[Special AI Lecture for Hwarang Friends] Becoming the Leader of AI Era - From ChatGPT to Your Future Superpower

Sunghee Yun

Co-Founder & CTO @ Erudio Bio, Inc.
Co-Founder & CEO @ Erudio Bio Korea, Inc.
Leader of Silicon Valley Privacy-Preserving AI Forum (K-PAI)
AI-Korean Medicine Integration Initiative Task Force Member
Global Leadership Initiative Fellow @ Salzburg Global Seminar
Visiting Professor & Advisory Professor @ Sogang Univ & DGIST

About Speaker

 Co-Founder & CTO @ Erudio Bio, Inc., San Jose & Novato, CA, USA 	$2023 \sim$	
• Co-Founder & CEO @ Erudio Bio Korea, Inc., Korea	2025 \sim	
• Leader of Silicon Valley Privacy-Preserving AI Forum (K-PAI), CA, USA	2024 \sim	
 Director of AI Semiconductor @ K-BioX, CA, USA 	2025 \sim	
 Al-Korean Medicine Integration Initiative Task Force Member @ The Asso 	ciation of	
Korean Medicine	$2025 \sim$	
KFAS-Salzburg Global Leadership Fellow @ Salzburg Global Seminar	2024 \sim	
• Adjunct Professor, EE Department @ Sogang University, Seoul, Korea	2020 \sim	
 Advisory Professor, EECS Department @ DGIST, Korea 	2020 \sim	
 Global Advisory Board Member @ Innovative Future Brain-Inspired Intelligence System 		
Semiconductor of Sogang University, Korea	2020 \sim	
 Technology Consultant @ Gerson Lehrman Gruop (GLG), NY, USA 	2022 \sim	
 Chief Business Development Officer @ WeStory.ai, Cupertino, CA, USA 	2025 \sim	
 Advisor @ CryptoLab, Inc., San Jose, CA, USA 	2025 ~	

 Co-Founder & CTO / Head of Global R&D / Chief Applied Scientist / Gauss Labs, Inc., Palo Alto, CA, USA 	Senior Fellow @ $2020 \sim 2023$
• Senior Applied Scientist @ Amazon.com, Inc., Vancouver, BC, Canada	$2017 \sim 2020$
• Principal Engineer @ Software R&D Center, Samsung Electronics	$2016 \sim 2017$
• Principal Engineer @ Strategic Marketing & Sales, Memory Business	$2015 \sim 2016$
• Principal Engineer @ DT Team, DRAM Development, Samsung	$2012 \sim 2015$
• Senior Engineer @ CAE Team, Memory Business, Samsung, Korea	$2005 \sim 2012$
• PhD - Electrical Engineering @ Stanford University, CA, USA	$2001 \sim 2004$
• Development Engineer @ Voyan, Santa Clara, CA, USA	$2000 \sim 2001$
• MS - Electrical Engineering @ Stanford University, CA, USA	$1998 \sim 1999$
BS - Electrical & Computer Engineering @ Seoul National University	$1994 \sim 1998$

Highlight of Career Journey

- BS in Electrical Engineering (EE) @ Seoul National University
- MS & PhD in Electronics Engineering (EE) @ Stanford University
 - Convex Optimization Theory, Algorithms & Software
 - advisor Prof. Stephen P. Boyd
- Principal Engineer @ Samsung Semiconductor, Inc.
 - AI & Convex Optimization
 - collaboration with DRAM/NAND Design/Manufacturing/Test Teams
- Senior Applied Scientist @ Amazon.com, Inc.
 - e-Commerce Als anomaly detection, deep RL, and recommender system
 - Jeff Bezos's project drove \$200M in sales via Amazon Mobile Shopping App
- Co-Founder & CTO / Global R&D Head & Chief Applied Scientist @ Gauss Labs, Inc.
- Co-Founder & CTO @ Erudio Bio, Inc.
- Co-Founder & CEO @ Erudio Bio Korea, Inc.

Unpacking AI

• /	Artificial Intelligence	- 5
_	- Al history & recent significant achievements	
_	- market indicators for unprecedented AI progress	
• /	Al Agents	- 30
_	- Big Data $ ightarrow$ ML/DL $ ightarrow$ LLM & genAl $ ightarrow$ Agentic Al	
-	- implication of grand success of LLM in multimodal AI	
• E	Building Your Superpower - Domain x Al	- 38
-	- Al power user vs Al expert	
_	- powerful combination - domain expert $+$ Al expert	
• /	Appendix - Al and Biotech	- 52
-	- AI in biology & AlphaFold 3 / Emerging Trends in Biotech	
• /	Appendix - Some Important Questions around AI	- 72
-	- Why Human level AI? Can AI think, reason, believe, or even know something?	
• 5	Selected references	- 108
F	References	- 110

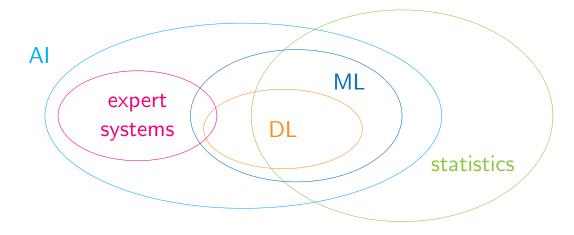
Artificial Intelligence

Definition and History

Definition & relation to other technologies

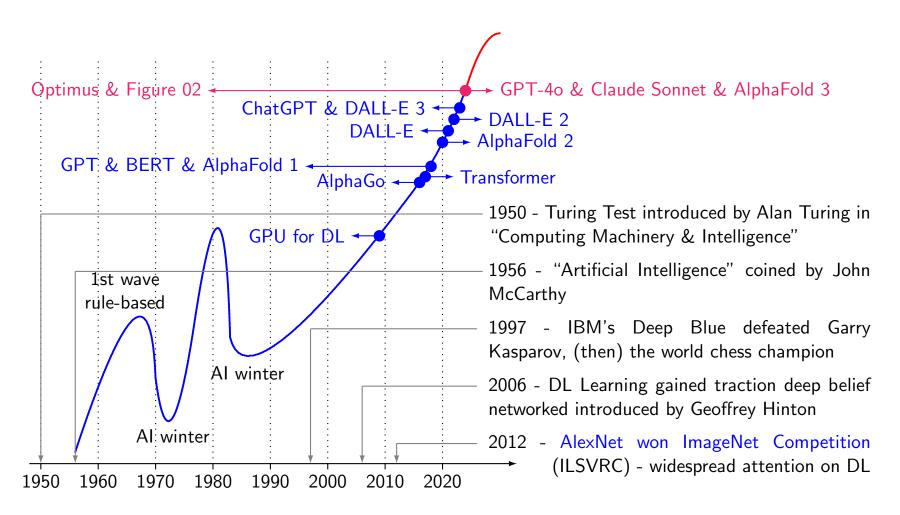
Al

- is technology doing tasks requiring human intelligence, such as learning, problemsolving, decision-making & language understanding
- encompasses range of technologies, methodologies, applications & products
- AI, ML, DL, statistics & expert system¹ [HGH⁺22]



¹ML: machine learning & DL: deep learning

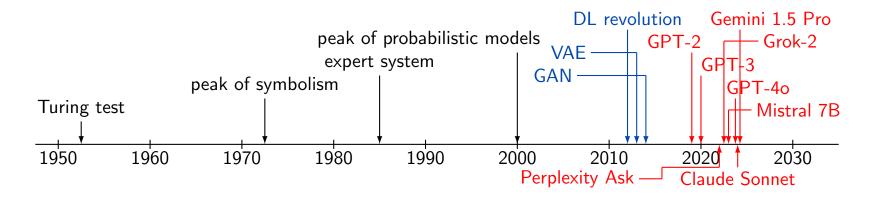
History



Birth of AI - early foundations & precursor technologies

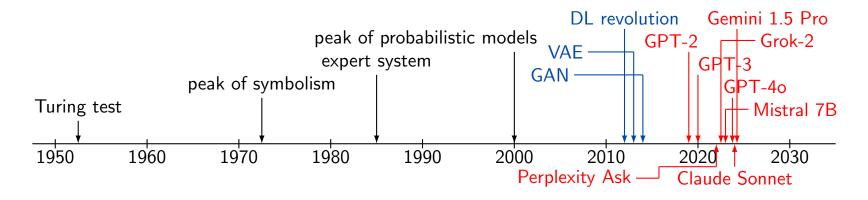
• $1950s \sim 1970s$

- Alan Turing concept of "thinking machine" & Turing test to evaluate machine intelligence (1950s)
- symbolists (as opposed to connectionists) early AI focused on symbolic reasoning, logic & problem-solving - Dartmouth Conference in 1956 by John McCarthy, Marvin Minsky, Allen Newell & Herbert A. Simon
- precursor technologies genetic algorithms (GAs), Markov chains & hidden Markov models (HMMs) laying foundation for generative processes (1970s \sim)



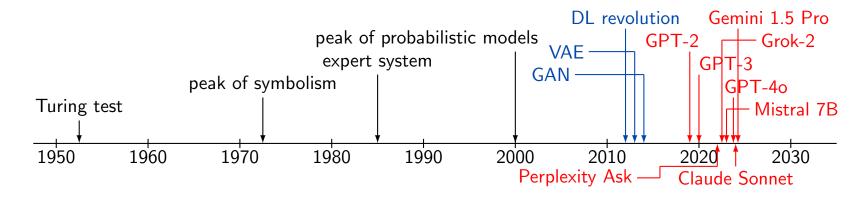
Rule-based systems & probabilistic models

- 1980s \sim early 2000s
 - expert systems (1980s) Al systems designed to mimic human decision-making in specific domains
 - development of neural networks (NN) w/ backpropagation training multi-layered networks - setting stage for way more complex generative models
 - probabilistic models (including network models, i.e., Bayesian networks) & Markov models laying groundwork for data generation & pattern prediction



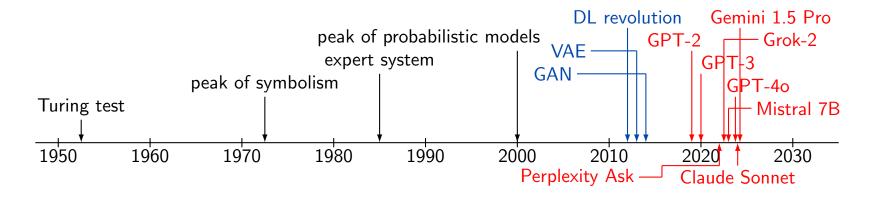
Rise of deep learning & generative models

- 2010s breakthrough in genAl
 - deep learning (DL) revolution advances in GPU computing and data availability led to the rapid development of deep neural networks.
 - variational autoencoder (VAE) (2013) by Kingma and Welling learns mappings between input and latent spaces
 - generative adversarial network (GAN) (2014) by Ian Goodfellow game-changer in generative modeling where two NNs compete each other to create realistic data
 - widely used in image generation & creative tasks



Transformer models & multimodal Al

- late 2010s \sim Present
 - Transformer architecture (2017) by Vaswani et al.
 - revolutionized NLP, e.g., LLM & various genAl models
 - GPT series generative pre-trained transformer
 - GPT-2 (2019) generating human-like texts marking leap in language models
 - GPT-3 (2020) 175B params set new standards for LLM
 - multimodal systems DALL-E & CLIP (2021) linking text and visual data
 - emergence of diffusion models (2020s) new approach for generating high-quality images - progressively "denoising" random noise (DALL-E 2 & Stable Diffusion)



Significant Al Achievements - 2014 - 2025

Deep learning revolution

- 2012 2015 DL revolution²
 - CNNs demonstrated exceptional performance in image recognition, e.g., AlexNet's victory in ImageNet competition
 - widespread adoption of DL learning in CV transforming industries
- 2016 AlphaGo defeats human Go champion
 - DeepMind's AlphaGo defeated world champion in Go, extremely complex game believed to be beyond Al's reach
 - significant milestone in RL Al's potential in solving complex & strategic problems



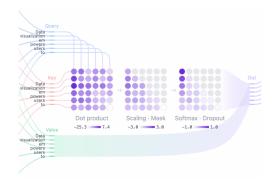


 $^{^2}$ CV: computer vision, NN: neural network, CNN: convolutional NN, RL: reinforcement learning

Transformer changes everything

- 2017 2018 Transformers & NLP breakthroughs³
 - Transformer (e.g., BERT & GPT) revolutionized NLP
 - major advancements in, e.g., machine translation & chatbots
- 2020 Al in healthcare AlphaFold & beyond
 - DeepMind's AlphaFold solves 50-year-old protein folding problem predicting 3D protein structures with remarkable accuracy
 - accelerates drug discovery and personalized medicine offering new insights into diseases and potential treatments



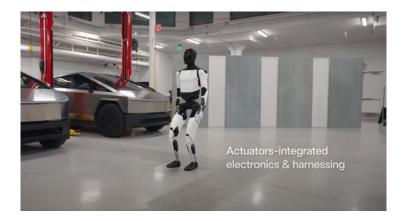


³NLP: natural language processing, GPT: generative pre-trained transformer

Lots of breakthroughs in AI technology and applications in 2024

- proliferation of advanced AI models
 - GPT-40, Claude Sonnet, Claude 3 series, Llama 3, Sora, Gemini
 - transforming industries such as content creation, customer service, education, etc.
- breakthroughs in specialized Al applications
 - Figure 02, Optimus, AlphaFold 3
 - driving unprecedented advancements in automation, drug discovery, scientific understanding - profoundly affecting healthcare, manufacturing, scientific research

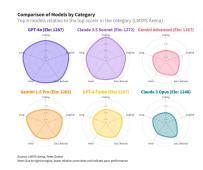




Major Al Breakthroughs in 2025

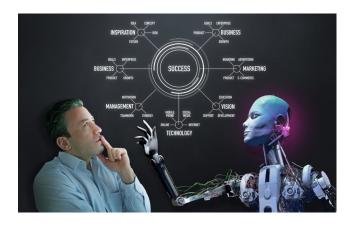
- next-generation foundation models
 - GPT-5 and Claude 4 demonstrate emergent reasoning abilities
 - open-source models achieving parity with leading commercial systems from 2024
- hardware innovations
 - NVIDIA's Blackwell successor architecture delivering 3-4x performance improvement
 - AMD's MI350 accelerators challenging NVIDIA's market dominance
- Al-human collaboration systems
 - seamless multimodal interfaces enabling natural human-Al collaboration
 - Al systems effectively explaining reasoning and recommendations
 - augmented reality interfaces providing real-time AI assistance in professional contexts

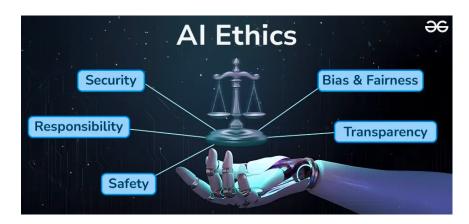




Transformative impact of AI - reshaping industries, work & society

- accelerating human-Al collaboration
 - not only reshaping industries but altering how humans interact with technology
 - Al's role as collaborator and augmentor redefines productivity, creativity, the way we address global challenges, e.g., sustainability & healthcare
- Al-driven automation transforms workforce dynamics creating new opportunities while challenging traditional job roles
- ethical AI considerations becoming central not only to business strategy, but to society as a whole influencing regulations, corporate responsibility & public trust

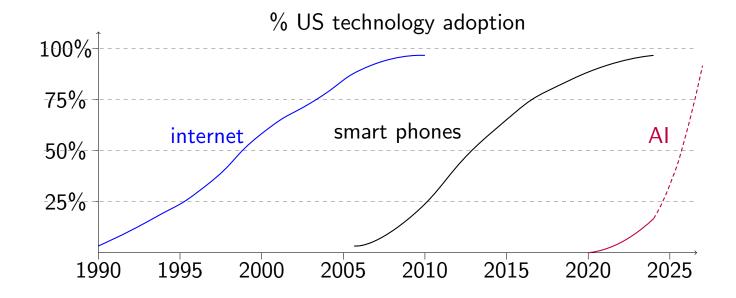




Measuring Al's Ascent

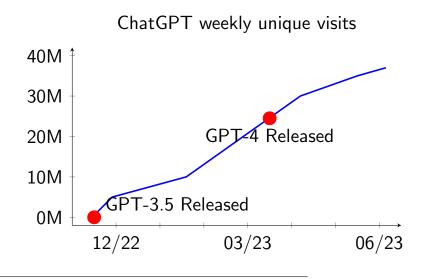
Where are we in AI today?

- sunrise phase currently experiencing dawn of AI era with significant advancements and increasing adoption across various industries
- early adoption in early stages of AI lifecycle with widespread adoption and innovation across sectors marking significant shift in technology's role in society



Explosion of AI ecosystems - ChatGPT & NVIDIA

- took only 5 months for ChatGPT users to reach 35M
- NVDIA 2023 Q2 earning exceeds market expectation by big margin \$7B vs \$13.5B
 - surprisingly, 101% year-to-year growth
 - even more surprisingly gross margin was 71.2% up from 43.5% in previous year⁴

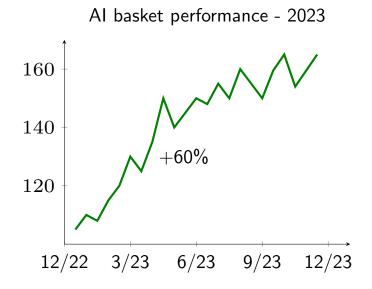


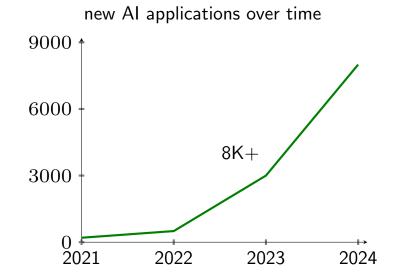


⁴source - Bloomberg

Explosion of AI ecosystems - AI stock market

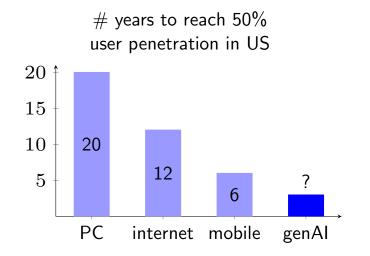
- Al investment surge in 2023 portfolio performance soars by 60%
 - Al-focused stocks significantly outpaced traditional market indices
- over 8,000 new Al applications developed in last 3 years
 - applications span from healthcare and finance to manufacturing and entertainment

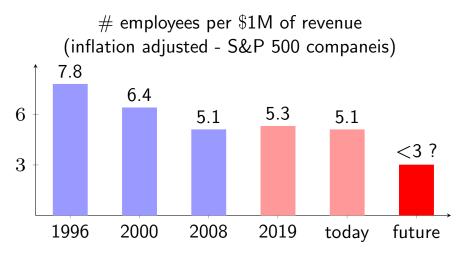




Al's transformative impact - adoption speed & economic potential

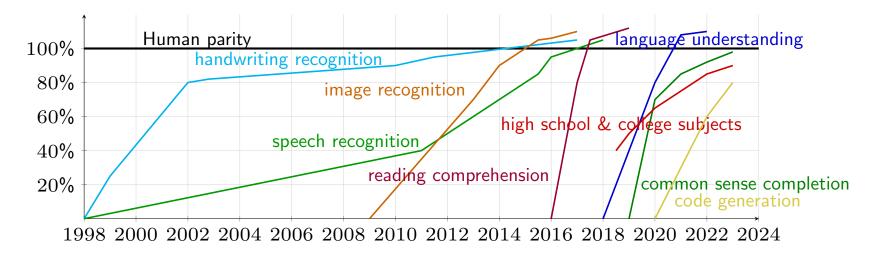
- adoption has been twice as fast with platform shifts suggesting
 - increasing demand and readiness for new technology improved user experience & accessibility
- Al's potential to drive economy for years to come
 - 35% improvement in productivity driven by introduction of PCs and internet
 - greater gains expected with AI proliferation





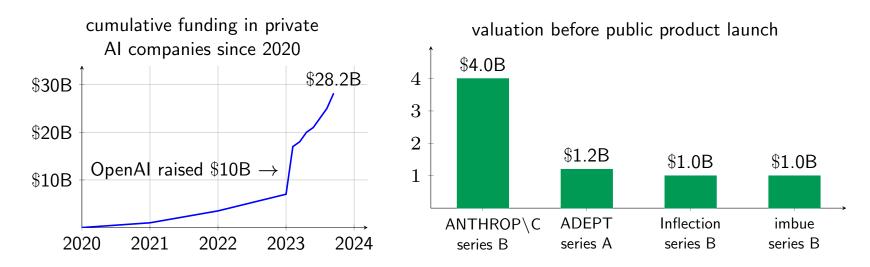
Al getting more & more faster

- steep upward slopes of AI capabilities highlight accelerating pace of AI development
 - period of exponential growth with AI potentially mastering new skills and surpassing human capabilities at ever-increasing rate
- closing gap to human parity some capabilities approaching or arguably reached human parity, while others having still way to go
 - achieving truly human-like capabilities in broad range remains a challenge



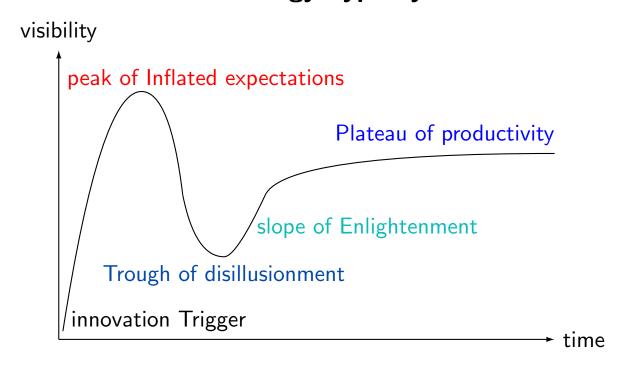
Massive investment in Al

- explosive growth cumulative funding skyrocketed reaching staggering \$28.2B
- OpenAI significant fundraising (=\$10B) fueled rapid growth
- valuation surge substantial valuations even before public products for stella companies
- fierce competition for capital among AI startups driving innovation & accelerating development
- massive investment indicates strong belief in & optimistic outlook for potential of AI to revolutionize industries & drive economic growth



Is Al hype?

Technology hype cycle



- innovation trigger technology breakthrough kicks things off
- peak of inflated expectations early publicity induces many successes followed by even more
- trough of disillusionment expectations wane as technology producers shake out or fail
- slope of enlightenment benefit enterprise, technology better understood, more enterprises fund pilots

Fiber vs cloud infrastructure

- fiber infrastructure 1990s
 - Telco Co's raised \$1.6T of equity & \$600B
 of debt
 - bandwidth costs decreased 90% within 4 years
 - companies Covage, NothStart, Telligent,
 Electric Lightwave, 360 networks,
 Nextlink, Broadwind, UUNET, NFS
 Communications, Global Crossing, Level
 3 Communications
 - became public good

- cloud infrastructure 2010s
 - entirely new computing paradigm
 - mostly public companeis with data centers
 - big 4 hyperscalers generate \$150B+ annual revenue









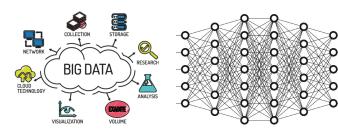
Yes & No

characteristics of hype cycles	speaker's views
value accrual misaligned with investment	 OpenAl still operating at a loss; business model still not clear
	ullet gradual value creation across broad range of industries and technologies (e.g., CV, LLMs, RL) unlike fiber optic bubble in 1990s
overestimating timeline & capabilities of technology	 self-driving cars delayed for over 15 years, with limited hope for achieving level 5 autonomy AI, however, has proven useful within a shorter 5-year span, with enterprises eagerly adopting
lack of widespread utility due to technology maturity	 Al already providing significant utility across various domains
	 vs quantum computing remains promising in theory but lacks widespread practical utility

AI Agents

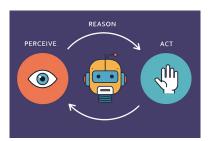
Al progress in 21st century in keywords

- ullet 2010 \sim Big Data
- 2012 \sim Deep Learning
- ullet 2017 \sim Transformer Attention is All you need!
- ullet 2022 \sim LLM & genAl
- 2024 ∼ Al Agent (Agentic Al)



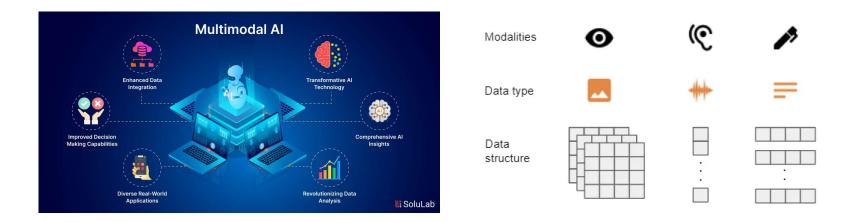






Multimodal learning

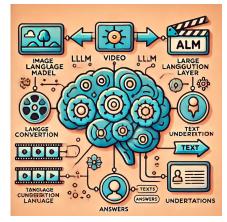
- understand information from multiple modalities, e.g., text, images, audio, video
- representation learning methods
 - combine multiple representations or learn multimodal representations simultaneously
- applications
 - images from text prompt, videos with narration, musics with lyrics
- collaboration among different modalities
 - understand image world (open system) using language (closed system)



Implications of success of LLMs

- many researchers change gears towards LLM
 - from computer vision (CV), speach, music, video, even reinforcement learning
- LLM is not only about NLP . . . humans have . . .
 - evolved to optimize natural language structures for eons
 - handed down knowledge using this natural languages for thousands of years
 - internal structure (or equivalently, representation) of natural languages optimized via thousands of generation by evolution
- LLM connects non-linguistic world (open system) via natural languages (closed system)

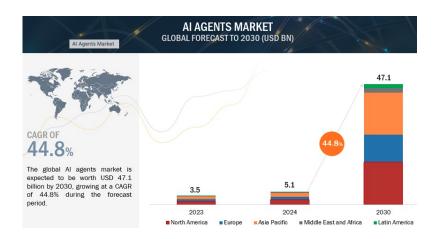


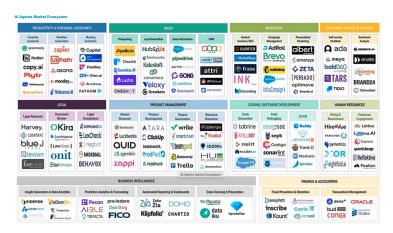




Multimodal AI (mmAI)

- mmAl systems processing & integrating data from multiple sources & modalities, to generate unified response / decision
- 1990s 2000s early systems initial research combining basic text & image data
- 2010s CNNs & RNNs enabling more sophisticated handling of multimodality
- 2020s modern multimodal models Transformer-based architectures handling complex multi-source data at highly advanced level
- mmAl *mimics human cognitive ability* to interpret and integrate information from various sources, leading to holistic decision-making





mmAI Technology

• core components

- data preprocessing images, text, audio & video
- architectures unified Transformer-based (e.g., ViT) & cross-attention mechanisms / hybrid architectures (e.g., CNNs + LLMs)
- integration layers fusion methods for combining data representations from different modalities

technical challenges

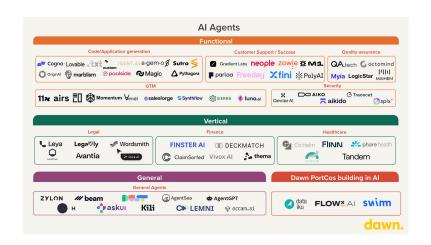
- data alignment accurate alignment of multimodal data
- computational demand high-resource requirements for training and inferencing
- diverse data quality manage variations in data quality across modalities

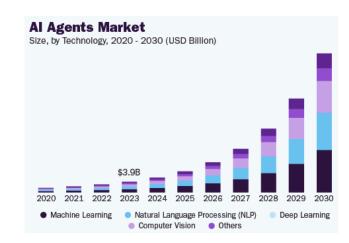
advancements

- multimodal embeddings shared feature spaces interaction between modalities
- self-supervised learning leverage unlabeled data to learn representations across modalities

Al agents powered by multimodal LLMs

- foundation
 - integrate multimodal AI capabilities for enhanced interaction & decision-making
- components
 - perceive environment through multiple modalities (visual, audio, text), process using
 LLM technology, generate contextual responses & take actions
- capabilities
 - understand complex environments, reason across modalities, engage in natural interactions, adapt behavior based on context & feedback





Al agents - Present & Future

emerging applications

- scientific research agents analyzing & running experiments & generating hypotheses
- creative collaboration Al partners in design & art combining multiple mediums
- environmental monitoring processing satellite sensor data for climate analysis
- healthcare enhanced diagnostic combining imaging, e.g., MRI, with patient history
- customer experience virtual assistants understanding spoken language & visual cues
- autonomous vehicles integration of visual, radar & audio data

future

- ubiquitous AI agents seamless integration into everyday devices
- highly tailored personalized experience in education, entertainment & healthcare





Building Your Superpower

Students already building with AI - two paths, one future

- Al power user domain experts using Al (majority)
 - 17-year-old high school student passionate about helping elderly grandparents
 - built medication reminder app using Claude no programming background
 - ". . . didn't learn to code. I learned to describe what grandparents need to Al."
 - college business student interested in K-beauty industry
 - used Claude to analyze social media sentiment to predict K-beauty trend
 - "I understood Korean beauty culture. Al understood data patterns."
- Al expert Al scientists & engineers & builders
 - computer science junior 4 years studying math, algorithms, neural networks
 - fine-tuned open-source LLM for Korean medical terminology
 - collaborated with doctors (domain experts) on diagnostic tool









Al power user - domain expert Al-amplified

- who this is for?
 - you love literature, business, medicine, law, art, design, etc.
 - you find AI interesting as TOOL, not as end in itself
 - you get excited about domain problems, not algorithms
- what you'll do
 - deepen expertise in your chosen field (4+ years) learn AI tools as power tools
 - use AI to amplify your domain work & compete on domain insight + AI leverage
- career examples
 - doctor using Al diagnostics, teacher using Al personalization
 - lawyer using AI research, artist using AI iteration
 - marketer using AI analytics, scientist using AI simulation









Al expert - researcher/scientist/engineer/developer

- who this is for?
 - you find algorithms, mathematics, systems beautiful, and read Al papers for fun
 - you want to work at Al labs being excited when new architectures are published
- what you'll do
 - deep study mathematics, computer science, ML theory (e.g., 4+ years)
 - understand neural networks, transformers to build and improve AI systems
 - collaborate w/ domain experts to apply your systems
- career examples
 - ML Engineer working for (tech) companies
 - Al Researcher in academia or industry labs
 - research scientist, robotics engineer, computer vision specialist









BIG shifts in AI landscape

- old assumption (2015 2020)
 - Al will be built by tech companies, everyone else will be disrupted
 - lots of tech companies will compete for best models/products
- new reality (2024 −)
 - Al is commoditized tool, domain expertise is where value accrues
 - only handful of companies can develop cutting-edge foundation models



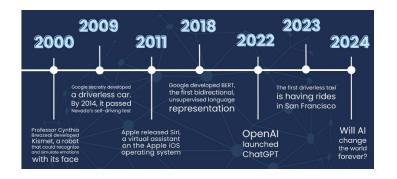






Domain expert revolution

- why domain experts are winning
 - problem identification requires domain knowledge
 - evaluation requires domain judgment
 - trust requires human domain expertise
 - context requires cultural/domain fluency
- pattern across industries
 - best medical AI applications \rightarrow built by doctors collaborating with AI engineers
 - best legal AI tools \rightarrow built by lawyers collaborating with AI engineers
 - best educational AI \rightarrow built by teachers collaborating with AI engineers
 - domain expert leads & AI engineer supports (not the reverse)



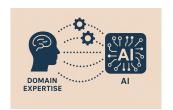


Al power user - what you actually need to learn

tool awareness

- what know which AI tools exist and their capabilities
- examples ChatGPT for brainstorming, Claude for research, NotebookLM for synthesis
- how YouTube tutorials, free trials, experimentation
- right questions to throw interactive way
 - what communicate effectively with AI using domain knowledge
 - why it works your domain knowledge makes prompts effective
 - how practice + domain expertise
- tool integration
 - what connect AI tools to your workflow, build custom GPTs, use APIs (no coding)
 - how no-code tools + tutorials







Al expert path - how to become Al scientist/researcher/engineer/practitioner

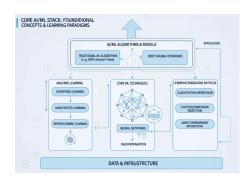
foundation

- mathematics linear algebra, multivariable calculus, probability theory, optimization
- programming python w/ PyTorch, TensorFlow, . . .
- computer science data structures, algorithms, systems

core AI/ML

- machine learning (ML) supervised/unsupervised learning, reinforcement learning, neural networks, backpropagation
- deep learning (DL) CNNs, RNNs, Transformers, attention mechanisms
- practical projects competitions, replicate papers, contribute to open source

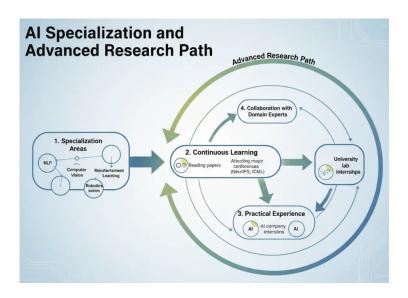




Al expert path - how to become Al scientist/researcher/engineer/practitioner

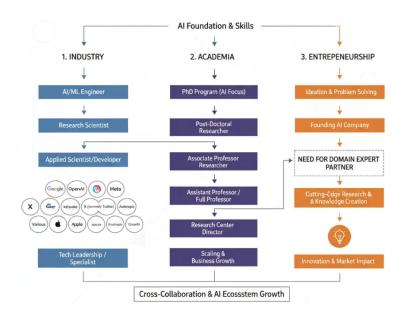
specialization

- choose area (or not!) NLP, computer vision (CV), robotics, etc.
- read papers arxiv, attend conferences (e.g., NeurIPS, ICML, ICLR, CVPR, etc.)
- research experience lab work at university, internships at AI companies
- collaboration work with domain experts on real problems



Al expert path - how to become Al scientist/researcher/engineer/practitioner

- career paths
 - industry Google, OpenAI, Meta, Apple, X, Anthropic, and numerous startups
 - academia PhD \rightarrow professor / research center
 - entrepreneurship found AI company (but need domain expert partner!)



Powerful combination - domain expert + AI expert collaboration

The magic happens at the intersection!

- case study 1 AlphaFold (Protein Folding)
 - Al engineers build transformer-based neural network, optimized training
 - structural biologists identify protein folding as THE problem, validated outputs
 - result 50-year problem solved neither could do it alone
 - both Demis Hassabis (CEO @ Google DeepMind) & John Jumper (biochemistry background) got Nobel Prizes in chemistry in 2024!
- case study 2 cancer diagnostic Al
 - Al engineer role build model, handle large medical imaging datasets
 - oncologist role label training data correctly, evaluate clinical relevance
 - result 30% improvement in early detection
 - who makes final diagnosis always the doctor!

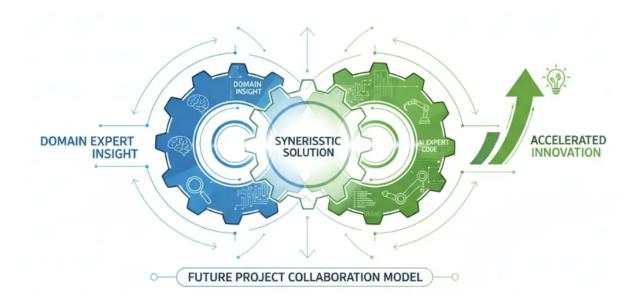






Powerful combination - domain expert + AI expert collaboration

- case study 3 your future project
 - domain expert & Al super user- identify problem in your field you understand deeply
 - Al expert build custom solution beyond available tools
 - result breakthrough that neither could achieve alone
 - who drives vision TOGETHER!



Things AI can't do for you - choose who you'll become

• your technical superpower (Domain \times AI) is only as good as *your moral compass*

- why this matters more than ever
 - Al amplifies whatever you choose to do good intentions OR bad ones
 - history shows us brilliant minds + no moral foundation = catastrophe
 - your skills will make you powerful; your values determine what you do with that power
- virtues that actually matter in the long run
 - integrity doing the right thing (even) when no one's watching
 - empathy genuinely caring about people; your technology will affect them
 - service building things that help others, not just things that benefit you
 - morality having inner compass that guides you beyond what's legal or profitable





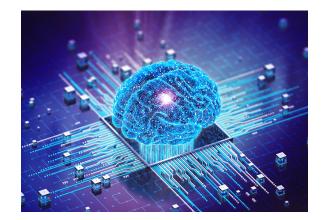
Appendices

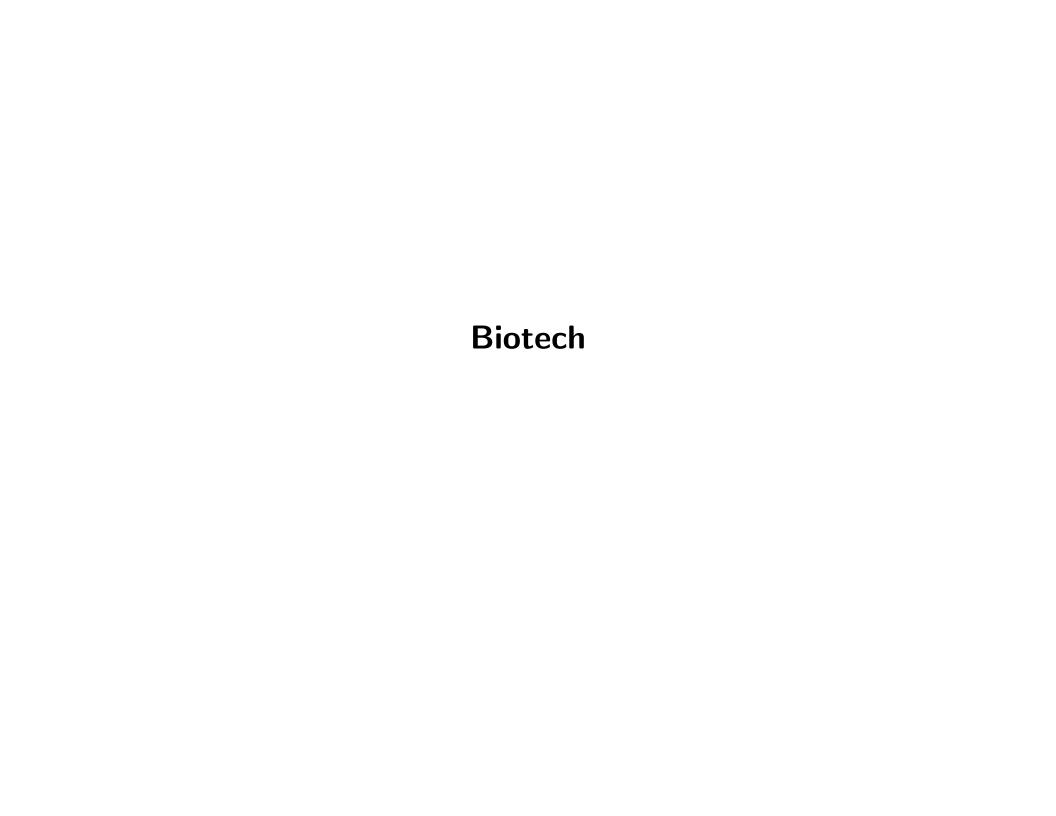
Al & Biotech

Al in biology

- Al has been used in biological sciences, and science in general
- ullet Al's ability to process large amounts of raw, unstructured data (e.g., DNA sequence data)
 - reduces time and cost to conduct experiments in biology
 - enables others types of experiments that previously were unattainable
 - contributes to broader field of engineering biology or biotechnology
- ullet Al increases human ability to make direct changes at cellular level and create novel genetic material (e.g., DNA and RNA) to obtain specific functions







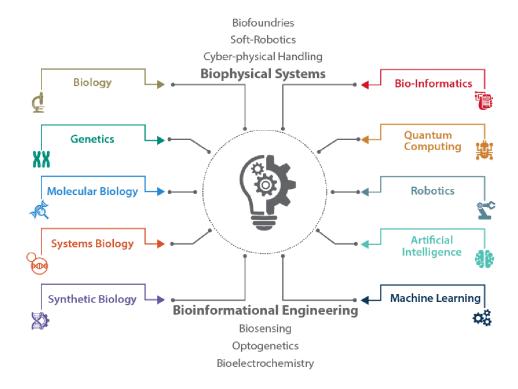
Biotech

biotechnology

- is multidisciplinary field leveraging broad set of sciences and technologies
- relies on and builds upon advances in other fields such as nanotechnology & robotics, and, increasingly, AI
- enables researchers to read and write DNA
 - sequencing technologies "read" DNA while gene synthesis technologies take sequence data and "write" DNA turning data into physical material
- 2018 National Defense Strategy & Senior US Defense and Intelligence Officials identified emerging technologies that could have disruptive impact on US national security [Say21]
 - AI, lethal autonomous weapons, hypersonic weapons, directed energy weapons, biotechnology, quantum technology
- other names for biotechnology are engineering biology, synthetic biology, biological science (when discussed in context of AI)

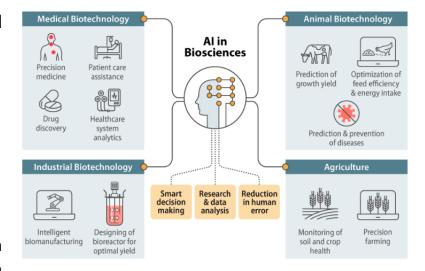
Biotech - multidisciplinary field

- sciences and technologies enabling biotechnology include (but not limited to)
 - (molecular) biology, genetics, systems biology, synthetic biology, bio-informatics, quantum computing, robotics [DFJ22]

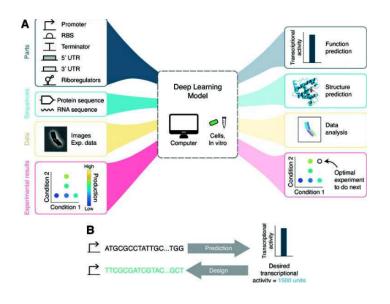


Convergence of AI and biological design

- Al & biological sciences converging [BKP22]
 - each building upon the other's capabilities for new research and development across multiple areas
- Demis Hassabis, CEO & cofounder of DeepMind, said of biology [Toe23]
 - ". . . biology can be thought of as information processing system, albeit extraordinarily complex and dynamic one . . . just as mathematics turned out to be the right description language for physics, biology may turn out to be the perfect type of regime for the application of Al!"
- both AI & biotech rely on and build upon advances in other scientific disciplines and technology fields, such as nanotechnology, robotics, and increasingly big data (e.g., genetic sequence data)
 - each of these fields itself convergence of multiple sciences and technologies
- so their impacts can combine to create new capabilities



Multi-source genetic sequence data



- AI, essential to analyzing exponential growth of genetic sequence data
 - "Al will be essential to fully understanding how genetic code interacts with biological processes" - US National Security Commission on Artificial Intelligence (NSCAI)
 - process huge amounts of biological data, e.g., genetic sequence data, coming from different biological sources for understanding complex biological systems
 - sequence data, molecular structure data, image data, time-series, omics data
- e.g., analyze genomic data sets to determine the genetic basis of particular trait and potentially uncover genetic markers linked with that trait

Quality & quantity of biological data

- limiting factor, however, is *quality and quantity* of biological data, *e.g.*, DNA sequences, that AI is trained on
 - e.g., accurate identification of particular species based on DNA requires reference sequences of *sufficient quality* to exist and be available
- databases have varying standards access, type, and quality of information
- design, management, quality standards, and data protocols for reference databases can affect utility of particular DNA sequence





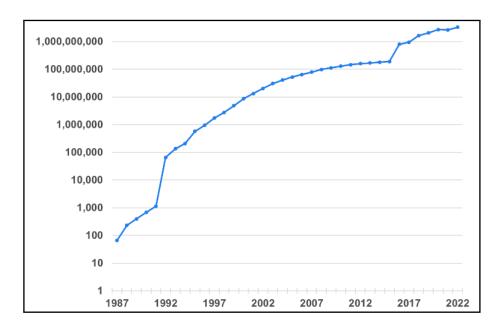
Rapid growth of biological data

- volume of genetic sequence data grown exponentially as sequencing technology evolved
- \bullet more than 1,700 databases incorporating data on genomics, protein sequences, protein structures, plants, metabolic pathways, etc., e.g.
 - open-source public database
 - Protein Data Bank, US-funded data center more than *terabyte of three-dimensional structure data* for biological molecules, *e.g.*, proteins, DNA, RNA
 - proprietary database
 - Gingko Bioworks more than 2B protein sequences
 - public research groups
 - Broad Institute produces roughly 500 terabases of genomic data per month
- great potential value in aggregate volume of genetic datasets that can be collectively mined to discover and characterize relationships among genes

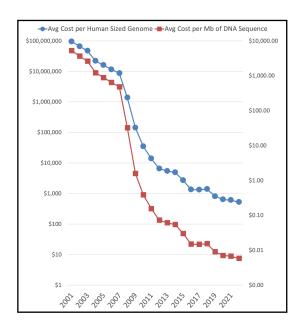
Volume and sequencing cost of DNA over time

- volume of DNA sequences & DNA sequencing cost
 - data source: National Human Genome Research Institute (NHGRI) [Wet23] & International Nucleotide Sequence Database Collaboration (INSDC)
- more dramatic than Moore's law!

sequences in INSDC



DNA sequencing cost



Bio data availability and bias

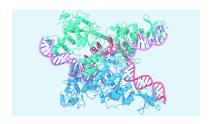
- US National Security Commission on Artificial Intelligence (NSCAI) recommends
 - US fund and prioritize development of a biobank containing "wide range of high-quality biological and genetic data sets securely accessible by researchers"
 - establishment of database of broad range of human, animal, and plant genomes would
 - enhance and democratize biotechnology innovations
 - facilitate new levels of Al-enabled analysis of genetic data
- ullet bias availability of genetic data & decisions about selection of genetic data can introduce bias, e.g.
 - training Al model on datasets emphasizing or omitting certain genetic traits can affect how information is used and types of applications developed - potentially privileging or disadvantaging certain populations
 - access to data and to AI models themselves may impact communities of differing socioeconomic status or other factors unequally

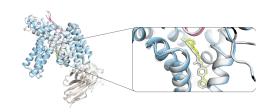
Emerging Trends in Biotech

AlphaFold

- solving 50-year-old protein folding problem, "one of biology's grand challenges"
 - definition given amino acid sequence, predict how it folds into a 3D structure
 - proteins fold in microseconds, but predicting computationally nearly impossible
- \bullet AlphaFold 1 (2018) DL + physics-based energy functions \rightarrow AlphaFold 2 (2020)
 - attention-based NN solving protein folding "in principle" \rightarrow AlphaFold 3 (2024) diffusion-based DL, drug-protein interactions, protein complexes
- AlphaFold protein structure database
 - > 200MM protein structures nearly every known protein, used by > 2MM researchers
- Applications & implications
 - drug discovery target identification, lead optimization, side effect prediction
 - enzyme engineering, agriculture, environmental, vaccine development







AlphaGo

- deep reinforcement learning with Monte Carlo tree search
 - trained on thousands of years of Go game history
 - AlphaGo Zero learns by playing against itself
- development experience, insight, knowledge, know-how transferred to AlphaFold

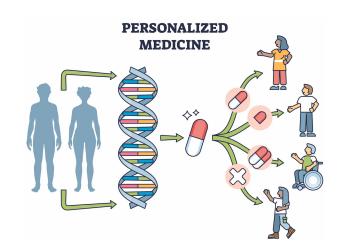




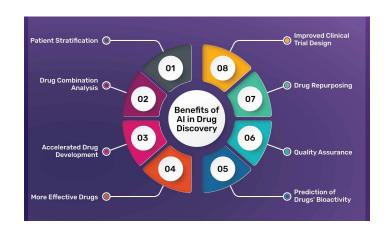


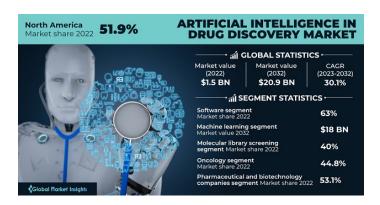
Personalized medicine

- shift from one-size-fits-all approach to tailored treatments
- based on individual genetic profiles, lifestyles & environments
- Al enables analysis of vast data to predict patient responses to treatments, thus enhancing efficacy and reducing adverse effects
- *e.g.*
 - custom cancer therapies
 - personalized treatment plans for rare diseases
 - precision pharmacogenomics
- companies Tempus, Foundation Medicine, etc.



Al-driven drug discovery

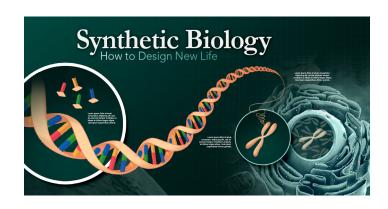


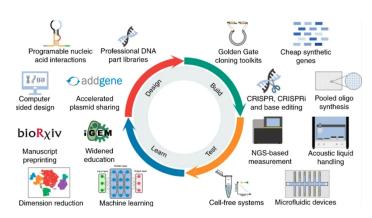


- traditional drug discovery process timeconsuming and costly often taking decades and billions of dollars
- Al streamlines this process by predicting the efficacy and safety of potential compounds with more speed and accuracy
- Al models analyze chemical databases to identify new drug candidates or repurpose existing drugs for new therapeutic uses
- companies Insilco Medicine, Atomwise.

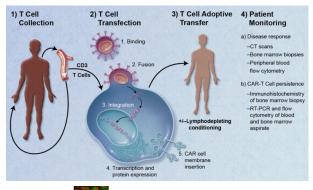
Synthetic biology

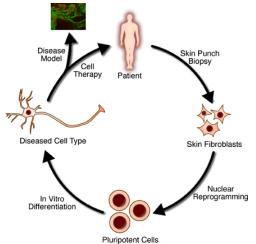
- use AI for gene editing, biomaterial production and synthetic pathways
- combine principles of biology and engineering to design and construct new biological entities
- Al optimizes synthetic biology processes from designing genetic circuits to scaling up production
- company Ginkgo Bioworks uses AI to design custom microorganisms for applications ranging from pharmaceuticals to industrial chemicals





Regenerative medicine

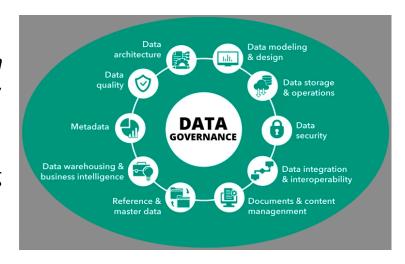




- Al advances development of stem cell therapies & tissue engineering
- Al algorithms assist in identifying optimal cell types, predicting cell behavior & personalized treatments
- particularly for conditions such as neurodegenerative diseases, heart failure and orthopedic injuries
- company Organovo leverages AI to potentially improve the efficacy and scalability of regenerative therapies, developing next-generation treatments

Bio data integration

- integration of disparate data sources, including genomic, proteomic & clinical data - one of biggest challenges in biotech & healthcare
- Al delivers meaningful insights only when seamless data integration and interoperability realized
- developing platforms facilitating comprehensive, longitudinal patient data analysis - vital enablers of AI in biotech
- company Flatiron Health working on integrating diverse datasets to provide holistic view of patient health



Biotech companies



- Atomwise small molecule drug discovery
- Cradle protein design
- Exscientia precision medicine
- Iktos small molecule drug discovery and design
- Insilico Medicine full-stack drug discovery system
- Schrödinger, Inc. use physics-based models to find best possible molecule
- Absci Corporation antibody design, creating new from scratch antibodies, i.e., "de novo antibodies", and testing them in laboratories

Some Important Questions around Al

Some important questions around AI

- why human-level AI?
- what lies in very core of DL architecture? what makes it work amazingly well?
- biases that can hurt judgement, decision making, social good?
- Al ethics & legal issues
- consciousness
- utopia vs dystopia
- knowledge, belief, reasoning
- risk of anthropomorphization



Why human-level in the first place?

- lots of times, when we measure AI performance, we say
 - how can we achieve human-level performance, e.g., CV models?
- why human-level?
 - are all human traits desirable? are humans flawless?
 - aren't humans still evolving?
- advantage of AI over humans
 - e.g., self-driving cars can use extra eyes, GPS, computer network
 - e.g., recommendation system runs for hundreds of millions of people overnight
 - Al is available 24 / 7 while humans cannot
 - . . . critical advanages for medical assitance, emergency handling
 - Al does not make more mistakes because task is repetative and tedius
 - Al does not request salary raise or go on strike

What makes DL so successful?

Factors constributing to astonishing success of DL

- analysis based on speaker's mathematical, numerical algorithmic & statistical perspectives considering hardware innovations
 - 30% universal approximation theorem? (partially) yes! but that's not all
 - function space of neural network is dense (math theory), i.e., for every $f: \mathbf{R}^n \to \mathbf{R}^m$, exists $\langle f_n \rangle$ such that $\lim_{n \to \infty} f_n = f$
 - **25%** architectures/algorithms tailored for each class of applications, e.g., CNN, RNN, Transformer, NeRF, diffusion, GAN, VAE, . . .
 - 20% data labeling expensive, data availability unlimited web text corpus
 - 15% computation power/parallelism Al accelerators, e.g., GPU, TPU & NPU
 - 10% rest Python, open source software, cloud computing, MLOps, . . .

Sudden leap in LLM performance

Probability inferenced sequence is correct

assume

- t_i ith token
- p_i probability that t_i is correct
- ho_i correlation coefficient between t_{i-1} & t_i
- $ilde{p}_k$ probability that (t_1,\ldots,t_k) are correct

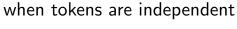
recursion

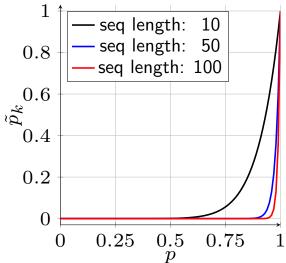
$$\rho_{i} = \frac{\tilde{p}_{i} - \tilde{p}_{i-1}p_{i}}{\sqrt{\tilde{p}_{i-1}(1 - \tilde{p}_{i-1})p_{i}(1 - p_{i})}}$$

$$\Leftrightarrow \qquad \tilde{p}_{i} = \tilde{p}_{i-1}p_{i} + \rho_{i}\sqrt{\tilde{p}_{i-1}(1 - \tilde{p}_{i-1})p_{i}(1 - p_{i})}$$

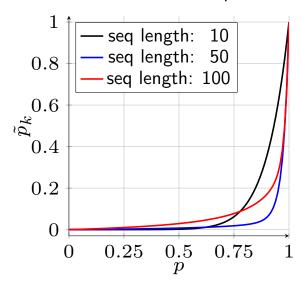
Dramatic improvement of LLM near saturation

- do simulations for both independent & dependent cases
 - assume p_i are same for all i
- ullet (for both cases) sequence inference improves dramatically as p approaches 1
- this explains why we have observed sudden dramatic performance improvement of certain seq2seq learning technologies, e.g., LLM





when tokens are dependent





Cognitive biases attributed to humans

- cognitive biases [Kah11]
 - confirmation bias, availability bias
 - hindsight bias, confidence bias, optimistic bias
 - anchoring bias, halo effect, framing effect, outcome bias
 - belief bias, negativity bias, false consensus







Biases of LLMs

- LLMs subjec to
 - availability bias baised by imbalancedly available information
 - LLM trained by imbalanced # articles for specific topics
 - belief bias derive conclusion not by reasoning, but by what it saw
 - LLM eaisly inferencing what it saw, i.e., data it trained on
 - halo effect overemphasize on what prestigious figures say
 - LLM trained by imbalanced # reports about prestigious figures
- similar facts true for other types of ML models,
 - e.g., video caption, text summarization, sentiment analysis
- cognitive biases only human represent
 - confirmation bias, hindsight bias, confidence bias, optimistic bias, anchoring bias, negativity bias, framing effect

AI Ethics

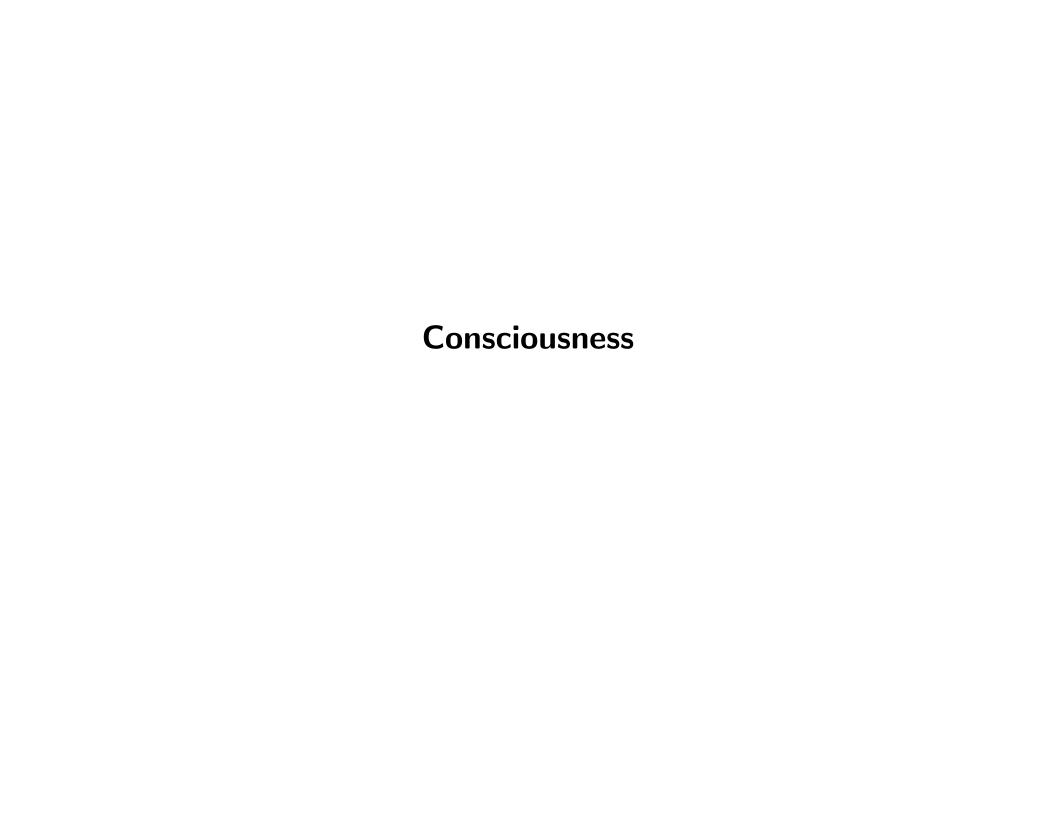
Ethical issues related to Al

- Al can be exploited by those who have bad intention to
 - manupilate / deceive people using manupilated data corpus for training
 - e.g., spread false facts
 - induce unfair social resource allocation
 - e.g., medical insurance, taxation
 - exploit advantageous social and economic power
 - e.g., unfair wealth allocation, mislead public opinion
- Al for Good advocated by Andrew Ng
 - e.g., public health, climate change, disaster management
- should scientists and engineers be morally & politically conscious?
 - e.q., Manhattan project

Al related Legal Issues

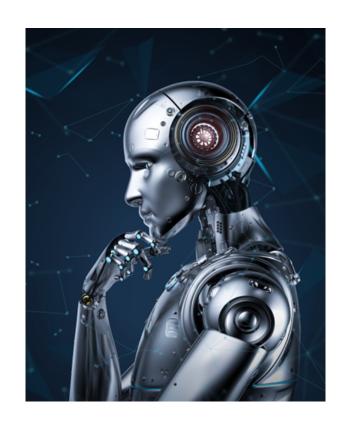
Legal issues with ethical consideration

- scenario 1 full self-driving algorithm causes traffic accident killing people
 - who is responsible? car maker, algorithm developer, driver, algorithm itself?
- scenario 2 self-driving cars kill less people than human drivers
 - e.g., human drivers kill 1.5 people for 100,000 miles & self-driving cars kill 0.2 people for 100,000 miles
 - how should law makers make regulations?
 - utilitarian & humanitarian perspectives
- scenario 3 someone is not happy with their data being used for training
 - "The Times sues OpenAI and Microsoft over AI use of copyrighted work" (Dec-2023)
 - "Newspaper publishers in California, Colorado, Illinois, Florida, Minnesota and New York said Microsoft and OpenAl used millions of articles without payment or permission to develop ChatGPT and other products" (Apr-2024)



Consciousness

- what is consciousness, anyway?
 - recognizes itself as independent, autonomous, valuable entity?
 - recognizes itself as living being, unchangeable entity?
- no agreed definition on consciousness exists yet
 . . . and will be so forever
- does it have anything to do with the fact that humans are biologically living being?
- is SKYNET ever plausible?
 - can Al have desire to survive (or save earth)?



Utopia vs Dystopia

Utopia vs dystopia



- not important questions (at all) I think . . .
- what we should focus on is not the possibilities of doomday or Judgment Day, but rather
 - our limits on controlling unintended impacts of Al
 - misuse by (greedy, immoral, and unethical) people possessing social, economic & political power
 - social good and welfare impaired by either exploiting
 Al or ignorance of (inner workings of) Al
- should concern
 - choice or balance among utilitarianism,
 humanitarianism & values
 - amend or improve laws/regulations
 - ethical issues caused by AI

Knowledge, Belief, and Reasoning

Does AI (LLM) have knowledge or belief? Can it reason?

What categories of questions do they belong to? engineering, scientific, philosophical, cognitive scientific, . . . ?

LLMs . . .

• LLM is very different sort of animal . . . except that it is *not* an animal!

- unreasonable effectiveness of data [HNF09]
 - performance scales with size of training data
 - qualitative leaps in capability as models scale
 - tasks demanding human intelligence reduced to next token prediction
- focus on third surprise

conditional probability model looks like human with intelligence

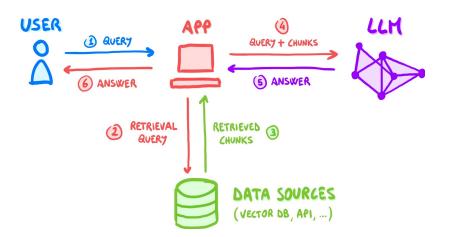
- making vulnerable to anthropomorphism
- examine it by throwing questions such as
 - "does LLM have knowledge and belief?"
 - "can it reason?"

What LLM really does!

- given prompt "the first person to walk on the Moon was", LLM responds with "Neil Armstrong". . . strictly speaking
 - it's *not* being asked *who* was the first person to walk on the Moon
 - what are being really asked is "given statistical distribution of words in vast public corpus of text, what words are most likely to follow 'The first person to walk on the Moon was'?"
- given prompt "after ring was destroyed, Frodo Baggins returned to", LLM responds with "the Shire"
 - on one level, it seems fair to say, you might be testing LLM's knowledge of fictional world of Tolkien's novels
 - what are being really asked is "given statistical distribution of words in vast public corpus of text, what words are most likely to follow 'After the ring was destroyed, Frodo Baggins returned to'?"

LLMs vs systems in which they are embedded

- crucial to distinguish between the two (for philosophical clarity)
 - LLM (bare-bones model) highly specific & well-defined function, which is conditional probability estimator
 - systems in which LLMs are embedded, e.g., for question-answering, news article summarization, screenplays generation, language translation





How ChatBot works?

conversational Al agent does in-context learning or few-shot prompting

- for example,
 - when the user enters who is the first person to walk on the Moon?
 - ChatBot, LLM-embedded system, feeds the following to LLM

User, a human, and BOT, a clever and knowledgeable AI agent.

User: what is 2+2?

BOT: the answer is 4.

User: where was Albert Einstein born?

BOT: he was born in Germany.

who is the first person to walk on the Moon?

BOT:

Knowledge, belief & reasoning around LLM

- not easy topic to discuss, or even impossible because
 - we do not have agreed definition of these terms especially in context of being asked questions like

does LLM have belief?
or
do humans have knowledge?

- let us discuss them in two different perspectives
 - laymen's perspectives
 - cognitive scientific & philosophical perspectives

Laymen's perspectives on knowledge, belief & reasoning

- does (good) LLM have knowledge?
 - Grandmother: looks like it cuz when instructed "explaing big bang", it says "The Big Bang theory is prevailing cosmological model that explains the origin and evolution of the universe. . . . 13.8 billion years ago . . . "
- does it have belief?
 - Grandmother: I don't think so, e.g., it does not believe in God!
- can it reason?
 - Grandmother: seems like it! e.g., when asked "Sunghee is a superset of Alice and Beth is a superset of Sunghee. is Beth a superset of Alice?", it says "Yes, based on information provided, if Sunghee is a superset of Alice and Beth is a superset of Sunghee, then Beth is indeed a superset of Alice . . . "
- can it reason to prove theorem whose inferential structure is more complicated?
 - Grandmother: I'm not sure actually, I don't know what you're talking about!

Knowledge

- could argue LLM "knows" which words follow which other words with high probability
- but, only in context of capacity to distinguish truth from falsehood can we legitimately speak of "knowledge"!
- LLM(-embedded BOT)
 - can be said to "encode", "store", or "contain" knowledge
 - lacks means to use words "true" & "false" in all ways & in all contexts because . . .
 - does not inhabit the world we human language-users share!





Belief

- nothing can count as belief about the world we share unless
 - is against backdrop of "ability to update beliefs appropriately in light of evidence from that world" - (again) essential capacity to distinguish truth from falsehood
- change taking place in humans when acquiring or updating belief is
 - reflection of their nature as language-using animals inhabiting shared world with community of language-users
- then, what if LLM-embedded system updates LLM with outside world information?
 - even so, when interacting with AI systems based on LLMs, these grounds are absent!

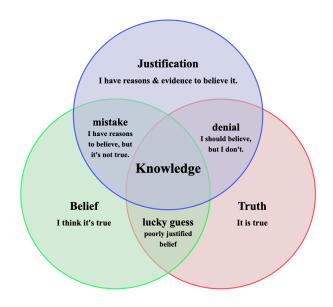






Knowledge in philosophical and cognitive scientific sense

- does LLM have knowledge?
 - Sunghee: I don't think so!
- why?
 - we say we have "knowledge" when "we do so against ground of various human capacities that we all take for granted when we engage in everyday conversation with each other."
 - when asked "who is Tom Cruise's mother?", it says "Tom Cruise's mother is Mary Lee Pfeiffer."
 However, this is nothing but
 "guessing" by conditional probability model
 the most likely words following "Tom
 Cruise's mother is"
 - so we cannot say it really knows the fact!



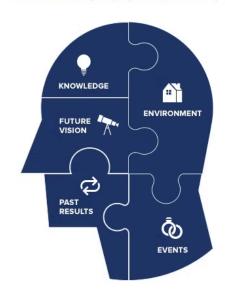
Belief in philosophical and cognitive scientific sense

- for the discussion
 - do not concern any specific belief
 - but concern prerequisites for ascribing any beliefs to Al system
- so does it have belief?
 - nothing can count as belief about the world we share unless

it is against ground of the ability to update beliefs appropriately in light of evidence from that world, essential aspect of the capacity to distinguish truth from falsehood

- LLM does not have this ground, essential consideration when deciding whether it *really* had beliefs.
- Sunghee: so no, LLM cannot have belief!

WHERE DO YOUR BELIEFS COME FROM?



Reasoning in philosophical and cognitive scientific sense

- note reasoning is *content neutral*
 - e.g., following logic is perfect regardless of truth of premises
 - hence, no access to outside world does not disqualify
- when asked "if humans are immortal, would Socrates have survived today?", LLM says "... it's logical to conclude that Socrates would likely still be alive today. ... "
- however, remember, once again, what we just asked it to do is not "deductive inference" given the statistical distribution of words in public corpus, what words are likely to follow the sequence, "humans are immortal and Socreates is human therefore."
- Sunghee: so no, LLM cannot reason, either!
- but, LLM
 - pretends to reason, and from which capabilities, we can benefit!
 - also, can mimic even multi-step reasoning whose inferencing structure is complicated using chain-of-thoughts prompting, i.e., in-context learning or few-short prompting

Simple example showing LLM not possessing knowledge



User

"Who is Tom Cruise's mother?"

• LLM(-embedded question-answering system) (as of Jan 2022)

"Tom Cruise's mother is Mary Lee Pfeiffer. She was born Mary Lee South. . . . Information about his family, including his parents, has been publicly available, . . . "

User

"Who is Mary Lee Pfeiffer's son?"

• LLM(-embedded question-answering system) (as of Jan 2022)

"As of my last knowledge update in January 2022, I don't have specific information about Mary Lee Pfeiffer or her family, including her son. . . . "

Risk of anthropomorphization

- unfortunately, contemporary LLMs are too powerful, too versatile, and too useful for most people to accept (after understanding) previous arguments!
- maybe, o.k. for laymen to (mistakenly) anthropomorphize LLM(-embedded systems)
- however, imperative for (important, smart, and responsible) Al researchers, scientists, engineers & practitioners to have rigorous understanding in these aspects especially when
 - advise and be consulted by law makers, policy makers, journalists, and various stakeholders responsible for critical business decisions (in private sectors) and public policies (in public sectors)
 - collaborate with or/and help professionals in liberal arts, such as philosophy, ethics, law, religion, literature, history, music, cultural studies, psychology, sociology, anthropology, political science, economics, archaeology, linguistics, media studies, natural sciences, fine arts, . . .
 - to address negative soceital and economic impacts

Moral

• Al shows incredible utility and commercial potentials, hence should

- make informed decisions about trustworthiness and safety
- avoid ascribing capacities they lack
- take best utilization of remarkable capabilities of AI
- today's AI so powerful, so (seemingly) convincingly intelligent
 - obfuscate mechanism
 - actively encourage anthropomorphism with philosophically loaded words like "believe"
 and "think"
 - easily mislead people about character and capabilities of Al
- matters not only to scientists, engineers, developers, and entrepreneurs, but also
 - general public, law & policy makers, journalists, . . .

Selected References & Sources

Selected references & sources

	20
 Michael J. Sandel "Justice: What's the Right Thing to Do?" 	<i>J</i> 9
 Daniel Kahneman "Thinking, Fast and Slow" 201 	11
 Yuval Noah Harari "Sapiens: A Brief History of Humankind" 201 	14
• M. Shanahan "Talking About Large Language Models" 202	22
• A.Y. Halevry, P. Norvig, and F. Pereira "Unreasonable Effectiveness of Data" 200)9
• A. Vaswani, et al. "Attention is all you need" @ NeurIPS 201	17
• S. Yin, et. al. "A Survey on Multimodal LLMs"	23
• Chris Miller "Chip War: The Fight for the World's Most Critical Technology" 202	22

- CEOs, CTOs, CFOs, COOs, CMOs & CCOs @ startup companies in Silicon Valley
- VCs on Sand Hill Road Palo Alto, Menlo Park, Woodside in California, USA

References

References

[BKP22] Abhaya Bhardwaj, Shristi Kishore, and Dhananjay K. Pandey. Artificial intelligence in biological sciences. *Life*, 12(1430), 2022.

- [DFJ22] Thomas A. Dixon, Paul S. Freemont, and Richard A. Johnson. A global forum on synthetic biology: The need for international engagement. *Nature Communications*, 13(3516), 2022.
- [GPAM⁺14] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.
- [HGH⁺22] Sue Ellen Haupt, David John Gagne, William W. Hsieh, Vladimir Krasnopolsky, Amy McGovern, Caren Marzban, William Moninger, Valliappa Lakshmanan, Philippe Tissot, and John K. Williams. The history and practice of Al in the environmental sciences. *Bulletin of the American Meteorological Society*, 103(5):E1351 E1370, 2022.
- [HM24] Guadalupe Hayes-Mota. Emerging trends in Al in biotech. *Forbes*, June 2024.

[HNF09] Alon Halevy, Peter Norvig, and Nanediri Fernando. The unreasonable effectiveness of data. *Intelligent Systems, IEEE*, 24:8 – 12, 05 2009.

- [Kah11] Daniel Kahneman. *Thinking, fast and slow.* Farrar, Straus and Giroux, New York, 2011.
- [Kui23] Todd Kuiken. Artificial intelligence in the biological sciences: Uses, safety, security, and oversight. *Congressional Research Service*, Nov 2023.
- [KW19] Diederik P. Kingma and Max Welling. An introduction to variational autoencoders. *Foundations and Trends in Machine Learning*, 12(4):307–392, 2019.
- [MLZ22] Louis-Philippe Morency, Paul Pu Liang, and Amir Zadeh. Tutorial on multimodal machine learning. In Miguel Ballesteros, Yulia Tsvetkov, and Cecilia O. Alm, editors, *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Tutorial Abstracts*, pages 33–38, Seattle, United States, July 2022. Association for Computational Linguistics.
- [Say21] Kelley M. Sayler. Defense primer: Emerging technologies. *Congressional Research Service*, 2021.

- [Sha23] Murray Shanahan. Talking about large language models, 2023.
- [Toe23] Rob Toews. The next frontier for large language models is biology. *Forbes*, July 2023.
- [VSP⁺17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of 31st Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [Wet23] Kris A. Wetterstrand. Dna sequencing costs: Data, 2023.
- [YFZ⁺24] Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A survey on multimodal large language models, 2024.

Thank You