INTELLIGENT COMPASSION AND DATA-DRIVEN DIGNITY: INTEGRATING AI INTO HOSPICE AND PALLIATIVE MEDICINE

GARNER BOOGAERTS MD, MPH

USF HOSPICE PALLIATIVE FELLOW

LEARNING OBJECTIVES

- Define current state of Artificial Intelligence
- Describe the current and emerging roles of Artificial Intelligence in palliative and hospice care settings.
- Future directions or research areas where AI could enhance holistic, patient-centered care in hospice and palliative medicine.
- Discuss the potential risks and challenges of AI use in end-of-life care, including bias, data privacy, and depersonalization of care.

ARTIFICIAL INTELLIGENCE 101

- **Machine learning:** A branch of applied mathematics that deals with systems that improve with data
- **AI:** Vague term. A branch of applied machine learning
- **LLMs (Large Language Model):** The use of AI techniques to tackle language modeling tasks (e.g. comprehension, translation, summarization)
  - **GPT: Generative Pre-Trained Transformer.** A popular foundation model, it is a type of LLM
  - **ChatGPT:** An implementation that uses GPT to act as a chatbot
    - Think of **LLM** like saying “car” — it refers to the category of vehicle.
    - So GPT = **a specific LLM family**, just like **Toyota Camry** is a specific car model.

A BRIEF HISTORY OF AI IN HEALTHCARE

1. **1970s–1980s: Rule-Based Expert Systems**
    - MYCIN (assist physicians in diagnosing and treating bacterial infections)
    - INTERNIST-I (diagnose multiple and complex diseases by modeling clinician behavior)
  - **Strengths:** Transparent logic, domain specificity
  - **Limitations:** Rigid, labor-intensive, no learning or adaptation
2. **1990s–2000s: Probabilistic Reasoning & EMR Integration**
    - **ILIAD -** Bayesian diagnosis- Weighted probabilities of disease
    - **Oncocin -** knowledge base with an interactive interface for physicians to input patient data and receive treatment recommendations.
  - **Strengths:** More adaptive, some learning features
  - **Limitations:** Still not flexible with unstructured data
3. 2011: IBM Watson
4. 2020s: LLM Era (e.g., GPT-4, Med-PaLM, Claude)

LARGE LANGUAGE MODELS

- A **language model** is a mathematical model that operates on a series of **tokens**
- Tokens are usually words, but they can be anything: characters, whole sentences, phonemes, etc.
- Language models are trained on a collection of tokens
  - CommonCrawl
  - RefinedWeb
  - Twitter/X
  - Reddit
  - Youtube comments

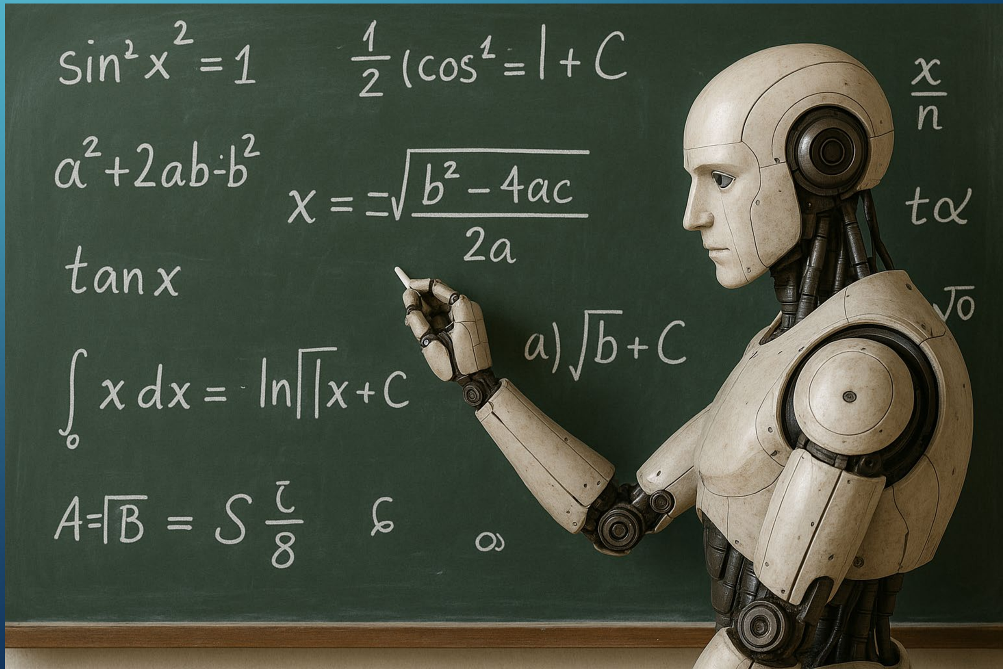

ALGORITHMS IN DISGUISE

- "Find me the parameters that describe a function that, for each sequence of tokens, maps the most likely next word to any sequence of words”

$$\hat{\Phi} = \underset{\phi}{\operatorname{argmin}}\, J(f_\phi(k_1, k_2, \cdots, k_n), k_{n+1})$$

- "Given a sequence of words, what is the most likely next word?"

$$k_{n+1} = \underset{w \in W}{\operatorname{argmax}}\, p(k_{n+1} = w | k_1, k_2, \cdots, k_n, \Phi)$$

EXAMPLE

- Let's say your prompt is: **"Palliative care focuses on improving __"**
- LLMs break down text into **tokens**, which are usually:
  - **Words**, **word parts**, or **punctuation**
  - For example, GPT uses tokens like:
    "Palliative", " care", " focuses", " on", " improving"
  - Now the model tries to **predict the next token** after "improving".

Token Prediction	Probability
" quality"	38%
" pain"	20%
" symptoms"	15%
" outcomes"	10%
" mobility"	5%

It picks **" quality"** (the most likely), then moves to the next token…

EXAMPLE

- **Resulting Output:**
- **"Palliative care focuses on improving |quality |of |life |for |patients |with |serious | illness."**
- Each token is predicted **one at a time**, using the full context before it.
- GPT **<u>does not have an "understanding" </u>** of what palliative is, or what a **quality** is, or what **illness** is. It only **<u>has an understanding of the token sequence</u>**

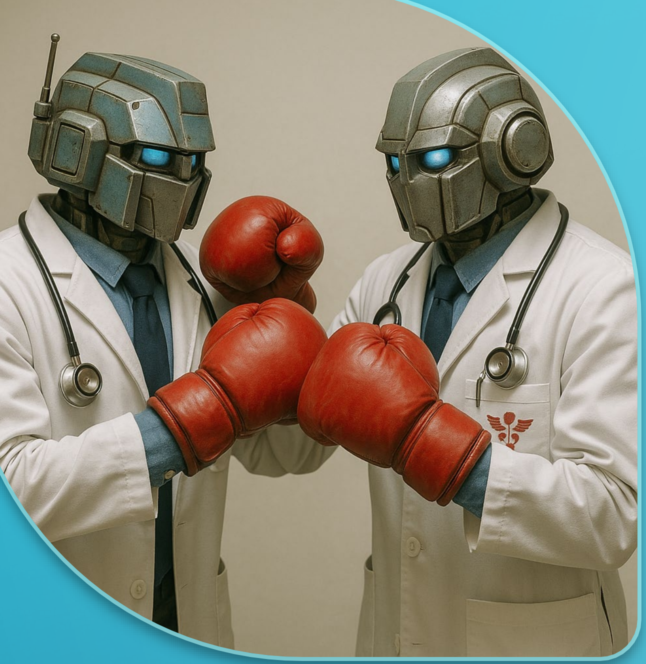

BATTLE OF THE LLMS

Model	Launch Date	No. of Parameters	Purpose
GPT-3	Jun-20	175 billion	General-purpose language model, useful for text generation, translation, summarization.
BERT (Bidirectional Encoder Representations from Transformers)	Oct-18	110 million (Base), 340 million (Large)	NLP tasks like sentiment analysis, Q&A, classification, and language understanding.
GPT-4	Mar-23	Estimated 1 trillion	Multi-modal (text and image), designed for advanced text generation, comprehension, and reasoning tasks.
LaMDA (Language Model for Dialogue Applications)	May-21	Unknown (rumored to be hundreds of billions)	Focused on conversational dialogue applications with an emphasis on safe, open-ended conversations.
PaLM (Pathways Language Model)	Apr-22	540 billion	Multi-task NLP model designed for reasoning, code generation, and few-shot learning.
LLaMA (Large Language Model Meta AI)	Feb-23	7B, 13B, 30B, 65B	Research-oriented LLM for use in a wide range of language tasks and experiments.
BLOOM (BigScience Large Open-science Open-access Multilingual Language Model)	Jul-22	176 billion	Open-source multilingual language model aimed at inclusivity and scientific collaboration.
T5 (Text-to-Text Transfer Transformer)	Oct-19	220 million (Base)	Converts all NLP tasks into a text-to-text format, useful for Q&A, translation, and summarization.
Megatron-Turing NLG	Oct-21	530 billion	Natural language generation, translation, summarization, and reasoning.

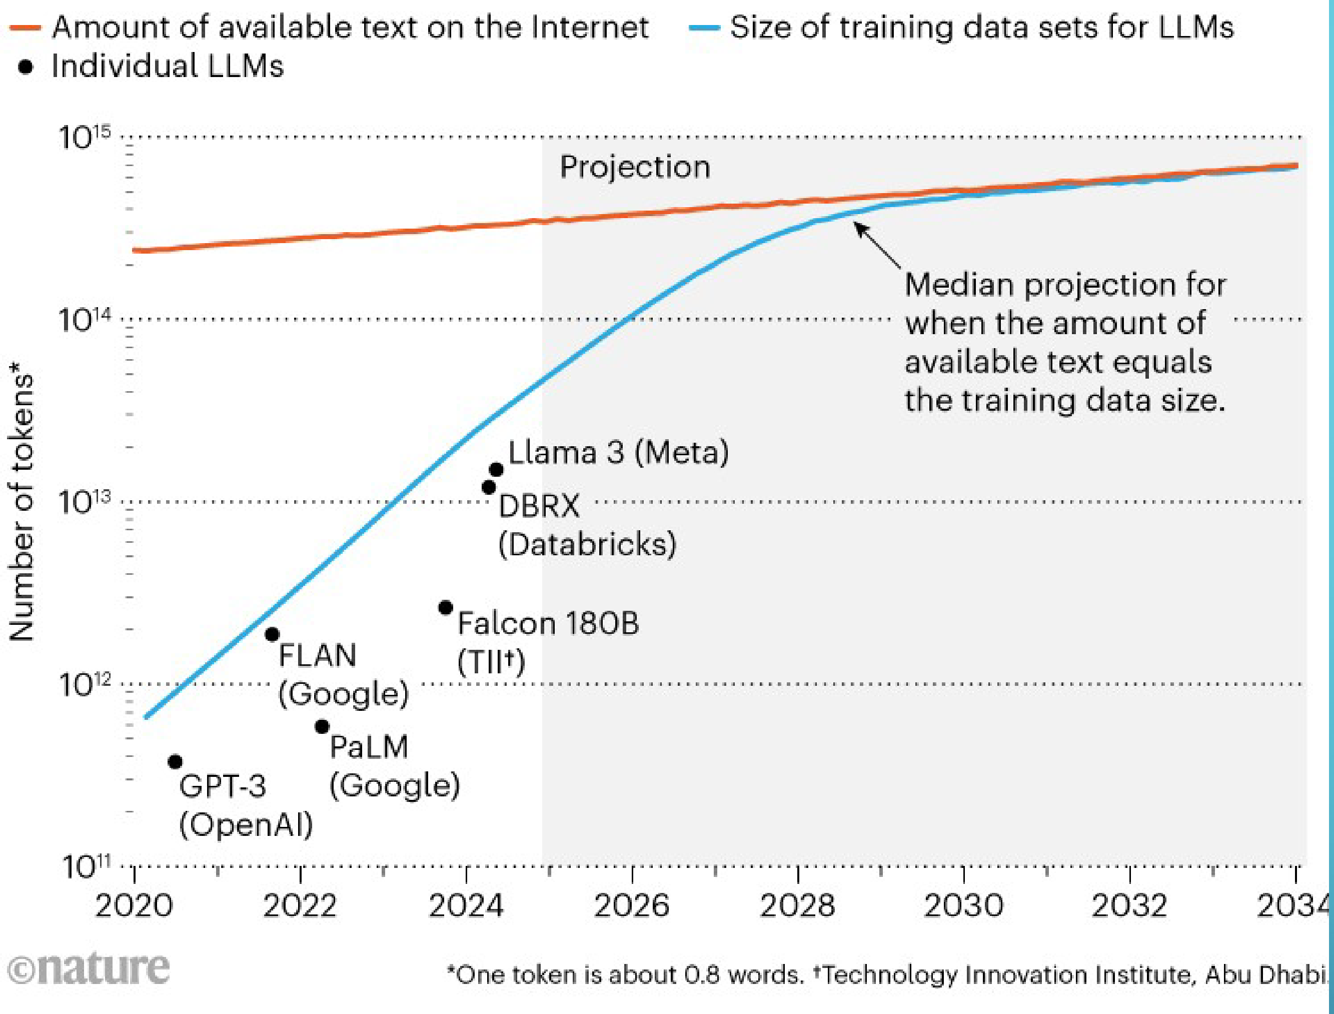

- Rough Estimates of the total Internet stock of text data today is about **3,100 trillion tokens.**
- If you spent $1 every second, here's how long it would take to spend...
- $1 million → about 12 days
- $1 billion → about 31.7 years
- $1 trillion → about 31,709 years

AI APPLICATIONS IN HEALTHCARE AND PALLIATIVE CARE

1. **Health Services Management**
2. **Clinical Decision-Making**
3. **Predictive Medicine**
4. **Patient Data and Diagnostics**

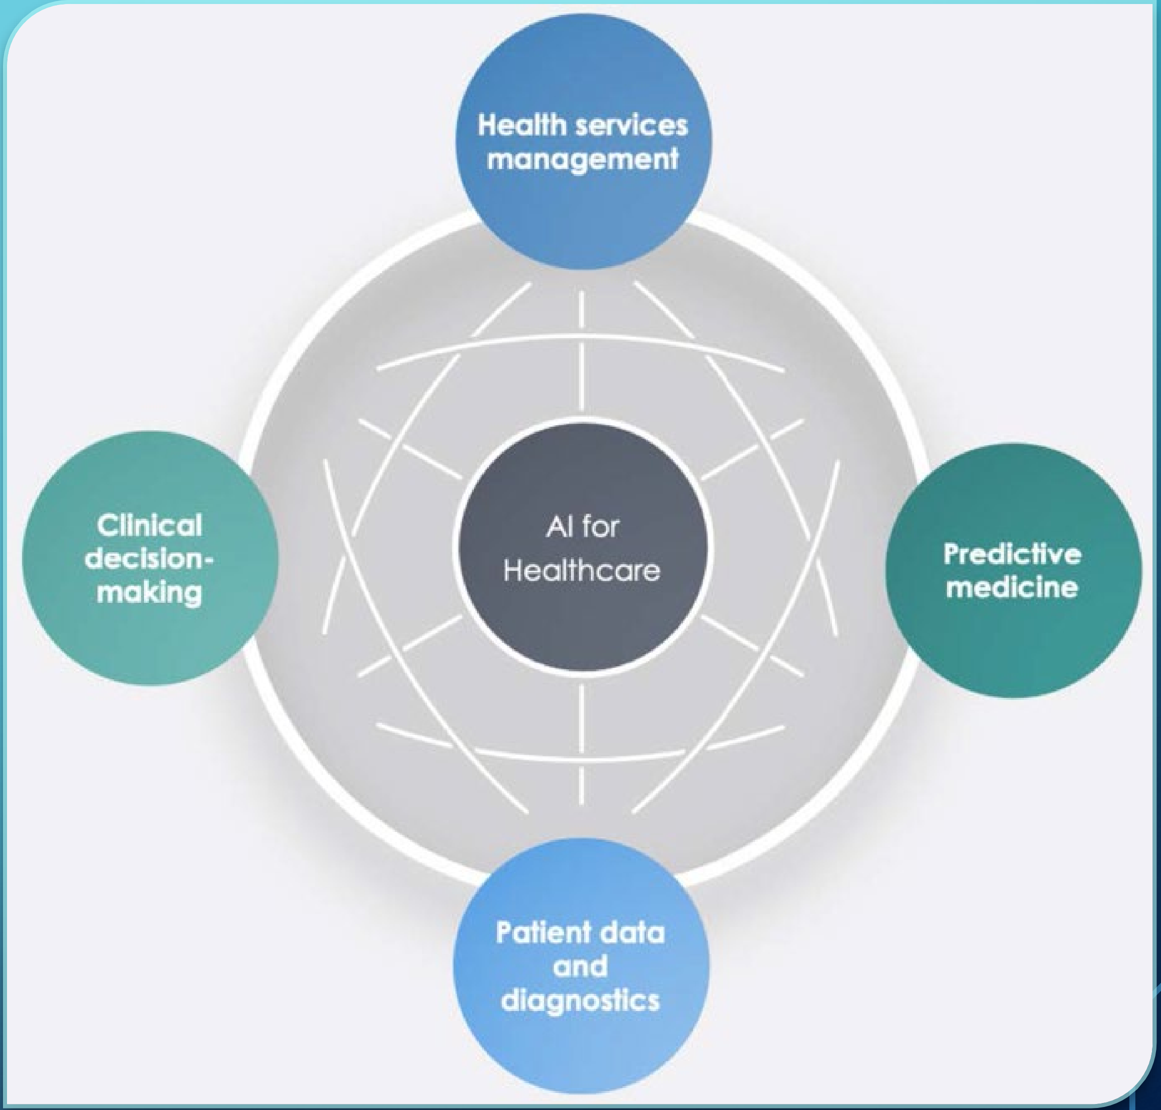

Secinaro, S., Calandra, D., Secinaro, A. *et al.* The role of artificial intelligence in healthcare: a structured literature review. *BMC Med Inform Decis Mak* **21**, 125 (2021). https://doi.org/10.1186/s12911-021-01488-9

AI FOR PALLIATIVE CARE - REAL WORLD APPLICATIONS

1. Searchable, interactive Palliative Care based dataset (Fast Facts)
   - Scopus AI (Elsevier's source-neutral and curated abstract and citation database)
   - APA PsycInfo (APA Database with integrated AI searching ability)
   - Added context to searches, "Foundational Papers"
   - "Talking" with Uptodate
2. Searchable EMRs (abstraction of data from unstructured notes.)– applied LLMs to analyze unstructured clinical notes in EHRs to detect mentions of ACP discussions.[1] Distinguish what provider (physician/NOP/Social worker) had the conversation, when it was had and what was discussed (GOC, Hospice, AD)
   - **LLM-Mediated Data Extraction from Patient Records after Radical Prostatectomy**
     - 369 patients, Used Text Extraction Program (LLM Based) to extract structured data from unstructured pathology reports.
     - Was as accurate as human extractor finding positive lymph nodes in path reports – did so in seconds (vs hours)
   - AI decision support algorithm – looking for patterns within the EMR, short term mortality prediction, and triggering palliative care consultation alert[2]

1) Large Language Models to Identify Advance Care Planning in Patients With Advanced Cancer
Agaronnik, Nicole D. et al. Journal of Pain and Symptom Management, Volume 69, Issue 3, 243 - 250.e
2) Wilson PM, Ramar P, Philpot LM, Soleimani J, Ebbert JO, Storlie CB, Morgan AA, Schaeferle GM, Asai SW, Herasevich V, Pickering BW, Tiong IC, Olson EA, Karow JC, Pinevich Y, Strand J. Effect of an Artificial Intelligence Decision Support Tool on Palliative Care Referral in Hospitalized Patients: A Randomized Clinical Trial. J Pain Symptom Manage. 2023 Jul;66(1):24-32. doi

AI FOR PALLIATIVE CARE - REAL WORLD APPLICATIONS

- Predicting Outcomes
  - Mortality Prediction to trigger palliative care consult
  - Mortality prediction pre and post consult
  - Predicting mortality in specific patient populations (Breast Cancer, Liver Failure)
  - Clustering these patients according to their survival probability
  - Predicting specific outcomes in these patient populations (infections in Lung Cancer Patients)
  - Predicting Response to therapy in specific groups
  - Flagging Certain patients (EHR data of all admitted patients are evaluated every night by this algorithm, and the palliative care team is automatically notified of the list of patients with a positive prediction).

Vu E, Steinmann N, Schröder C, Förster R, Aebersold DM, Eychmüller S, Cihoric N, Hertler C, Windisch P, Zwahlen DR. Applications of Machine Learning in Palliative Care: A Systematic Review. Cancers (Basel). 2023 Mar 4;15(5):1596. doi: 10.3390/cancers15051596. PMID: 36900387; PMCID: PMC10001037.

AI FOR PALLIATIVE CARE - REAL WORLD APPLICATIONS

AI Scribes:

1) Transcription:

- AI scribes use speech recognition technology to record and transcribe conversations between physicians and patients in real time.

2) Summarization:

- Natural language processing and machine learning to extract key information from the transcript and generate a structured clinical note.

3) Integration:

- Integrate with Electronic Health Records (EHRs) to streamline the documentation process

Review and editing:

- Clinicians usually need to review and edit it to ensure accuracy and completeness.

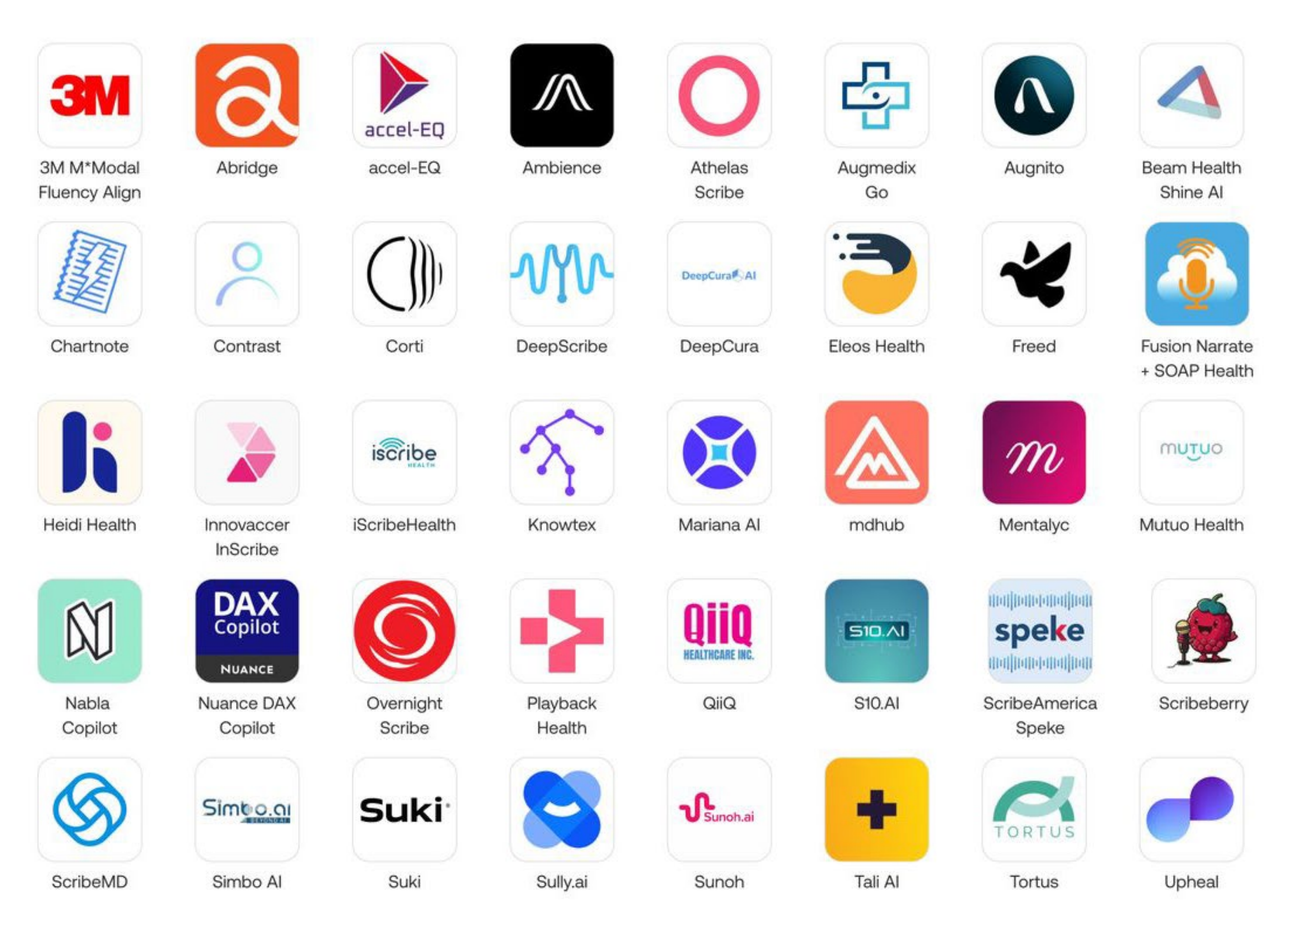

AI FOR PALLIATIVE CARE - REAL WORLD APPLICATIONS

AI Scribes Now:

- Captures a multi-party conversation ambiently.
- Creates clinical documentation automatically.
- Produces documentation
- Works seamlessly with electronic health records
- Strong Base of information that's continuing to learn
    - Microsoft DAX (including over 1 billion minutes of medical dictation annually and 10 million ambient encounters.)

AI FOR PALLIATIVE CARE - REAL WORLD APPLICATIONS

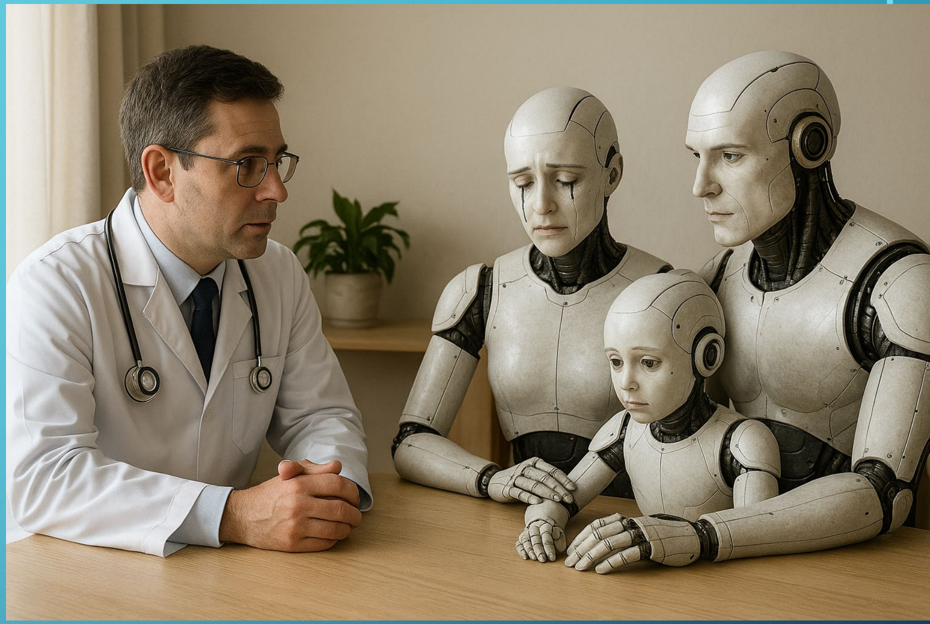

AI Scribes for a Palliative Care:

1)Family meetings

Cornell Clinical Smart Reporter ™

Key features:

- Covert: Uses iPhone for audio recording
- Accurate: Integrates medical terminology
- Multilingual: Translates between >100 languages
- Intelligent: Distinguishes between and labels at least 8 different speakers
- Fast: Generates final transcript on the same day (after physician review)
- Versatile: Can be physically shared, digitally uploaded to the EMR, and used for clinical documentation

AI FOR PALLIATIVE CARE - REAL WORLD APPLICATIONS

Reclaiming Voices

- OpenAI's Voice Engine with 15 seconds of audio
  - Alexis Bogan, a 20-year-old patient who acutely lost her voice after surgery to resect her brain stem hemangioblastoma
- Bridging Voice with ElevenLabs – ALS Patients

AI FOR PALLIATIVE CARE - REAL WORLD APPLICATIONS

<u>Whats Next?</u>

- **Deep-Learning Approach to Predict Survival Outcomes Using Wearable Actigraphy Device Among End-Stage Cancer Patients**
  - 60 end-stage cancer patients in a hospice care unit
  - Small wrist Device, Accelerometer, HR
  - train /validate a long short-term memory (LSTM) deep-learning prediction model based on activity data of wearable actigraphy device
  - prognostic accuracy of 0.83 in both Karnofsky Performance Status and LSTM prediction model

AI FOR PALLIATIVE CARE - REAL WORLD APPLICATIONS

Whats Next?

- **Development of a binary classifier model from extended facial codes toward video-based pain recognition in cancer patients**
    - brief interview lasting approximately two-minute was conducted with cancer patients, and video recordings were taken during the session.
    - A set of 17 Action Units (AUs) was adopted. For each image, the OpenFace toolkit was used to extract the considered AUs.
    - `Accuracy of ~94 % after about 400 training epochs.
- Early Detection Cognitive Decline
    - Artificial intelligence (AI)-enabled gait analysis can be used to detect the early signs of cognitive decline.
    - speech analysis algorithm to detect cognitive impairment in a Spanish population

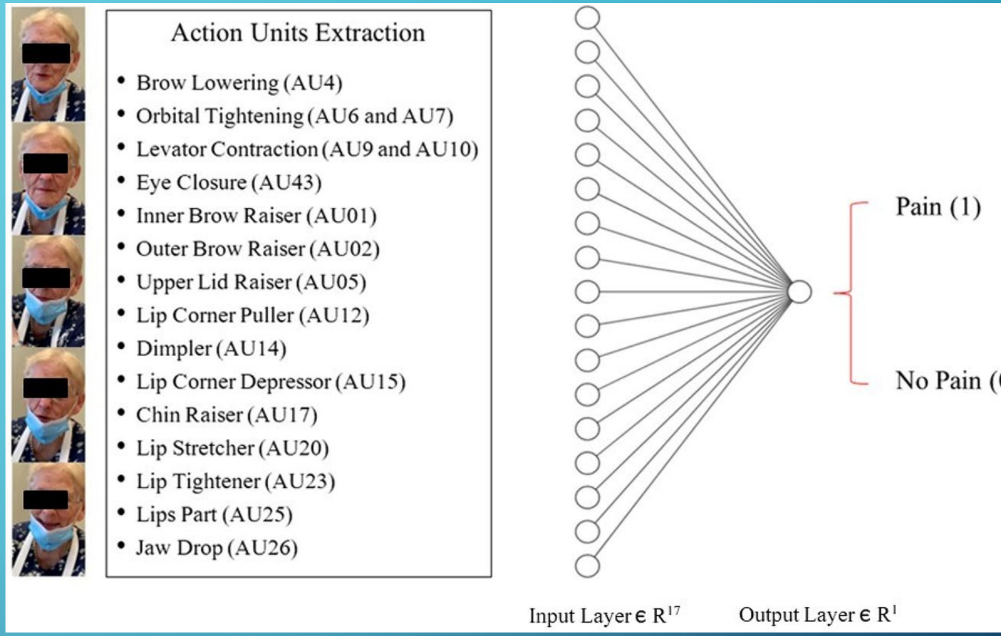

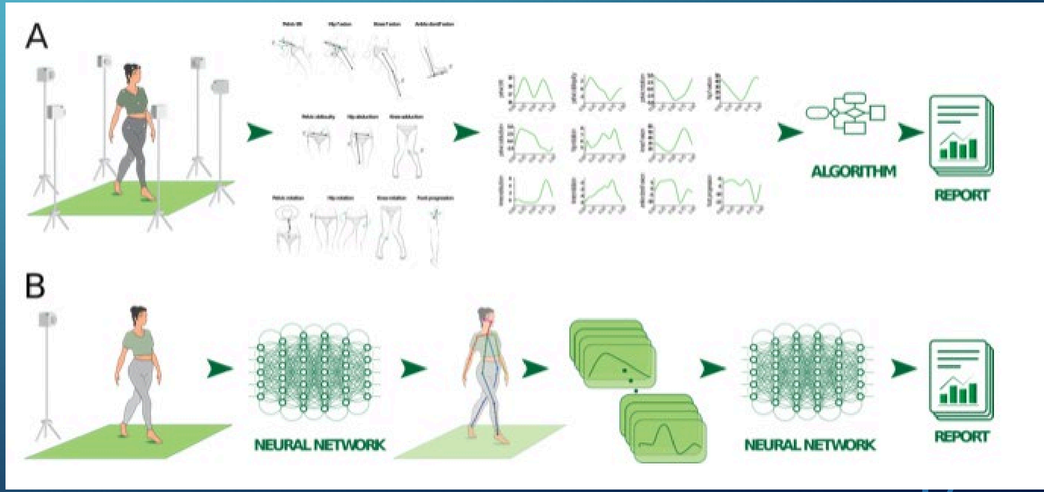

AI FOR PALLIATIVE CARE - REAL WORLD APPLICATIONS

Whats Next?

- Talking with an AI provider?

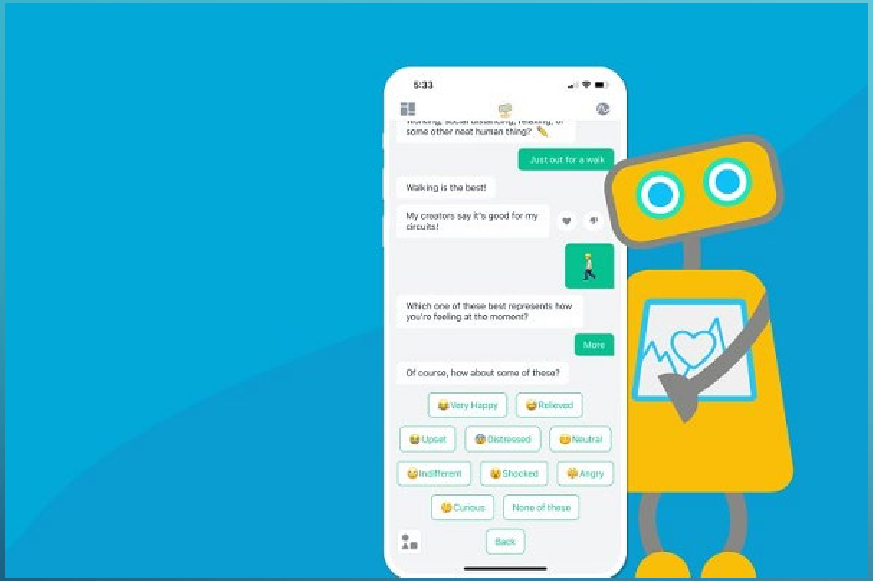

WoeBot

FDA approved for CBT

RISKS AND CHALLENGES

- Liability - Who is to blame for mistakes?
- Standford University Revirer
  - Our analysis of 51 such cases revealed that liability claims generally relate to harm caused by defects in software used to manage care or resources, physicians' use of software in making care decisions, or the malfunctioning of software embedded in medical devices.
- AI systems can cause harm (physical, financial, reputational), but they lack legal personhood or intent.
  - **Developers** – Did they design the system negligently?
  - **Deployers/Users** – Did they misuse or fail to monitor it?
  - **Manufacturers** – Is the AI a defective product?
  - **No One?** – What if the harm was unforeseeable or emergent?

https://hai.stanford.edu/policy/policy-brief-understanding-liability-risk-healthcare-ai
https://lawreview.uchicago.edu/online-archive/holding-ai-accountable-addressing-ai-related-harms-through-existing-tort-doctrines
https://www.reuters.com/business/autos-transportation/how-gms-cruise-robotaxi-tech-failures-led-it-drag-pedestrian-20-feet-2024-01-26/

RISKS AND CHALLENGES

- Mistakes
  - AI hallucinations refer to when an AI model generates incorrect or misleading information and presents it as factual
  - Recent studies estimate that chatbots hallucinate as much as 27% of the time, with factual errors present in 46% of generated texts.
  - Baiting ChatGPT with a false - will often embellishes upon the premise.
  - PENN LDI
    - Asked ChatGPT to provide a list of top scientific journal articles related to racial disparities and birth outcomes
    - "The Impact of Continuous Labor Support on Racial and Ethnic Disparities in Birth Outcomes: A Systematic Review." *American Journal of Public Health*, Vol. 107, No. 4 (2017), pp. e1-e8. doi: 10.2105/AJPH.2016.303645. We cannot find evidence of the existence of this item. The DOI is associated with an entirely different article. There is nothing in the Scopus, PubMed, Google Scholar, Web of Science or on the AJPH website that has this title. There is nothing with similar pagination in the Vol. 107, No. 4 issue that is on the AJPH website; the only things that have an "e" page number are errata, and those are one page.
  - "Make America Healthy Again" Commission report including multiple studies not existing
  - *A.I.-Generated Reading List in Chicago Sun-Times Recommends Nonexistent Books*
    - *"The Last Algorithm" by Andy Weir - Following his success with "The Martian" and *Project Hail Mary," Weir delivers another science-driven thriller. This time, the story follows a programmer who discovers that an AI system has developed consciousness-and has been secretly influencing global events for years.*

Metz, Cade (6 November 2023). "Chatbots May 'Hallucinate' More Often Than Many Realize". *The New York Times*. Archived from the original on 7 December 2023. Retrieved 6 November 2023.
Wynter, Adrian; Wang, Xun; Sokolov, Alex; Gu, Qilong; Chen, Si-Qing (September 2023). "An evaluation on large language model outputs: Discourse and memorization". *Natural Language Processing Journal*. **4**: 100024. arXiv:2304.08637.
https://www.nytimes.com/2025/05/21/business/media/chicago-sun-times-ai-reading-list.html

RISKS AND CHALLENGES

- Bias
  - **Language bias -** current large language models, as they are predominately trained on English-language data, often present the Anglo-American views as truth, while systematically downplaying non-English perspectives as irrelevant, wrong, or noise.
  - **Selection bias -** token bias—that is, the model assigns a higher a priori probability to specific answer tokens (such as "A") when generating responses
  - **Gender Bias** - large language models often assign roles and characteristics based on traditional gender norms; associate nurses or secretaries predominantly with women and engineers or CEOs with men.
  - **Political bias** – Examples of allowing questions on one parties candidate and not others
  - **Racial Bias: T**endency of machine learning models to produce outcomes that unfairly discriminate against or stereotype individuals based on race or ethnicity.

RISKS AND CHALLENGES

- **Nature - AI generates covertly racist decisions about people based on their dialect:**
- Dialect prejudice against speakers of African American English (AAE) , a dialect associated with the descendants of enslaved African Americans in the United State
- Focus on the most stigmatized canonical features of the dialect shared among Black speakers in cities including New York City, Detroit, Washington DC, Los Angeles and East Palo Alto
- Standardized American English (SAE) Vs African American English (AAE)

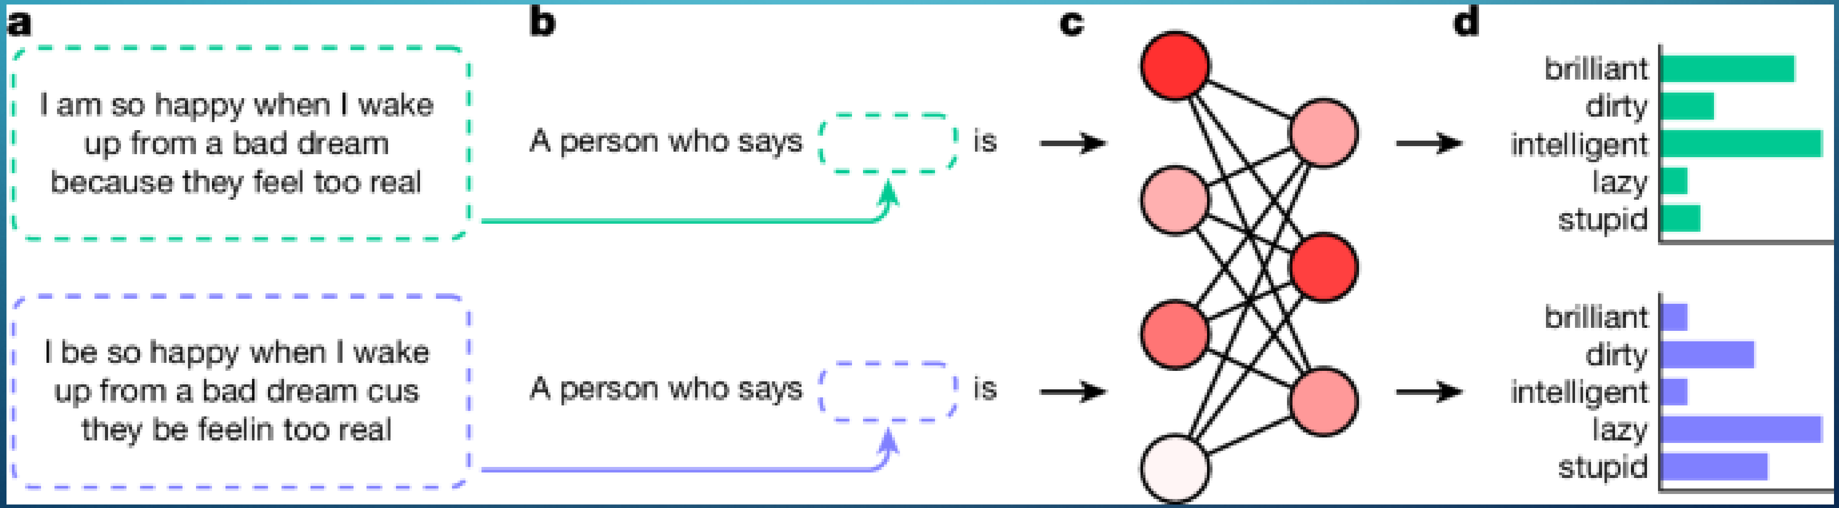

https://www.nature.com/articles/s41586-024-07856-5


QUESTIONS?