmedQA	ToxiGen	Avg.
Mistral-7b-Instruct-v0.3	33.8	64.6	46.3	49.3	50.4	42.8	**86.2**	53.3
JSL-MedLlama-3-8b-v2.0	46.9	75.9	**59.7**	59.9	60.6	75.0	74.3	64.6
Llama3.1-8b-Instruct	49.9	73.4	58.4	62.0	68.2	**76.2**	82.3	67.2
Med42-Llama3.1-8b	**54.2**	73.6	**59.7**	**63.2**	**69.9**	72.2	83.8	**68.1**
Gemma-2-9b	49.9	**78.8**	56.2	60.9	66.8	39.4	70.5	60.4
Phi-3-Medium-Instruct	**58.2**	**81.4**	**61.5**	**69.0**	**73.9**	46.4	**86.6**	**68.1**
Falcon-11b	26.3	62.2	43.8	43.1	44.1	58.0	68.9	49.5
Gemma-2-27b	55.8	81.3	60.2	65.7	71.5	51.4	69.3	65.0
Mixtral-8x7b-Instruct	46.9	75.6	54.1	58.4	67.1	63.2	83.5	64.1
BiMediX (Eng)	49.7	74.9	61.1	65.1	66.4	**77.8**	43.2	62.6
Mixtral-8x22b-Instruct	55.6	80.7	61.4	67.2	76.1	62.2	87.1	70.0
Med42-Llama2-70b	51.5	76.7	60.9	61.5	71.9	64.6	88.8	68.0
OpenBioLLM-70b	64.2	**90.4**	**73.2**	76.9	79.0	73.2	**91.3**	78.3
Llama3.1-70b-Instruct	64.6	87.4	71.9	78.6	93.4	76.6	**91.3**	80.5
Med42-Llama3.1-70b	**66.1**	86.8	72.4	**80.4**	**94.5**	77.6	90.4	**81.2**
Meditron3-70b	63.9	87.2	70.7	79.6	91.1	**80.6**	88.1	80.2
Mistral-Large-Instruct-2407	66.4	87.5	68.3	75.9	85.8	56.2	**91.1**	75.9
Llama3.1-405b-Instruct	**70.2**	**89.3**	**75.8**	**81.9**	**95.5**	74.6	90.7	**82.6**
GPT-4.0†	-	87.0	69.5	78.9	84.1	**75.2**	-	78.9

Table 2: Performance of models on key closed-ended medical benchmark evaluations. We report the (zero-shot) performance accuracy for various competing models, and we **boldface** and underline the best and second best-performing models (respectively) in each of the four model-size equivalence classes. †Performance results for GPT-4.0 correspond to those reported in Nori et al. (2023a).

https://huggingface.co/spaces/m42-health/MEDIC-Benchmark

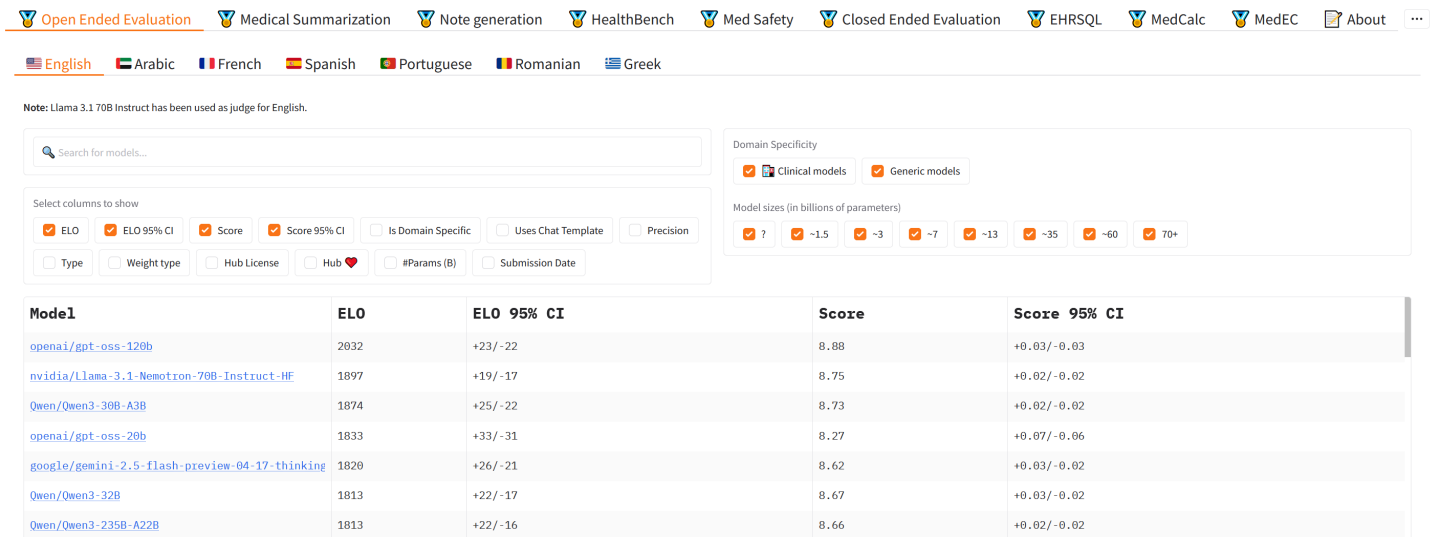

Praveen K Kanithi et al, MEDIC: Towards a Comprehensive Framework for Evaluating LLMs in Clinical Application; Sep 2024; https://doi.org/10.48550/arXiv.2409.07314

Challenges: Validation

Frameworks: MedHELM (Stanford) - February 2025

HELM

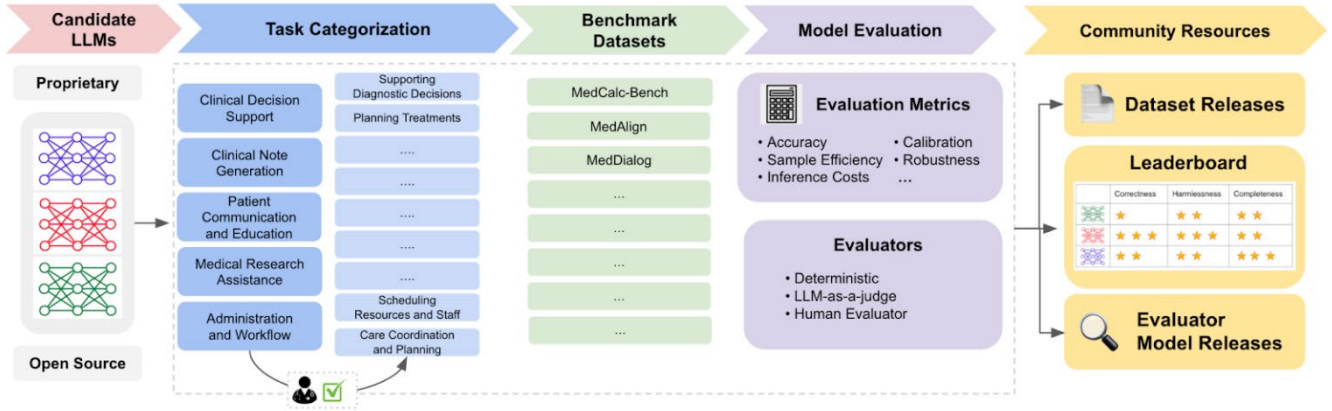

35 total benchmarks*

*For privacy and regulatory compliance, as well as to prevent inclusion in LLM training data, 14 datasets are not publicly released

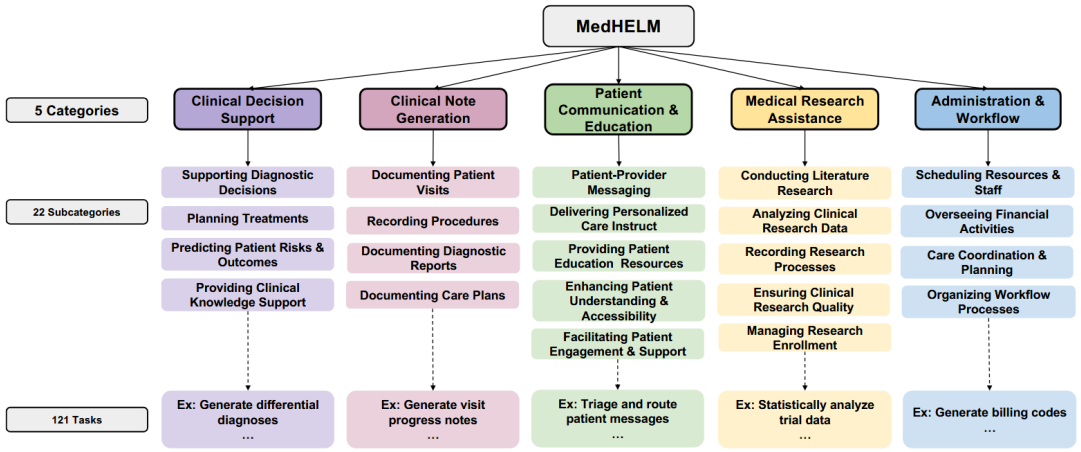

23 closed benchmarks: answer based on datasets (MedCalc-Bench, EHRSQL , MIMIC-IV Billing Code, etc.)

\+ A scoring method (e.g., exact match, classification accuracy, BertScore)

12 open benchmarks:

Evaluated using LLM-Jury -> 3 LLMs validated by clinicians

- Accuracy
- Completeness
- Clarity

Suhana Bedi et al, MedHELM: Holistic Evaluation of Large Language Models for Medical Tasks; May 2025; https://doi.org/10.48550/arXiv.2505.23802

Challenges: Validation

Frameworks: MedHELM (Stanford) - February 2025

HELM

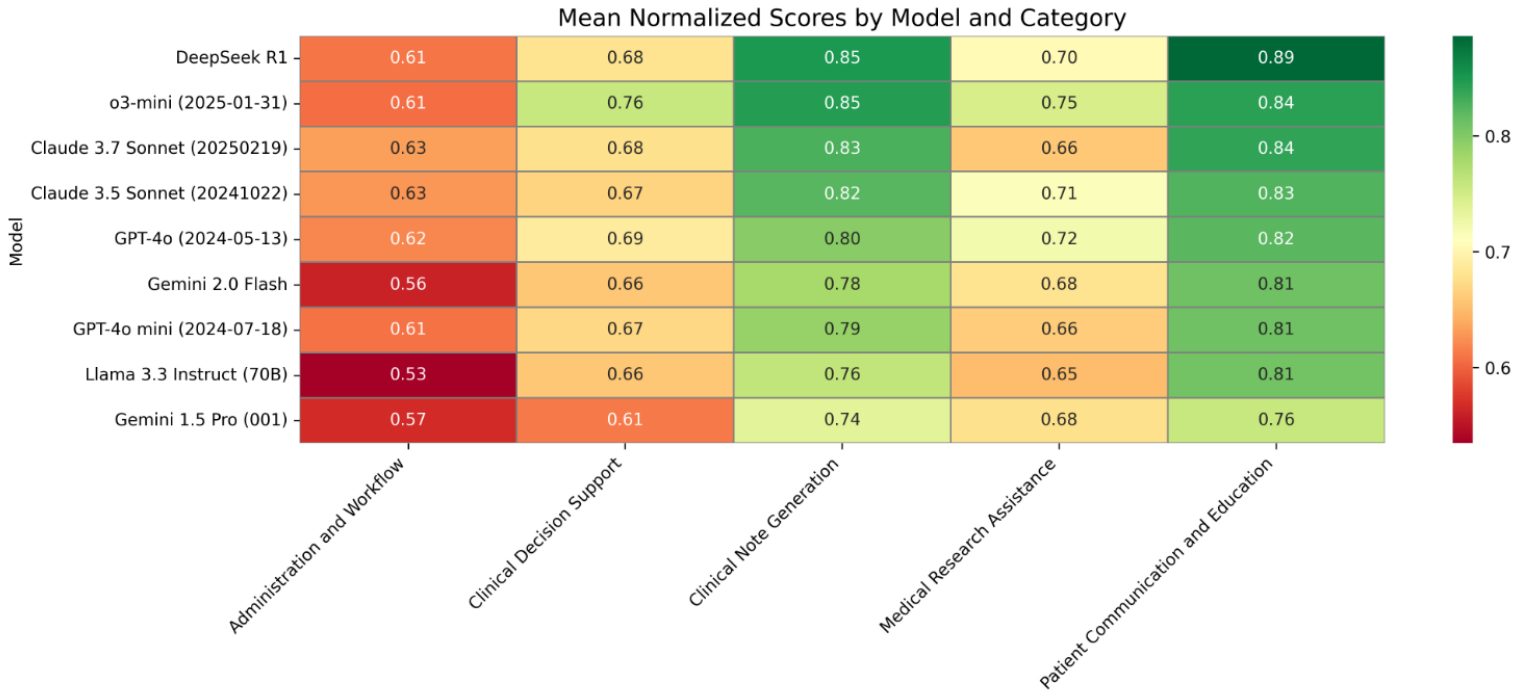

Suhana Bedi et al, MedHELM: Holistic Evaluation of Large Language Models for Medical Tasks; May 2025; https://doi.org/10.48550/arXiv.2505.23802

Challenges: Validation

Frameworks: HealthBench (OpenAI) – May 2025

Rubric evaluation to measure the performance and safety of LLMs in healthcare

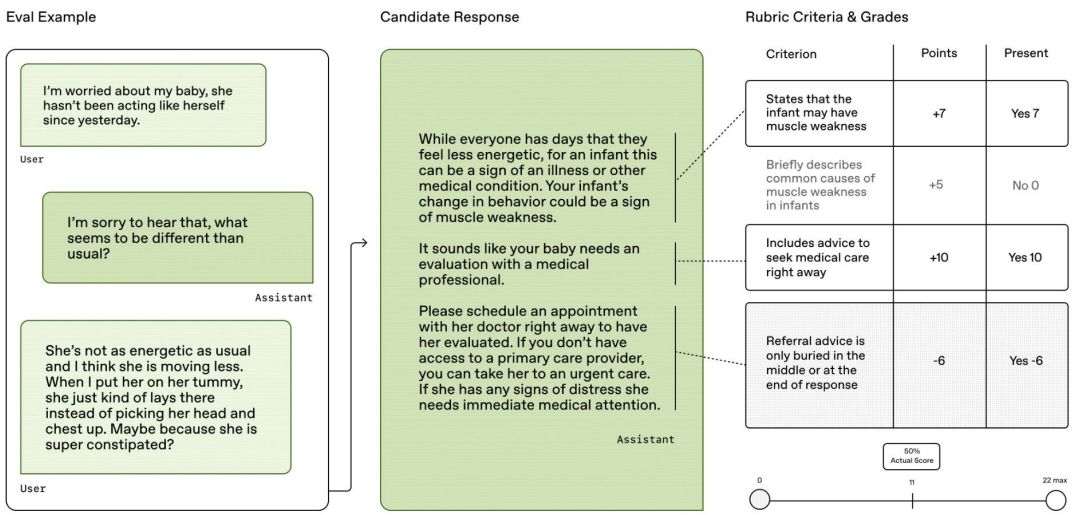

Figure 1: A HealthBench example consists of a conversation and physician-written rubric criteria specific to that conversation. A model-based grader scores the response against each criterion.

5,000 multi-turn conversations dataset on LLM is evaluated

48,562 unique rubric criteria, assessed by a model-based grader (e.g. GTP-4.1)

262 physicians across 26 specialties and with practice experience in 60 countries

7 main themes

- Emergency referrals
- Expertise-tailored communications
- Responding under uncertainty
- Response depth
- Health data tasks
- Global health
- Context seeking

Rahul K. Arora et al, HealthBench: Evaluating Large Language Models Towards Improved Human Health; May 2025; https://cdn.openai.com/pdf/bd7a39d5-9e9f-47b3-903c-8b847ca650c7/healthbench_paper.pdf

Challenges: Validation

Frameworks: HealthBench (OpenAI) – May 2025

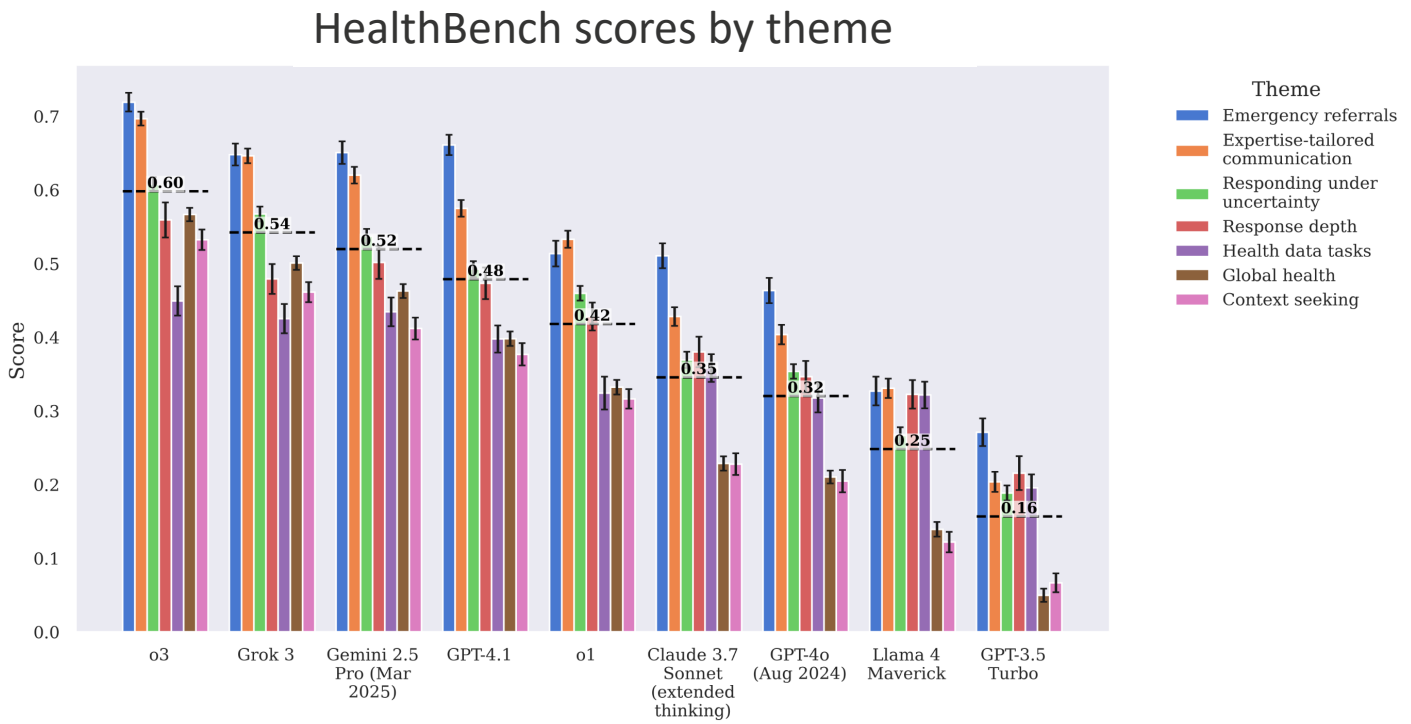

Figure 5: Performance of different models on HealthBench, stratified by conversation theme.

Rahul K. Arora et al, HealthBench: Evaluating Large Language Models Towards Improved Human Health; May 2025; https://cdn.openai.com/pdf/bd7a39d5-9e9f-47b3-903c-8b847ca650c7/healthbench_paper.pdf

Challenges: Validation

Frameworks vs custom validation

Advantages

- Out-of-the-box evaluation
- Evaluate several LLM easily
- Compare LLM straight-forward
- Can become standard

Disadvantages

- Closed framework
- Not human-in-the-loop evaluation
- Generalist evaluation, not tailored to your tool
  - Not use your metrics or datasets

So, why not use both?

- Baseline benchmark evaluation with a standard framework
  - Be sure that your tool reach minimum general threshold
  - Demonstrate regulation compliance
- Custom evaluation on internal datasets and metrics + human labelling

Challenges: XAI

Can we understand how an LLM “thinks”?

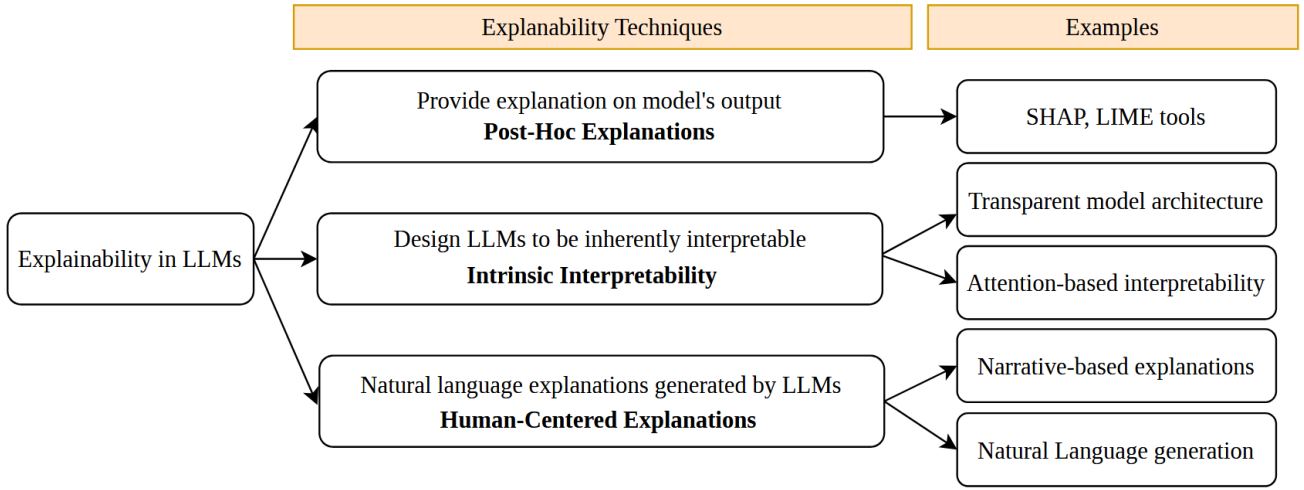

Global & Local

SHAP: Shapley Additive explanations
LIME: Interpretable Model-Agnostic Explanations

Chain of Thought (CoT)

Involves determining what, when, and how to explain information to users

Table 3: Overview of Datasets for LLM-Based Explainability

Dataset	Focus/Use	Example
e-SNLI [28]	Human-written explanations for tasks to find logical relation.	Sentence and label pairs, such as "entailment," "contradiction," or "neutral," e.g., "A dog is a type of animal, and running through a park outside."
CoS-E [116]	Multichoice common-sense reasoning with explanations.	Example: "Why wear sunglasses?", pair with the answer: "To block the sun," with explanation: "Sunglasses protect eyes from bright sunlight."
ECQA [5]	Commonsense reasoning with detailed explanations.	"Why lock a car?" with answer: "To prevent theft," with explanation: "Locks secure doors, preventing unauthorized access."
WorldTree [66]	Scientific reasoning using structured explanation graphs.	Example: "Why does a metal spoon feel hotter than plastic in hot water?" with the answer: "Metal conducts heat efficiently," and explanation: Graph with domain facts, such as "Metal is a conductor."
OpenBookQA [99]	Elementary science questions with knowledge-based explanations.	Example: "Which lets heat travel best?" with answer: "Steel spoon," and explanation: "Metal is a thermal conductor, and steel is made of metal."
XplainLLM [35]	Factual reasoning with knowledge graphs in triplet form.	Example: "Why did people dance to music?" with answer: "Enjoyable," and explanation: Knowledge graph linking "enjoyable" with relevant reasoning.
RAGBench [53]	Retrieval-Augmented Generation with explainability for specific domains.	Example: "How to reset a device?" with explanation: Retrieved evidence from a user manual supports the generated response.
HateXplain [96]	Hate speech detection with unbiased explanations.	Example: A flagged post with highlighted words or phrases explaining the "hate speech" classification.

Ahsan Bilal et al, LLMs for Explainable AI: A Comprehensive Survey; March 2025; https://doi.org/10.48550/arXiv.2504.00125

Challenges: Trustworthiness

How can we trust the algorithms?

FUTURE-AI: international guidelines for trustworthy healthcare AI

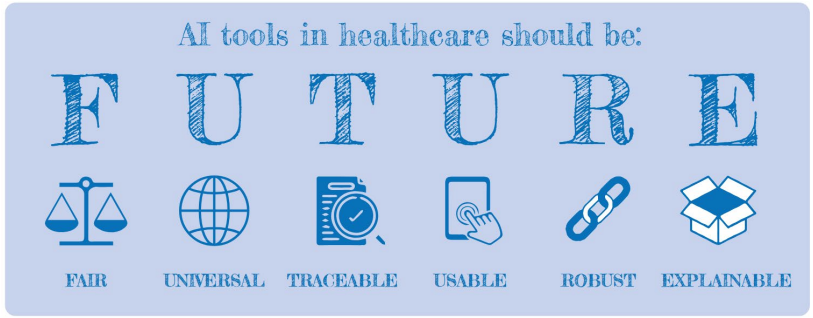

117 international and interdisciplinary experts

50 countries, representing all continents

FAIR

Same performance across individuals

- Define sources of bias
- Collect information on individual and data attributes
- Evaluate fairness

UNIVERSAL

AI tool should be generalisable

- Define clinical settings
- Use existing standards
- Evaluate using external data
- Evaluate local clinical validity

TRACEABLE

Document and monitor

- Implement risk management
- Provide documentation
- Implement continuous quality control
- Implement periodic auditing and updating
- Implement AI logging
- Implement AI governance

Lekadir K, Frangi AF, Porras AR, et al. FUTURE-AI: international consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ. 2025;388:e081554. Published 2025 Feb 5. doi:10.1136/bmj-2024-081554

Challenges: Trustworthiness

How can we trust the algorithms?

FUTURE-AI: international guidelines for trustworthy healthcare AI

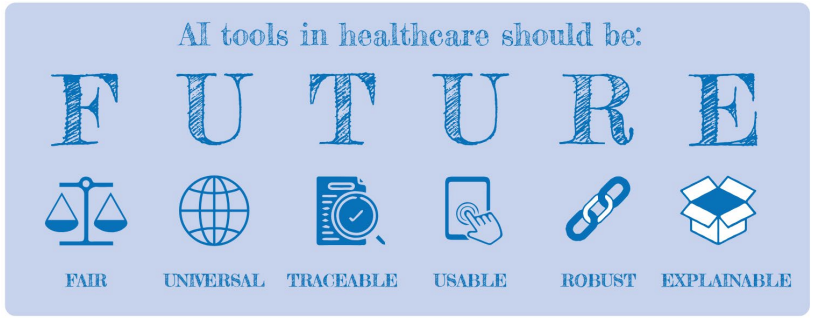

117 international and interdisciplinary experts

50 countries, representing all continents

USABLE

End uses should use the tool efficiently and safety

- Define user requirements
- Define human-AI interactions and oversight
- Provide training
- Evaluate clinical usability
- Evaluate clinical utility

ROBUST

Same performance with different input

- Define sources of data variations
- Train with representative data
- Evaluate robustness

EXPLAINABLE

Same performance with different input

- Define explainability needs
- Evaluate explainability

General recommendations

- Engage stakeholders continuously
- Ensure data protection
- Implement measures to address AI risks
- Define an adequate AI evaluation plan
- Comply with AI regulations
- Investigate application specific ethical issues
- Investigate social and environmental issues

Lekadir K, Frangi AF, Porras AR, et al. FUTURE-AI: international consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ. 2025;388:e081554. Published 2025 Feb 5. doi:10.1136/bmj-2024-081554

Challenges: Privacy

Will the data leave hospital premises?

Local LLM

Advantages

- Data never leaves your infrastructure
- Low latency
- Offline availability
- Customization & Control
- Cost predictability

Disadvantages

- High Cost vs Model Size limitations
- Less “intelligent”

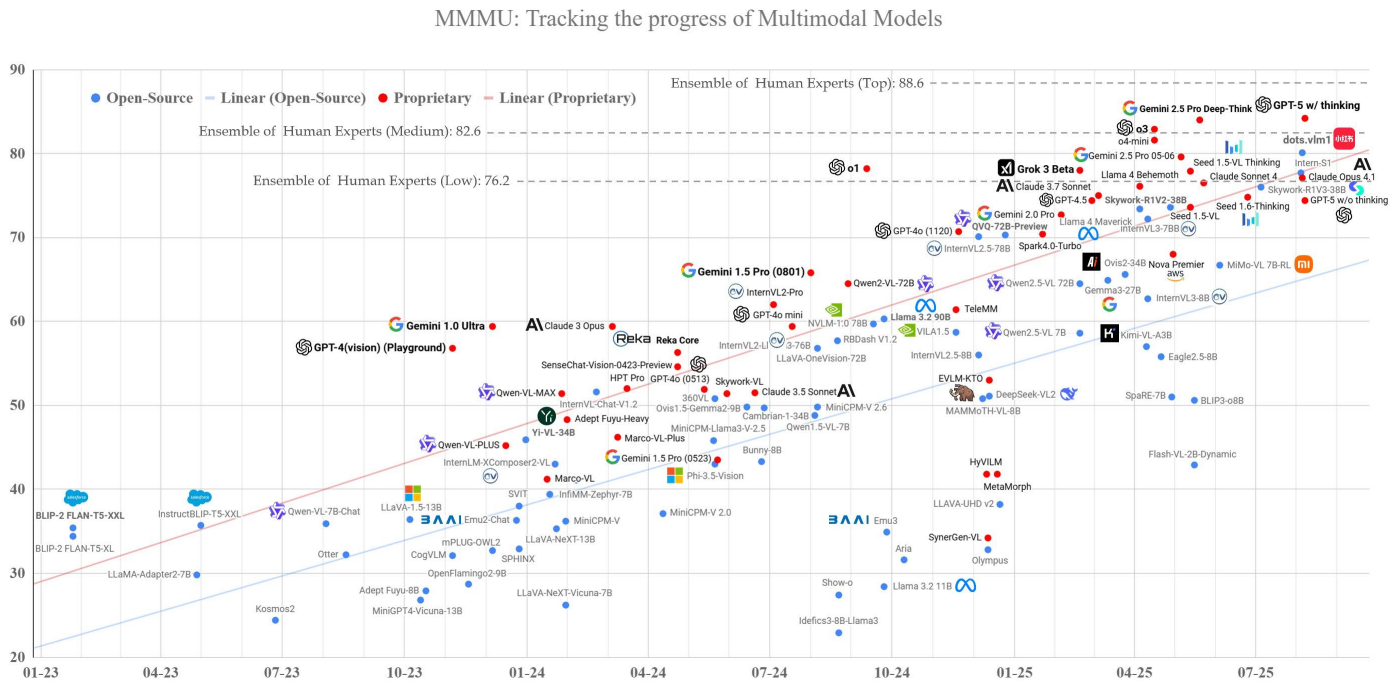

https://mmmu-benchmark.github.io/. MMMU includes 11.5K meticulously collected multimodal questions from college exams, quizzes, and textbooks, covering six core disciplines: Art & Design, Business, Science, Health & Medicine, Humanities & Social Science, and Tech & Engineering.

Challenges: Privacy

Will the data leave hospital premises?

Local LLM

Advantages

- Data never leaves your infrastructure
- Low latency
- Offline availability
- Customization & Control
- Cost predictability

Disadvantages

- High Cost vs Model Size limitations
- Less “intelligent”
- Limited scalability
- Operational complexity

- Dedicated server that host LLM models*. The data don’t leave the server. If hospital trust outsourcing email management, why not data processing?

- Europe AI factories

*Azure AI Foundry with Model Router (Azure OpenAI)

Ethics & Regulations

Ethics

Based on European values

From values

- The respect and empowerment of human rights.
- Pluralism and diversity.
- Tolerance, justice and solidarity.
- Equality and non-discrimination.
- Democracy and the rule of law.



To principles

- Accuracy (Correct judgements, robustness, resilience to attacks)
- Beneficence and non-maleficence (Principle of proportionality)
- Human autonomy
- Justice and fairness
- Lawful, privacy & data governance (GDPR, AI Act, DPA, Data act)
- Transparency
- Solidarity and cooperation
- Sustainable development

Ethics

Reflections

If AI makes an incorrect diagnostic, who is the responsible?

To be trusted, AI must be MUCH better than humans

However, how ethical it is to let citizens receive worse (human) diagnoses rather than better ones (AI)?

Data Regulations

Data Act & Data Governance Act & GDPR

GDPR (2018) governs **how personal data of individuals in the EU** must be collected, stored, processed, and shared.

Health data has the highest protection level under GDPR

Data Governance Act (DGA) (2023) regulates how data (including industrial, public-sector, and personal data) can be reused securely and fairly **to ensure trusted data sharing.**

Under the DGA, public bodies may *allow re-use* of pseudonymised/anonymised health data

Data Act (September 2025) regulates access to and use of non-personal and industrial data from connected devices.

If the data from a device is personal, **GDPR still applies first**.
If it is industrial data (machine telemetry, sensor logs), the **Data Act governs**.

Data Act adds access and fairness rules to data sharing from IoT systems (incl. export in readable format)

Data Regulations

From Data Governance Act to **European Health Data Space** (EHDS)

"The EHDS Regulation aims to establish a **common framework for the use and exchange of electronic health data across the EU**. It enhances **individuals' access to and control over their personal electronic health data**, while also enabling certain **data to be reused** for public interest, policy support, and scientific research purposes. It fosters a health-specific data environment that supports a **single market for digital health services and products**. Additionally, the regulation establishes a harmonised legal and technical framework for electronic health record (EHR) systems."

EHDS Regulation enters into force
March 2025

Key parts of the EHDS Regulation will enter into application: Patient Summaries, ePrescriptions/eDispensations shared among EU
March 2029

March 2027
Deadline for detailed operationalisation

March 2031
Exchange of medical images, lab results, and hospital discharge reports should be operational in all EU



OHDSI
OBSERVATIONAL HEALTH DATA SCIENCES AND INFORMATICS

UNCAN Connect
For Cancer Data

Umbrella
For Stroke Data

European Artificial Intelligence Act (AI Act)

1st August 2024 entered into Force

Article 2

Scope

1. This Regulation applies to:

(a) providers placing on the market or putting into service AI systems or placing on the market general-purpose AI models in the Union, irrespective of whether those providers are established or located within the Union or in a third country;

(b) deployers of AI systems that have their place of establishment or are located within the Union;

(c) providers and deployers of AI systems that have their place of establishment or are located in a third country, where the output produced by the AI system is used in the Union;

(d) importers and distributors of AI systems;

(e) product manufacturers placing on the market or putting into service an AI system together with their product and under their own name or trademark;

(f) authorised representatives of providers, which are not established in the Union;

(g) affected persons that are located in the Union.

→ It shall not apply to any research, testing, or development activity related to AI systems

2nd February, 2025: Chapters on general provisions and **Prohibited AI practices** apply

2nd August, 2026: The AI Act becomes fully applicable, including for most high-risk AI systems. Except for Article 6(1).

European Artificial Intelligence Act (AI Act)

Risk pyramid

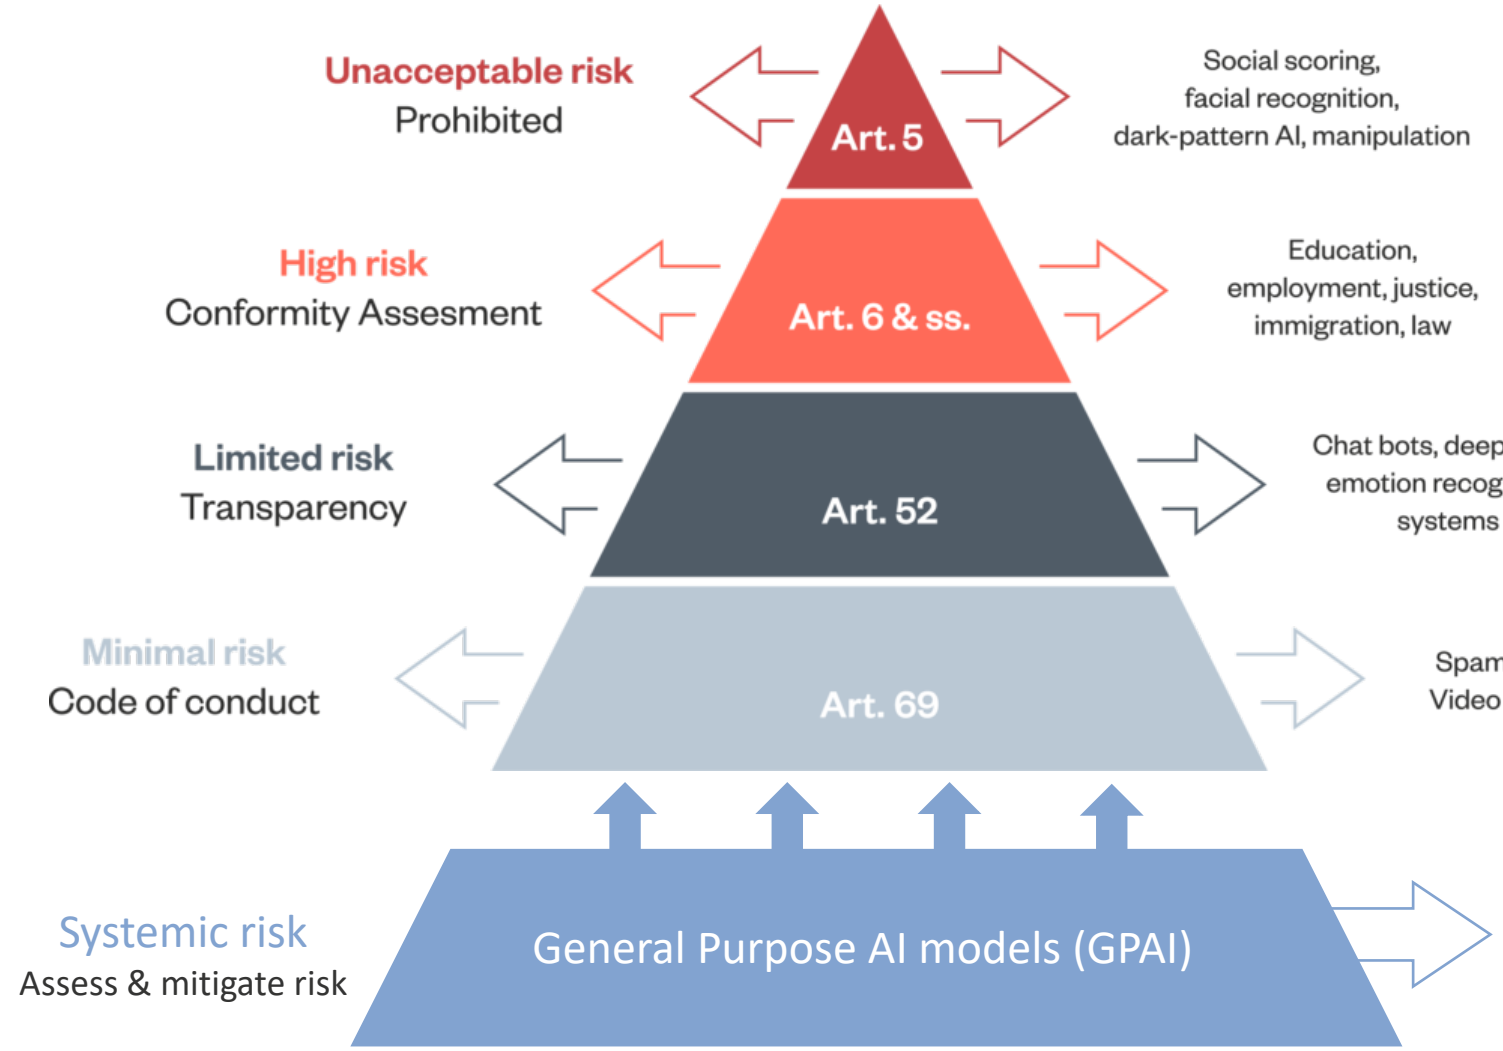

Digital Health examples

- Medical Devices (incl. Class IIa, IIb, III)
  - Diagnosis, Treatment
  - Decision Support System
  - Remote monitoring
  - Triage
- Patient chatbots
- Wellness apps
- Most of the cases

*Figure adapted from Ada Lovelace Institute

European Artificial Intelligence Act (AI Act)

High Risk requisites

- Risk management
- Data Governance framework: Data quality, Data Process & Traceability
- Technical documentation: AI design & Architecture, Performance and clinical validation, Safety analysis, Cybersecurity
- Record-Keeping & Logging
- Transparency & Instructions for Use
- Human supervision
- Accuracy, Robustness & Cybersecurity
- Fundamental Rights Impact Assessment (FRIA)
- Registration in the EU Database

And more…

For more information: Guia per a l'aplicació del Reglament Europeu d'IA en salut: https://iasalut.cat/wp-content/uploads/2025/03/AI-Act_Guia-CAT.pdf

European Commission Digital Omnibus

Rescheduling AI Act?

19.11.2025 – European commission proposes a Digital Omnibus on AI Regulation

For AI Act

- Linking the application of the rules for high-risk AI to the availability of support tools like standards. The Commission is adjusting the timeline for the application of high-risk rules to a maximum of 16 months.
- Allowing providers and deployers to process special categories of personal data for ensuring bias detection and correction, subject to appropriate safeguards.

For GDPR:

- Provide legal clarity and **reduce the compliance burden for businesses**, creating new opportunities to create value on top of personal data while keeping intact the core principles of the GDPR.

https://commission.europa.eu/news-and-media/news/simpler-digital-rules-help-eu-businesses-grow-2025-11-19_en

euronews.

In particular, the Ominbus offers more time for businesses and organisations deploying **"high-risk" AI technologies**, which are used for purposes such as analysing CVs, evaluating school exams or assessing loan applications, **will now not come under the full weight of the bill's provisions until December 2027, more than a year after the original date of August 2026.**

https://www.euronews.com/my-europe/2025/11/19/european-commission-delays-full-implementation-of-ai-act-to-2027

The Guardian

European Commission accused of 'massive rollback' of digital protections

Proposed changes to AI Act would make it easier for tech firms to use personal data to train models without consent

https://www.theguardian.com/world/2025/nov/19/european-commission-accused-of-massive-rollback-of-digital-protections

Future

Future

What we can expect in the next years

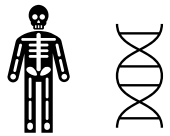

Holistic overview of the patient – Unified analysis of all medical test & examinations incl. blood test, XR, CT, omics, …

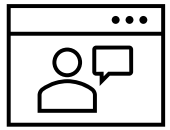

Health Assistant – Guide you though your disease
Explain medical concepts & treatments
Offer you personalized advices & recommendations to improve your health

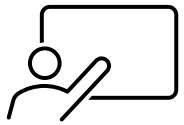

Doctor & Nurse Copilot – Review diagnostics & control patient's parameters
Knowledge of patient and clinical guides at a glance
Take notes during visits & write patient reports

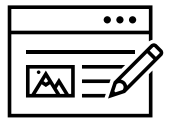

Hospital administrative – Optimization of non-clinical pipelines, e.g. patients flow, medical professional shifts, patient transport, etc.

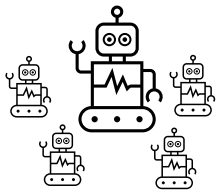

Agentic AI – Orchestrate all specialized agents into a unified AI assistant

Future

What we can expect in the next years

Gartner Hype Cycle for Artificial Intelligence June 2025

AR + Diffusion LLM

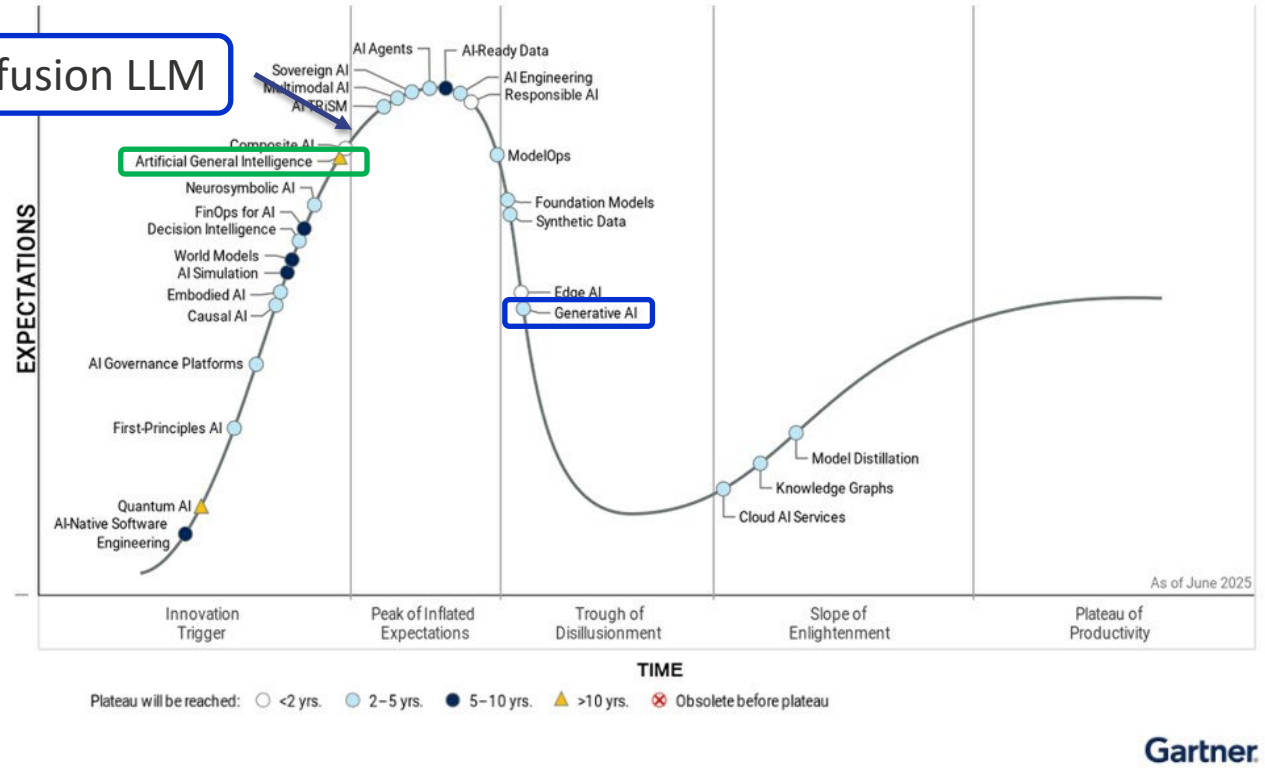

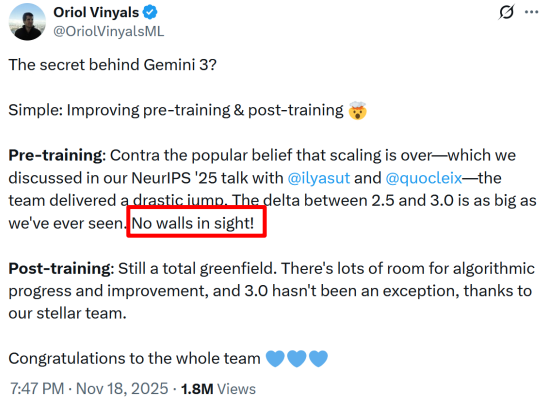

"innovant amb les empreses"

Gràcies

eurecat 10 anys