

The Shifting Landscape of Modern AI

Prof Truyen Tran

Head of AI, Health and Science

Applied AI Institute (A2I2), Deakin University

truyen.tran@deakin.edu.au | truyentran.github.io



1997: My journey - IPhO

**This was the first time I got
into airplane, and went
overseas!**

Silver medal! Yay!

Suit & shoes borrowed from
my Dad.



1956: The quest began




Image credit: Dickson Wu

Among the most challenging scientific questions of our time are the corresponding **analytic** and **synthetic** problems:

- How does the brain function?
- Can we design a machine which will simulate a brain?

-- *Automata Studies*

The quiet period (to the public)



2024



What will a digital **companion** look like?

How will we interact with “it”?

How can we build one?

(Hint: Multimodal, embodiment, theory of mind etc.).



AI is general purpose

Three kinds of AI


- **Cognitive automation**: encoding human abstractions → automate tasks normally performed by humans.
- **Cognitive assistance**: AI helps us make sense of the world (perceive, think, understand).
- **Cognitive autonomy**: Artificial minds thrive independently of us, exist for their own sake.



François Chollet

The quest for general purpose: **Learning** a Turing machine

Can we **design** a program that **learns** to program?




A mechanical Turing machine

2008: When things become over complex, time for a fresh air!



ELSEVIER



Hierarchical semi-Markov conditional random fields for deep recursive sequential data



Truyen Tran^{a,*}, Dinh Phung^a, Hung Bui^b, Svetha Venkatesh^a

^a Center for Pattern Recognition and Data Analytics, Deakin University Geelong, Australia

$$\Phi[\zeta, z] = \left[\prod_{d \in [1, D]} \prod_{i_k, i_{k+1} \in \tau^d} R_{i_{k+1}:i_k}^{d,s,z} \right] \prod_{d \in [1, D-1]} \left\{ \left[\prod_{i_k \in \tau^{d+1}, i_k \notin \tau^d} A_{u,v,i_k}^{d+1,s,z} \right] \left[\prod_{i_k \in \tau^{d+1}} \pi_{u,i_{k+1}}^{d,s,z} \right] \left[\prod_{i_k \in \tau^{d+1}} E_{u,i_k}^{d,s,z} \right] \right\} \quad (8)$$

$$\begin{aligned} \alpha_{i:j}^{d,s}(u) &= \sum_{t \in [i+1, j]} \sum_{v \in S^{d+1}} \sum_{\zeta_{i:t-1}^{d,s}(v)} \sum_{\zeta_{t:j}^{d+1,u}} \Phi \left[\hat{\zeta}_{i:t-1}^{d,s}(v) \right] \Phi \left[\hat{\zeta}_{t:j}^{d+1,u} \right] A_{v,u,t-1}^{d+1,s} R_{t:j}^{d+1,u} + \sum_{\zeta_{i:j}^{d+1,u}} \pi_{u,i}^{d,s} \Phi \left[\hat{\zeta}_{i:j}^{d+1,u} \right] R_{i:j}^{d+1,u} \\ &= \sum_{t \in [i+1, j]} \sum_{v \in S^{d+1}} \alpha_{i:t-1}^{d,s}(v) \hat{\Delta}_{t:j}^{d+1,u} A_{v,u,t-1}^{d+1,s} + \hat{\Delta}_{i:j}^{d+1,u} \pi_{u,i}^{d,s} \end{aligned} \quad (36)$$

2012: Deep learning took off

Machine learns as data speak

AI, machine learning, deep learning, data science and all those topics!

Monday 21 May 2012


Machine learning at its turning point: Non-convexity

Machine learning had enjoyed the last decade of specifying and optimising convex regularised risks. It seems that we have reached the point when we need to move on, accepting the fact that global optimum may not be optimal after all.

To me, learning based on convex functions is desirable as long as we know what we are looking for. In many cases, however, we do not, and we are more interested in discovery. And one of the most effective way is to loosely specify latent variables which, hopefully captures the generative process of the data. Whenever we introduce latent variables, it is likely that we hit a non-convex objective function. In fact, the goal now is not to find a optimal configuration given the known, well-behaved search space, but to search for *reasonably effective* configurations in a largely unknown space. It is actually evidenced in most powerful natural living systems, such as human brains: They are very effective and time-tested, but they are hardly optimal in any sense.


About Me



 **Truyen Tran**

I'm pushing the frontiers of AI by: ++ Unlocking intelligence & consciousness, ++ Designing competent

intelligent machines, and ++ Transforming physical and digital fields through AI. My team and I have made advances in deep neural networks, visual reasoning, machine learning for biomedicine, materials science, recommender system and software analytics. Currently we are building machines that can remember, reason about what they see, be mindful of what other think, and respect human values. More at: truyentran.github.io



2016: The
defining AI
moment for
the world

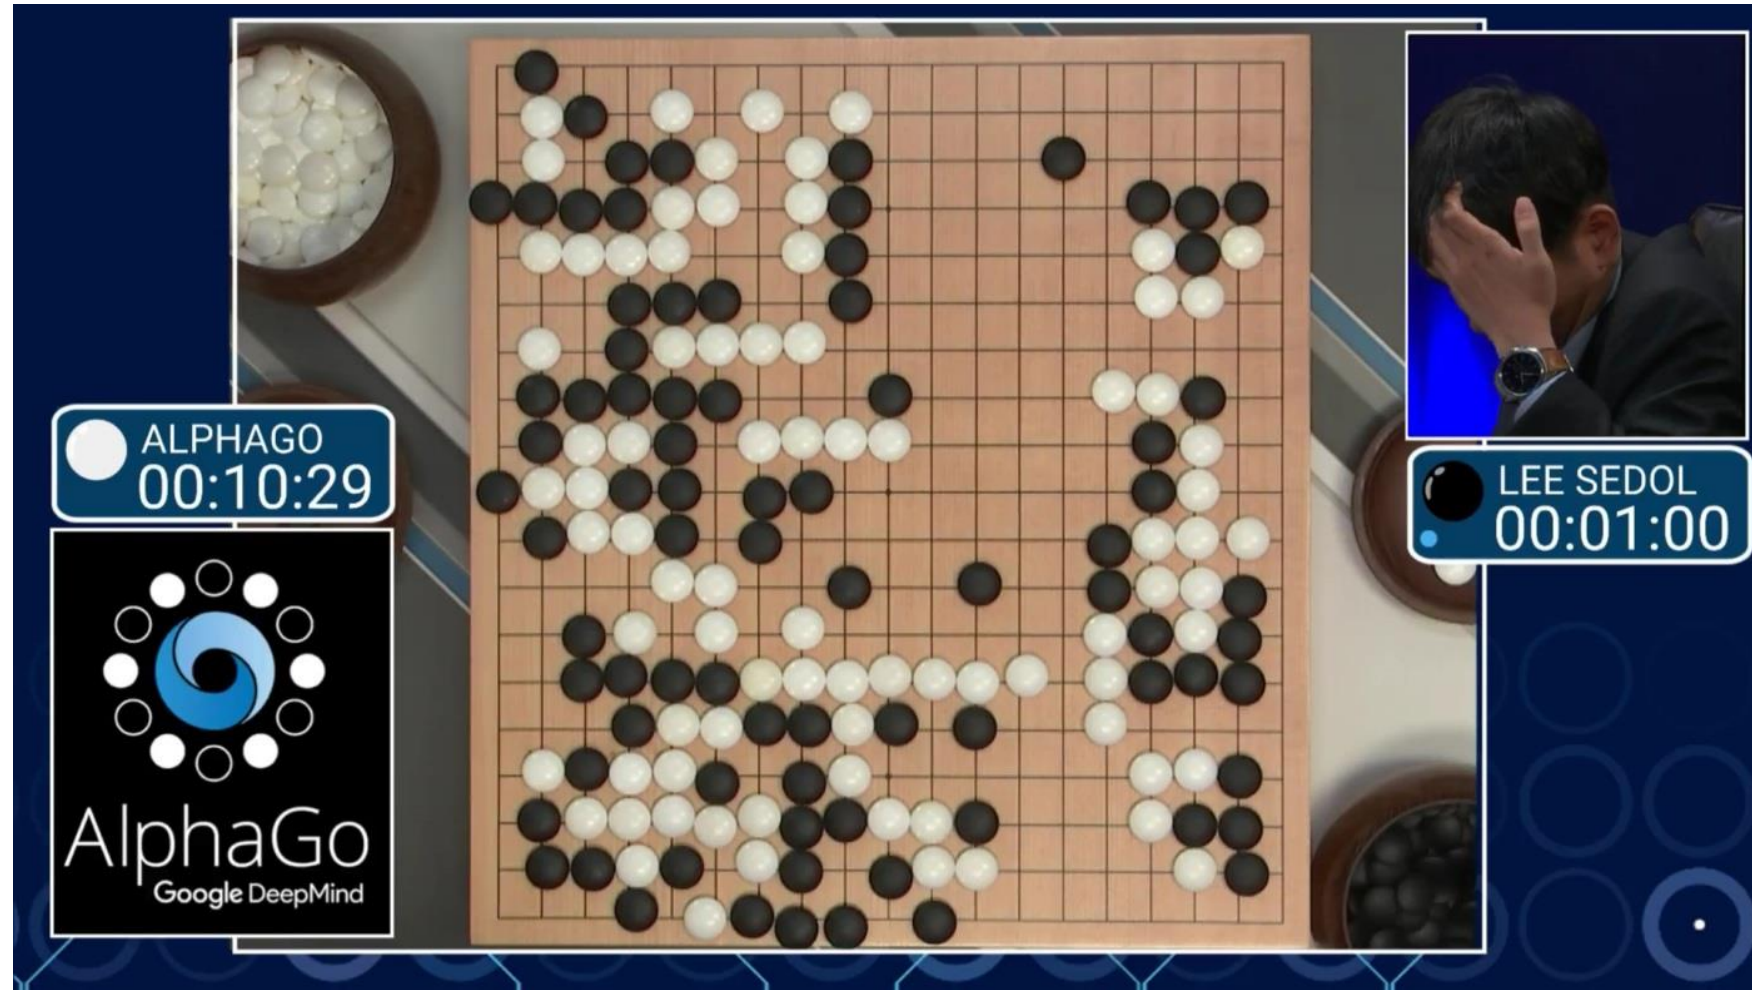


Photo credit: Shelly Palmer

2017

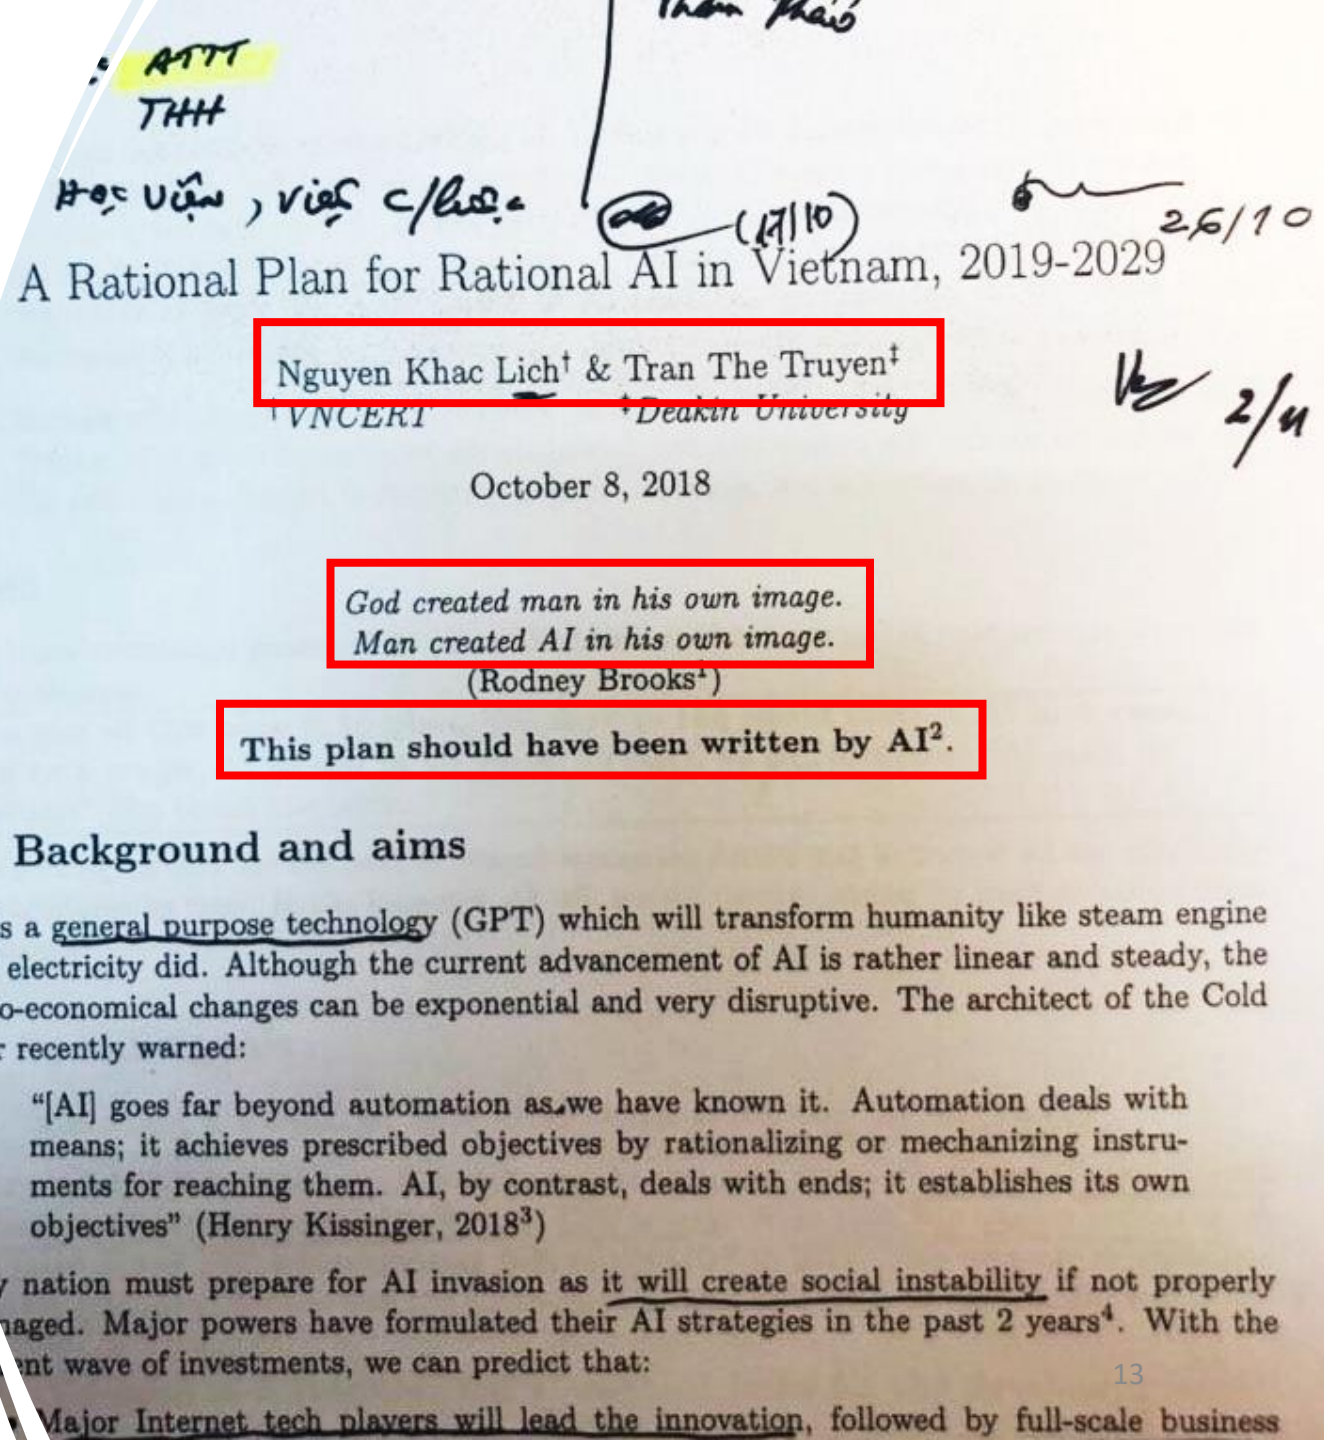
“Software 2.0 is written in
neural network weights”

Andrej Karpathy



2018: Vietnam

- “This plan should have been written by AI”
- “Any nation must prepare for AI invasion as it will create social instability if not properly managed”.
- “Major Internet tech players will lead the innovation”.



2019: The Bitter Lesson

By Rich Sutton

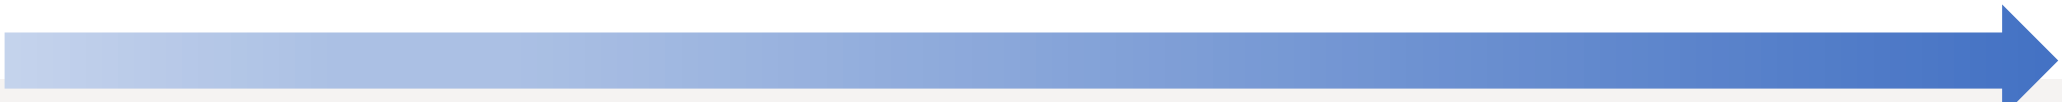
“The biggest lesson that can be read from 70 years of AI research is that **general methods that leverage computation** are ultimately the *most effective, and by a large margin.*”

“The two methods that seem to scale arbitrarily in this way are search and learning.”

2020: The sentiment

“Deep learning is going to be able to do everything”
(Geoff Hinton)

2020 onward: The landscape is shifting, rapidly!



	PRE - 2020	2020	2022	2023?	2025?	2030?
TEXT	Spam detection Translation Basic Q&A	Basic copy writing First drafts	Longer form Second drafts	Vertical fine tuning gets good (scientific papers, etc)	Final drafts better than the human average	Final drafts better than professional writers
CODE	1-line auto-complete	Multi-line generation	Longer form Better accuracy	More languages More verticals	Text to product (draft)	Text to product (final), better than full-time developers
IMAGES			Art Logos Photography	Mock-ups (product design, architecture, etc.)	Final drafts (product design, architecture, etc.)	Final drafts better than professional artists, designers, photographers)
VIDEO / 3D / GAMING			First attempts at 3D/video models	Basic / first draft videos and 3D files	Second drafts	AI Roblox Video games and movies are personalized dreams

Large model availability:



First attempts



Almost there



Ready for prime time

GenAI is now very powerful



GenAI are
compression
engine

Prompting is conditioning
for the (preference-
guided) decompression.



GenAI are
approximate
program database

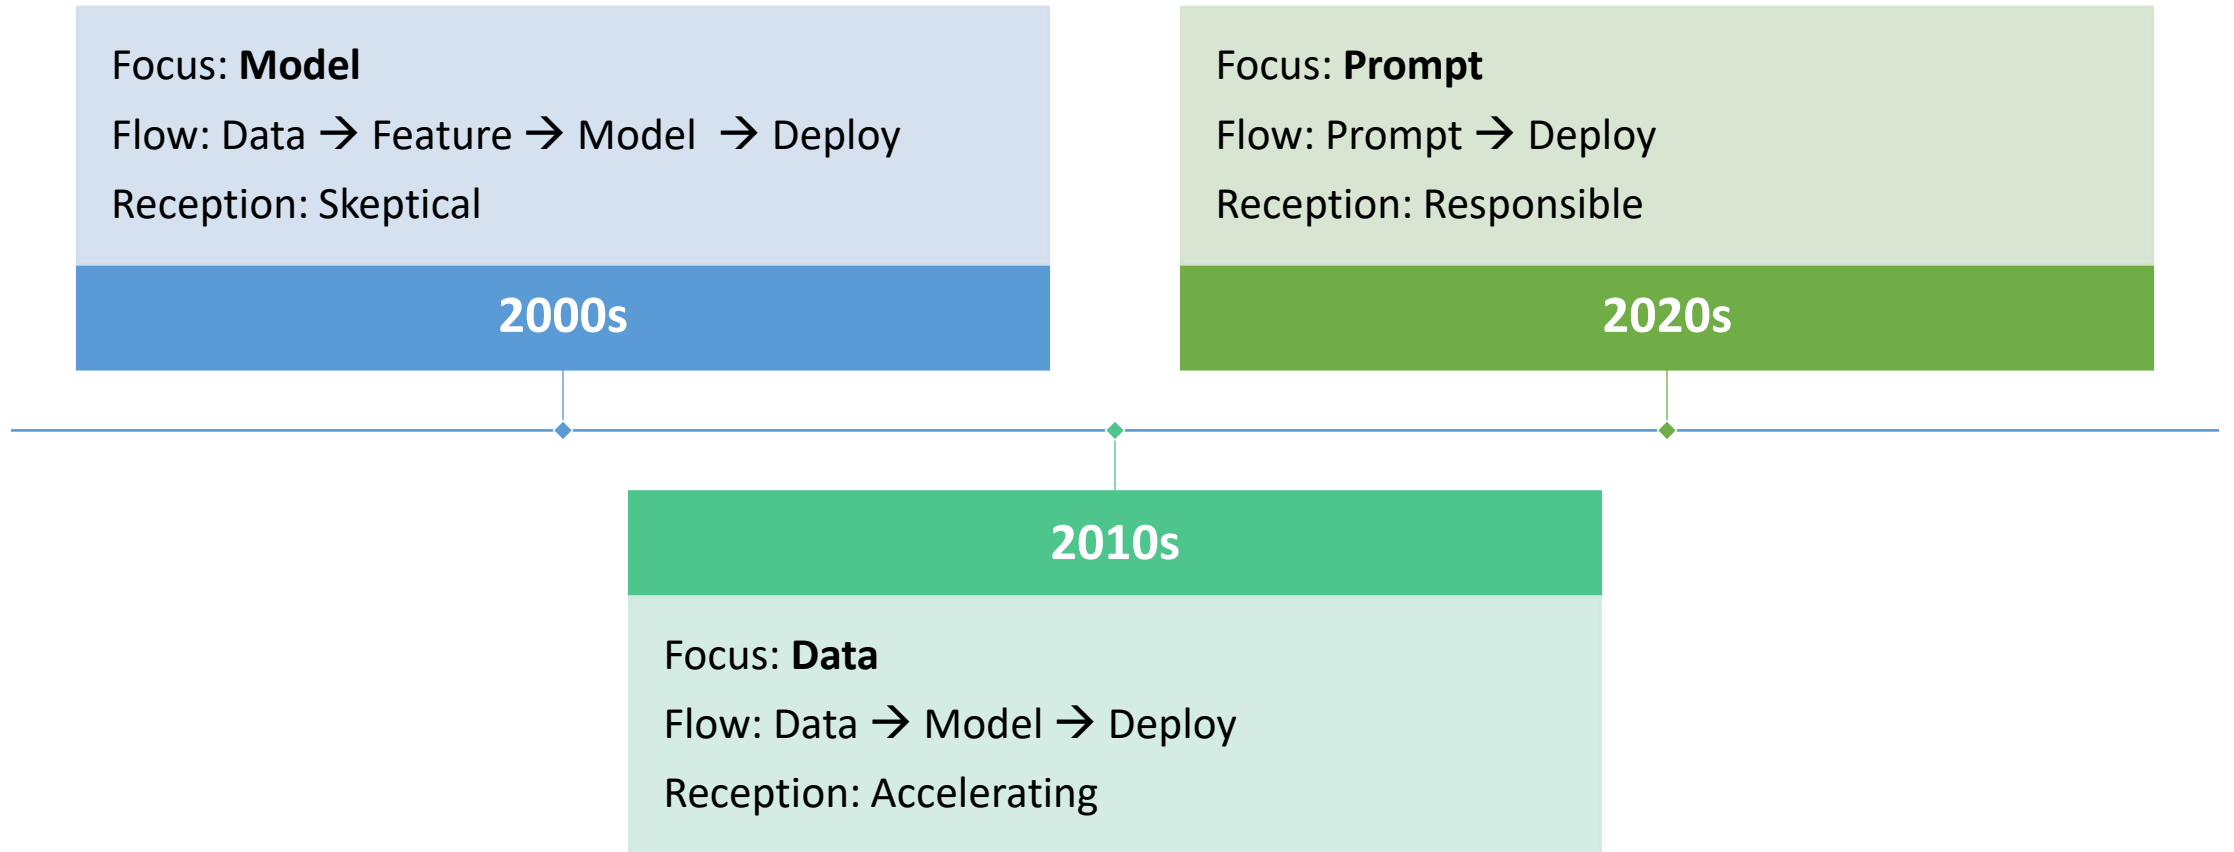
Prompting is retrieving an
approximate program that
takes input and delivers
output.




GenAI are
World Model

We can live entirely in
simulation!

The practice of AI is changing...



Generative AI as first-mile tech



AI has changed how we build, maintain and interact with systems.



Generative AI are the jack of all trades, master of none.



Last-mile AI apps with built on general-purpose Generative AI



The shifting AI research



Engineering

Design man-made systems



Generative AI





Discover emergent behaviours



Science

Discover laws in nature






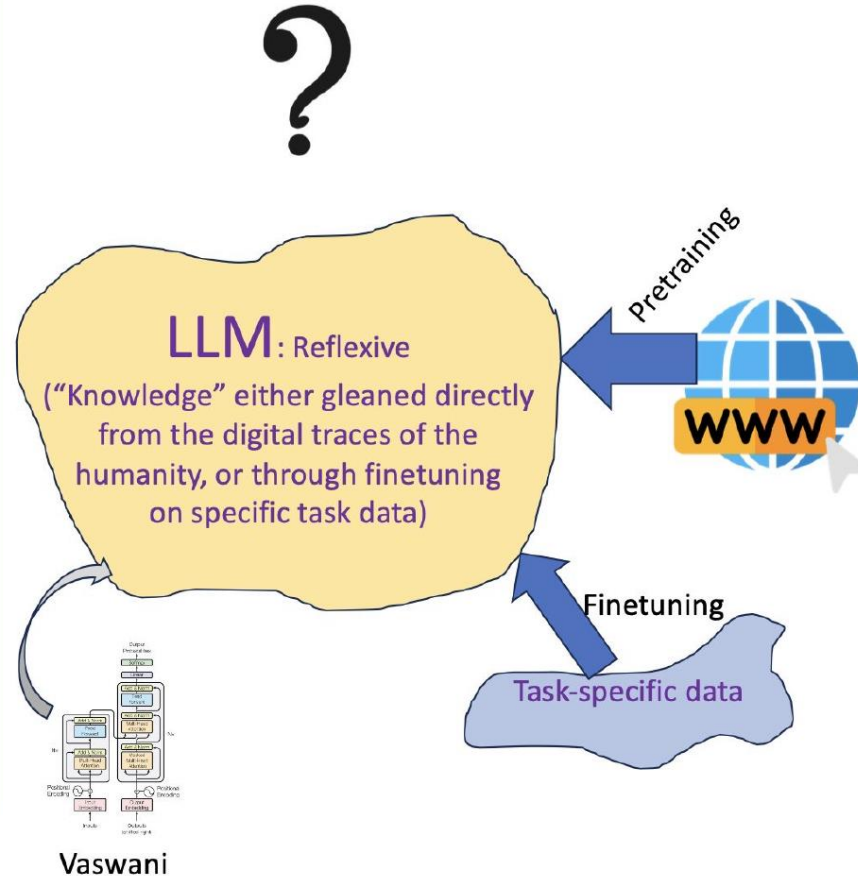
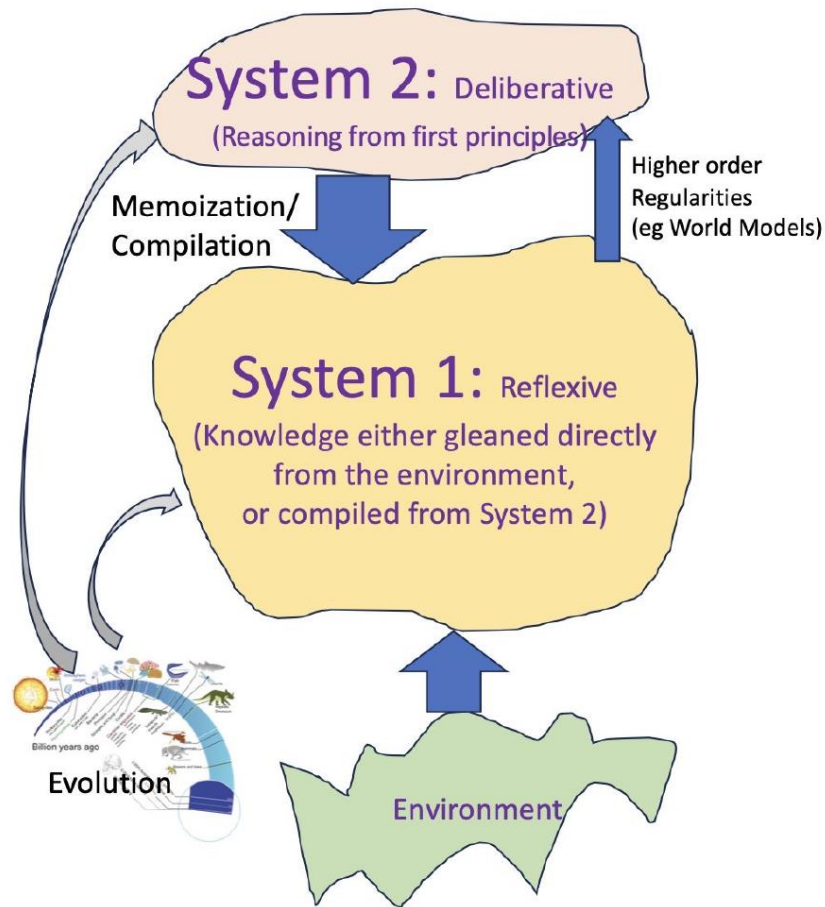
GenAI for Science:

Knowledge \rightarrow Reward \rightarrow Generation

Materials design at A2I2, Deakin Uni.



True reasoning is still missing



Machine reasoning: What



Leslie Valiant

Reasoning is concerned with arriving at a deduction about a **new combination of circumstances**.




Leon Bottou

Reasoning is to deduce **new** knowledge from previously acquired knowledge in response to a **query**.

Machine reasoning: How



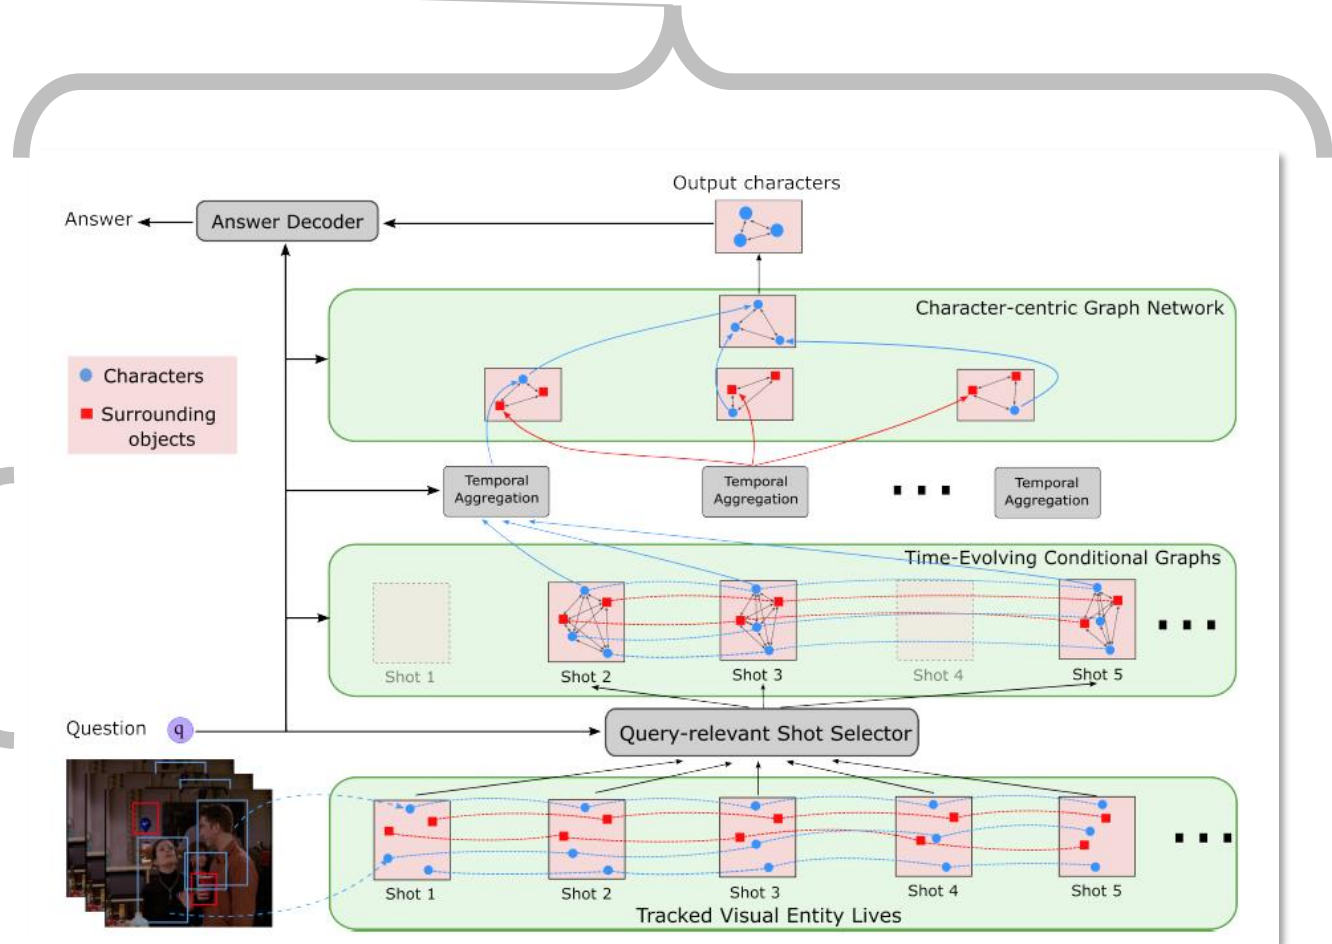
Reasoning as just-in-time **program synthesis**.



It employs **conditional computation**.

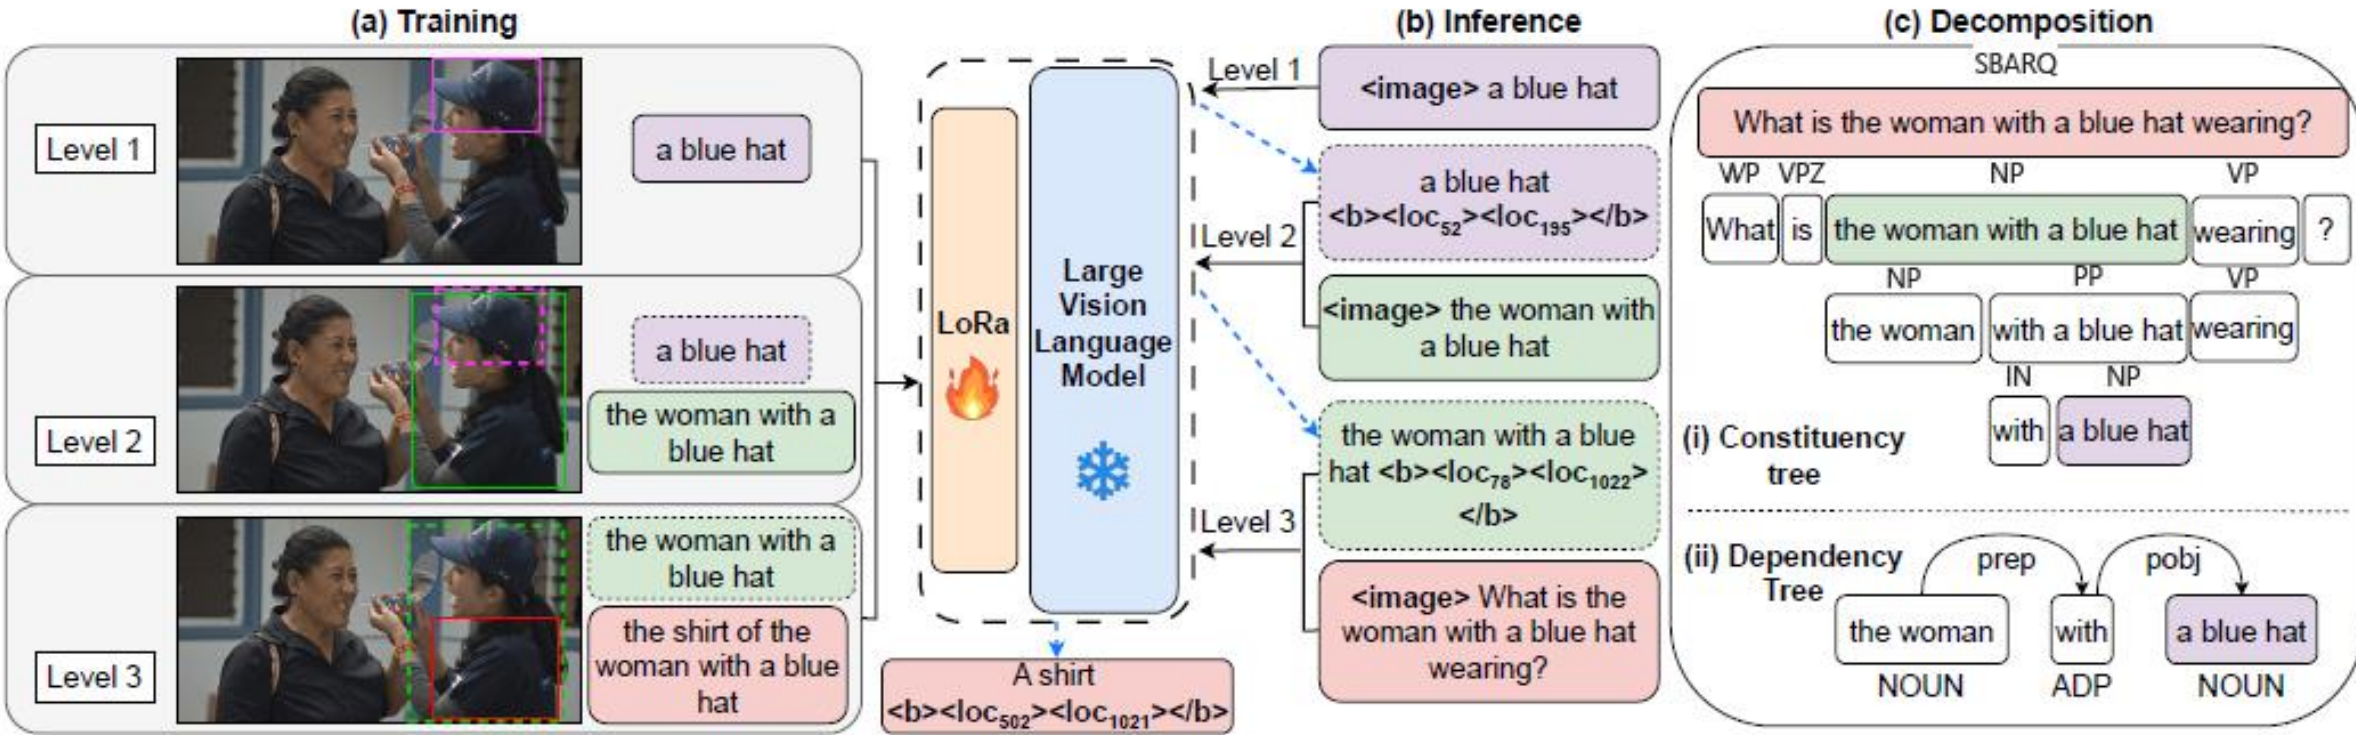


Reasoning is **recursive**,
e.g., mental travel.



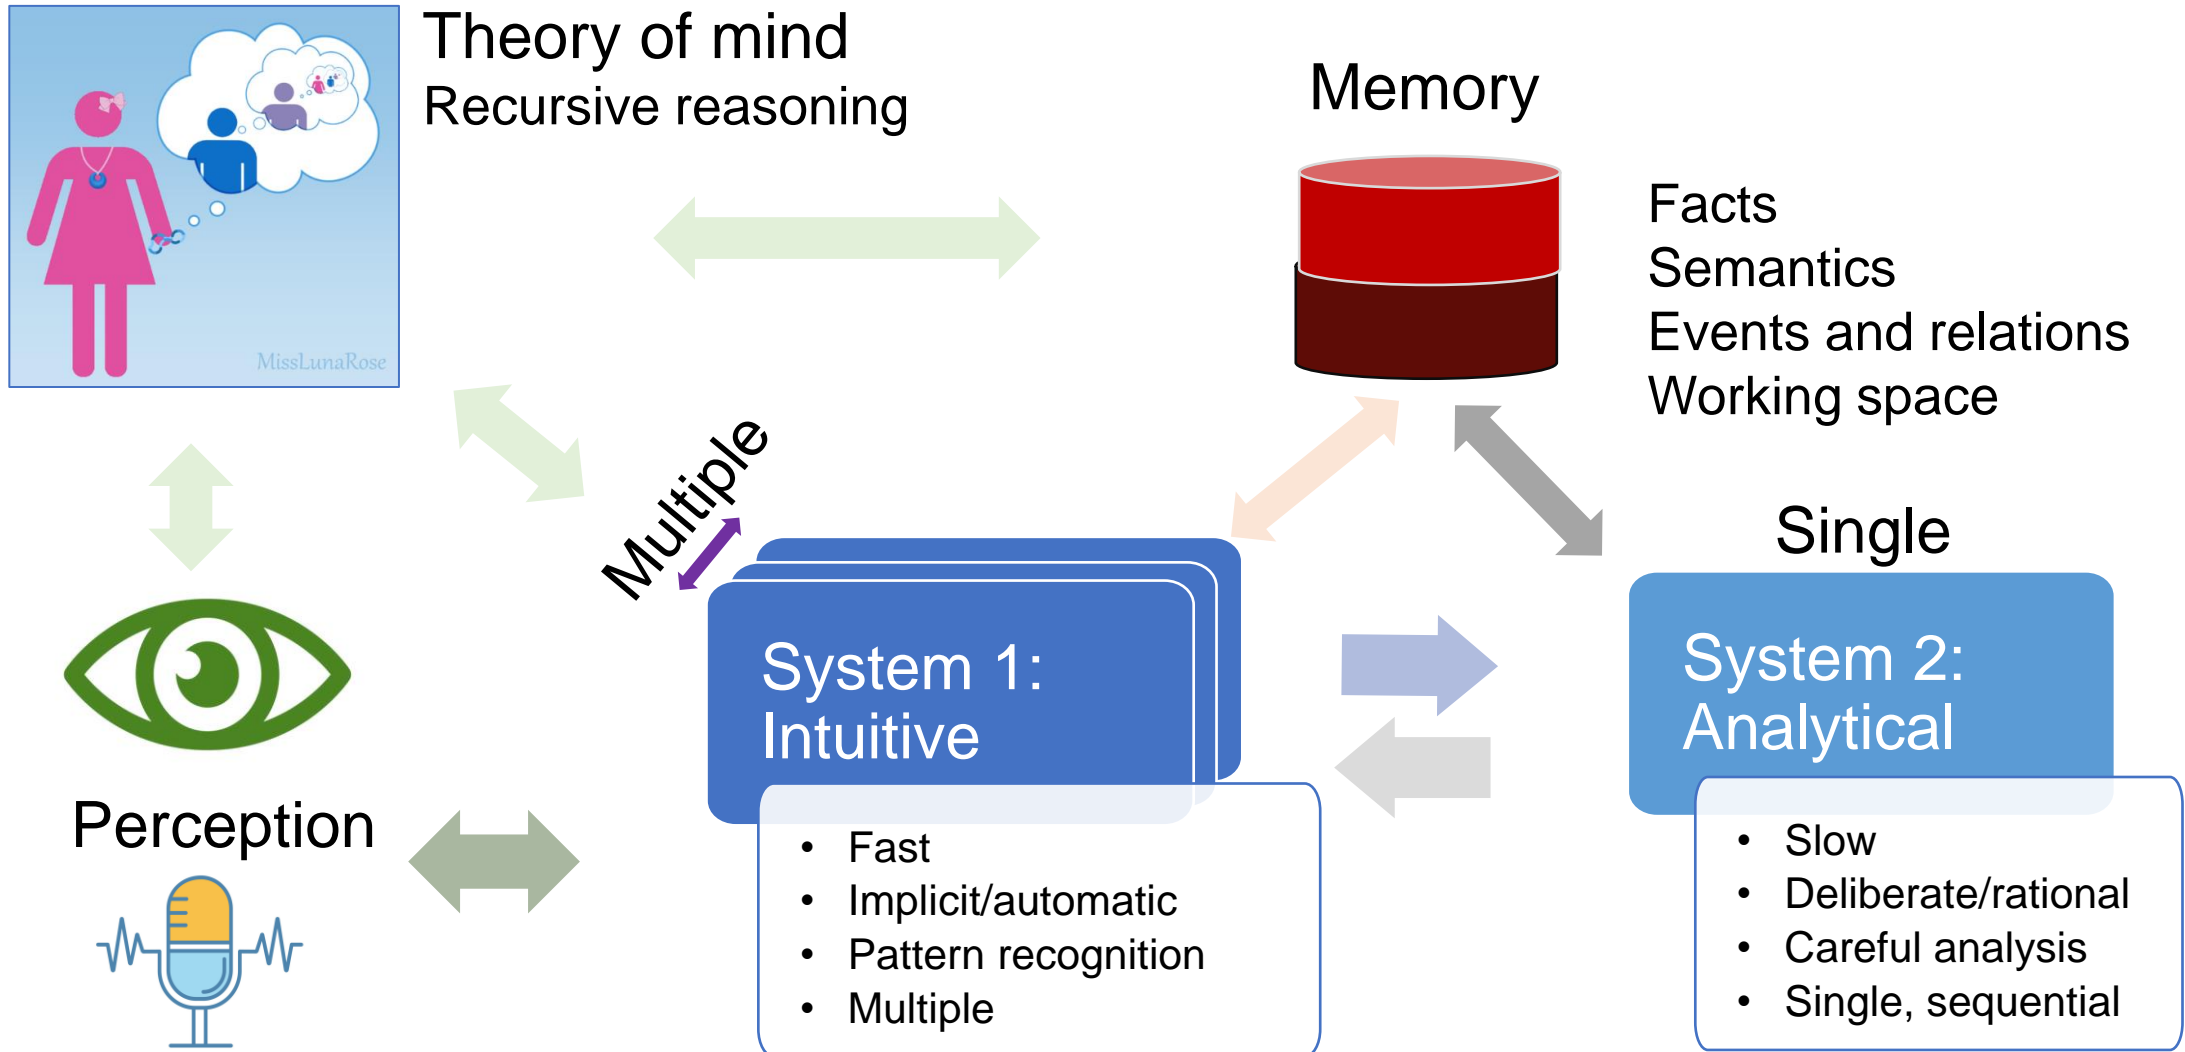
Dang, Long Hoang, Thao Minh Le, Vuong Le, Tu Minh Phuong, and Truyen Tran. "Dynamic Reasoning for Movie QA: A Character-Centric Approach." *IEEE Transactions on Multimedia* (2023).²⁴

Machine reasoning: Multimodality



Hung Le et al. Work in progress.

A possible cognitive architecture: Going social



AI enables thinking about BIG questions

- AGI is nearer. Even already here ... to some
- What if all modalities are connected?
- What will human be like if some of our brain is silicon-based, but not carbon-based?
- Is AI conscious? How is about the current AI infrastructure with billions of connected AI agents?
- Can science be automated?
- What will AI look like in the post-quantum era? Quantum AI?


THE SINGULARITY IS NEARER



When We Merge With AI

RAY KURZWEIL

'The best person I know at predicting the future of artificial intelligence' **BILL GATES**



Thank you!