AI4Science: The 5th Paradigm

Prof Truyen Tran
Head of AI, Health and Science

truyen.tran@deakin.edu.au
truyentran.github.io
@truyenoz
@tranthetruyen
shorturl.at/r3R6i
truyen-tran

My full-loop journey from physics to AI and back

Me and my teammates representing Vietnam at the International Physics Olympiad 1997, Canada.

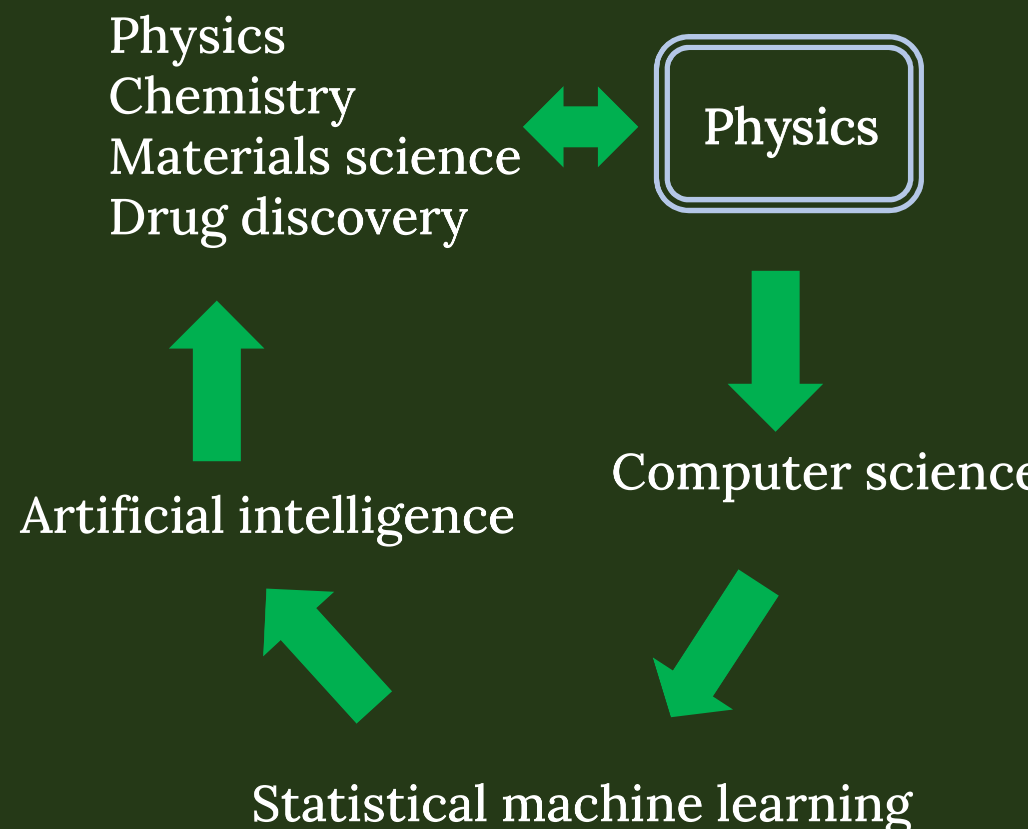

Collaborations

Topics

- Why AI for Science?
- Representation
- Prediction
- Optimization & Generalization
- Explanation
- Physics-informed ML
- AI Co-scientist
- Future & Risks

AI for Science

- "Progress in science depends on new techniques, new discoveries and new ideas, probably in that order."
  - => AI is a great set of techniques!
- Investments of AI4Science are on the rise
  - Governments
  - Big Tech (Google, Microsoft, Meta, IBM, NVIDIA, etc).
  - Startups
  - Universities

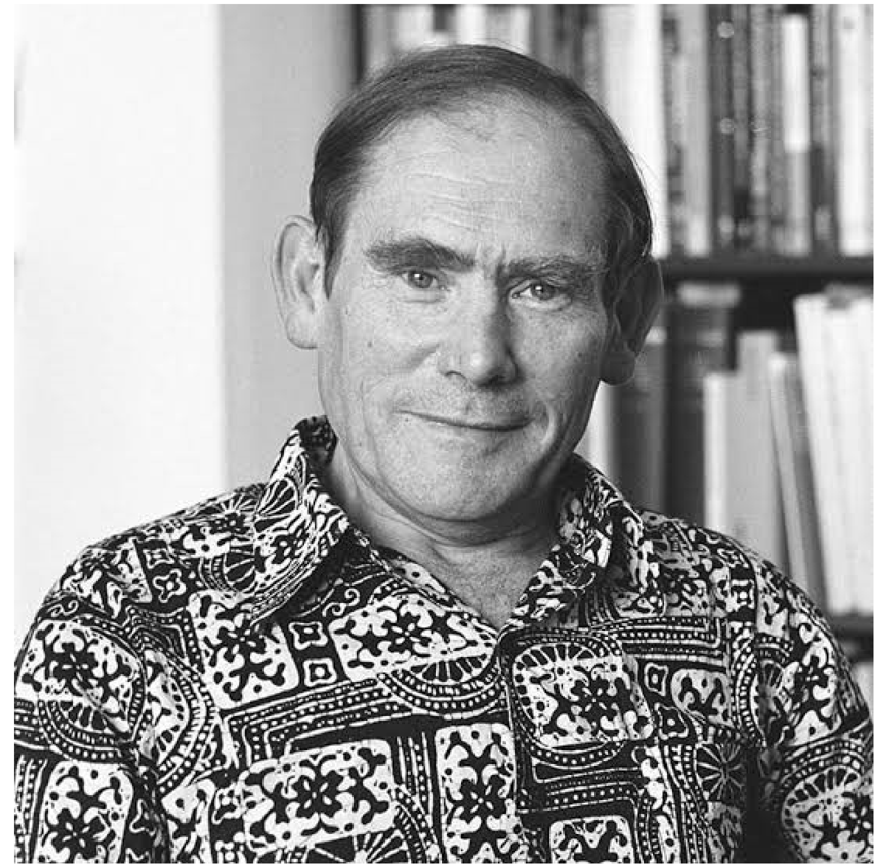

Sydney Brenner (1927 – 2019)
2002 Nobel Prize in Physiology

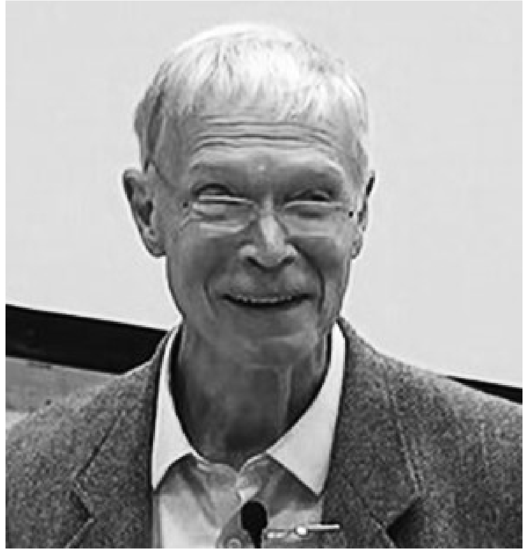

“Physics is a point of view that the world around us is, with effort, ingenuity, and adequate resources, understandable in a predictive and reasonably quantitative fashion."

(John Hopfield, 2024 Nobel Prize in Physics)

Academic recognition

- 2018 Turing Award (~Nobel Prize in Computing)
- 2024 Nobel Prize in Physics – foundations of AI
- 2024 Nobel Prize in Chemistry – AlphaFold, protein folding
- 2025 Turing Award (~Nobel Prize in Computing)

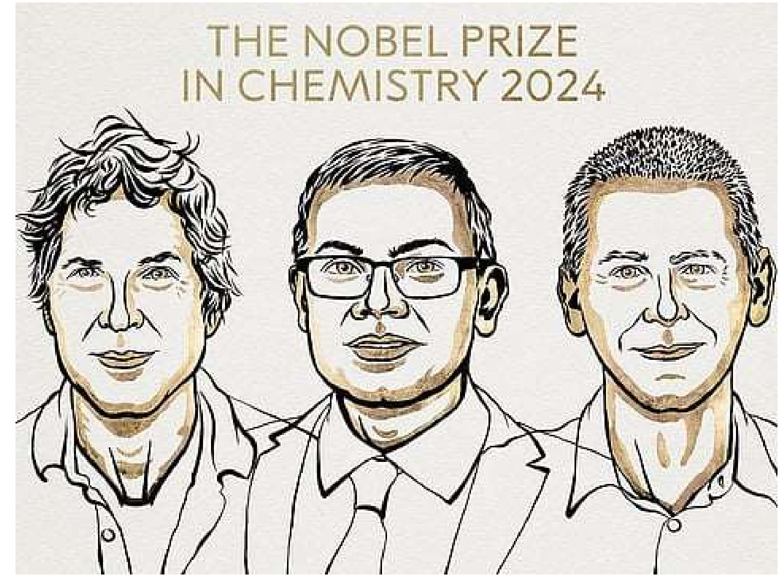

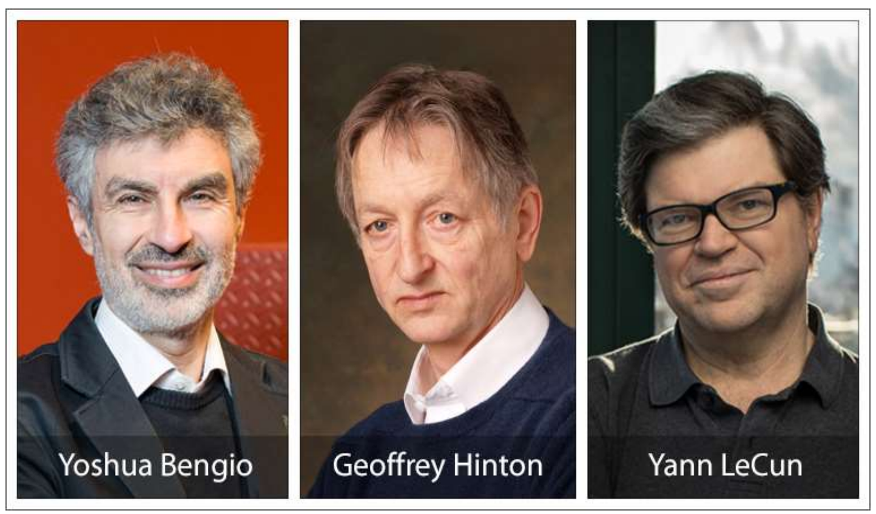

Turing Awards 2018

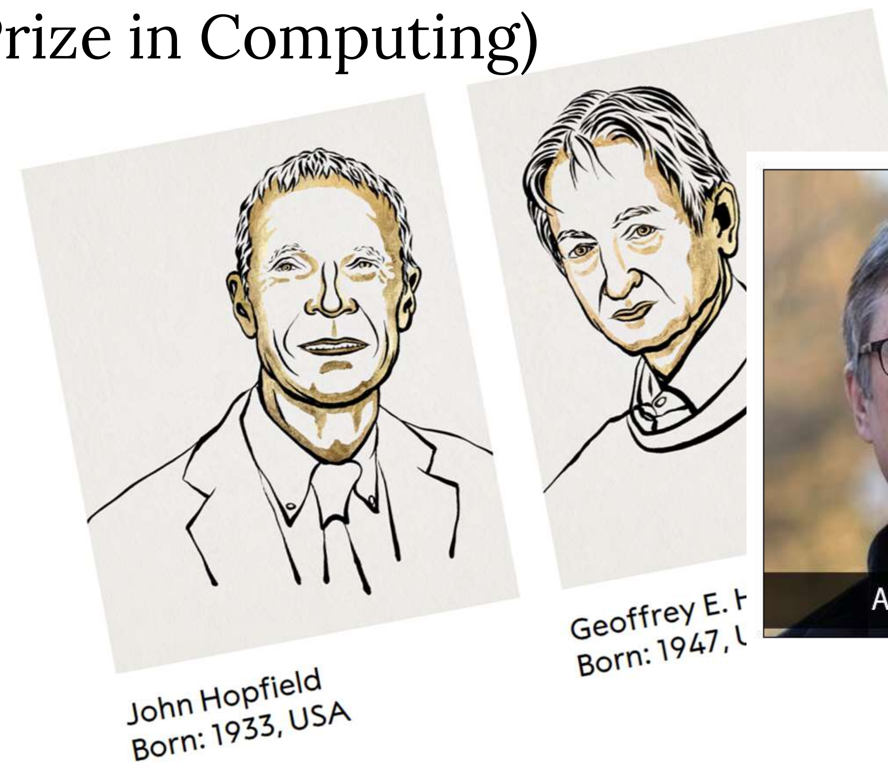

Turing Awards 2025

AI is general purpose technology (GPT)

Image source: inc42

The 100 years of AI making …

AGI = Artificial General Intelligence

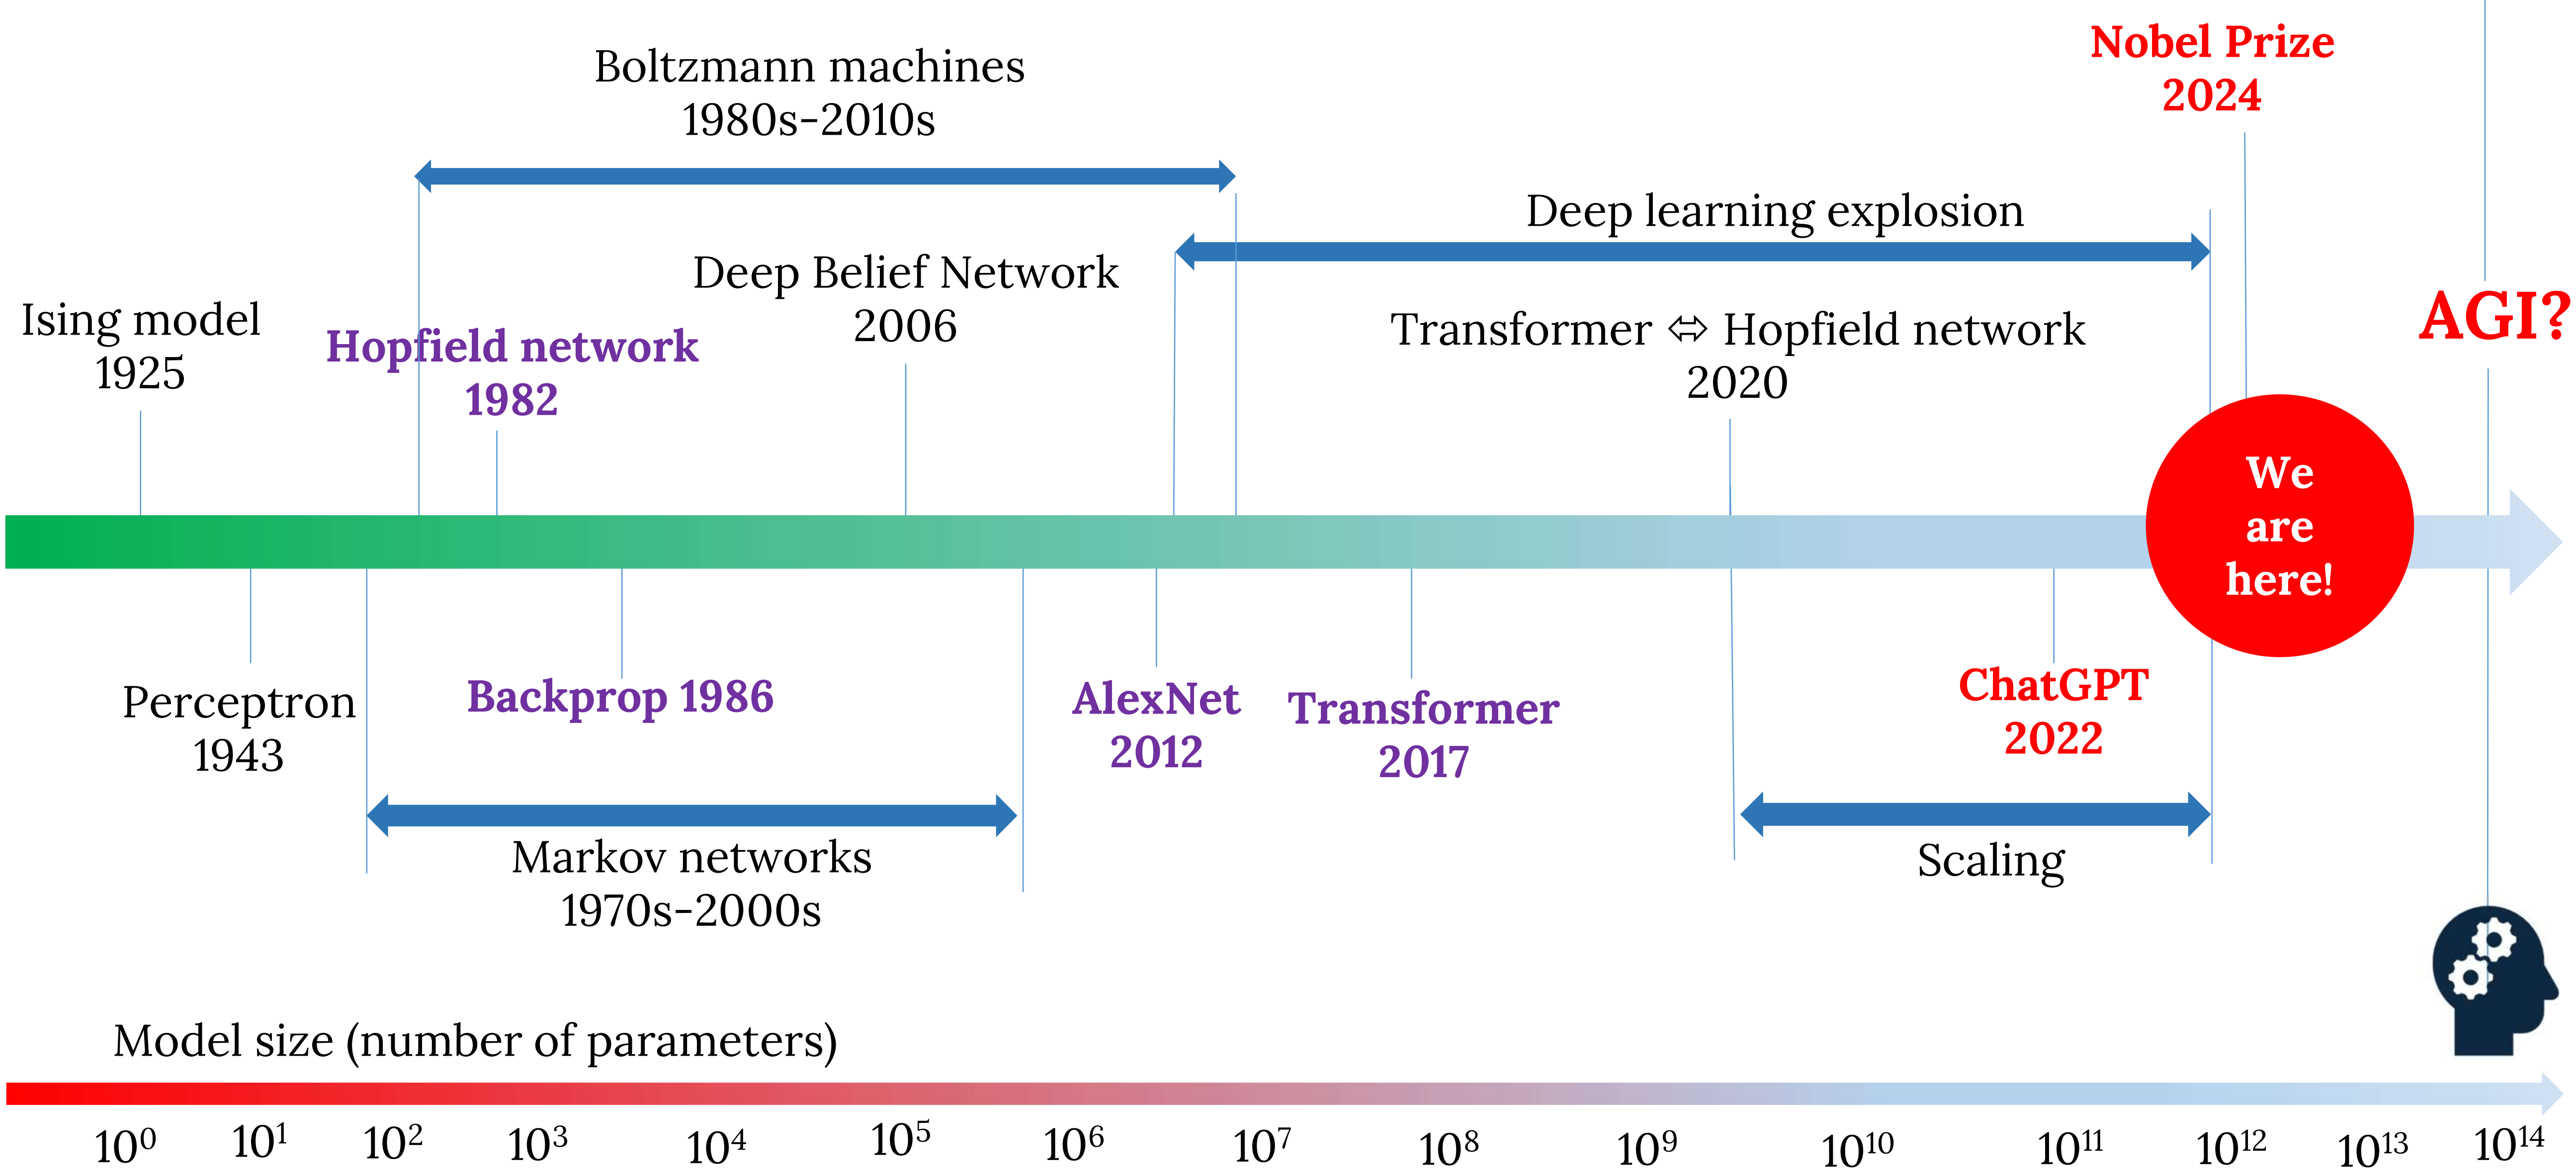

Image source: AssemblyAi

The (sudden) emergence

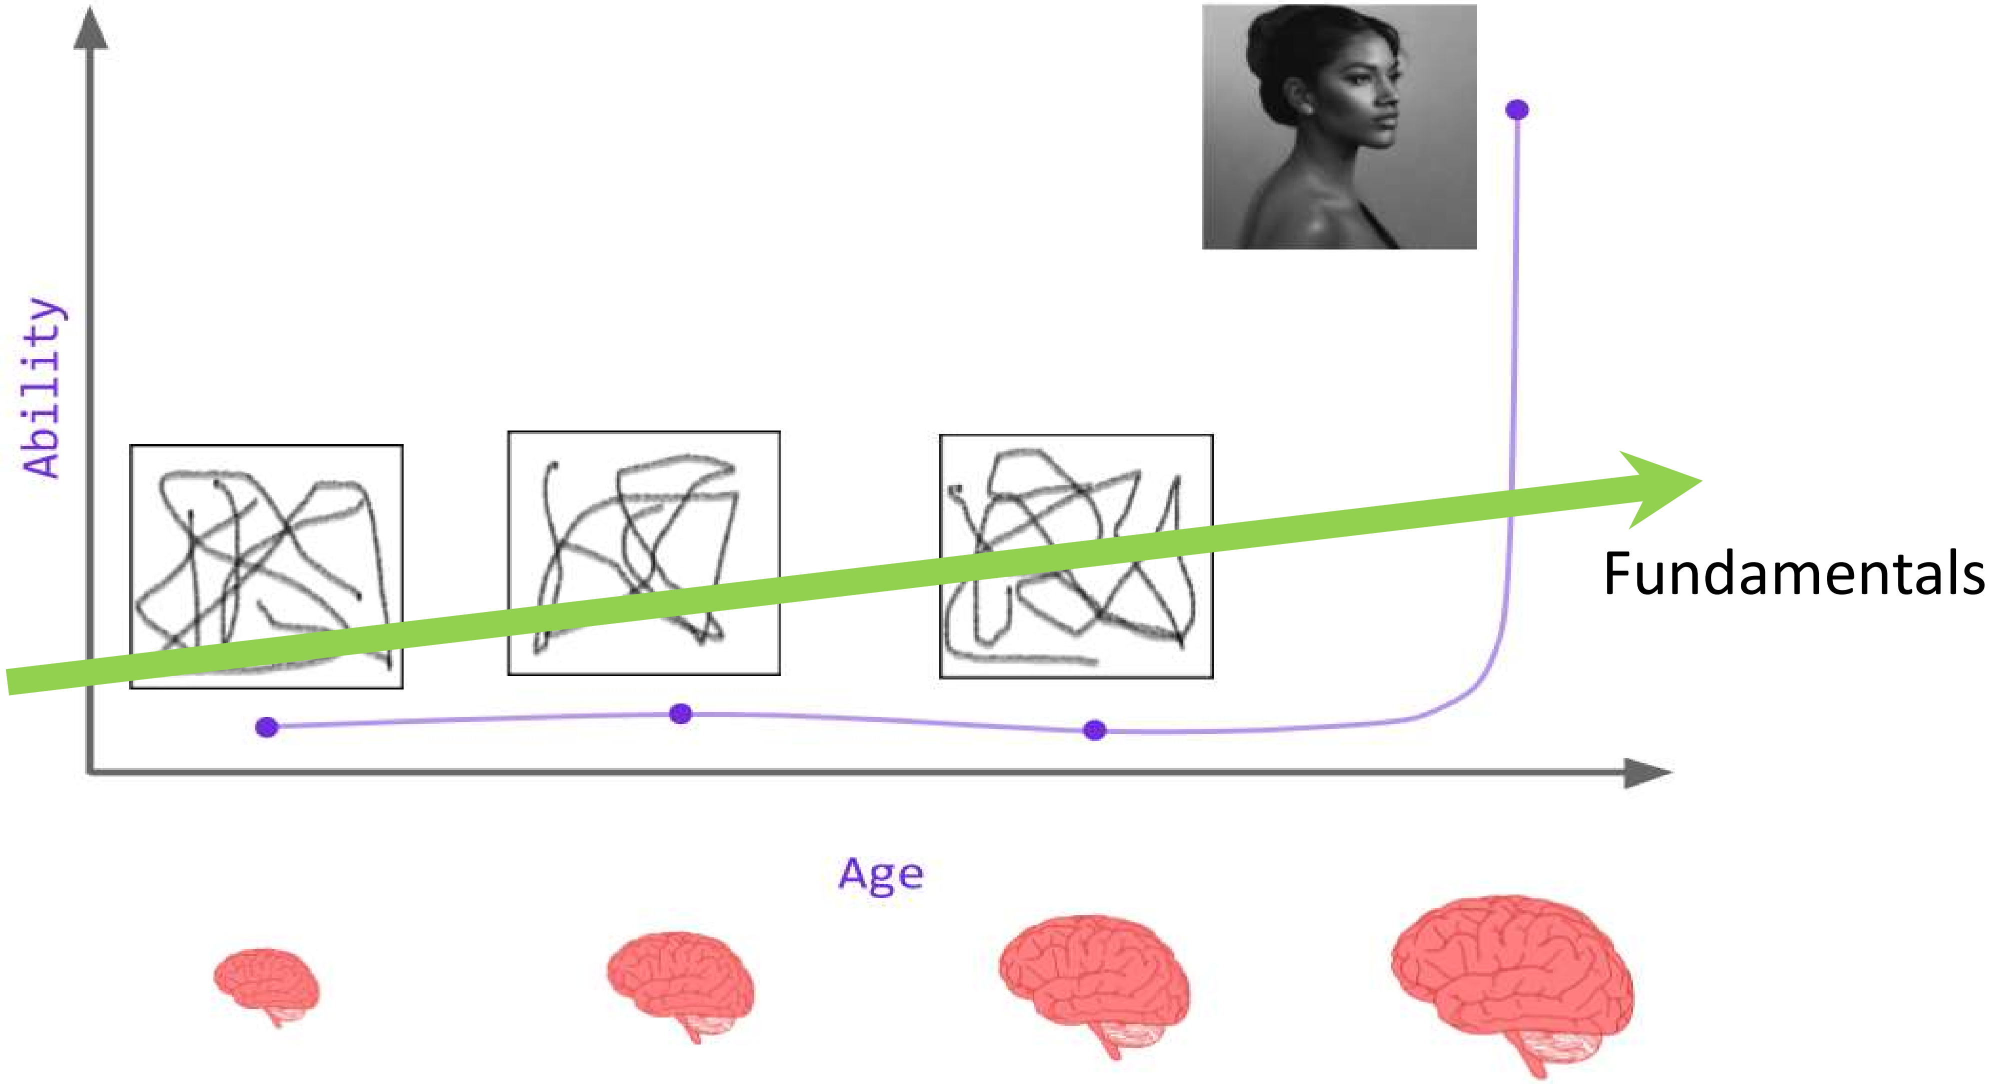

Source: assemblyai

A tipping point: Foundation models

- A foundation model is a model trained at broad scale that can adapted to a wide range of downstream tasks
- Scale and the ability to perform tasks beyond training
- Emergence of new properties
- Uniformity of architecture.

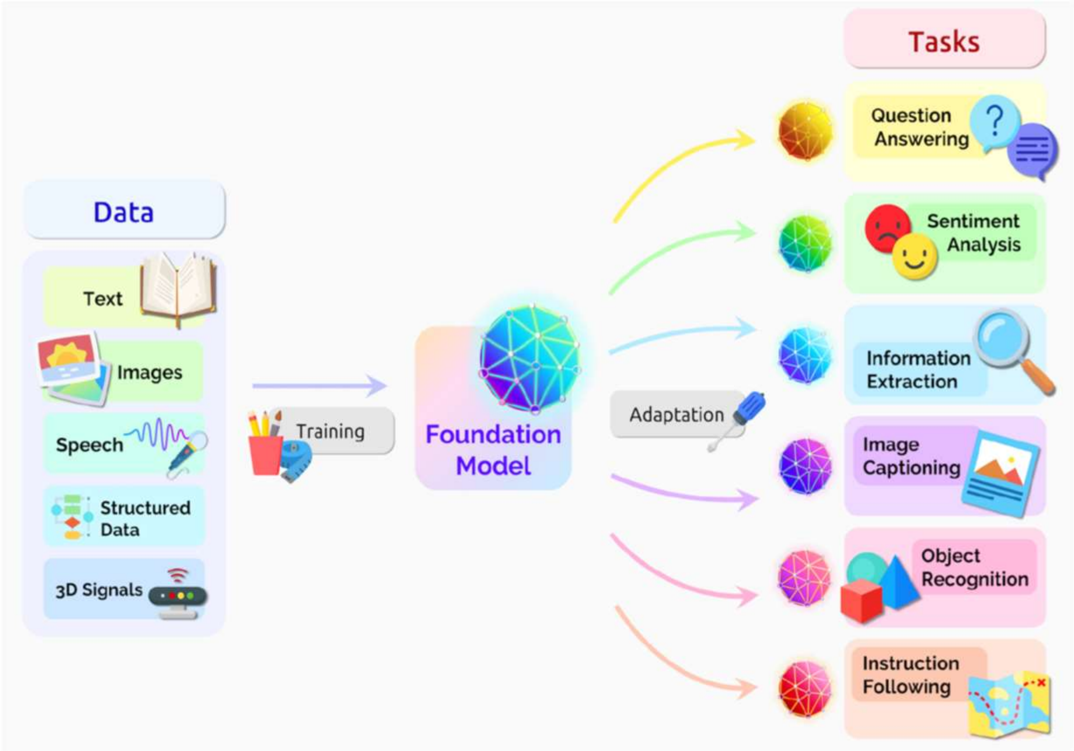

Picture taken from (Bommasani et al, 2021)

Slide credit: Samuel Albanie, 2022

Credit: DeepMind

One model for all – the case of Gato (2022)

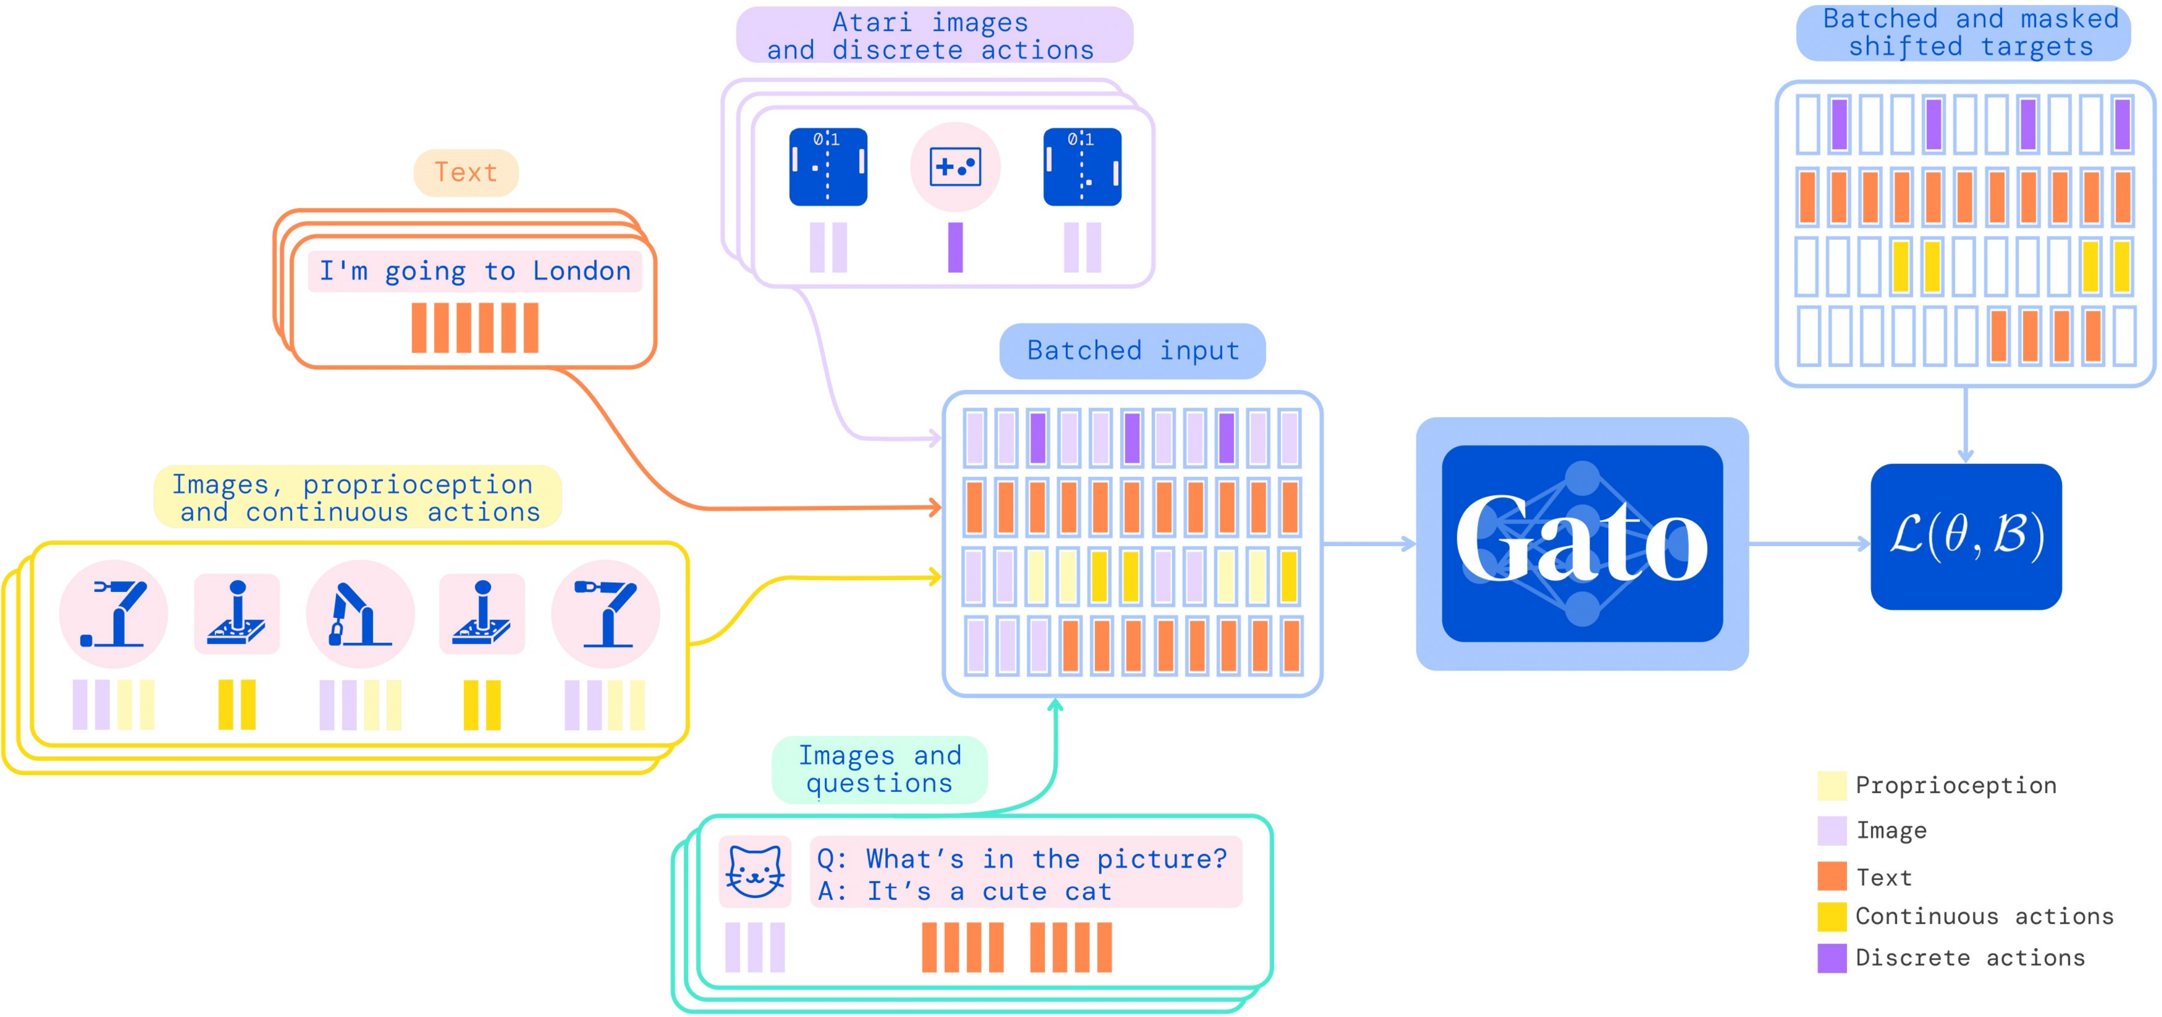

Reed, Scott, Konrad Zolna, Emilio Parisotto, Sergio Gomez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Gimenez et al. "A generalist agent." *arXiv preprint arXiv:2205.06175* (2022).

Why one-model-for-all possible?

The world is regular: Rules, patterns, motifs, grammars, recurrence

- World models are learnable from data!

Human brain gives an example

- One brain processes all modalities, doing plenty of tasks, and learning from different kind of training signals.
- Thinking at high level is independent of input modalities and task-specific skills.

Three kinds of AI use

- 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.

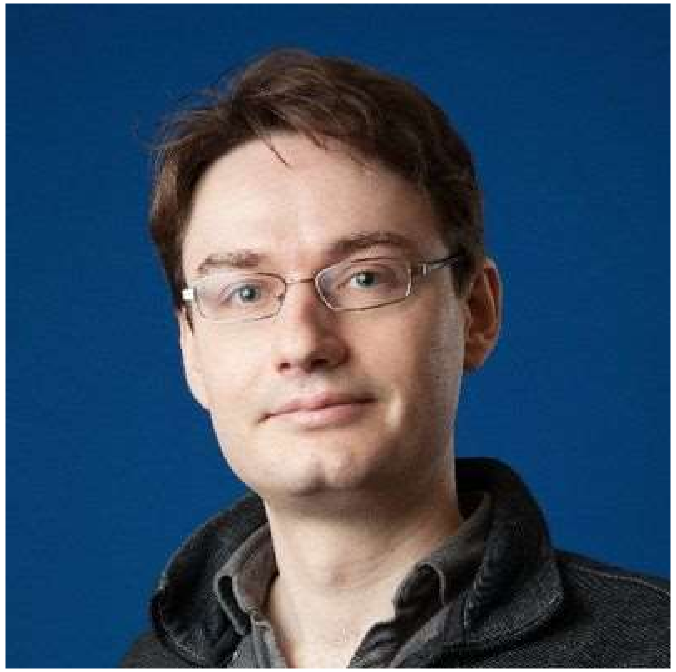

François Chollet

Technical capabilities

Predictive AI: Pattern recognition, out-of-distribution detection, prediction.

Generative AI: Generating new designs to meet performance criteria.

Agentic AI: Coordinating in TEAMS to achieve a goal by themselves.

Optimization: Refining the generated designs to optimise the performance.

Mental capabilities

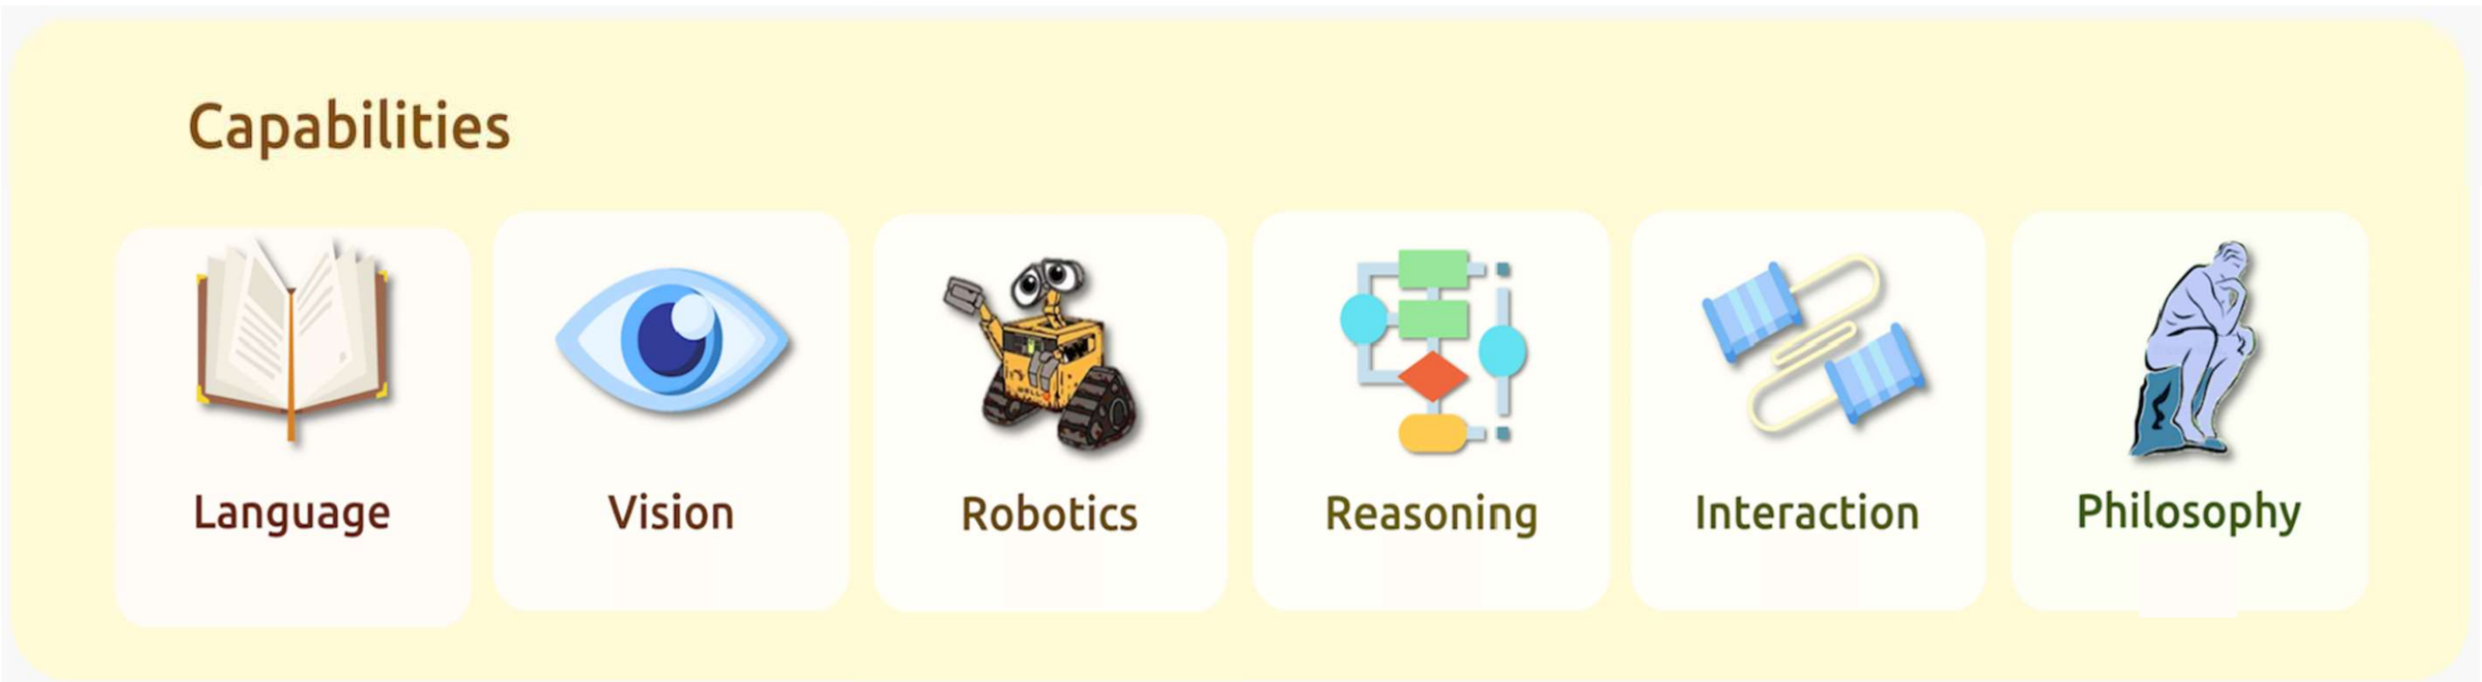

Image credits/References:
R. Bommasani et al., "On the opportunities and risks of foundation models", arxiv (2021)
(ImageNet) O. Russakovsky et al., "Imagenet large scale visual recognition challenge", IJCV (2015)
(CLIP) A. Radford et al., "Learning transferable visual models from natural language supervision", ICML (2021)
D. Silver et al., "Mastering the game of Go with deep neural networks and tree search", Nature (2016)

The shifting AI research

Engineering

Design man-made systems

Generative AI

Discover emergent behaviours

Science

Discover laws in nature

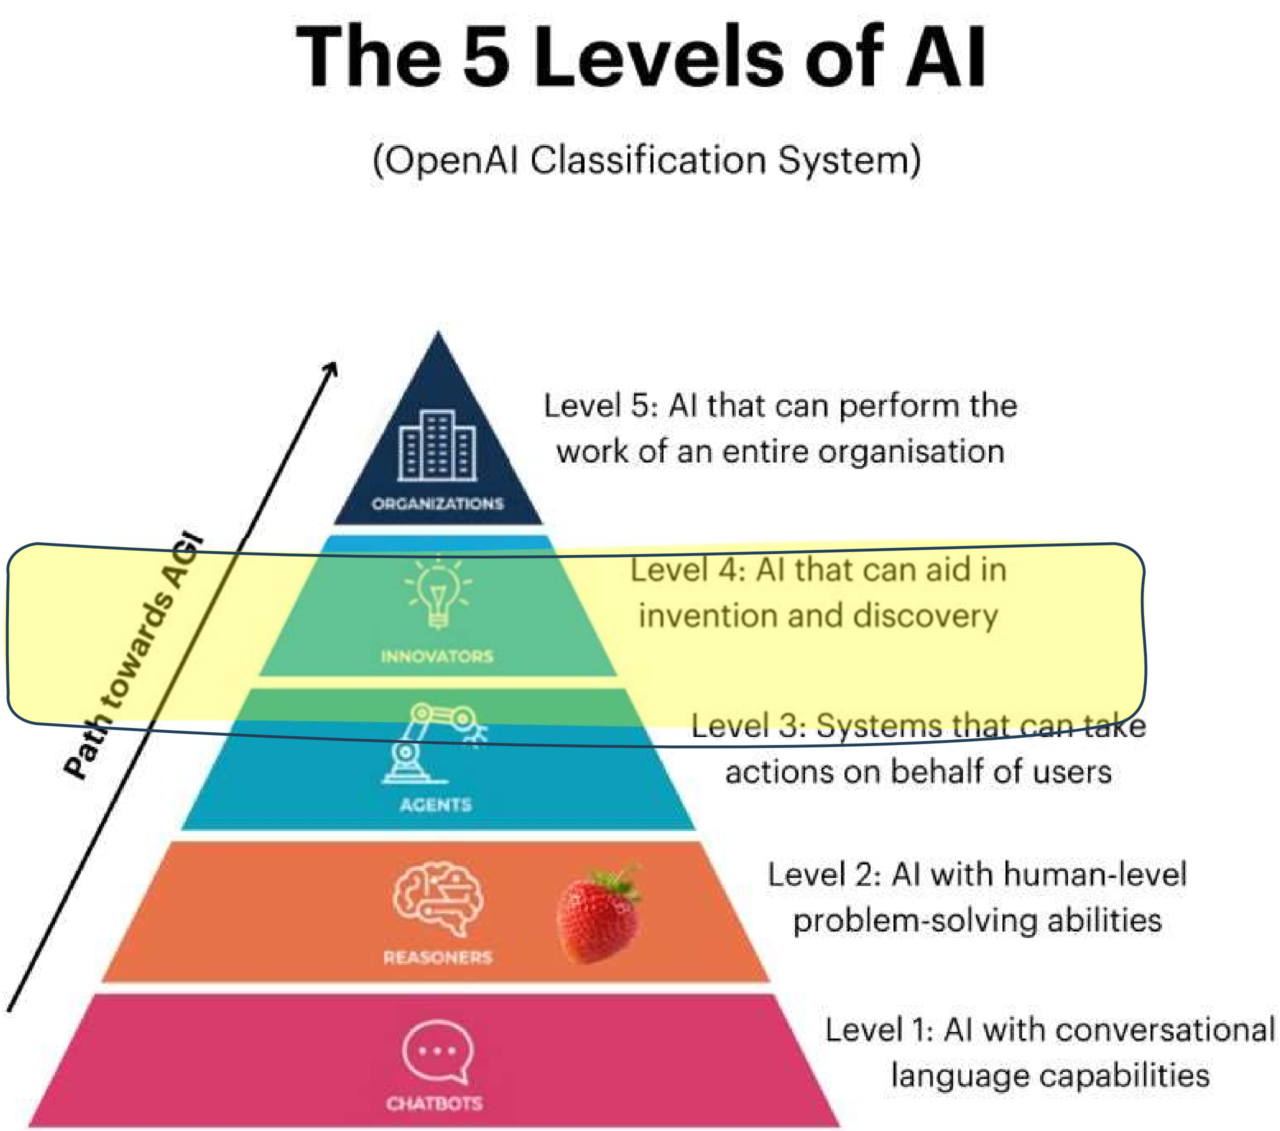

Source: QuestionPro

Credit: McLennan

Why is AI co-scientist possible?

The world is regular: Rules, patterns, motifs, grammars, recurrence

> World models are learnable from data (real or simulated)!

> "Any pattern that can be generated or found in nature can be efficiently discovered or modelled by a classical learning algorithm".

(***Demis Hassabis***, *World Chess Championship*, CEO *of Google DeepMind*, *Nobel Laureate in Chemistry*, 2024)

exposition ⇔ pretraining
(background knowledge)

worked problems ⇔ supervised finetuning
(problem + demonstrated solution, for imitation)

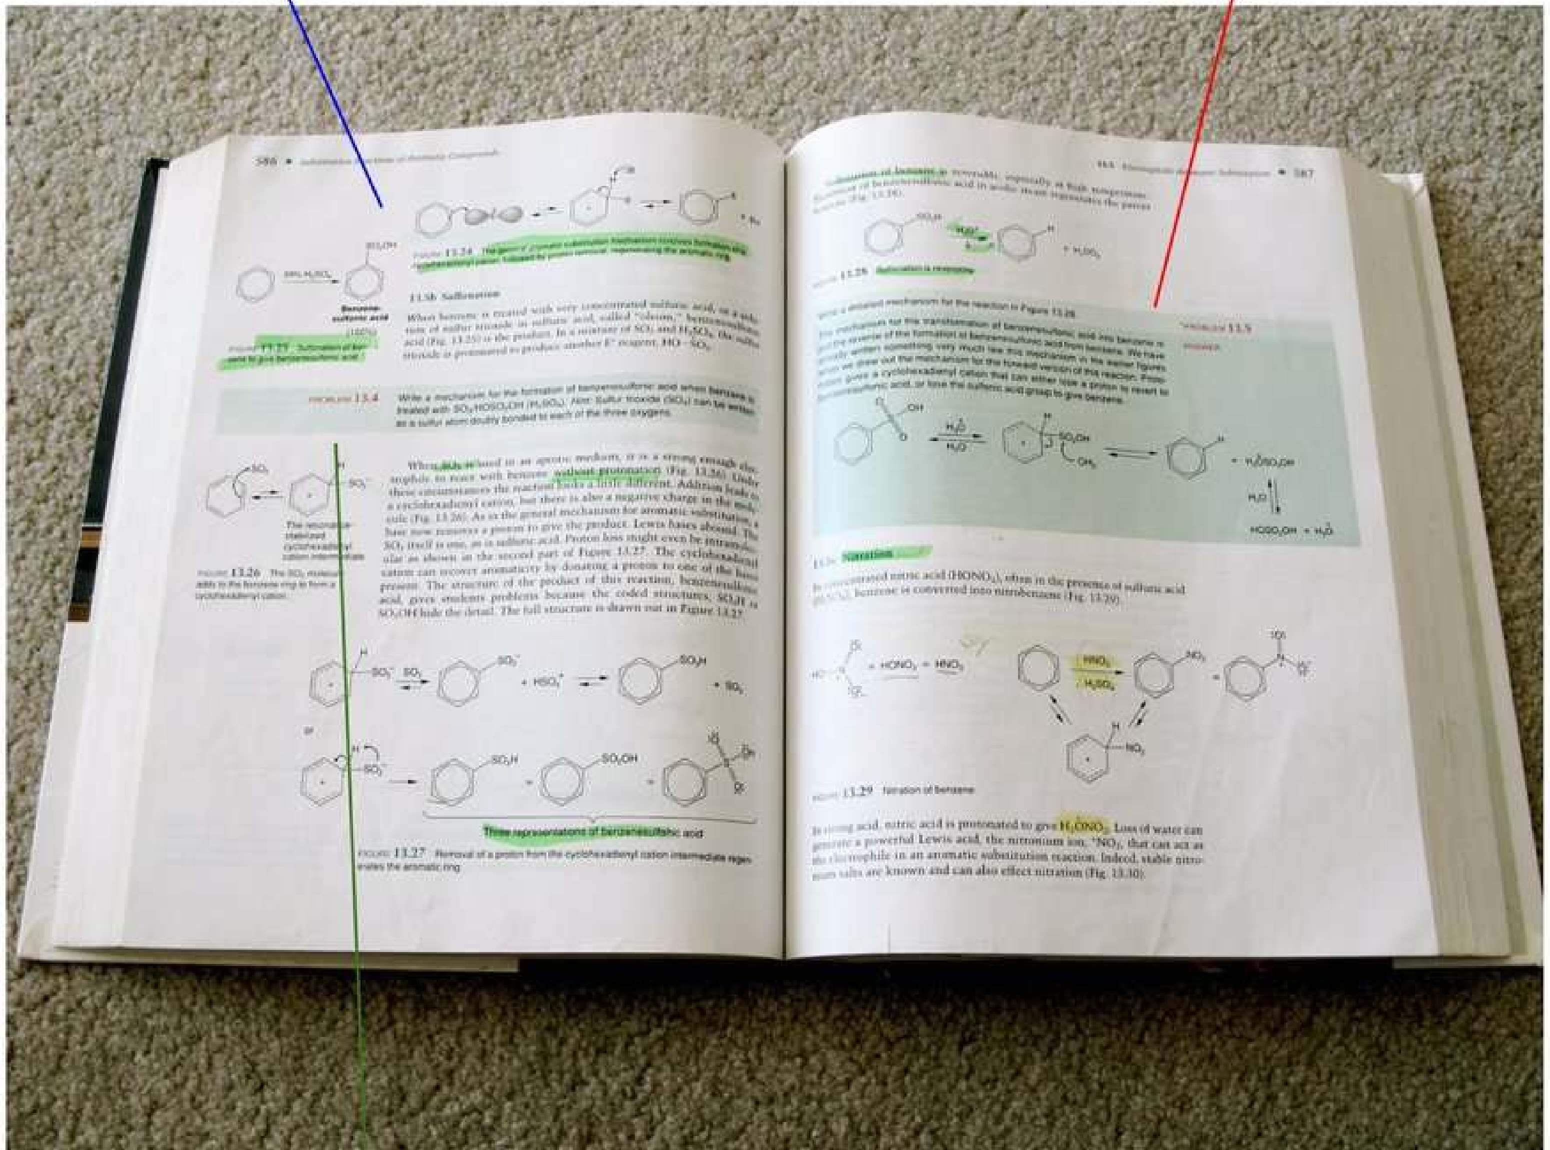

practice problems ⇔ reinforcement learning
(prompts to practice, trial & error until you reach the correct answer)

Modern AI training mimics human training

=> *Human scientists learn, so does AI*

Source: Andrej Karpathy

A simple plan

- Build generalist AI
- Use AI to solve science

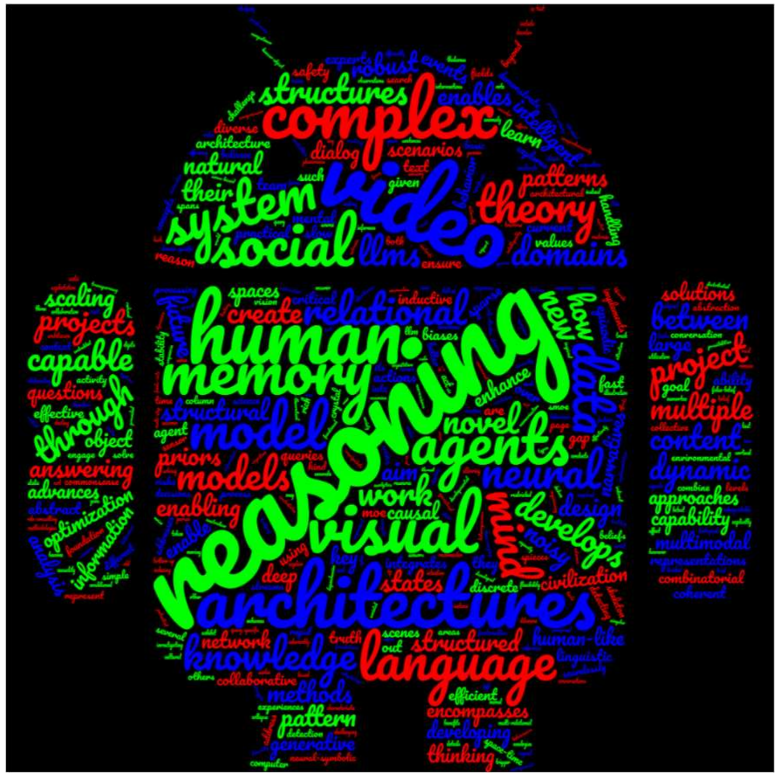

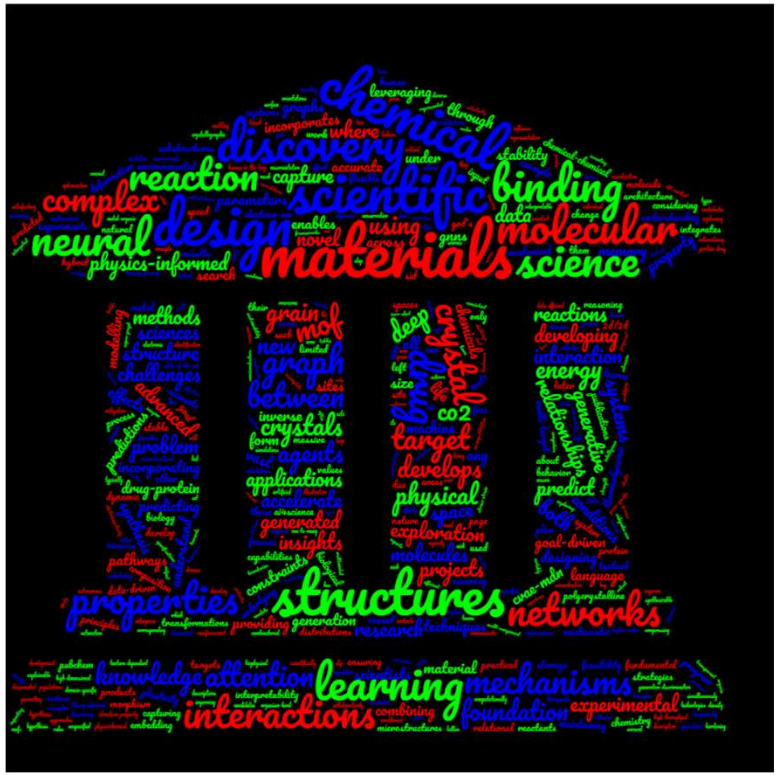

Generalist AI in the making, for proteomics

Protein Large Language Models: A Comprehensive Survey

Yijia Xiao[♠] *, Wanjia Zhao[♣], Junkai Zhang[♠], Yiqiao Jin[♥], Han Zhang[♠],
Zhicheng Ren[♠], Renliang Sun[♠], Haixin Wang[♠], Guancheng Wan[♠], Pan Lu[♣],
Xiao Luo[♠], Yu Zhang[♦], James Zou[♣], Yizhou Sun[♠], Wei Wang[♠]
[♠]UCLA, [♣]Stanford, [♥]Georgia Tech, [♦]Texas A&M

https://github.com/Yijia-Xiao/Protein-LLM-Survey

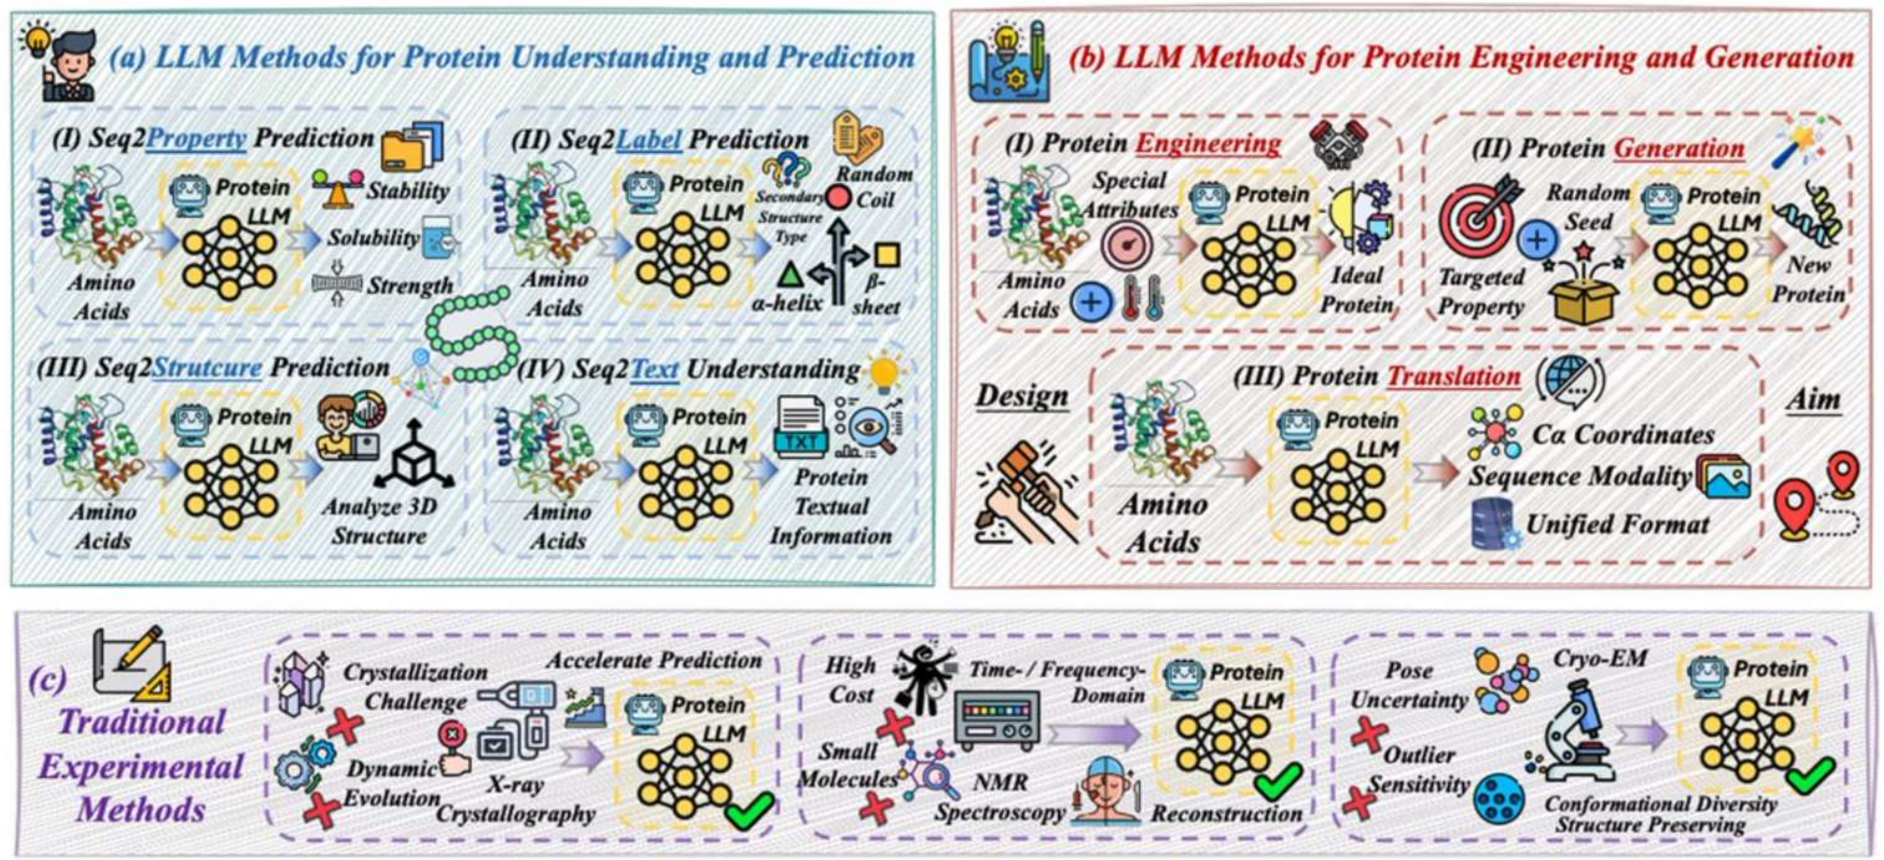

A not so simple plan

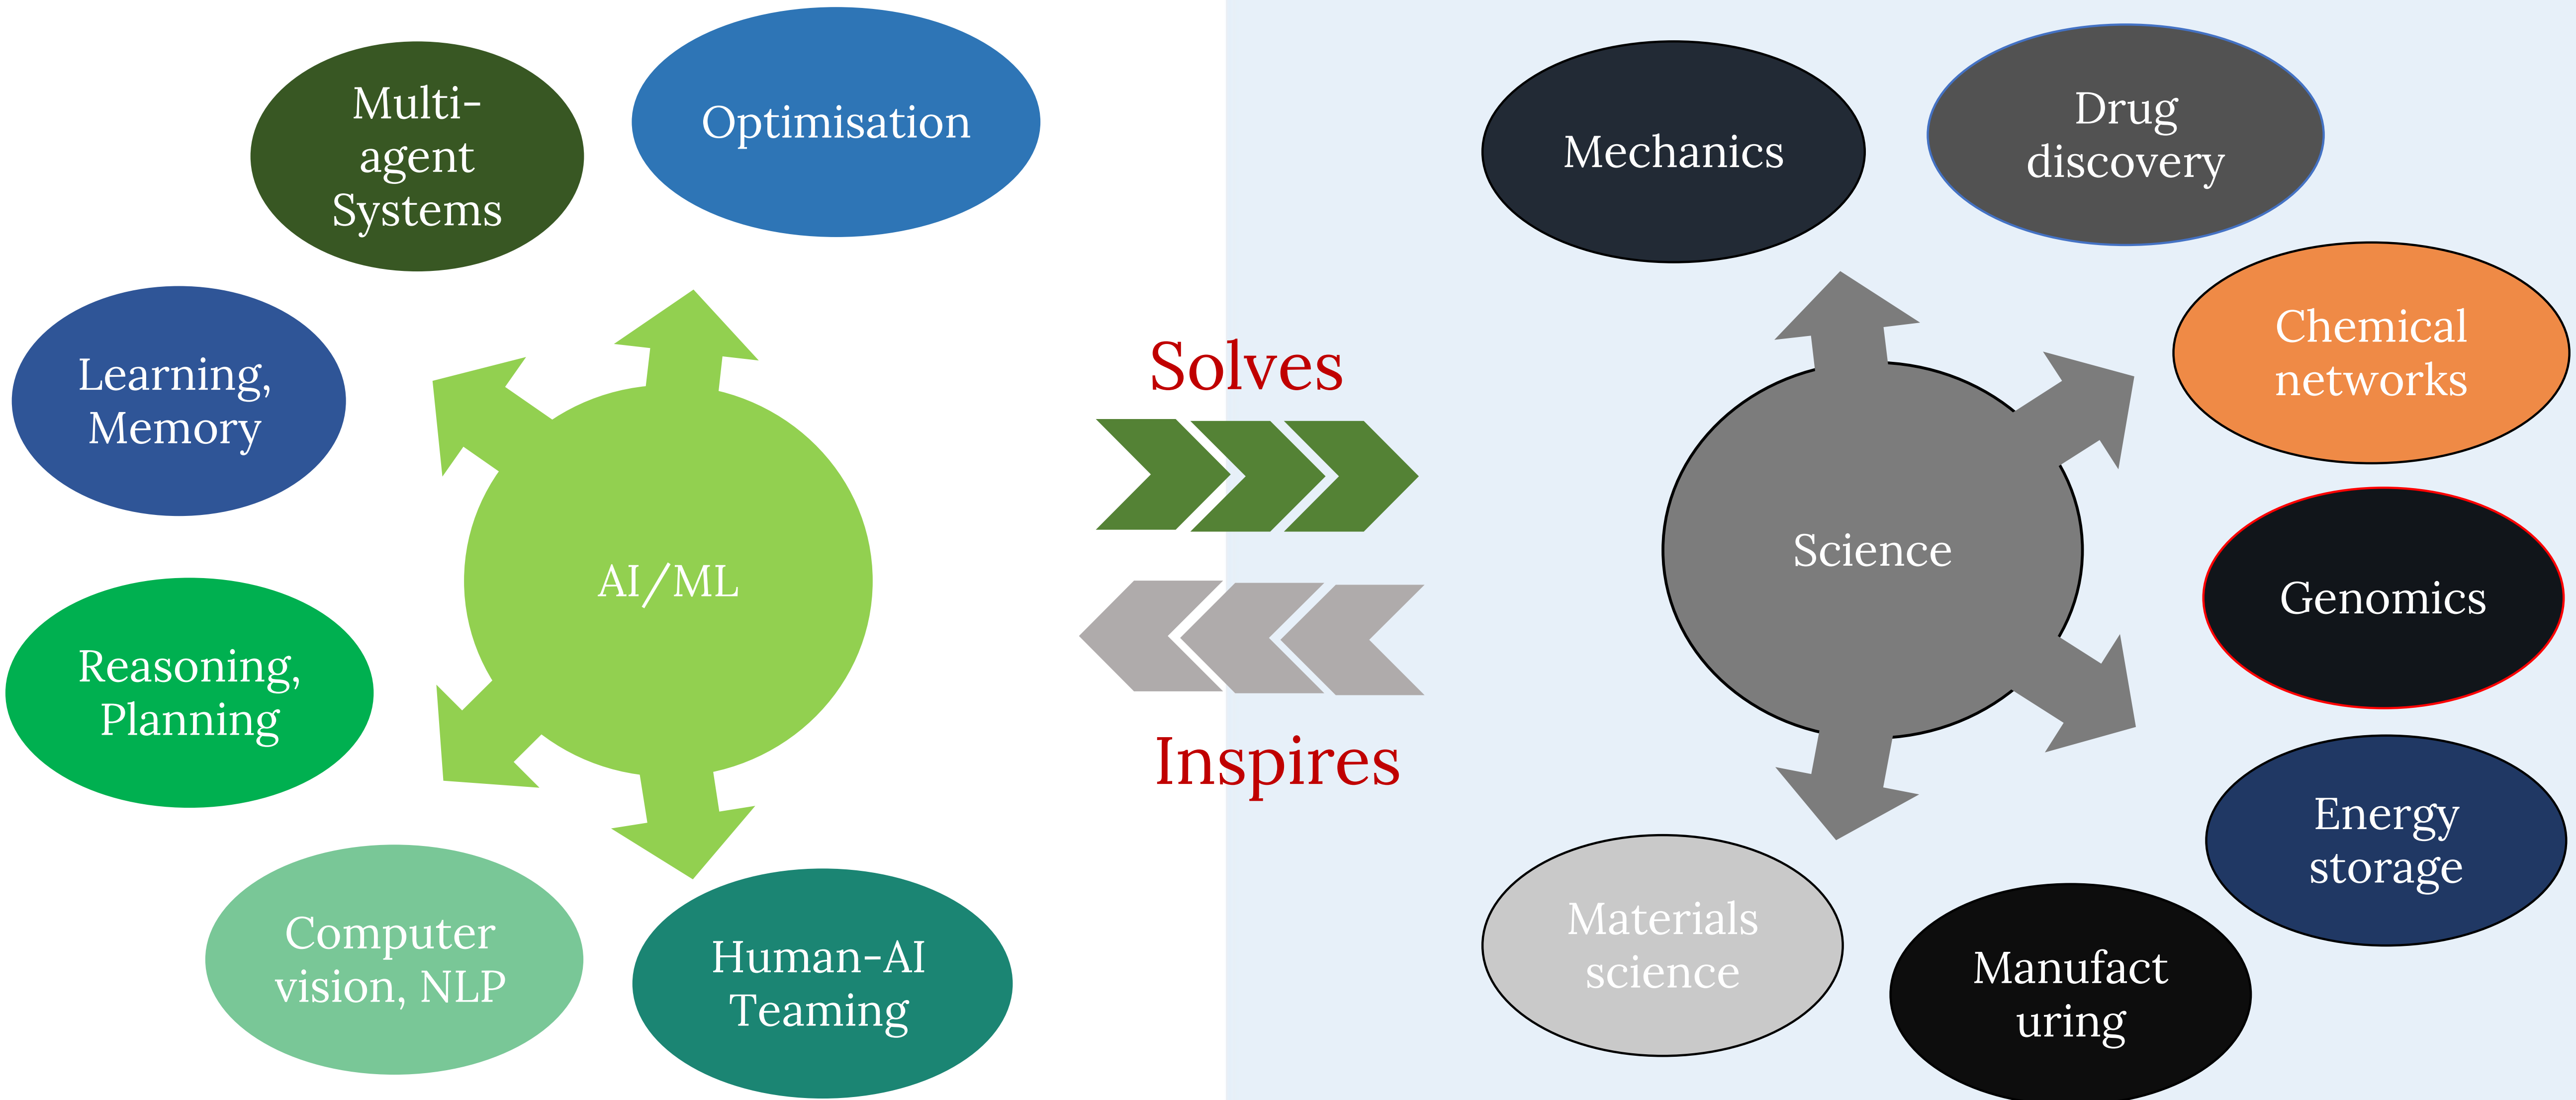

The fifth paradigm of science

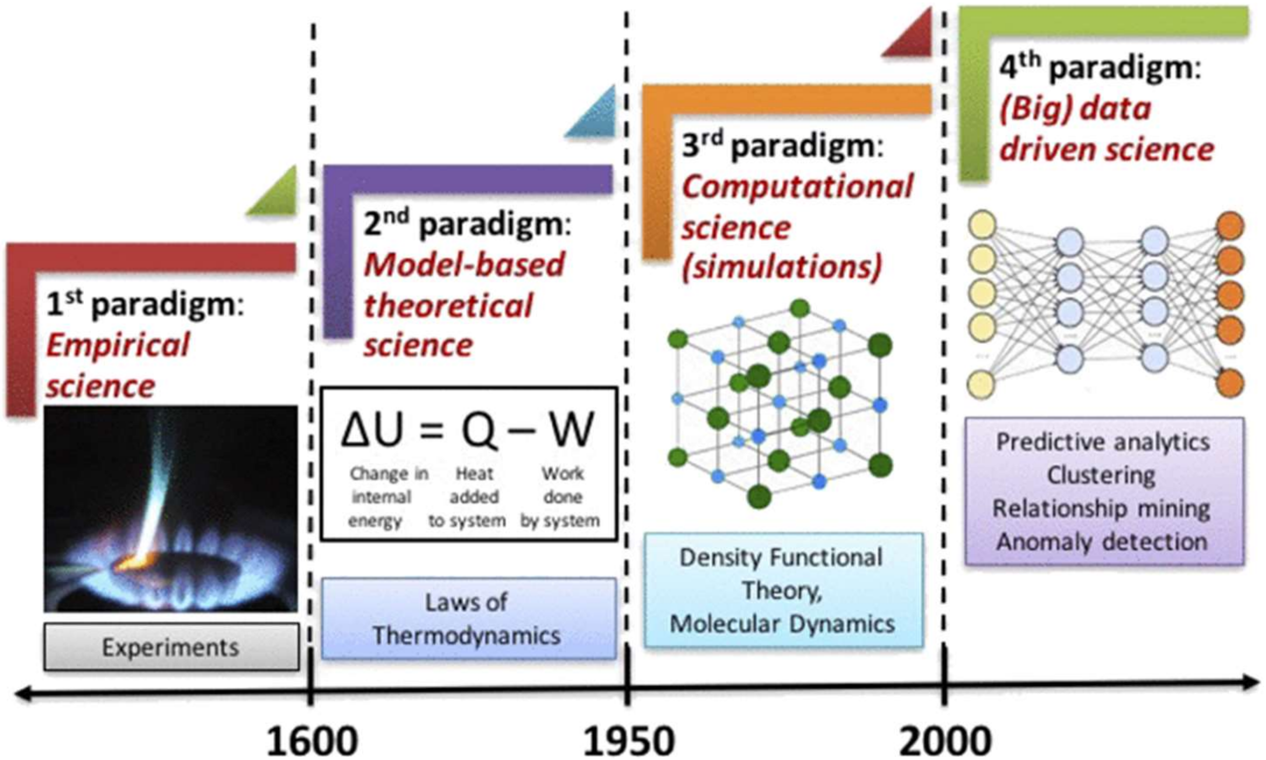

The 5th paradigm (2020-present)

- Advanced deep learning
- Massive data simulation
- Powerful Foundation Models

Agrawal, A., & Choudhary, A. (2016). Perspective: Materials informatics and big data: Realization of the "fourth paradigm" of science in materials science. *Apl Materials*, 4(5), 053208.

https://www.microsoft.com/en-us/research/blog/ai4science-to-empower-the-fifth-paradigm-of-scientific-discovery/

AI4Science: The "Law of Accelerating Returns" at work

- The entire history of human's invention: Tool that produces tools
  - For past 50 years: Software that writes software
  - Basis for the prediction of Singularity in 2045 by Ray Kurzweil
- Science & AI
  - Science invented computer
  - Computer helps build AI
  - AI accelerates science
- => This loops seem to help accelerate entropy increase in the 2nd Law of Thermodynamics!

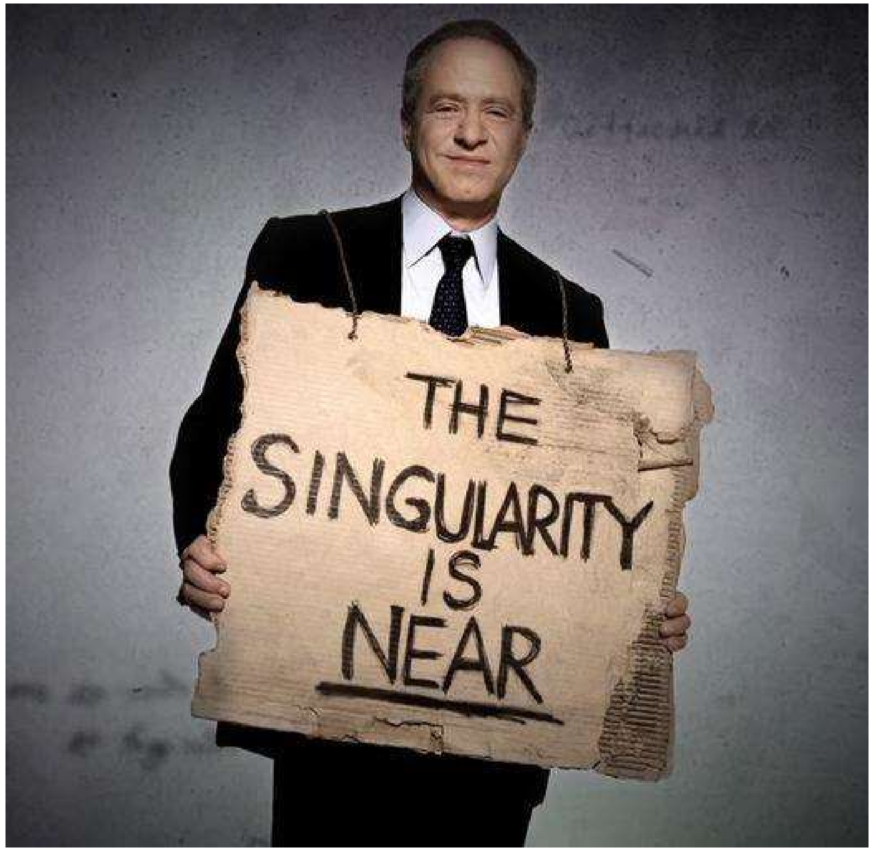

Image: Shutterstock

Space of innovation in materials science: An example

Materials

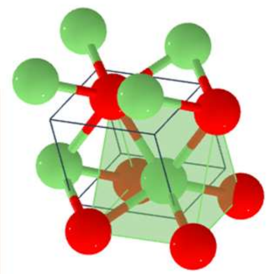

- Molecular space exploration
  - Small, medium, large, supra
- Molecular interaction
  - Network, docking
- Chemical reaction, retrosynthesis
  - Catalyst, yield, free-energy
- Crystal space exploration
- Alloy space exploration
- Microstructures
- Knowledge extraction, coding, expression, manipulation

AI/ML

- Representation
  - Graphs, geometry, periodicity, token
  - Materials manifold
- Learning, attention and memory
  - Self-supervised, supervised, reinforcement
  - Transfer, zero-shot, few-shot, adaptation learning
  - Learning to reason
- Reasoning & planning
  - Planning
  - Optimisation
  - Extrapolation, generation
  - Abductive, inductive, deductive reasoning
  - Tool use

Topics

- ~~Why AI for Science?~~
- Representation
- Prediction
- Optimization & Generalization
- Explanation
- Physics-informed ML
- AI Co-scientist
- Future & Risks

Representing molecules as graphs

- **No regular, fixed-size structures**
- Graphs are ***permutation invariant***:
  - #permutations are exponential function of #nodes
  - The probability of a generated graph G need to be marginalized over all possible permutations
- **Multiple objectives:**
  - **Diversity** of generated graphs
  - **Smoothness** of latent space
  - Agreement with or optimization of multiple "**drug-like**" objectives

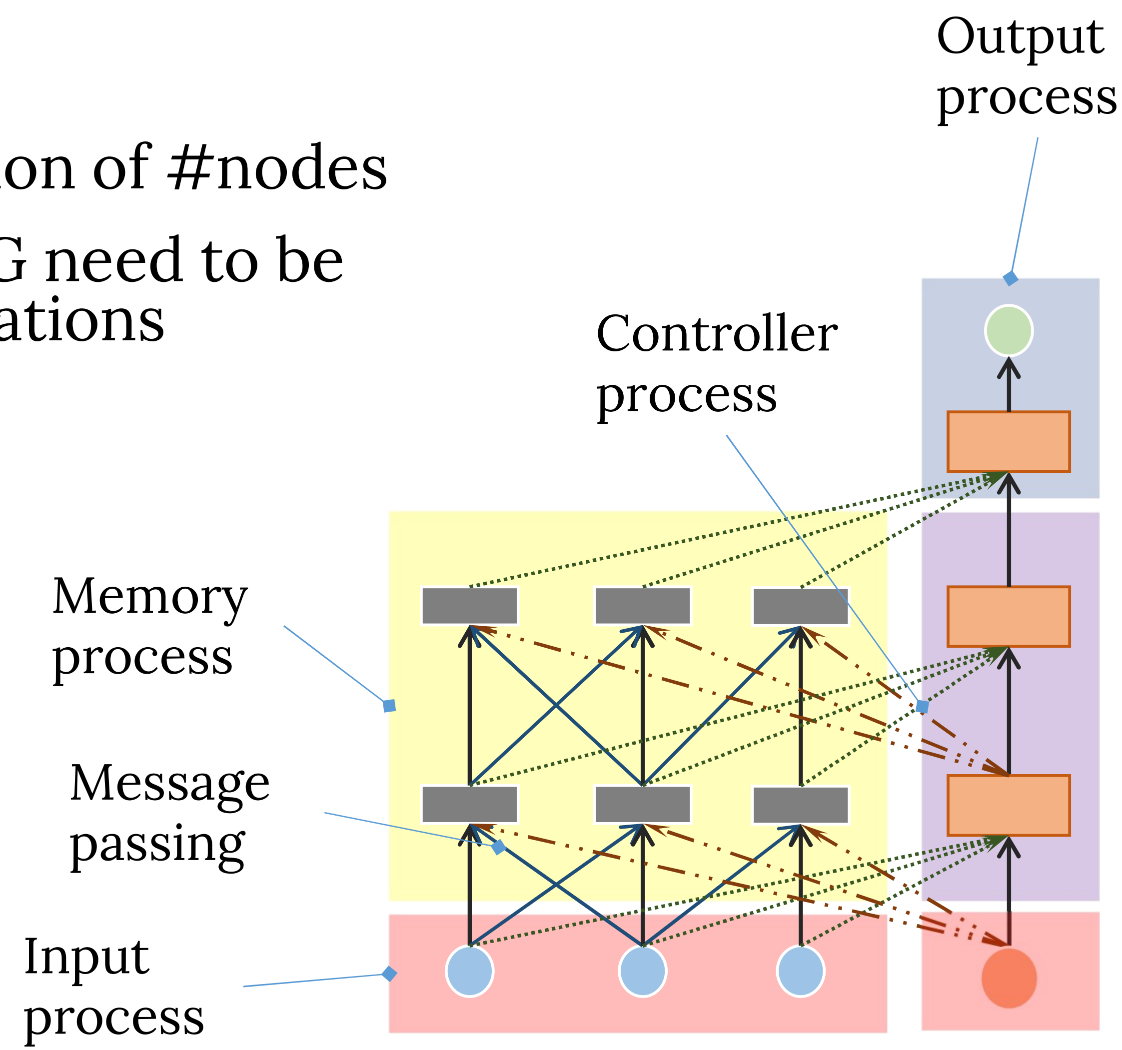

#REF: Pham, T., Tran, T., & Venkatesh, S. (2018). Relational dynamic memory networks. *arXiv preprint arXiv*:1808.04247.

Representing proteins

1D sequence (vocab of size 20) – hundreds to thousands in length

2D contact map – requires prediction

3D structure – requires folding information, either observed or predicted. Now available thanks to AlphaFold 2.

NLP-inspired embedding (word2vec, doc2vec, glove, seq2vec, ELMo, BERT, GPT).

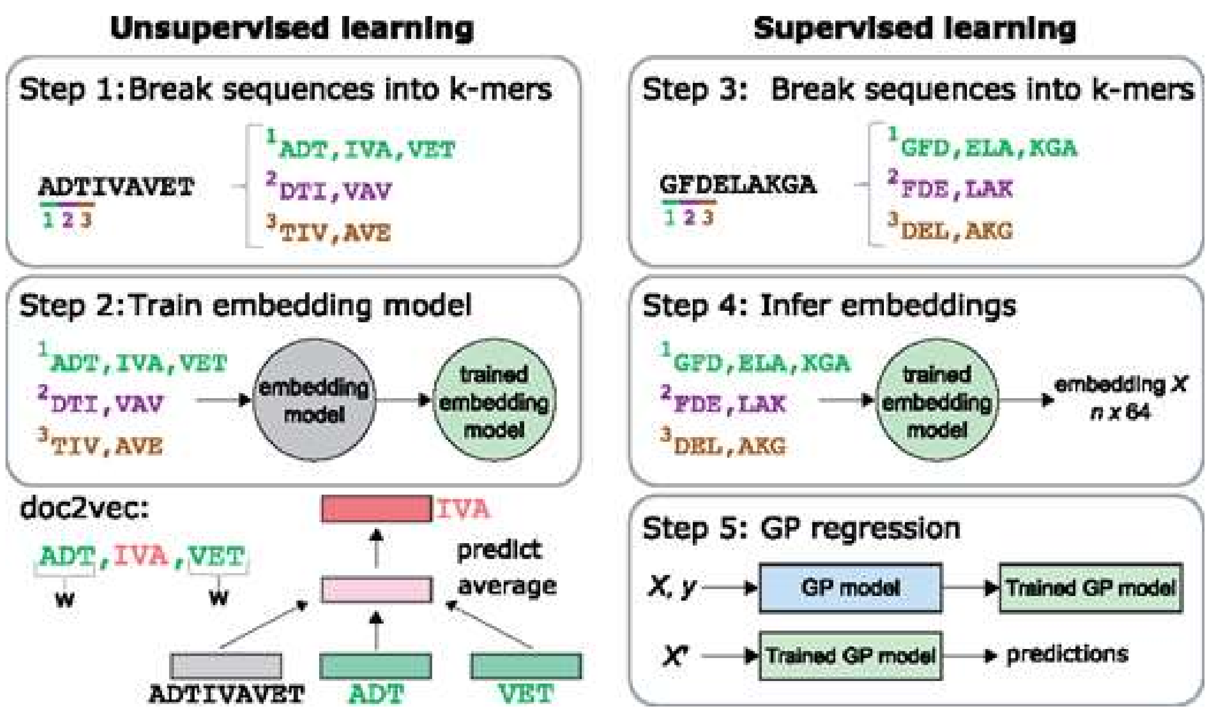

#REF: Yang, K. K., Wu, Z., Bedbrook, C. N., & Arnold, F. H. (2018). Learned protein embeddings for machine learning. *Bioinformatics*, *34*(15), 2642-2648.

Crystal structure representation

- Crystal structure input:
  - Atom type $\boldsymbol{A} = (a_0, \dots, a_N) \in \mathbb{A}^N$
  - Atom coordinates $\boldsymbol{X} = (\boldsymbol{x}_0, \dots, \boldsymbol{x}_N) \in \mathbb{R}^{N \times 3}$
  - Periodic lattice $\boldsymbol{L} = (\boldsymbol{l}_1, \boldsymbol{l}_2, \boldsymbol{l}_3) \in \mathbb{R}^{3 \times 3}$
- Multi-graph representation to model the periodic interaction

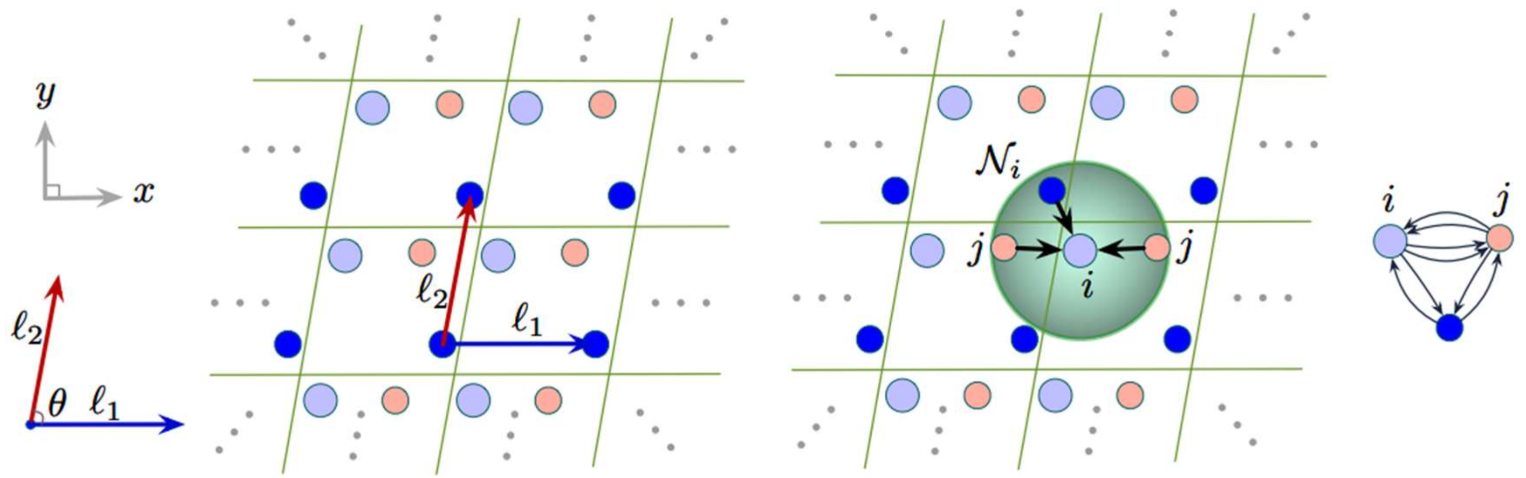

(a) (b) (c)

Representing microstructures of crystal mixture

Generate dual phase models

Volume domain: 10^6 voxels

Generate prior β grains

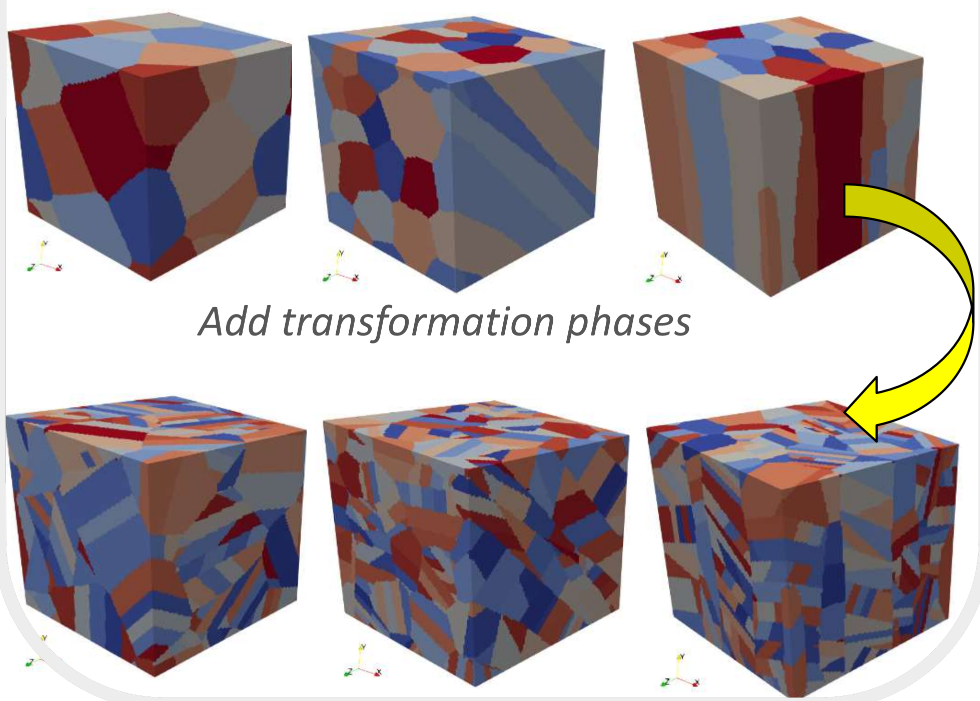

Feature information

- *Input information for each microstructure saved per voxel*
- *Saved data considers local environment*

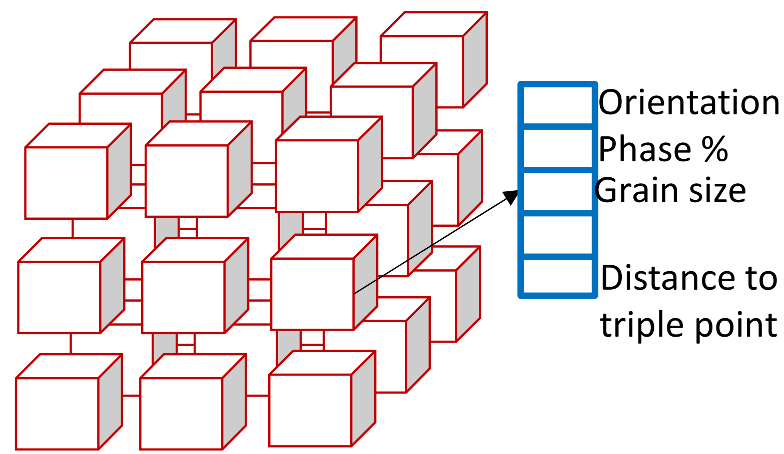

Slide credit: Sterjovski and Agius

Topics

- ~~Why AI for Science?~~
- ~~Representation~~
- Prediction
- Optimization & Generalization
- Explanation
- Physics-informed ML
- AI Co-scientist
- Future & Risks

A graph processing machine for molecular property prediction

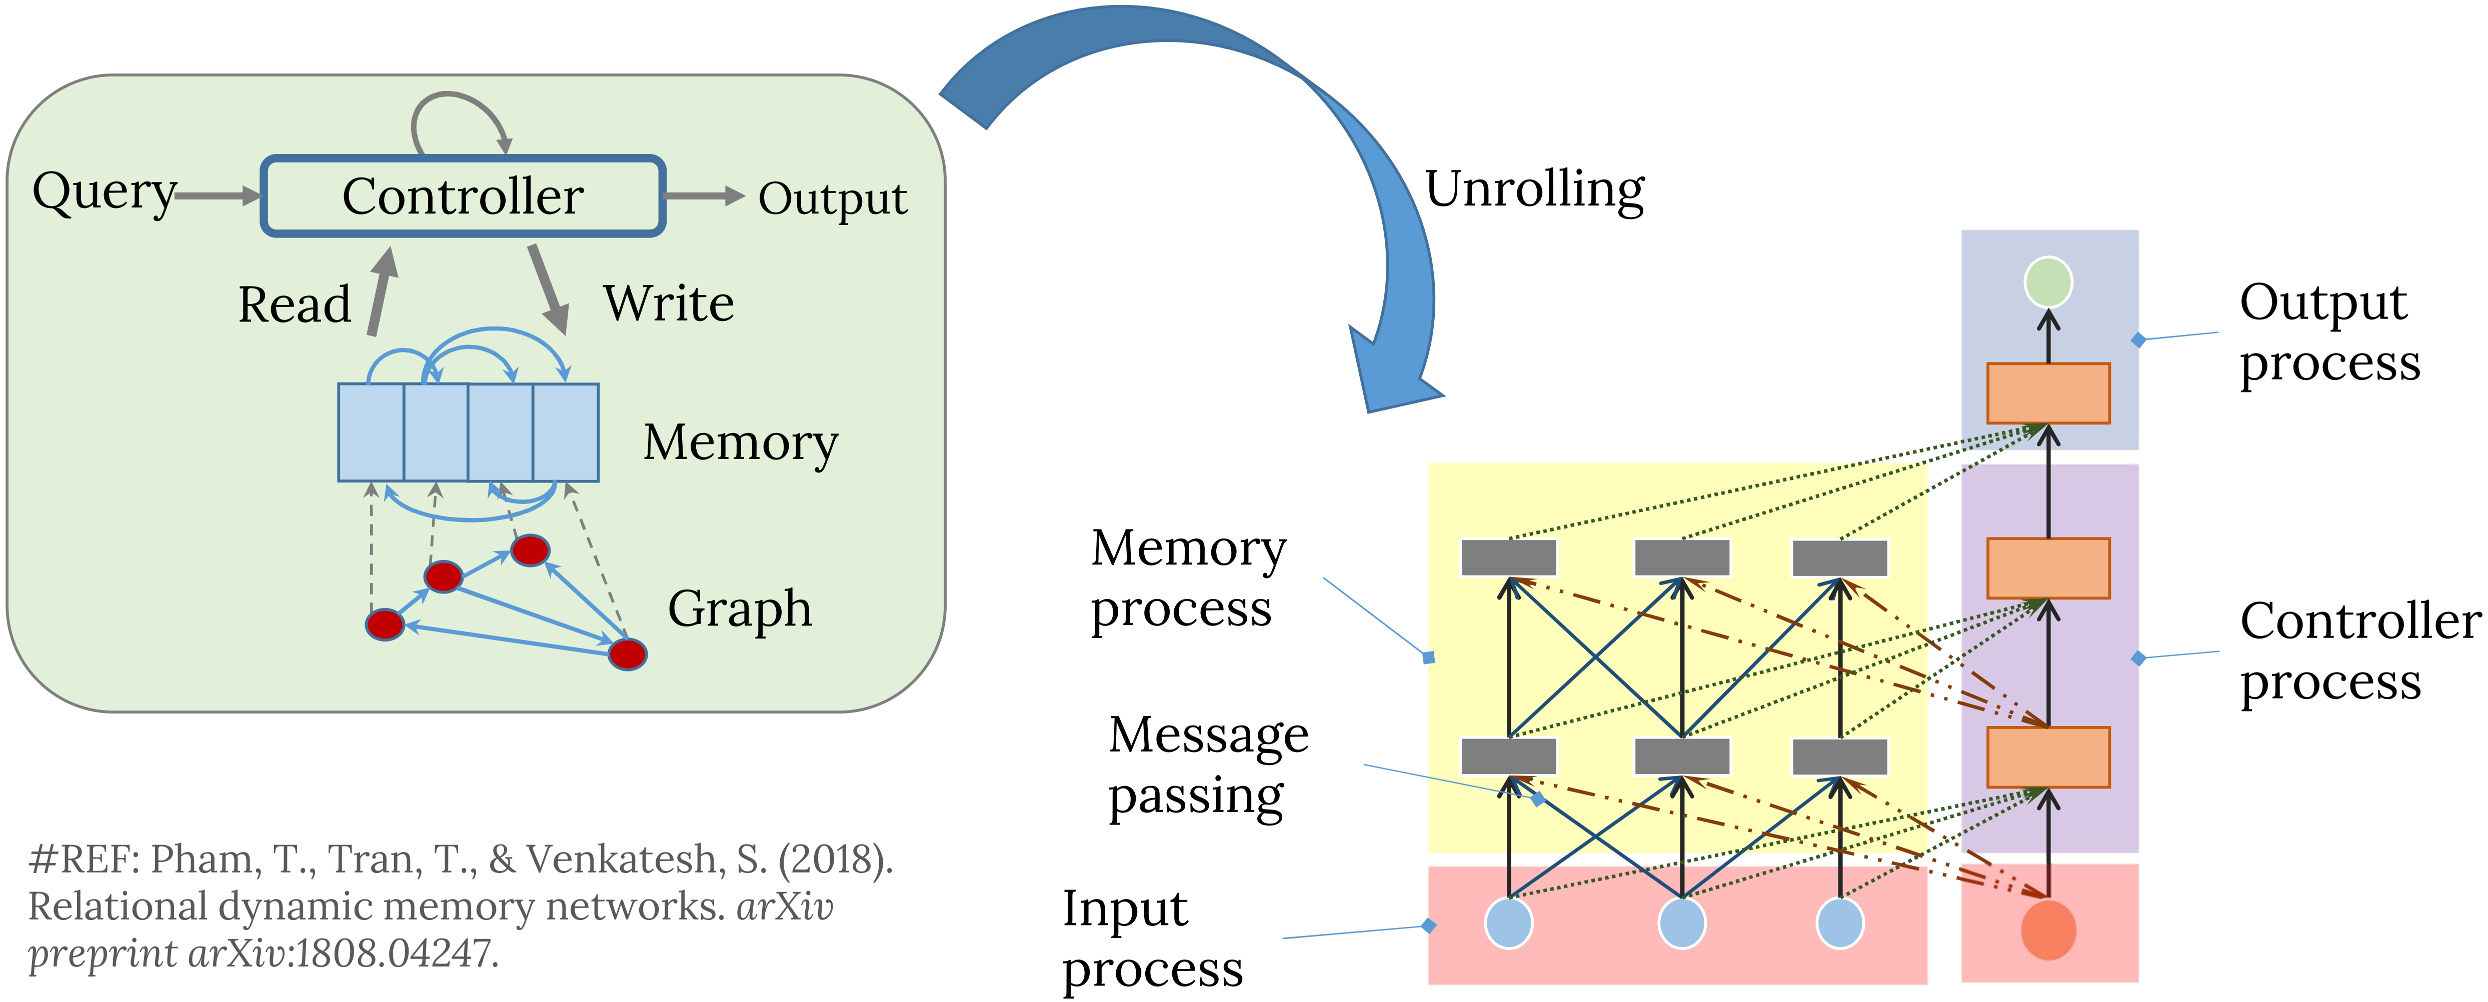

#REF: Pham, T., Tran, T., & Venkatesh, S. (2018). Relational dynamic memory networks. *arXiv preprint arXiv*:1808.04247.

Multi-target prediction

Possible targets

Molecular graph

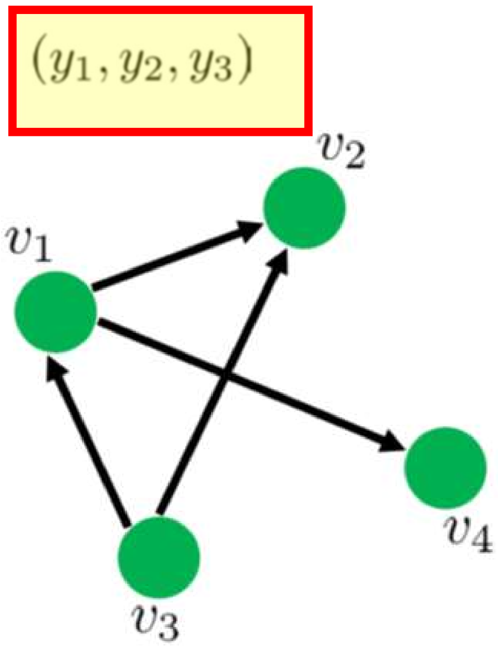

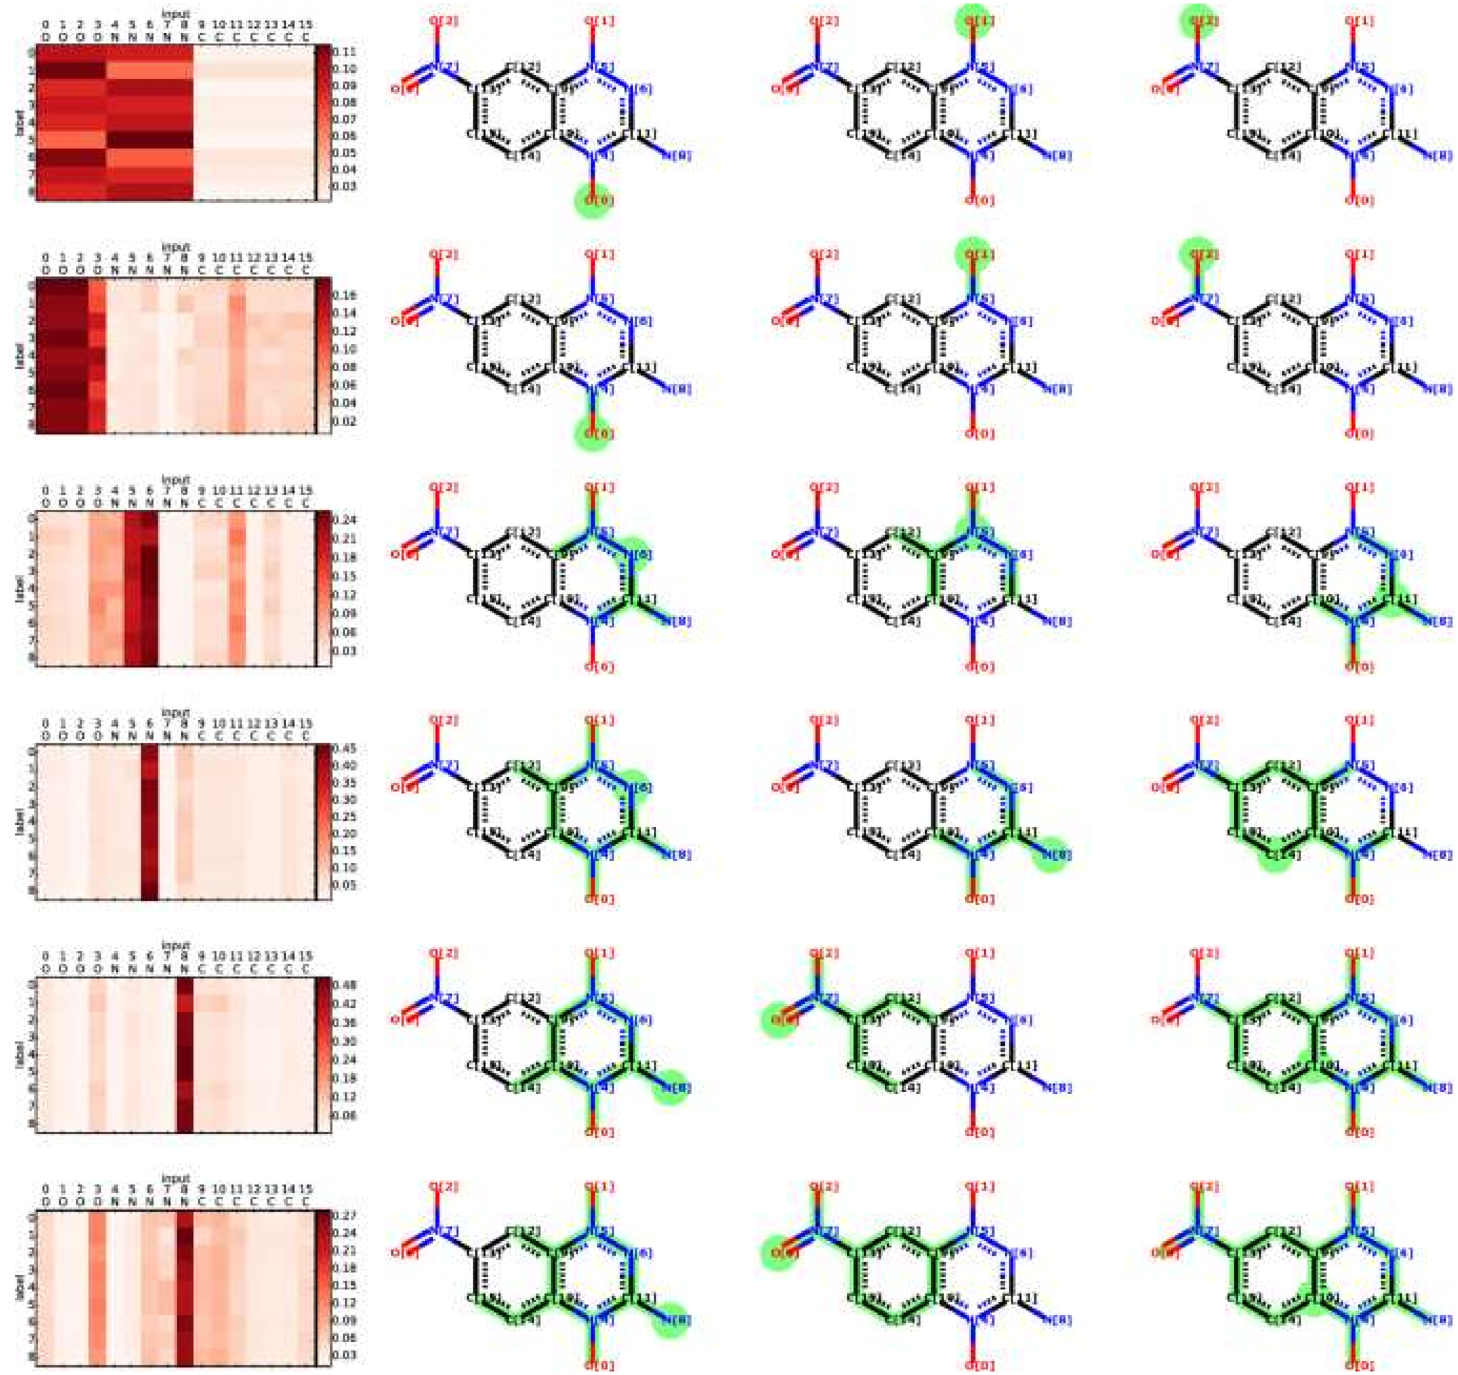

#REF: Do, Kien, et al. "Attentional Multilabel Learning over Graphs-A message passing approach." *Machine Learning, 2019*.

Predict multiple properties

- #REF: Do, Kien, et al. "Attentional Multilabel Learning over Graphs-A message passing approach." *Machine Learning, 2019*.

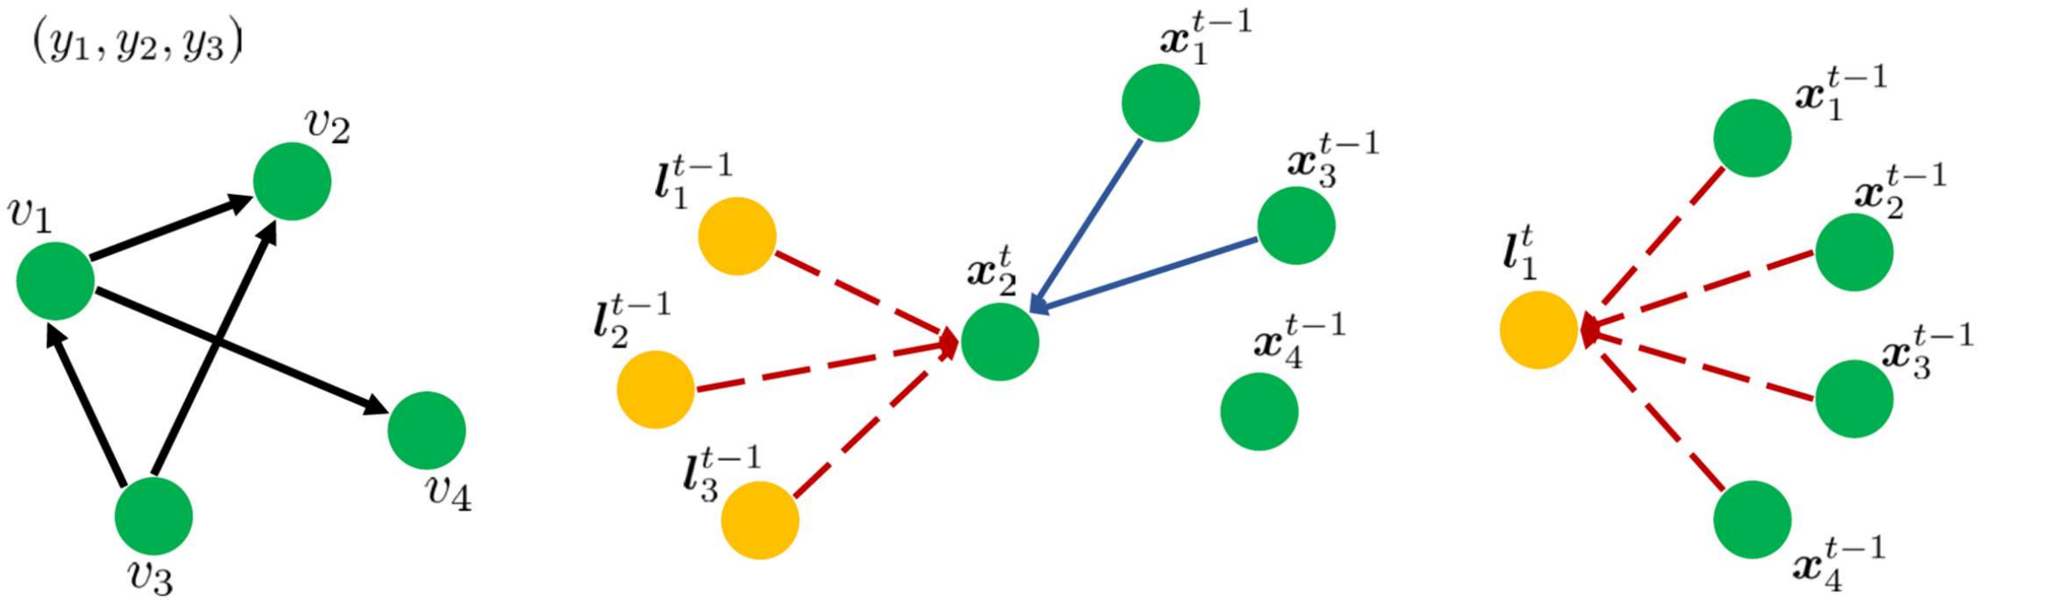

(a) A input graph with 4 nodes and 3 labels

(b) Input node update

(c) Label node update

Chemical-chemical interaction via *Relational Dynamic Memory Networks*

- #REF: Pham, Trang, Truyen Tran, and Svetha Venkatesh. "Relational dynamic memory networks." *arXiv preprint arXiv:1808.04247*(2018).


https://usaherald.com/6018-2/

USA Herald
The People's Voice

Home › Expert Advice

Expert Advice

2017 Drug Interactions Report: A Growing and Deadly Threat

By **Marina Noble MD, MPH** - November 7, 2017 1810 0 0 Comments

SHARE Share 3 Tweet Share 1

https://usaherald.com/wp-content/uploads/2017/11/Drug-Cocktail.png

Bupa Medical Certi....pdf Bupa claim form.pdf Show all

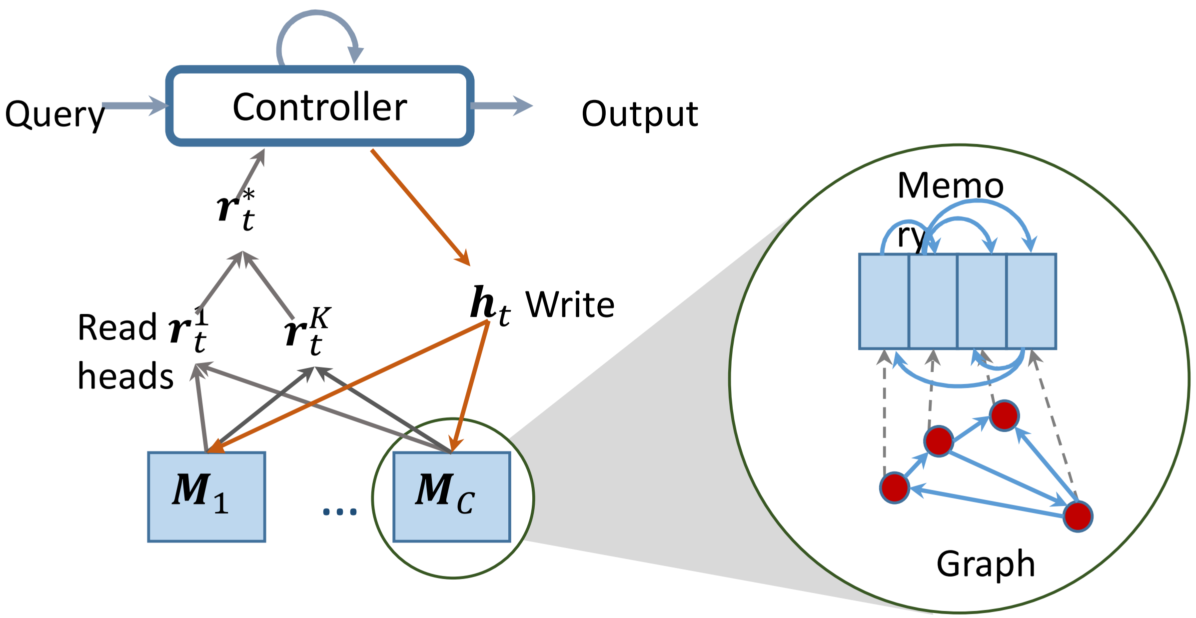

GEFA: Drug-protein binding as graph-in-graph interaction

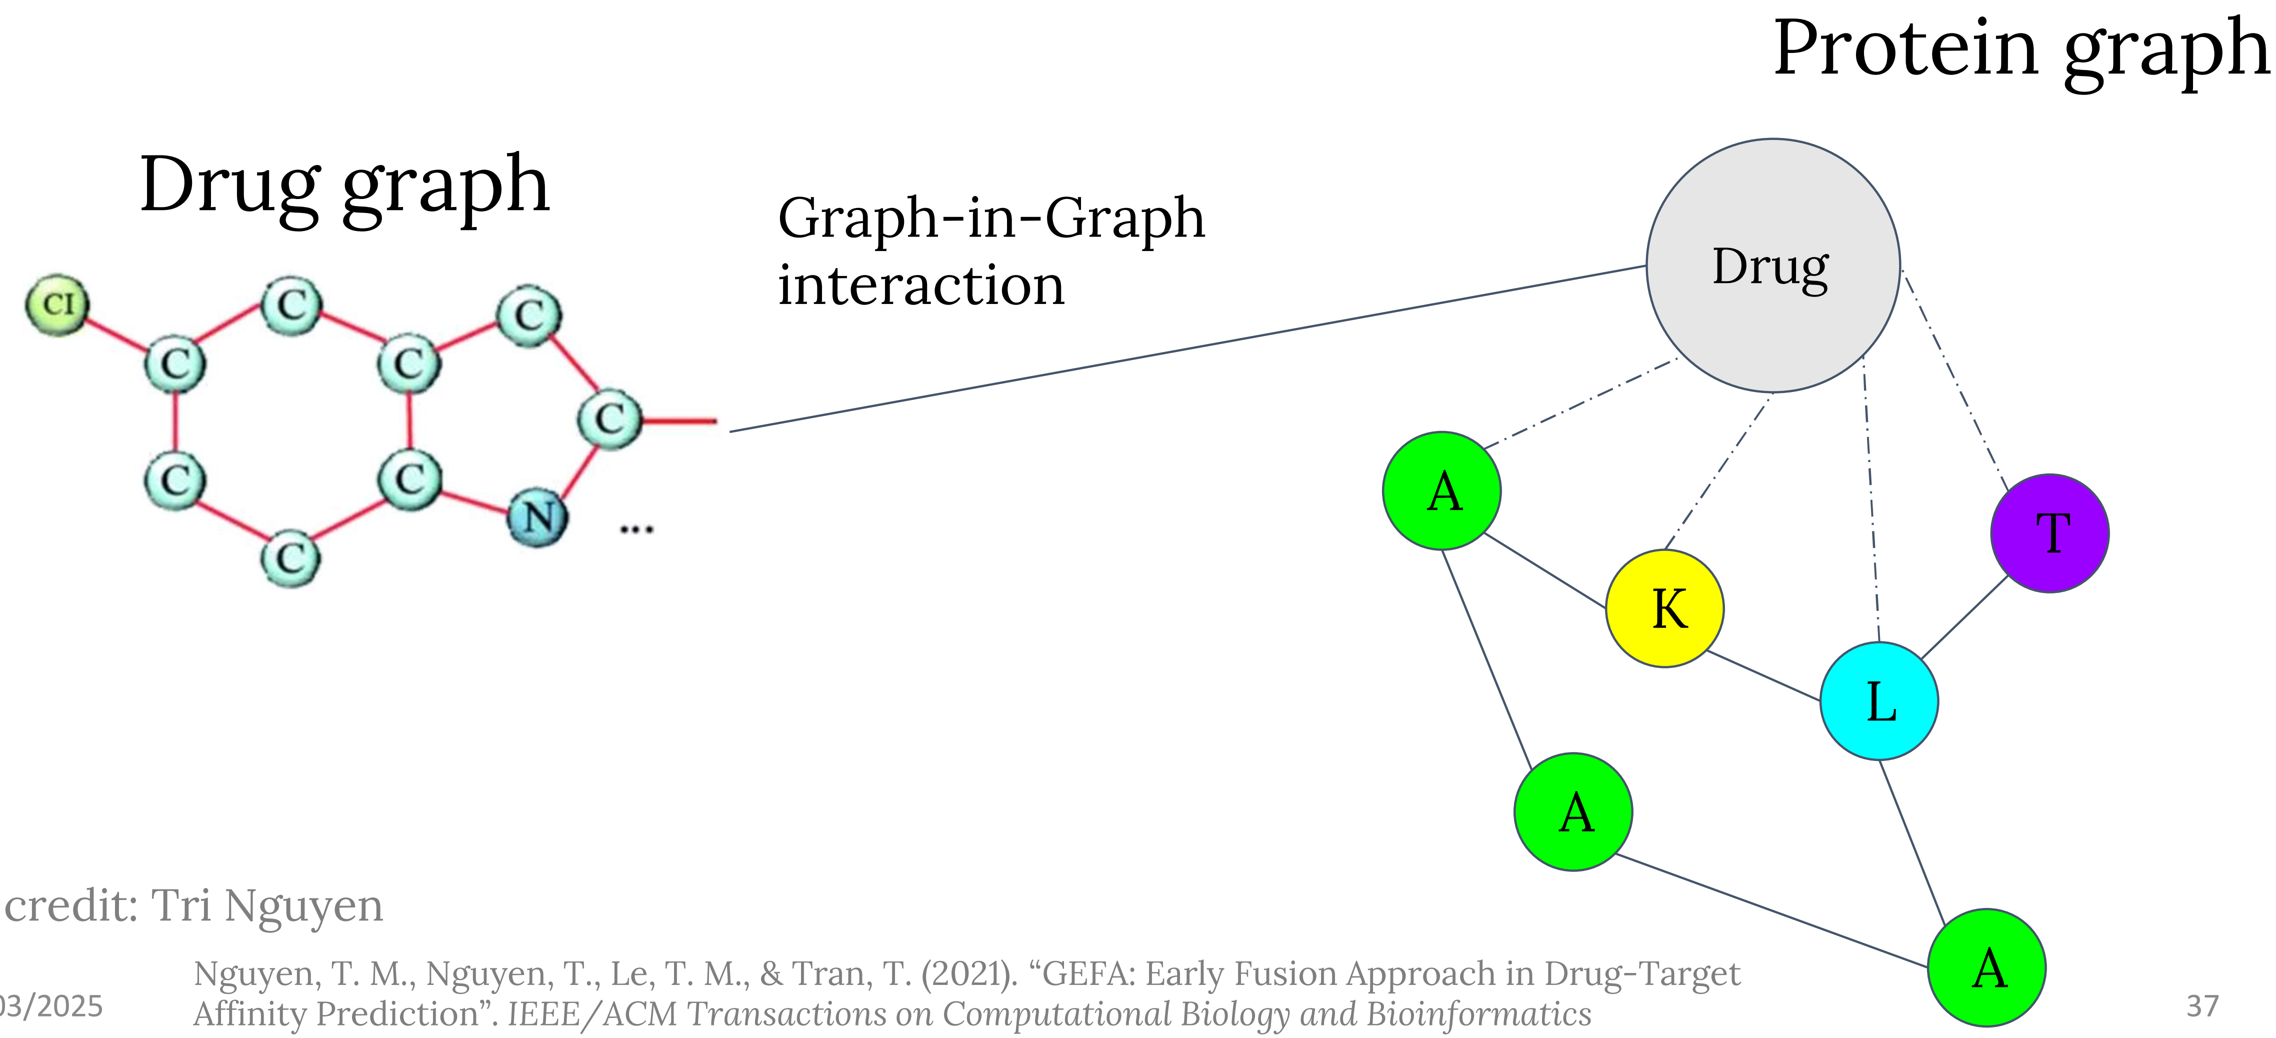

Slide credit: Tri Nguyen

Nguyen, T. M., Nguyen, T., Le, T. M., & Tran, T. (2021). “GEFA: Early Fusion Approach in Drug-Target Affinity Prediction”. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*

GEFA (cont.)

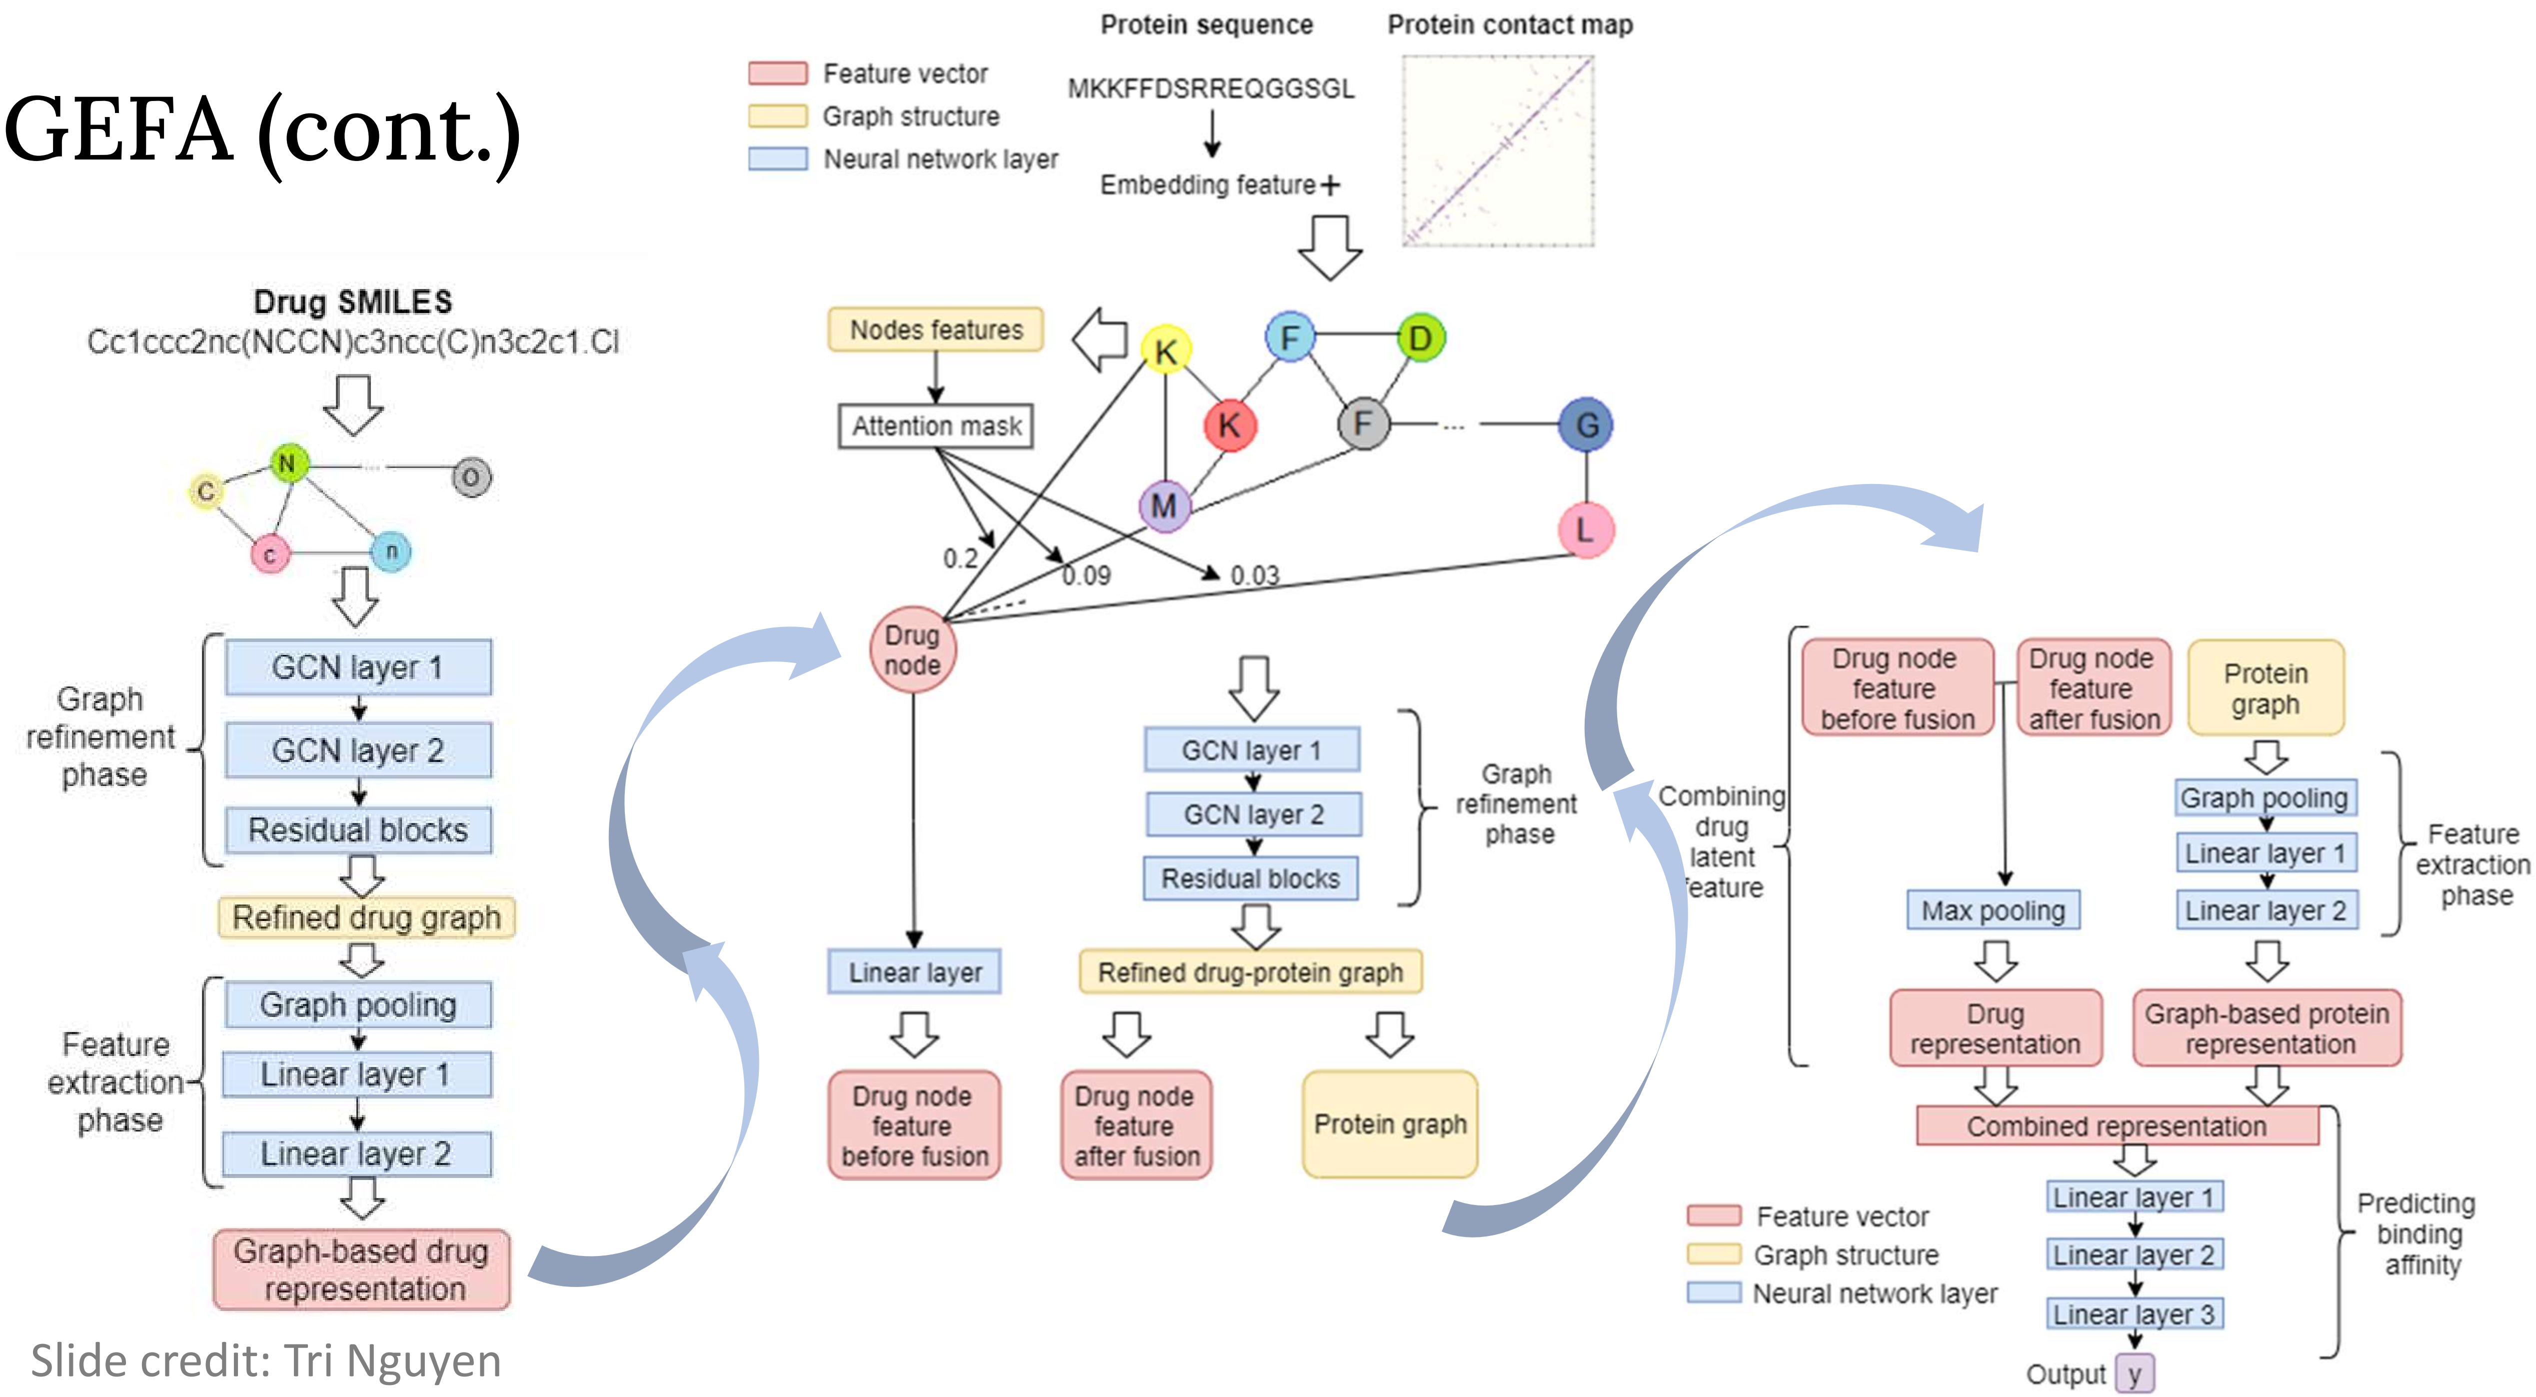

Slide credit: Tri Nguyen

Nguyen, T. M., Nguyen, T., Le, T. M., & Tran, T. (2021). "GEFA: Early Fusion Approach in Drug-Target Affinity Prediction". *IEEE/ACM Transactions on Computational Biology and Bioinformatics*

Predicting stress-strain curve from crystal mixture

- Transformer/3D-UNet to leverage long-range dependencies between voxels
- Input: Feature vectors per voxel.
- Output: Strain curve per voxel.

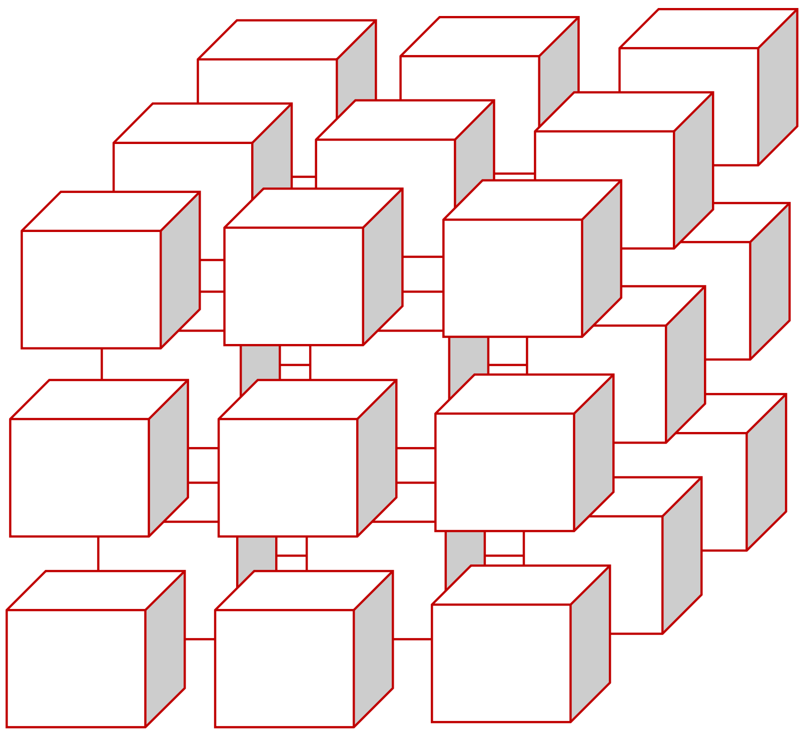

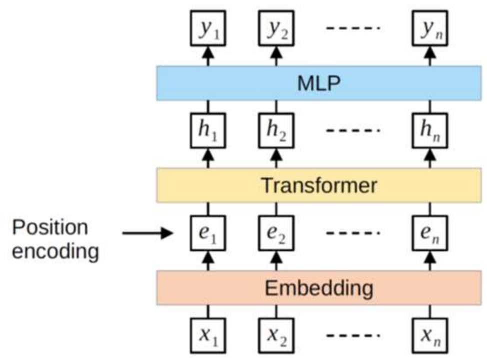

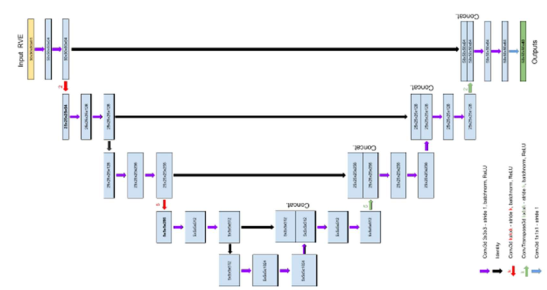

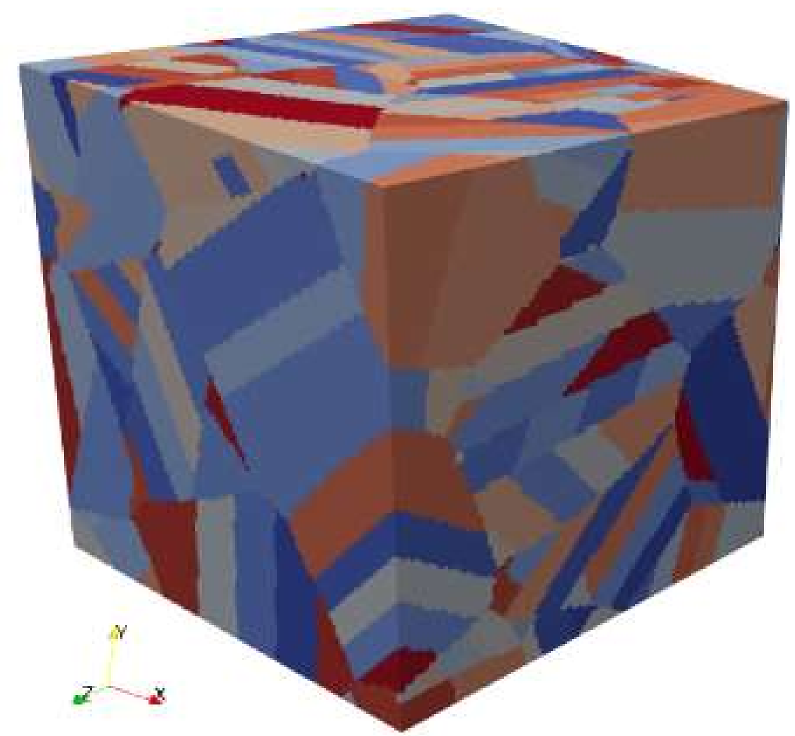

Topics

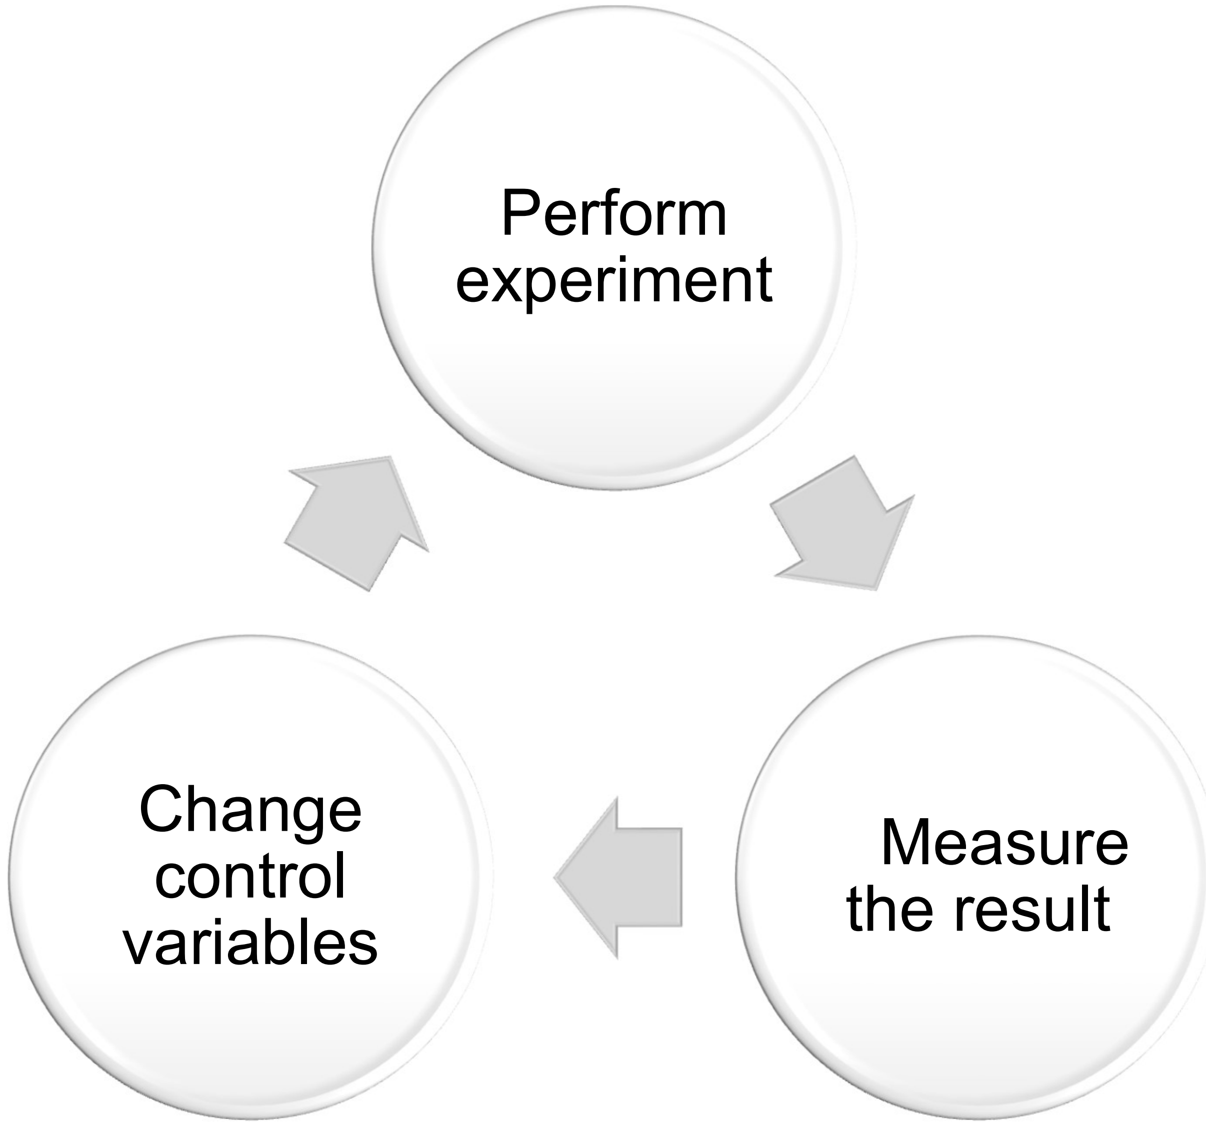

experimentation drives scientific enquiry

Crashworthiness test

10 variables

→

1024 experiments

(each takes 98 hours)

→

Requires 12 years to finish

the curse of complexity

That is why the space is open

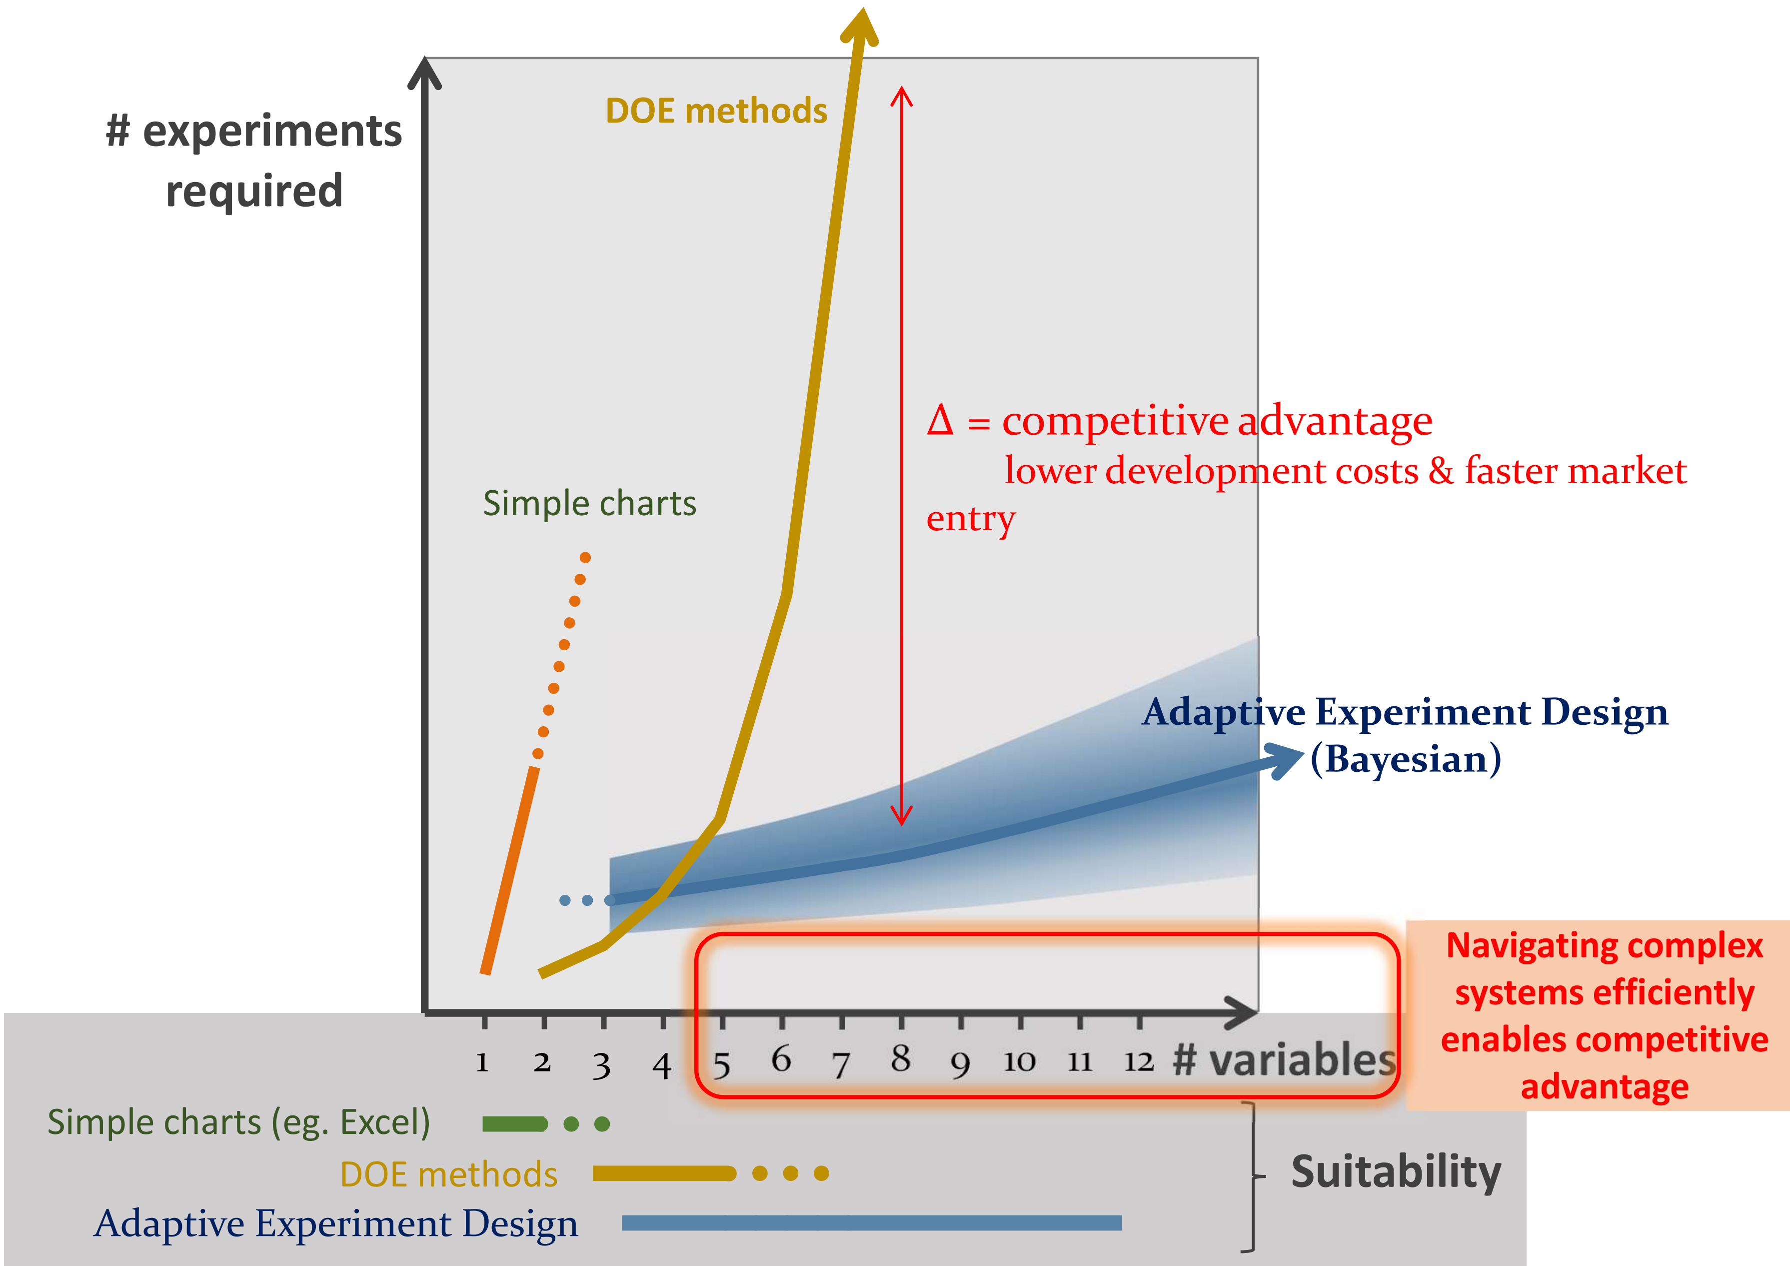

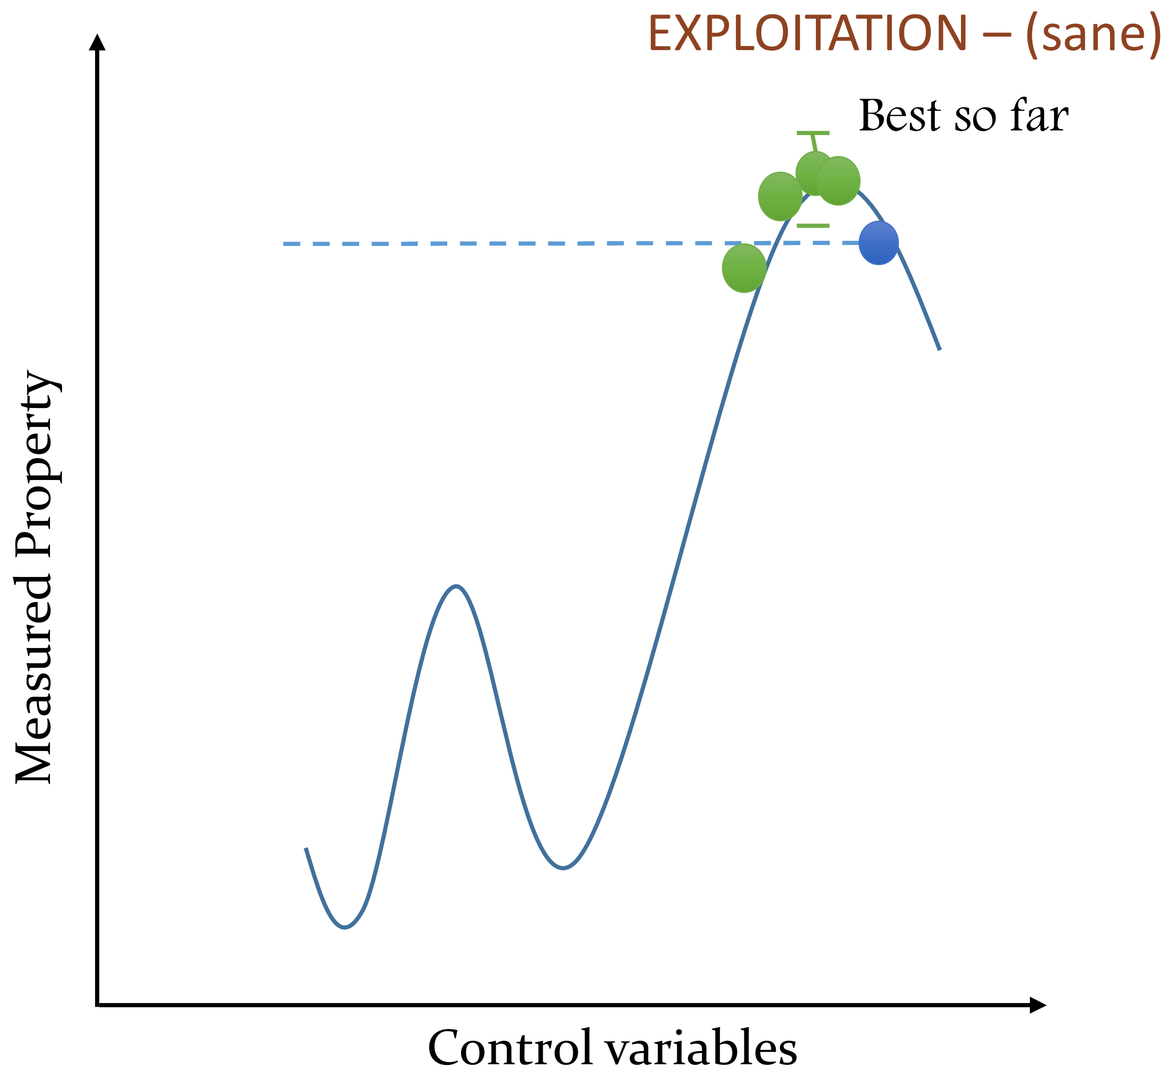

Underpinning theory – Bayesian optimisation

Target : Length (Quantitative)
Diameter (Quantitaive)
Quality (Qualitative)

- *Rapid Bayesian optimisation for synthesis of short polymer fiber materials". In: Scientific reports 7.1 (2017)*

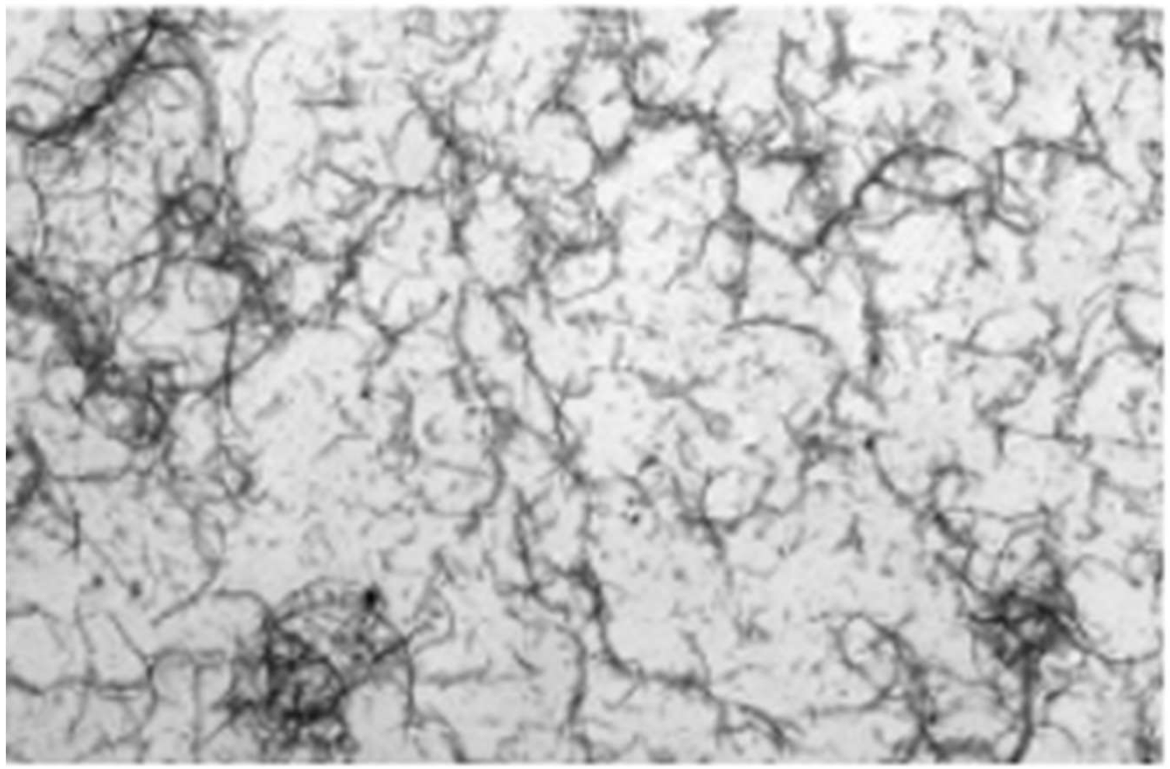

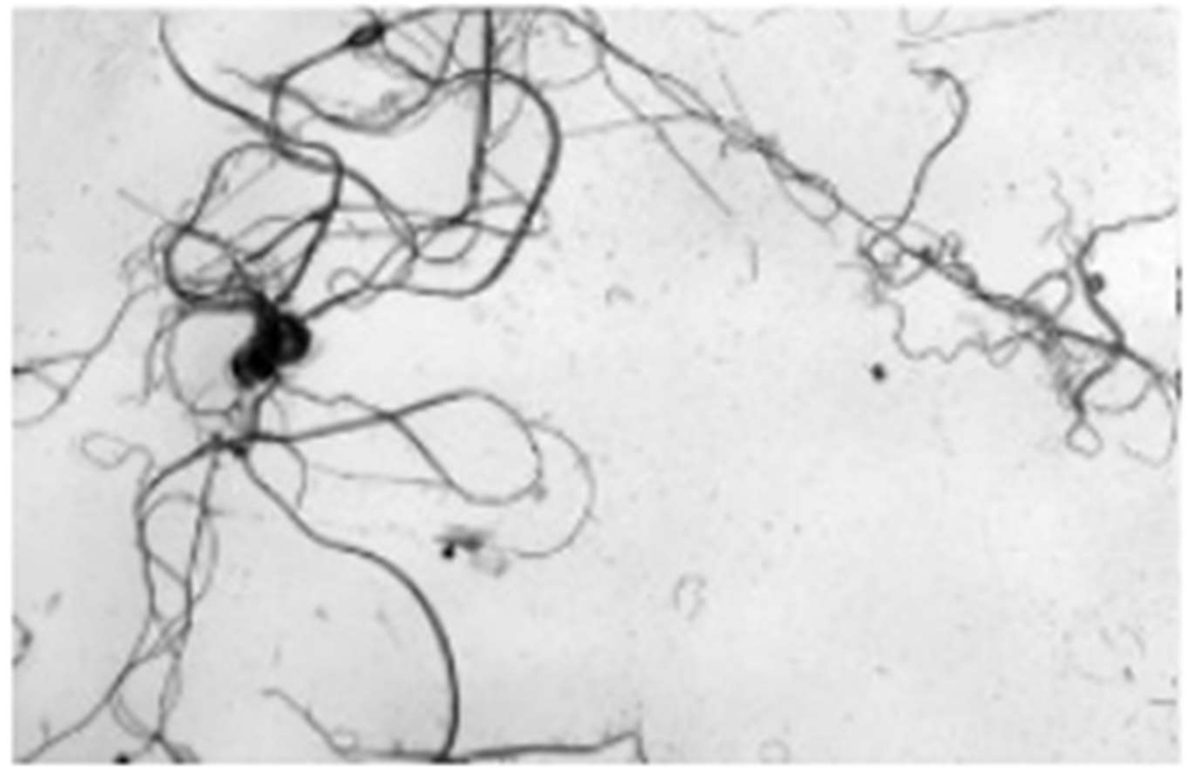

GTPN: Synthesis via reaction prediction as neural graph morphism

- **Input**: A set of graphs = a single big graph with disconnected components
- **Output**: A new set of graphs. Same nodes, different edges.
- **Model**: Graph morphism
- **Method**: Graph transformation policy network (GTPN)

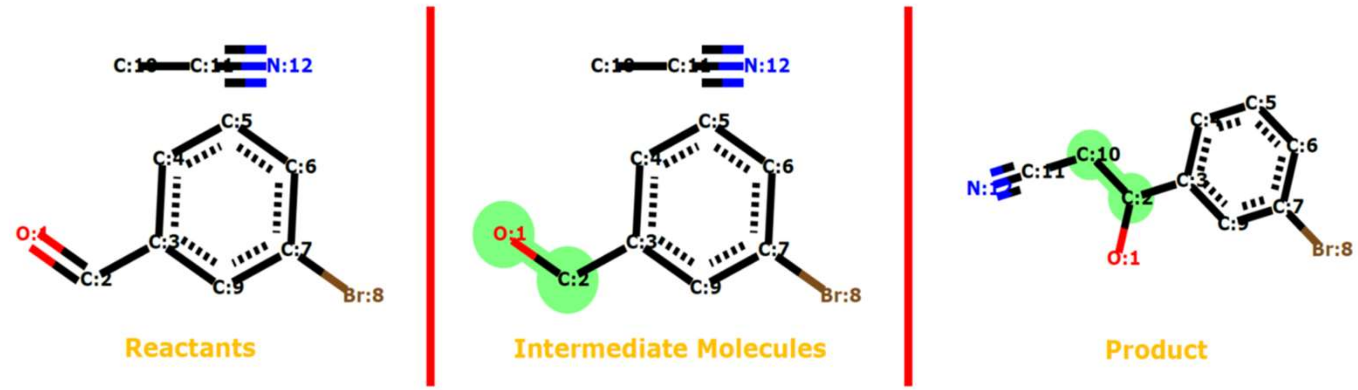

Figure 1: A sample reaction represented as a set of graph transformations from reactants (leftmost) to products (rightmost). Atoms are labeled with their type (Carbon, Oxygen,...) and their index (1, 2,...) in the molecular graph. The atom pairs that change connectivity and their new bonds (if existed) are highlighted in green. There are two bond changes in this case: 1) The double bond between O:1 and C:2 becomes single. 2) A new single bond between C:2 and C:10 is added.

Kien Do, Truyen Tran, and Svetha Venkatesh. "Graph Transformation Policy Network for Chemical Reaction Prediction." *KDD'19*.

Nguyen, P., Tran, T., Gupta, S., Rana, S., Barnett, M. and Venkatesh, S., 2019, May. Incomplete conditional density estimation for fast materials discovery. In *Proceedings of the 2019 SIAM International Conference on Data Mining* (pp. 549-557). Society for Industrial and Applied Mathematics.

Alloy design generation

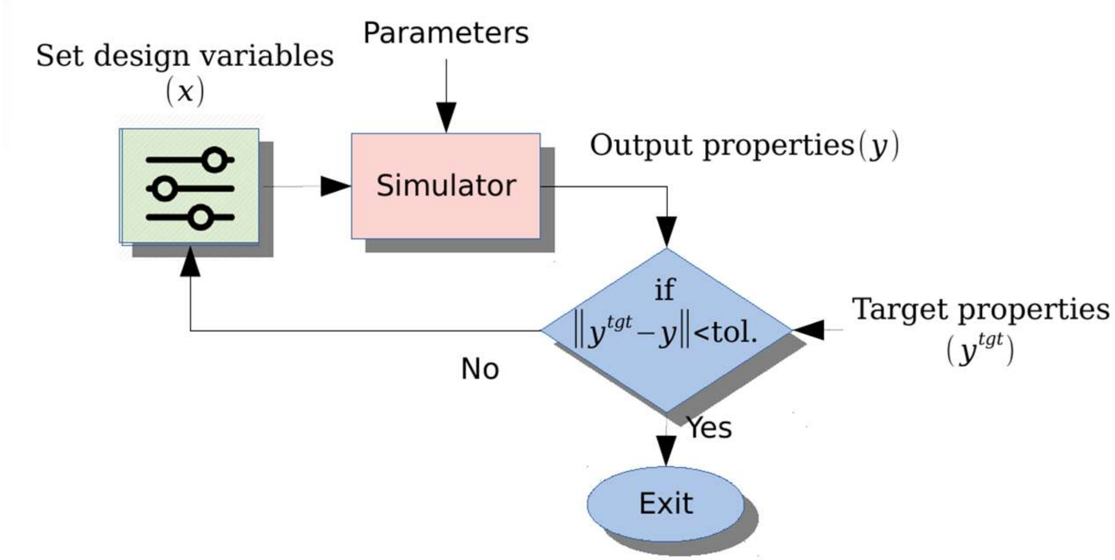

- Scientific innovations are expensive
- One search per specific target
- Availability of growing data

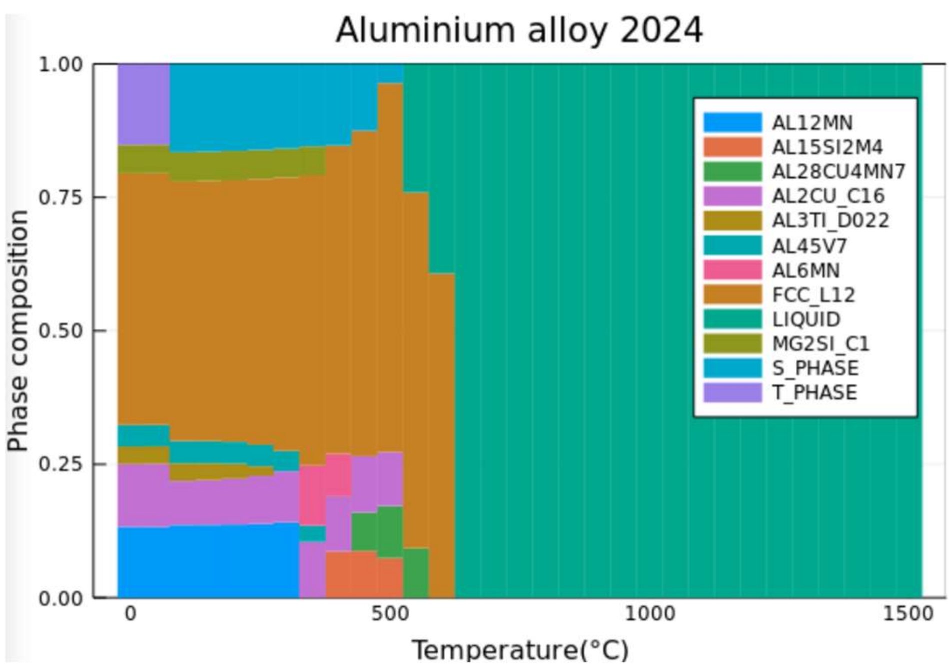

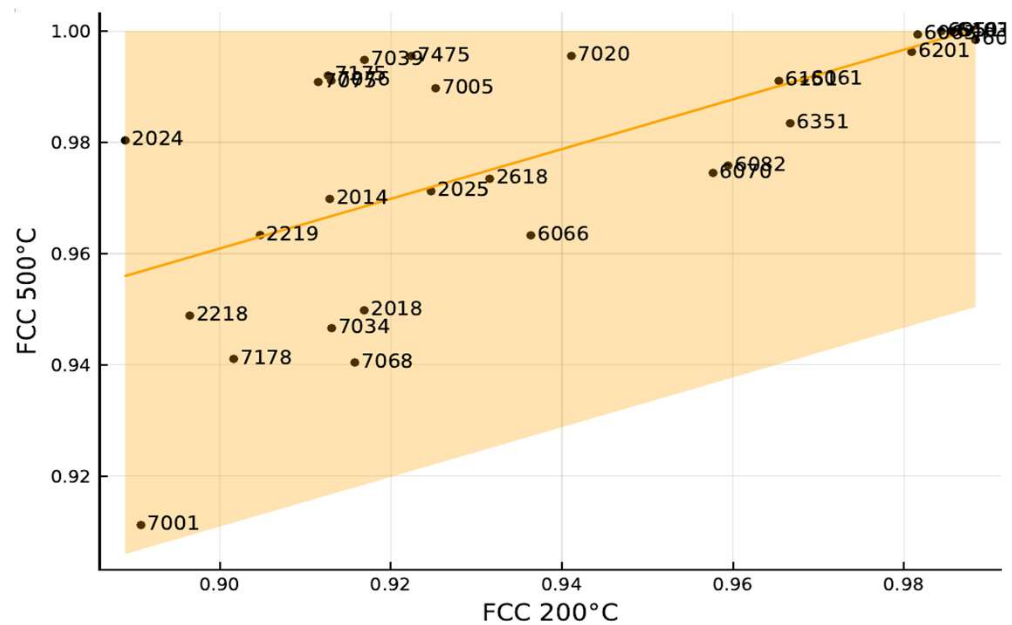

Inverse design

Leverage	Leverage the existing data and query the simulators in an offline mode
Avoid	Avoid the global optimization by learning the inverse design function f -1(y)
Predict	Predict design variables in a single step

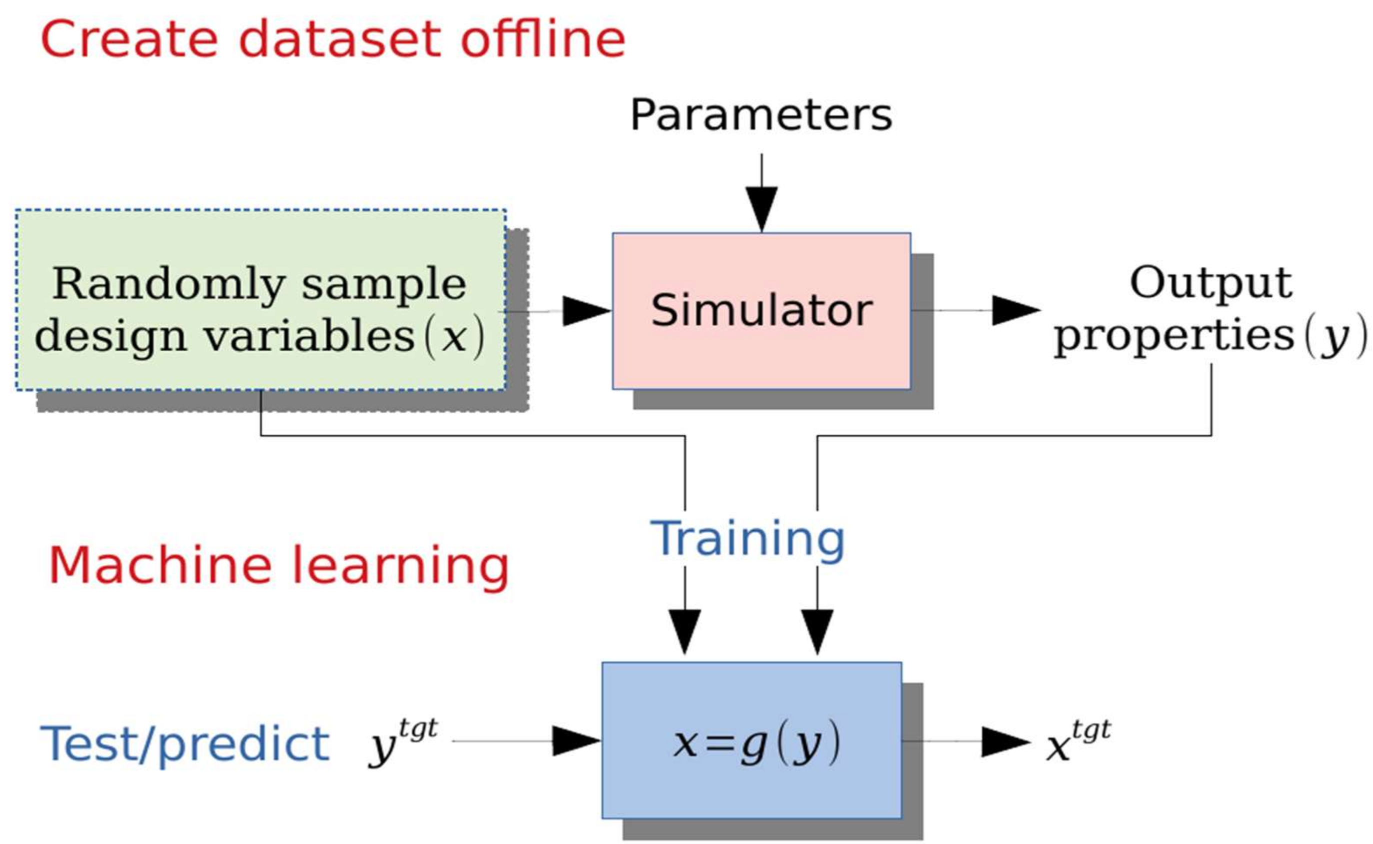

$$x^{\text{target}} = g(y^{\text{target}})$$

Incomplete conditional density estimation

- Multimodal density estimation given incomplete conditions

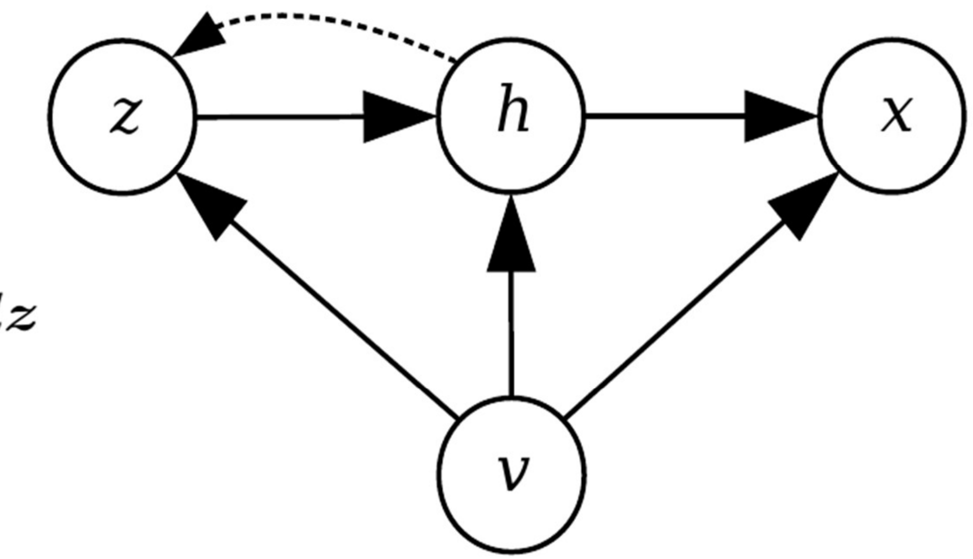

$$P(\boldsymbol{x} \mid \boldsymbol{v}) = \int_{\boldsymbol{h}} \int_{\boldsymbol{z}} P(\boldsymbol{x} \mid \boldsymbol{v}, \boldsymbol{h}) \, P(\boldsymbol{h} \mid \boldsymbol{v}, \boldsymbol{z}) \, P(\boldsymbol{z} \mid \boldsymbol{v}) d\boldsymbol{h} d\boldsymbol{z}$$
$$\approx \frac{1}{N} \sum_{i=1}^{N} \mathbb{E}_{P(\boldsymbol{h} \mid \boldsymbol{v}, \boldsymbol{z}^{(i)})} \left[P(\boldsymbol{x} \mid \boldsymbol{v}, \boldsymbol{h}) \right]$$

- However, integrating over h is still intractable, we approximate the expectation by a function evaluation at the mode

$$P(\boldsymbol{x} \mid \boldsymbol{v}) \approx \frac{1}{N} \sum_{i=1}^{N} P\left(\boldsymbol{x} \mid \boldsymbol{v}, \boldsymbol{\mu}\left(\boldsymbol{v}, \boldsymbol{z}^{(i)}\right)\right)$$

Generated alloys example

- Known-alloy dataset: 15,000 variations from 30 known series of Aluminum alloys
- BO-search dataset: 15,000 variations from 1000 found alloys by Bayesian optimization
- **Input**: phase diagram | **Output**: element composition

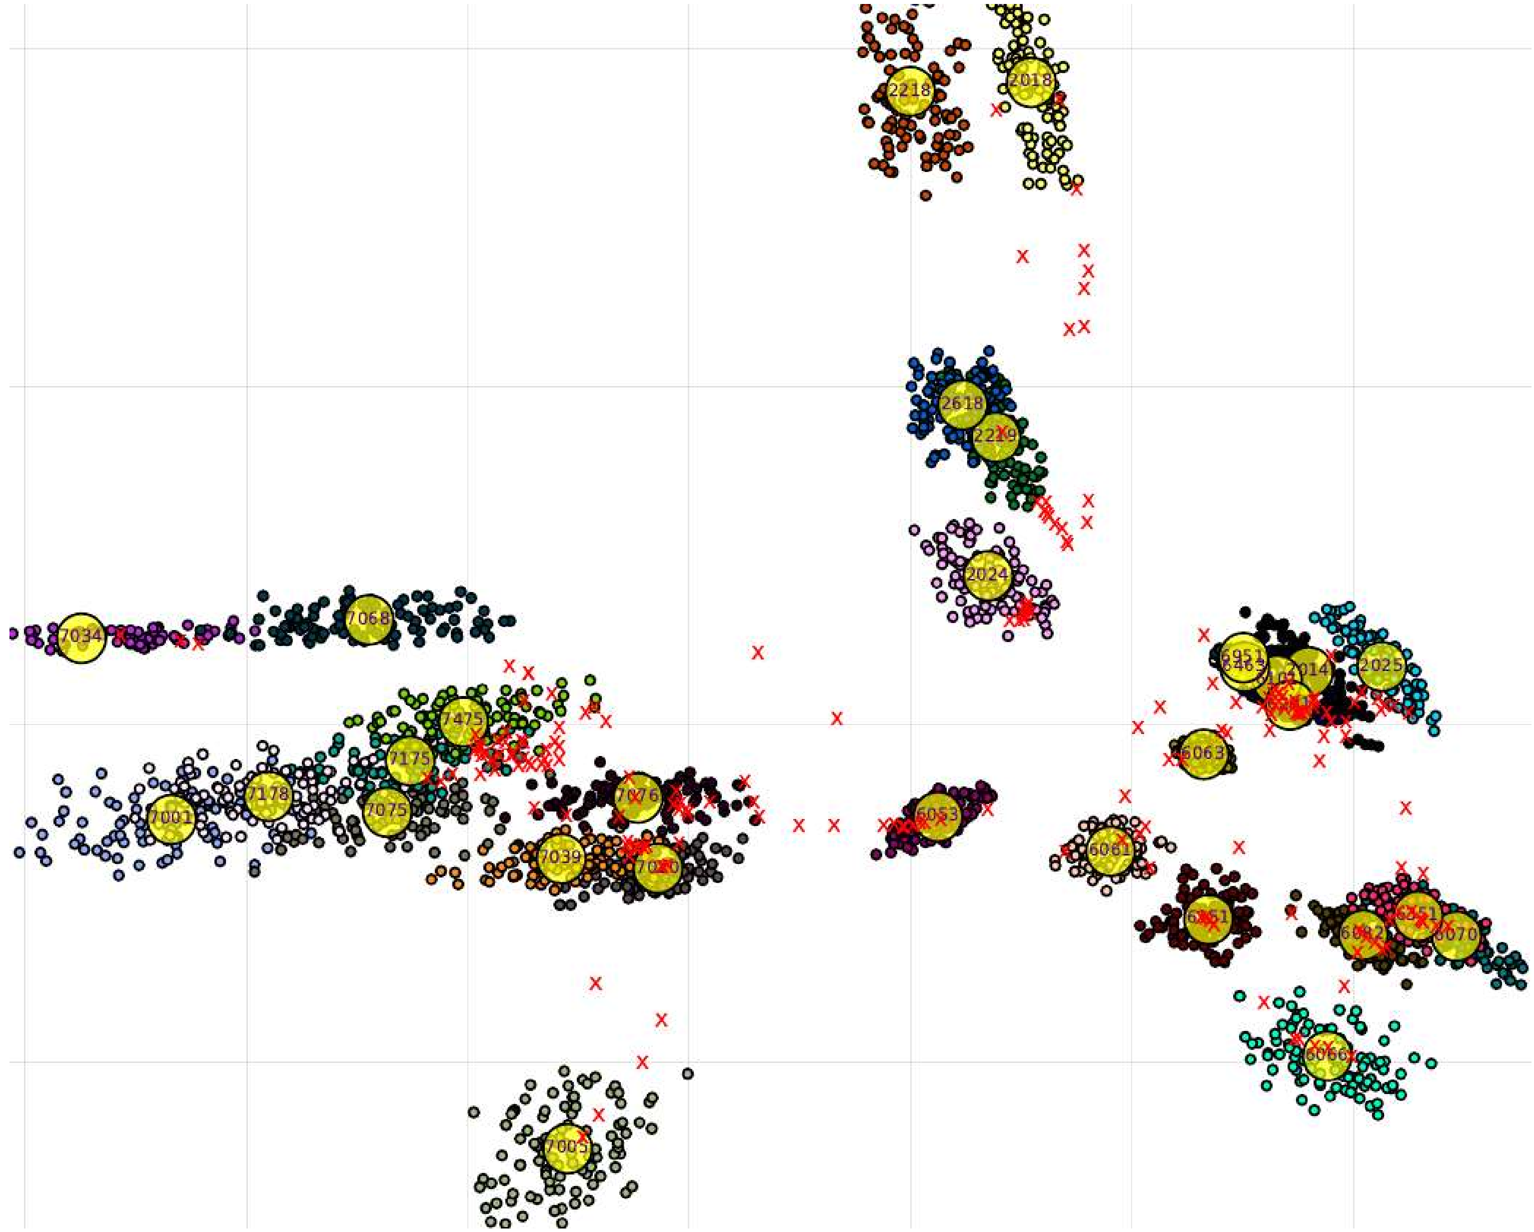

Crystal structure generation

Application in structure discovery: battery, aerospace materials, etc.

The stability of a solid-state crystal structure is connected to its formation energy

Target:

- Generate crystal-like stable structure
- Diversity set of crystal structure candidates for active learning

Slide credit: Tri Nguyen

GFlownet

- GFlownet learns to generate the composition object:
  - From the starting state, policy network output the probability distribution over building blocks
  - Select building blocks randomly based on the output probability distribution and create a new state → calculate the new probability distribution
  - Repeat until reaching the terminal state
  - Getting the reward from the environment (sparse reward)
- The complete set of actions from starting state to terminal state is a trajectory
- Flow is a non-negative function defined on the set of complete trajectory
- GFlownet is trained by matching the flow going through state: in-flow = out-flow

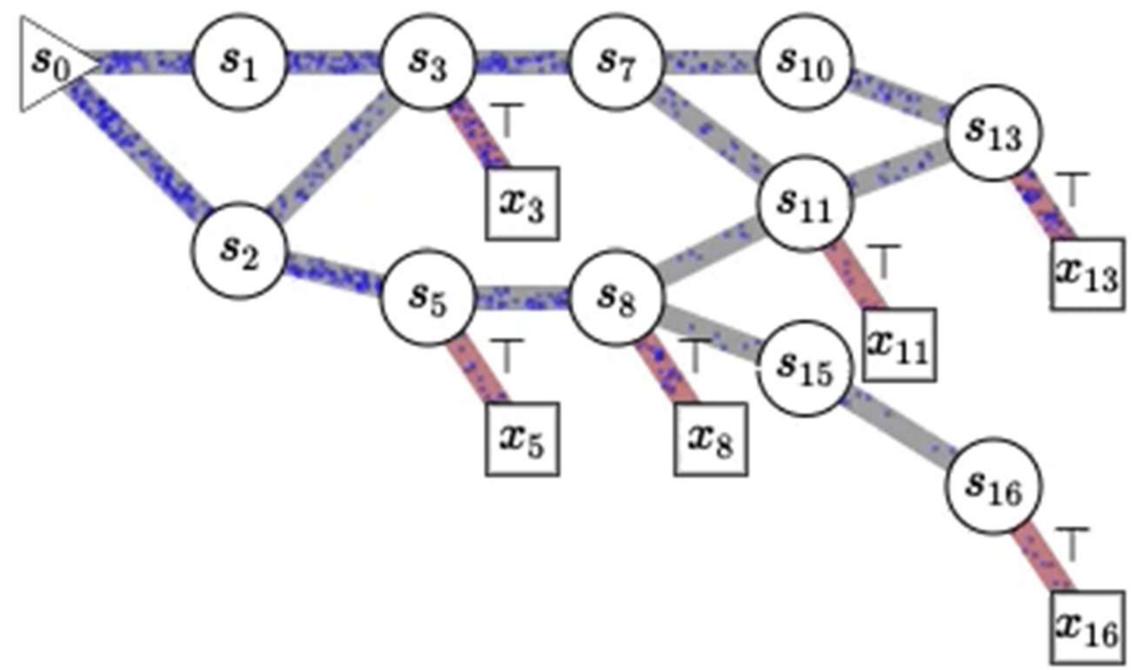

Slide credit: Tri Nguyen

Gflownet (cont.)

- Advantage of GFlownet
  - Diverse set of candidates → avoid getting stuck in multi-modal distribution (e.g stability/energy landscape of crystal structure)
  - Can sample in proportion to a given reward function (crystal structure generation: formation energy)

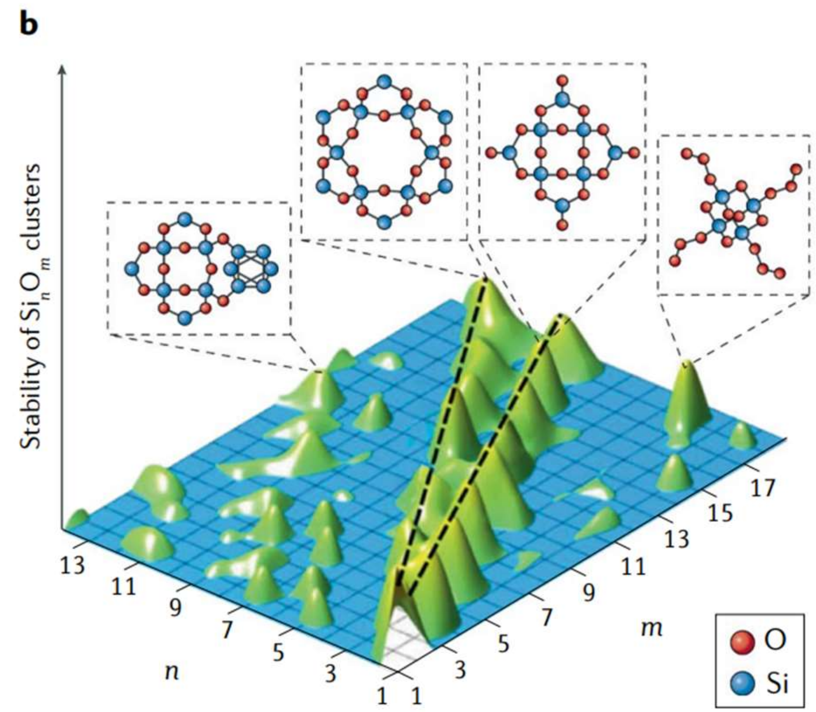

Slide credit: Tri Nguyen

Crystal structure generation with GFlownet

- State:
  - Multi-graph representation for structure:
    - Node: atoms
    - Node feature: element type, fraction coordinate
    - Edges: built using near-neighbor-based method CrystalNN with search cut-off starting from 13 and increasing to 20
    - Edge feature: cell-direction vector 'to_image', bond distance
  - 3D grid space: currently occupied and available position to insert new atom
- Action:
  - Available fraction coordinate on a 3D grid.
  - The chosen element

Slide credit: Tri Nguyen

GFlownet - Forward policy

- Policy network:
  calculate the probability distribution over actions

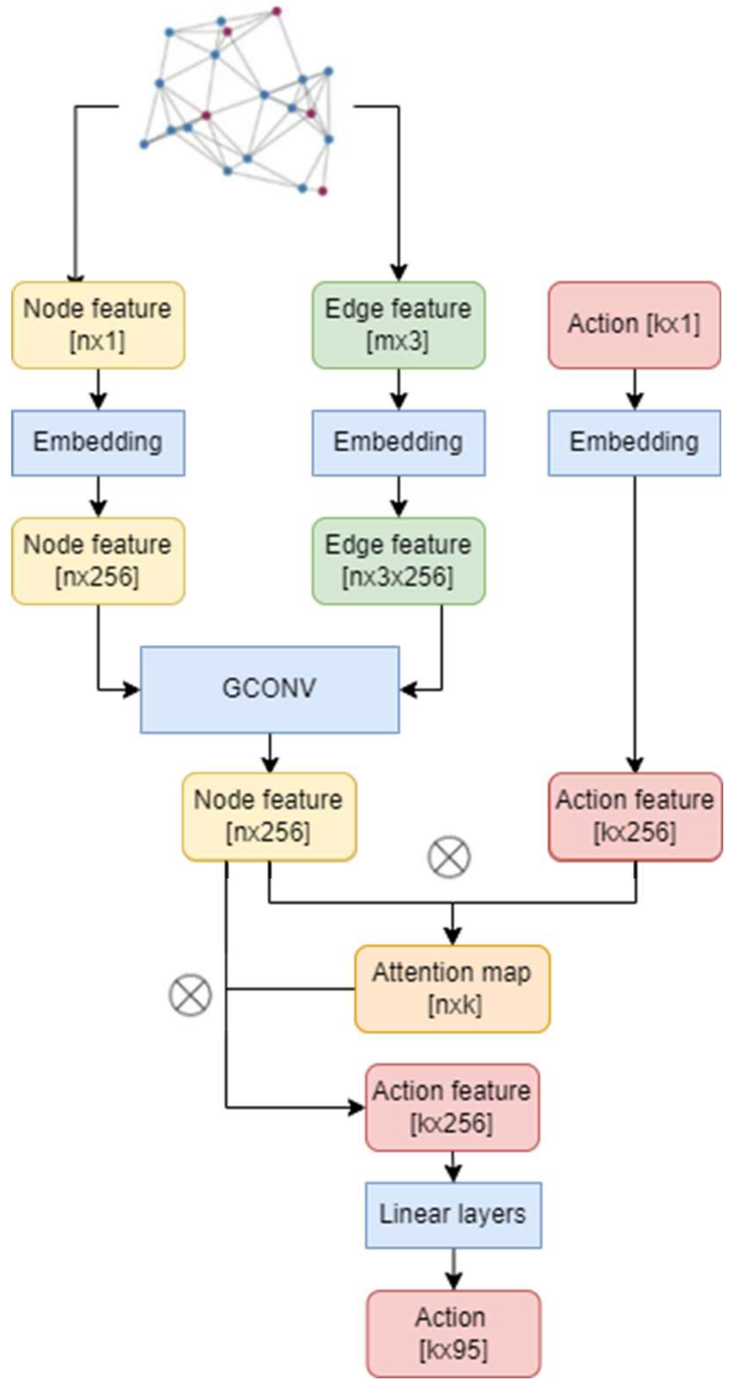

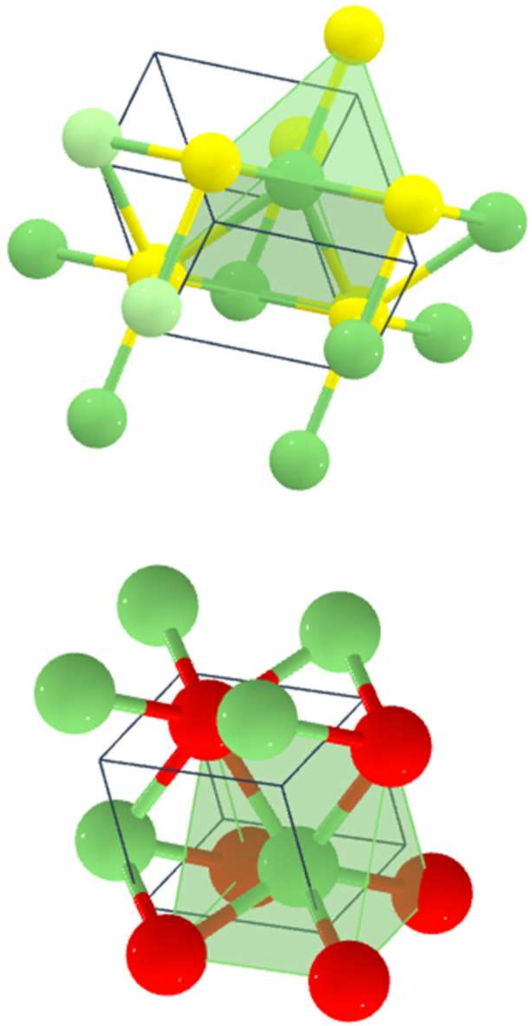

Examples of generated crystal structure

Slide credit: Tri Nguyen

SHAFT: Symmetry!

- Generalise GFlowNet
- Decompose materials space
- Exploit the symmetry of the materials
- -> Fast exploration!

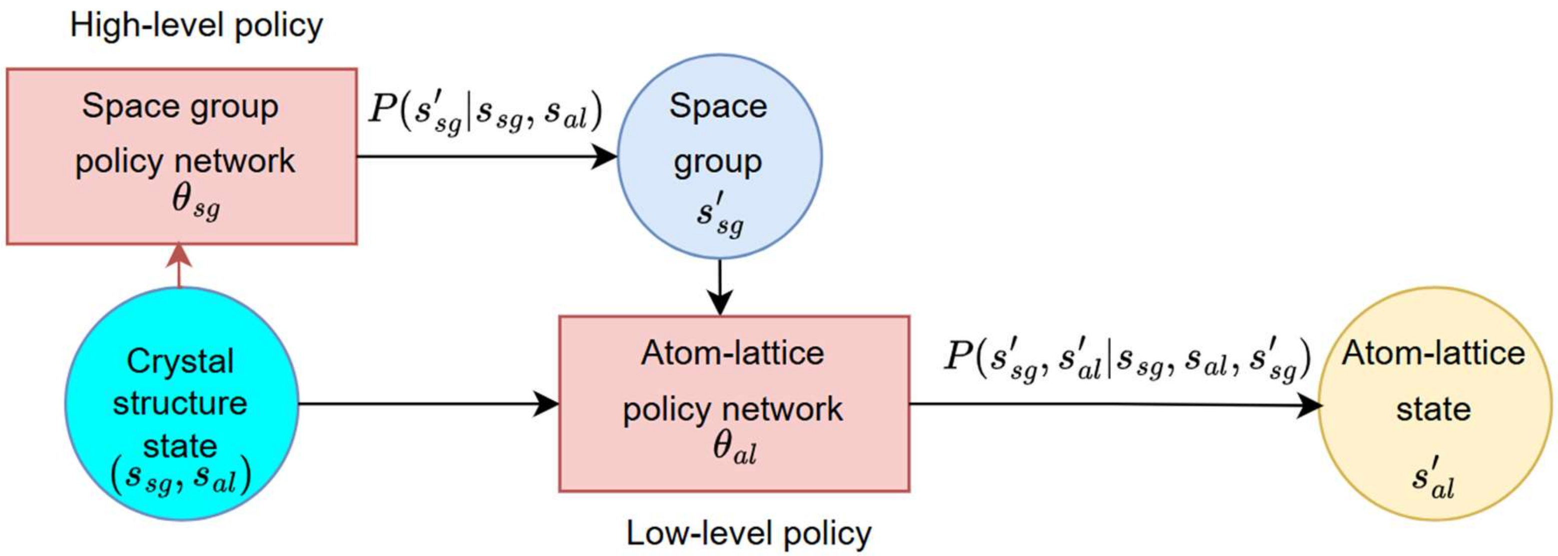

Nguyen, Tri Minh, Sherif Abdulkader Tawfik, Truyen Tran, Sunil Gupta, Santu Rana, and Svetha Venkatesh. "Efficient Symmetry-Aware Materials Generation via Hierarchical Generative Flow Networks." *arXiv preprint arXiv:2411.04323* (2024).

Example: Metal–organic framework(MOF) generated by SHAFT

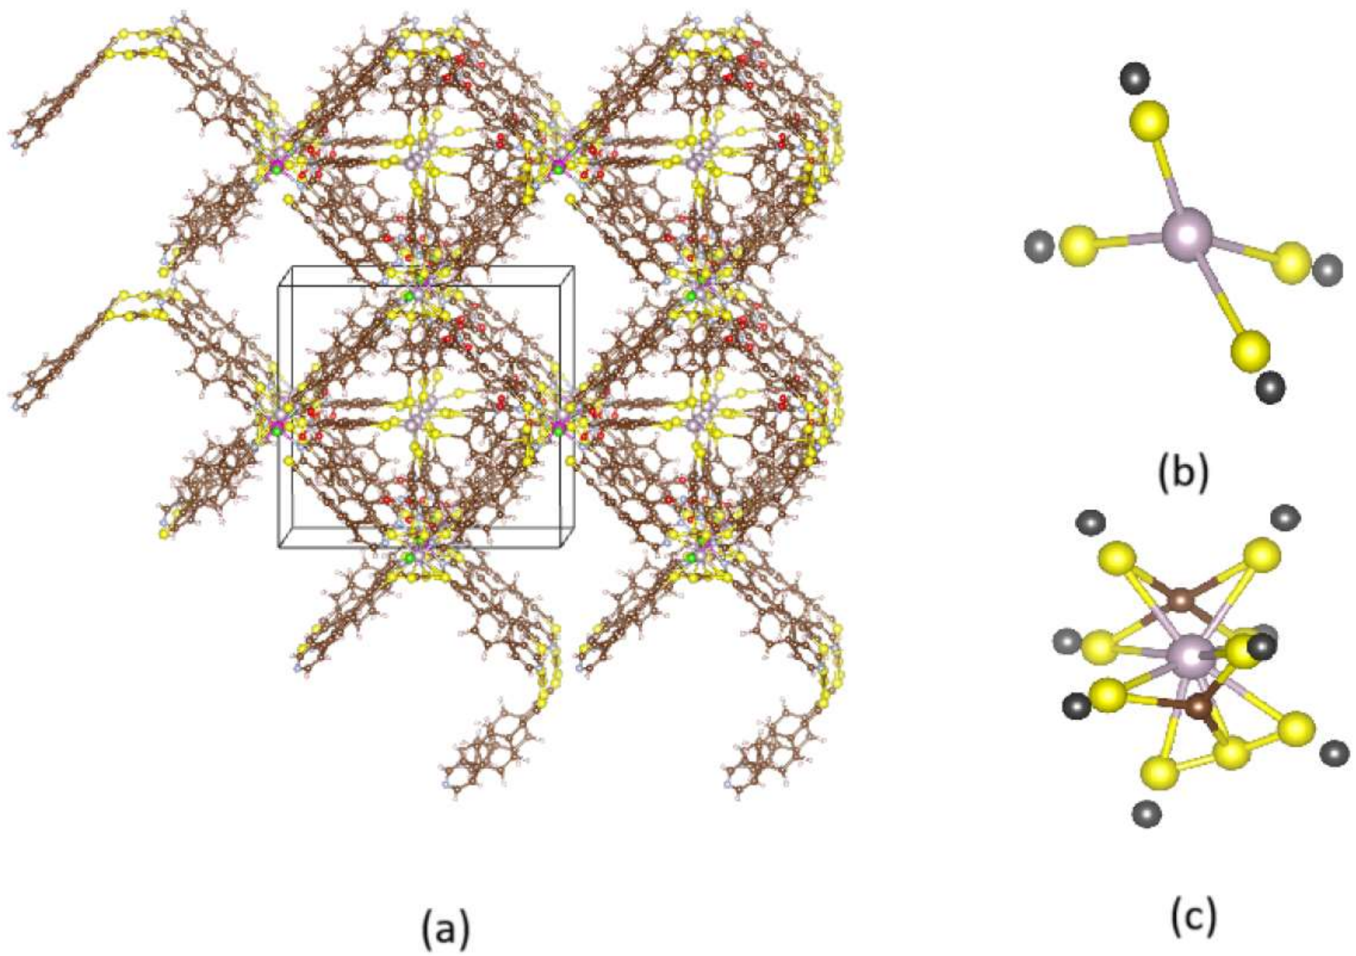

Example: Battery materials discovery process

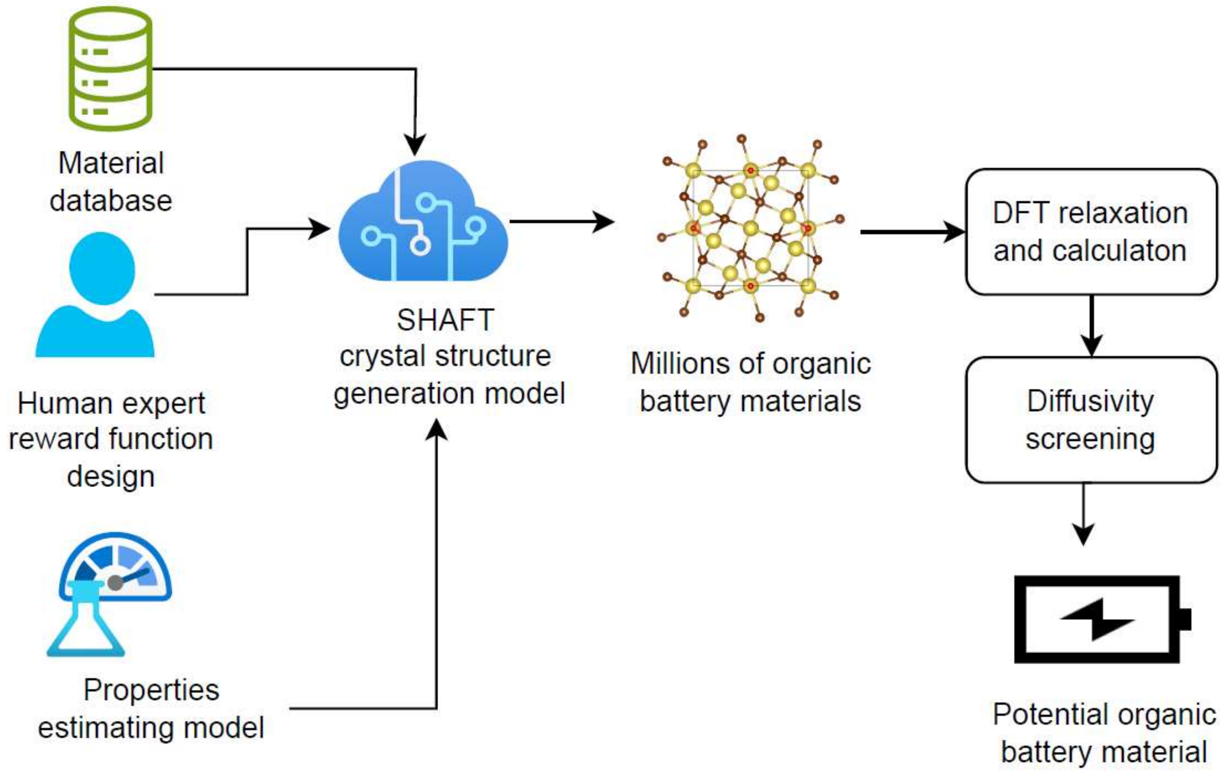

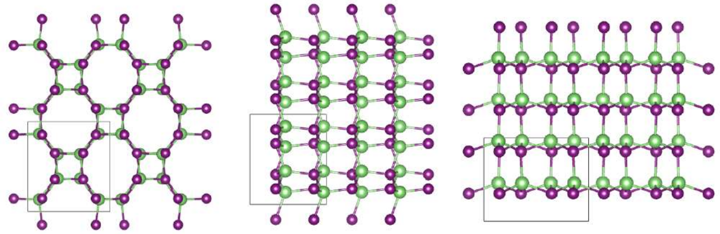

Li8I8

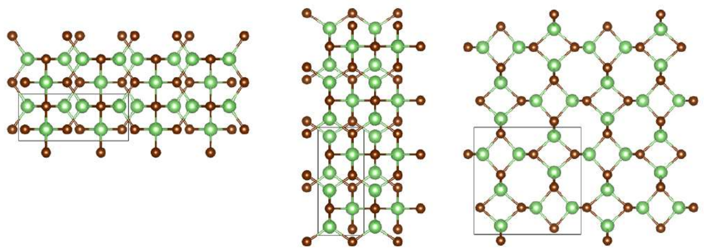

Li8Br8

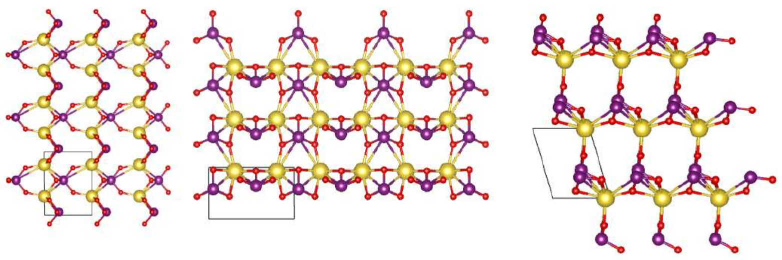

Na2I2O6

Topics

Explaining DTA deep learning model: feature attribution

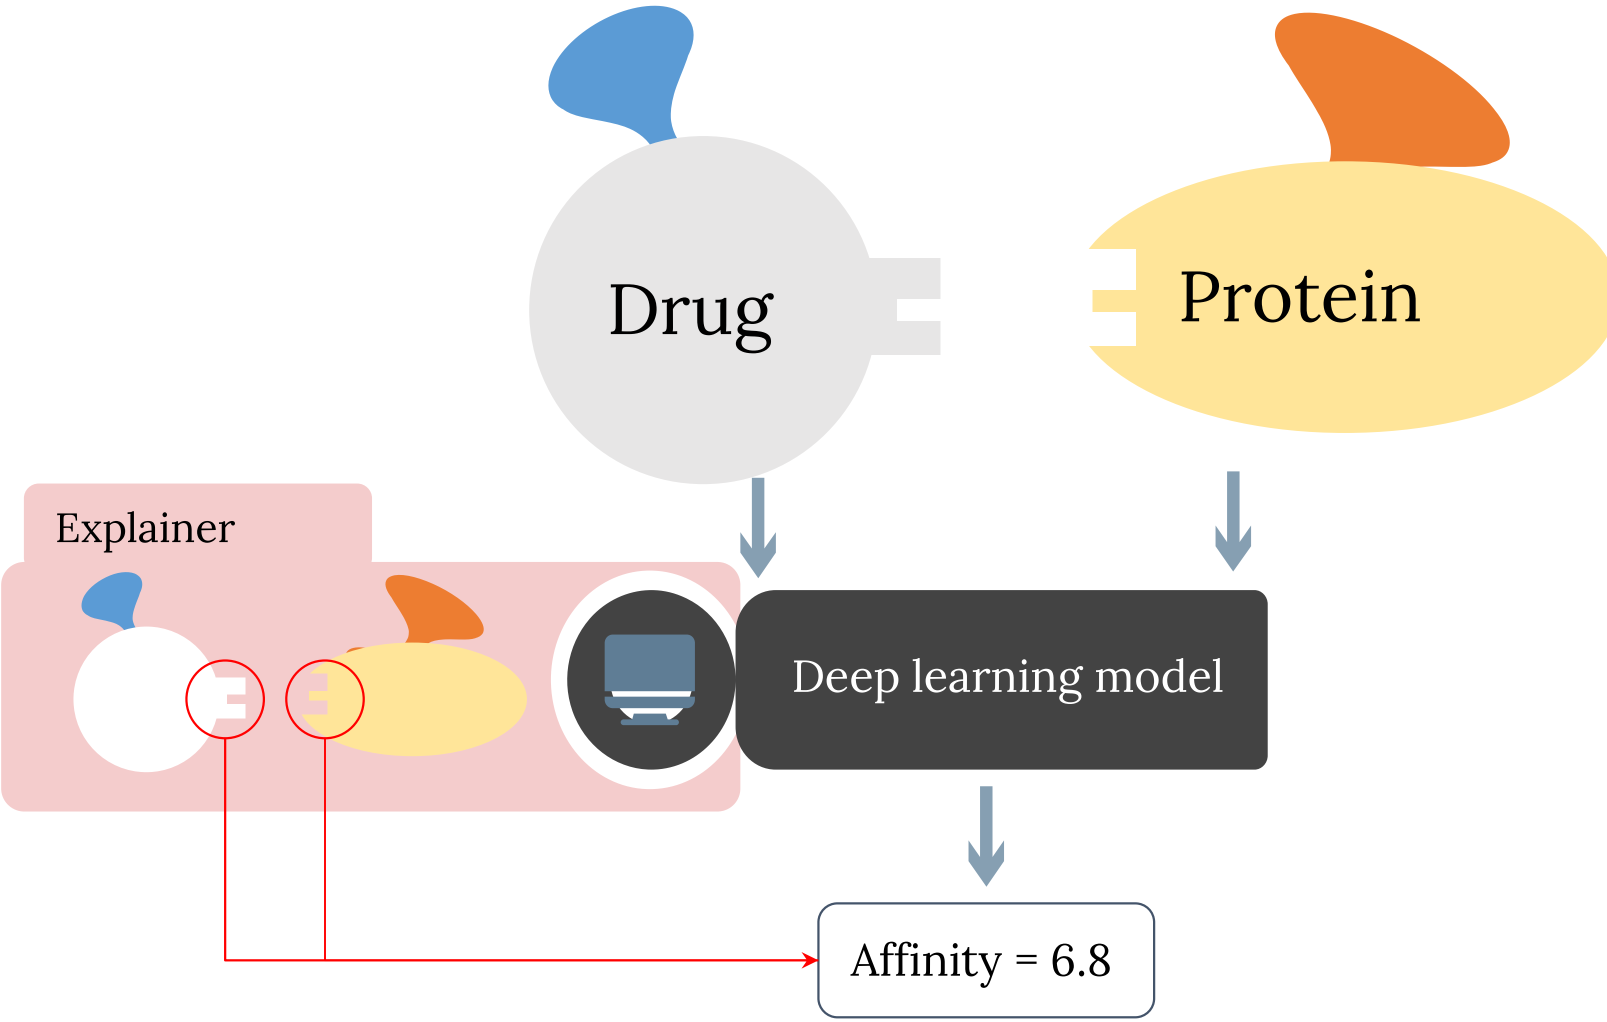

Show the contribution of each part of input to the model decision

Does not show the causal relationship between the input and the output of model

Slide credit: Tri Minh Nguyen

MACDA: MultiAgent Counterfactual Drug-target Affinity framework

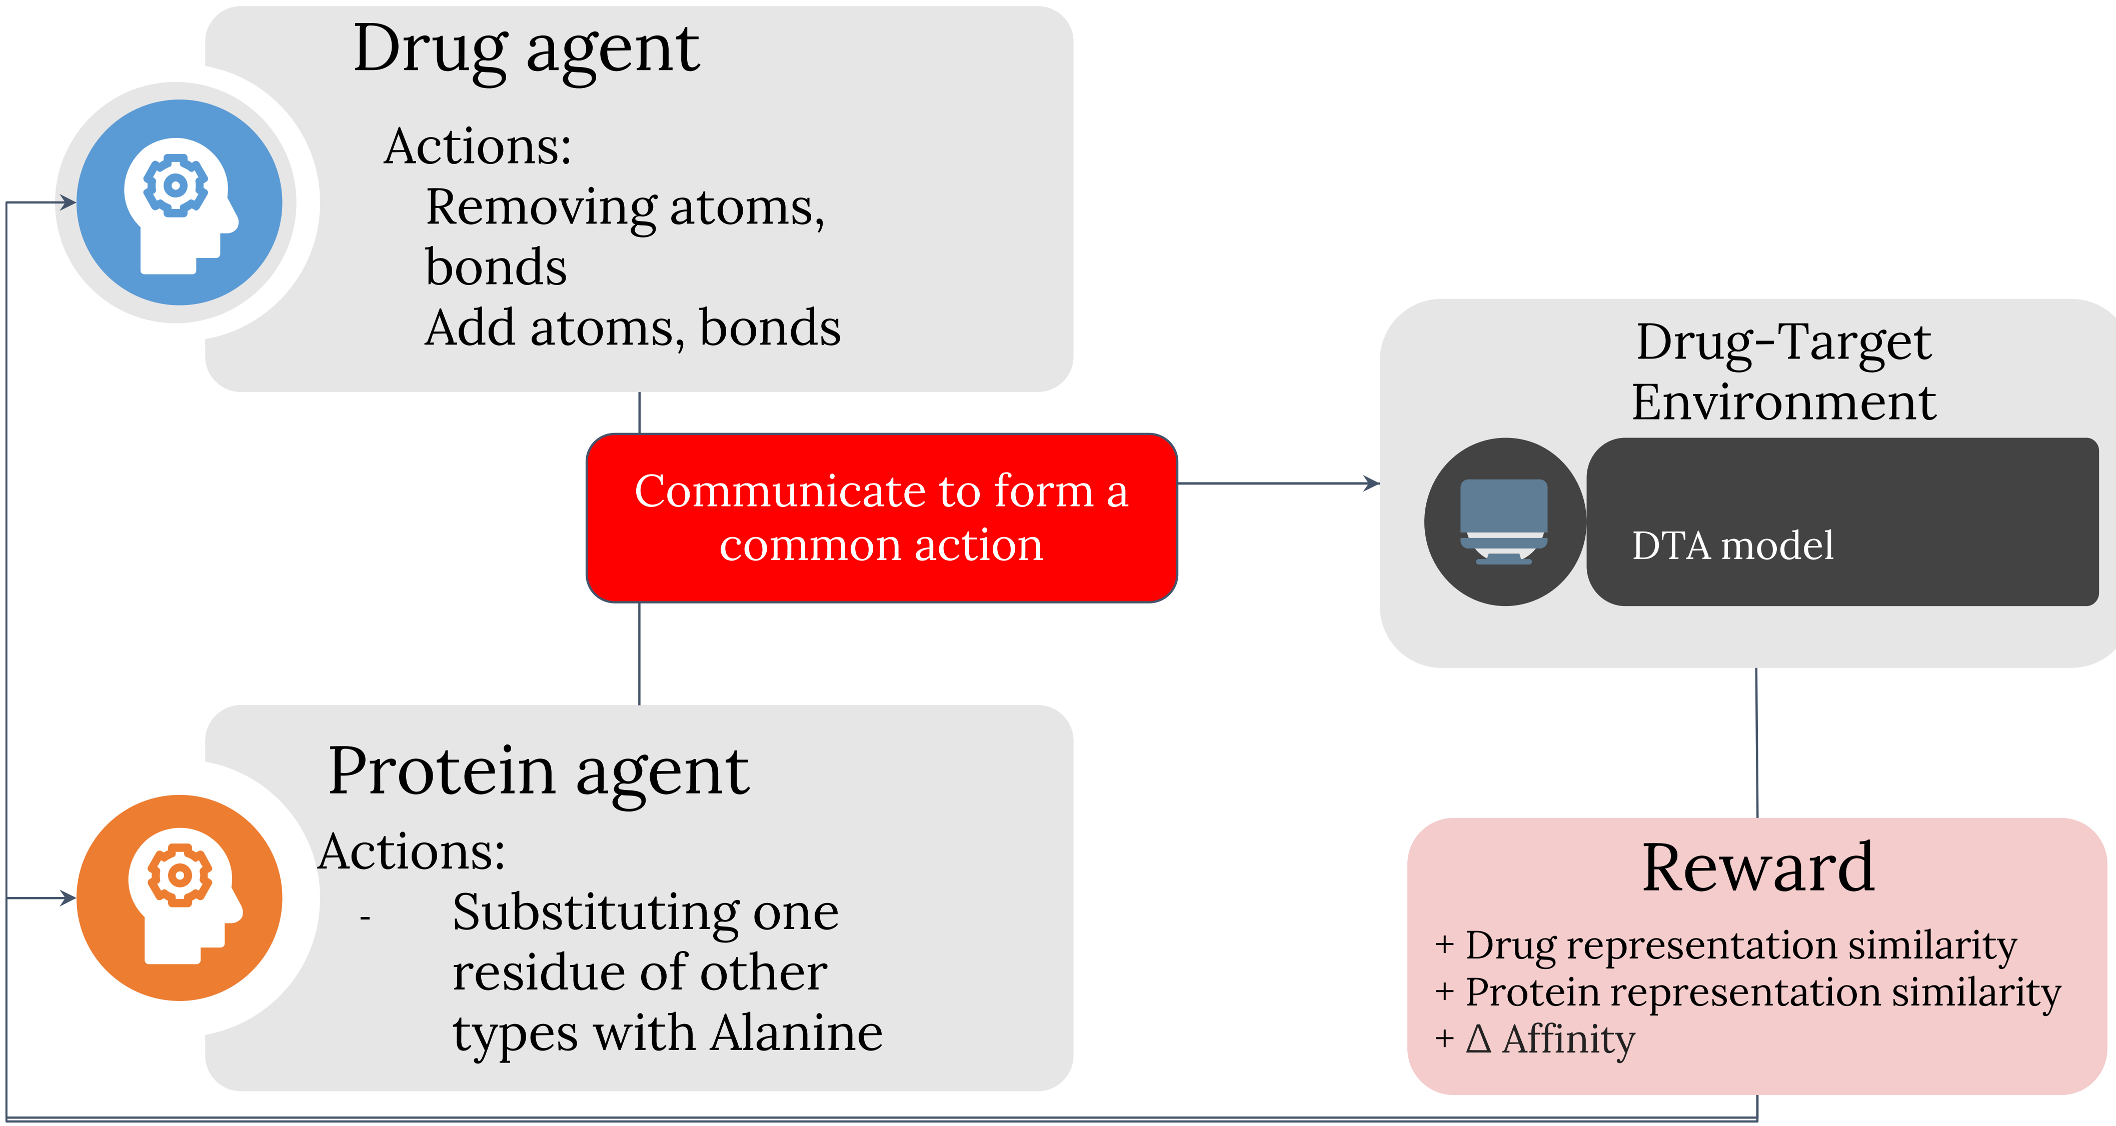

Nguyen, T.M., Quinn, T.P., Nguyen, T. and Tran, T., 2022. Explaining Black Box Drug Target Prediction through Model Agnostic Counterfactual Samples. *IEEE/ACM Transactions on Computational Biology and Bioinformatics.*

Slide credit: Tri Minh Nguyen

Topics

Embedding symbolic physics into ML

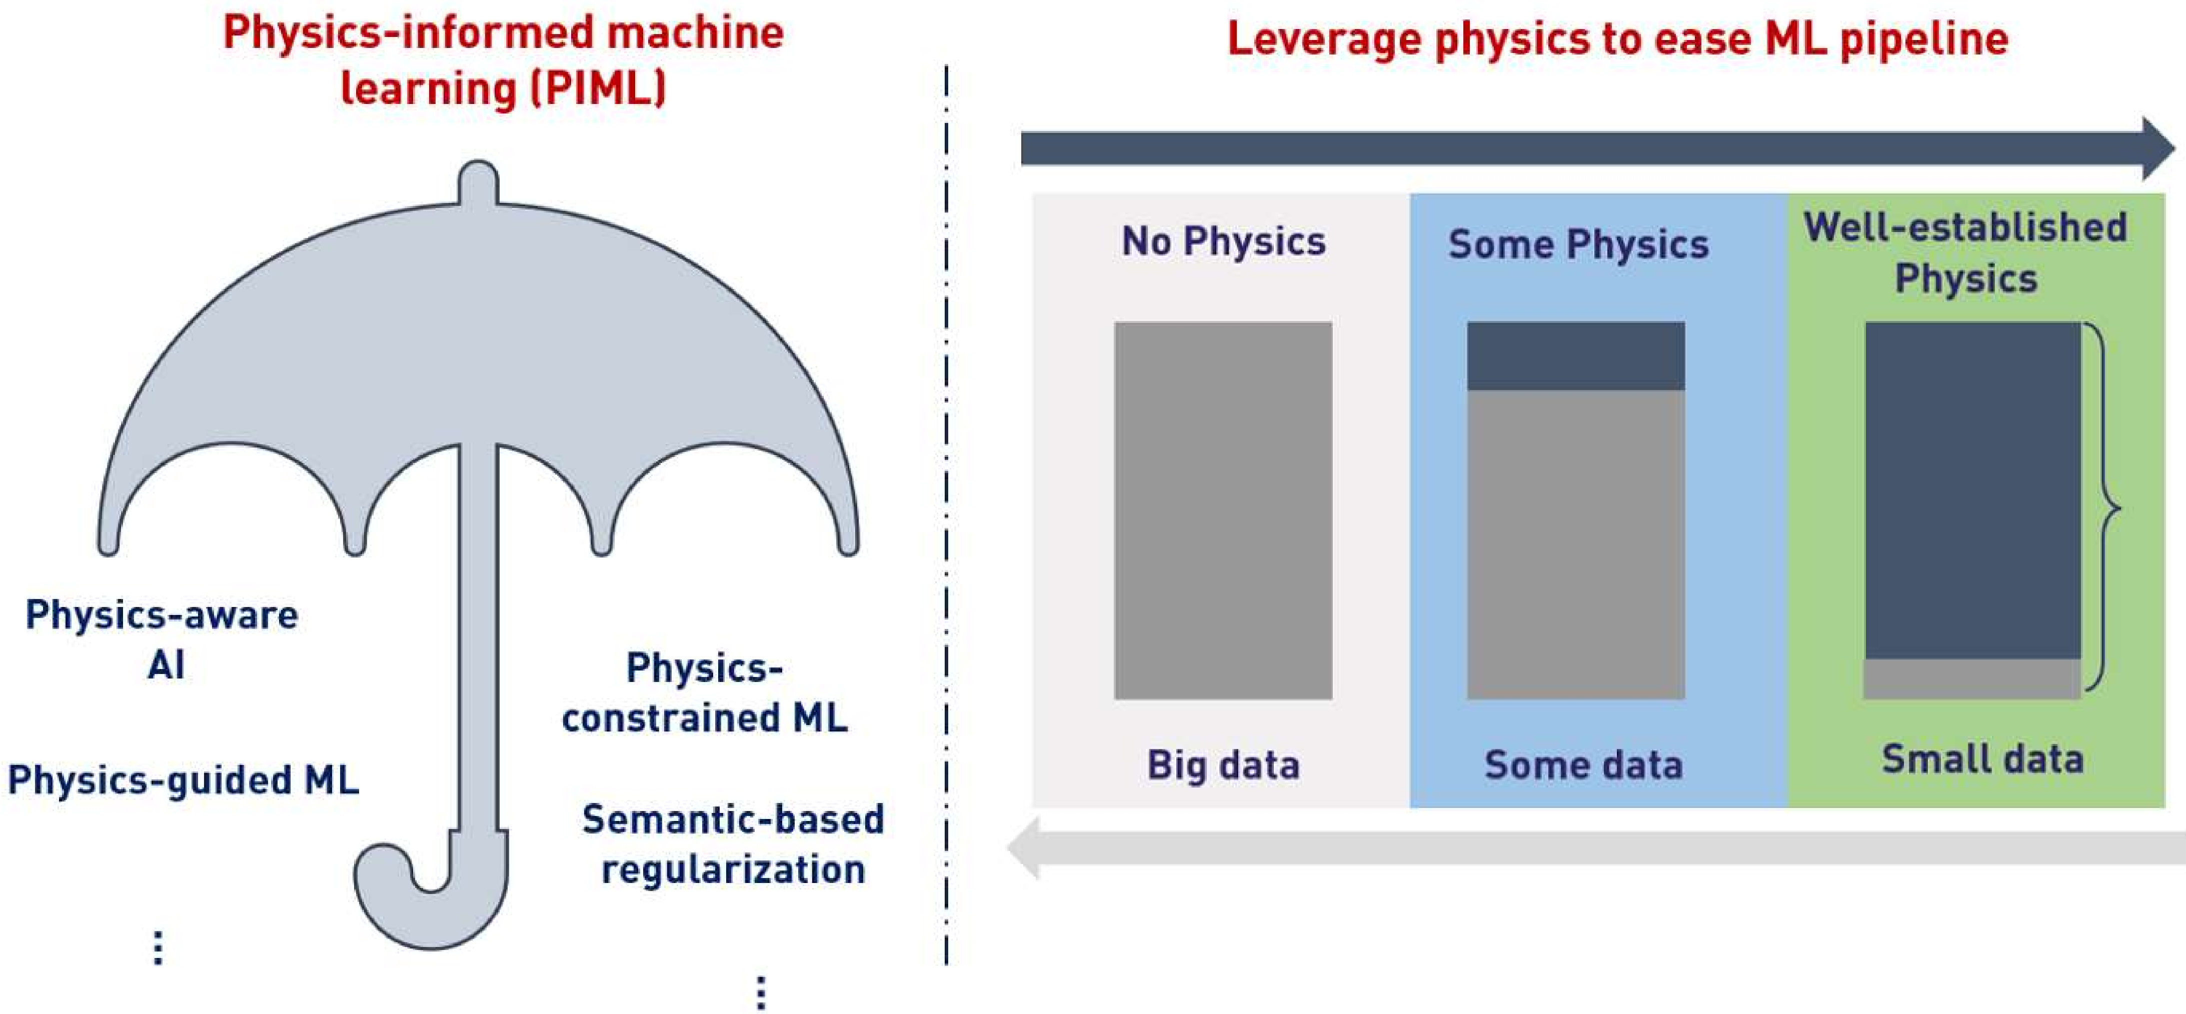

Source: @zhaoshuai1989

Leveraging physics invariance

- Newton laws
- Symmetry
- Conversation laws
- Noether’s Theorem linking symmetry and conservation.

Invariante Variationsprobleme.

(F. Klein zum fünfzigjährigen Doktorjubiläum.)

Von

Emmy Noether in Göttingen.

Vorgelegt von F. Klein in der Sitzung vom 26. Juli 1918[1].

Es handelt sich um Variationsprobleme, die eine kontinuierliche Gruppe (im Lieschen Sinne) gestatten; die daraus sich ergebenden Folgerungen für die zugehörigen Differentialgleichungen finden ihren allgemeinsten Ausdruck in den in § 1 formulierten, in den folgenden Paragraphen bewiesenen Sätzen. Über diese aus Variationsproblemen entspringenden Differentialgleichungen lassen sich viel präzisere Aussagen machen als über beliebige, eine Gruppe gestattende Differentialgleichungen, die den Gegenstand der Lieschen Untersuchungen bilden. Das folgende beruht also auf einer Verbindung der Methoden der formalen Variationsrechnung mit denen der Lieschen Gruppentheorie. Für spezielle Gruppen und Variationsprobleme ist diese Verbindung der Methoden nicht neu; ich erwähne Hamel und Herglotz für spezielle endliche, Lorentz und seine Schüler (z. B. Fokker), Weyl und Klein für spezielle unendliche Gruppen[2]. Insbesondere sind die zweite Kleinsche Note und die vorliegenden Ausführungen gegenseitig durch einander beein-

1) Die endgültige Fassung des Manuskriptes wurde erst Ende September eingereicht.

2) Hamel: Math. Ann. Bd. 59 und Zeitschrift f. Math. u. Phys. Bd. 50. Herglotz: Ann. d. Phys. (4) Bd. 36, bes. § 9, S. 511. Fokker, Verslag d. Amsterdamer Akad., 27./1. 1917. Für die weitere Litteratur vergl. die zweite Note von Klein: Göttinger Nachrichten 19. Juli 1918.

In einer eben erschienenen Arbeit von Kneser (Math. Zeitschrift Bd. 2) handelt es sich um Aufstellung von Invarianten nach ähnlicher Methode.

Kgl. Ges. d. Wiss. Nachrichten. Math.-phys. Klasse. 1918. Heft 2. 17

First page of **Emmy Noether**'s article "Invariante Variationsprobleme" (1918). Source: Wikipedia

Physics-informed neural networks

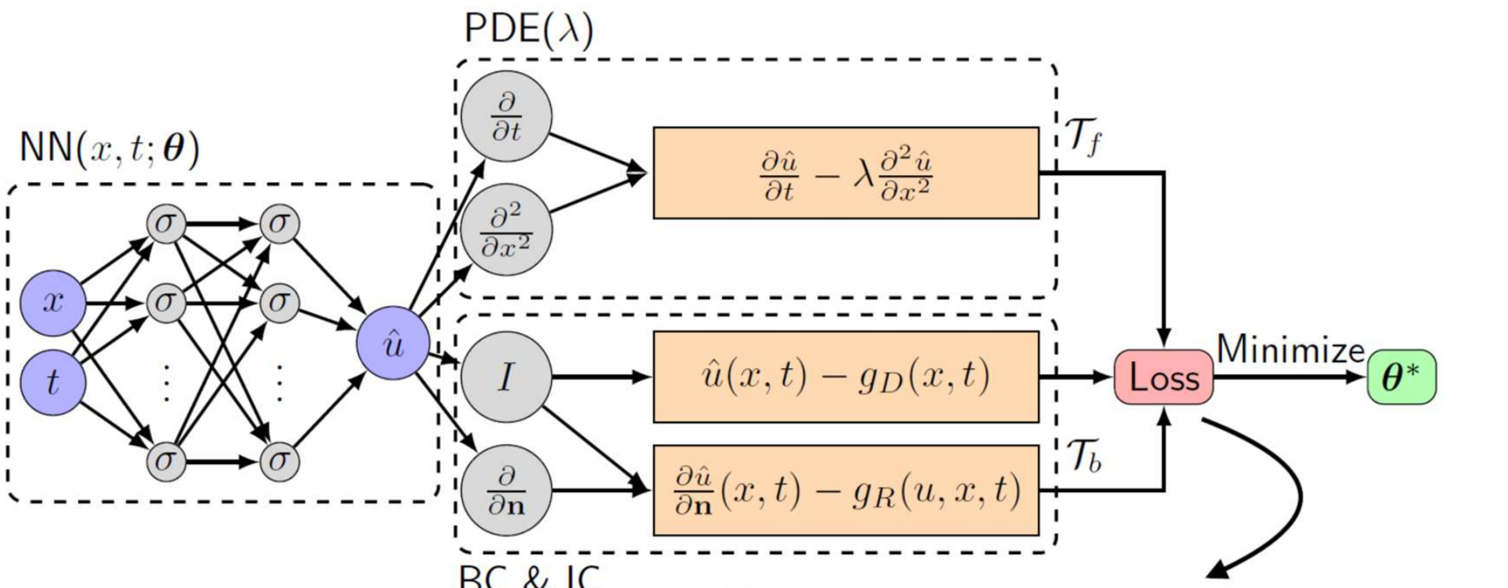

$$\mathcal{L}(\theta) := \underbrace{\frac{1}{N_u}\sum_{i=1}^{N_u}[\boldsymbol{u}_i - f_\theta(\boldsymbol{x}_i)]^2}_{\text{Data fit}} + \underbrace{\frac{1}{\lambda}\mathcal{R}[f_\theta(\boldsymbol{x})]}_{\text{Physics regularization}}$$

Figure from talk by Perdikaris & Wang, 2020.

Case study: Covid-19 infections in VN 2021

- Classic model SIR: Close-form solutions hard to calculate
  - Parameters change over time due to intervention → Need more flexible framework.
- Solution: Richard's ODE equation → Mixture of Gompertz curves
- Task: 10-20 data points → Extrapolate 150 more.

$$Y'(t) = \alpha\left(1 - \left(\frac{Y}{K}\right)^{\nu}\right)Y$$

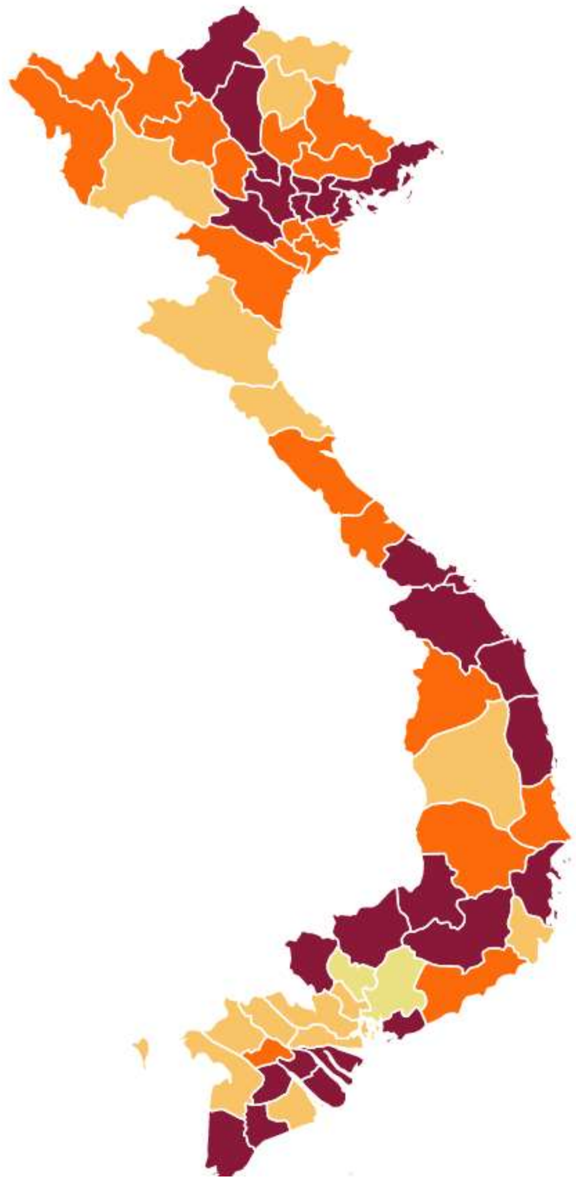

Case of HCM City

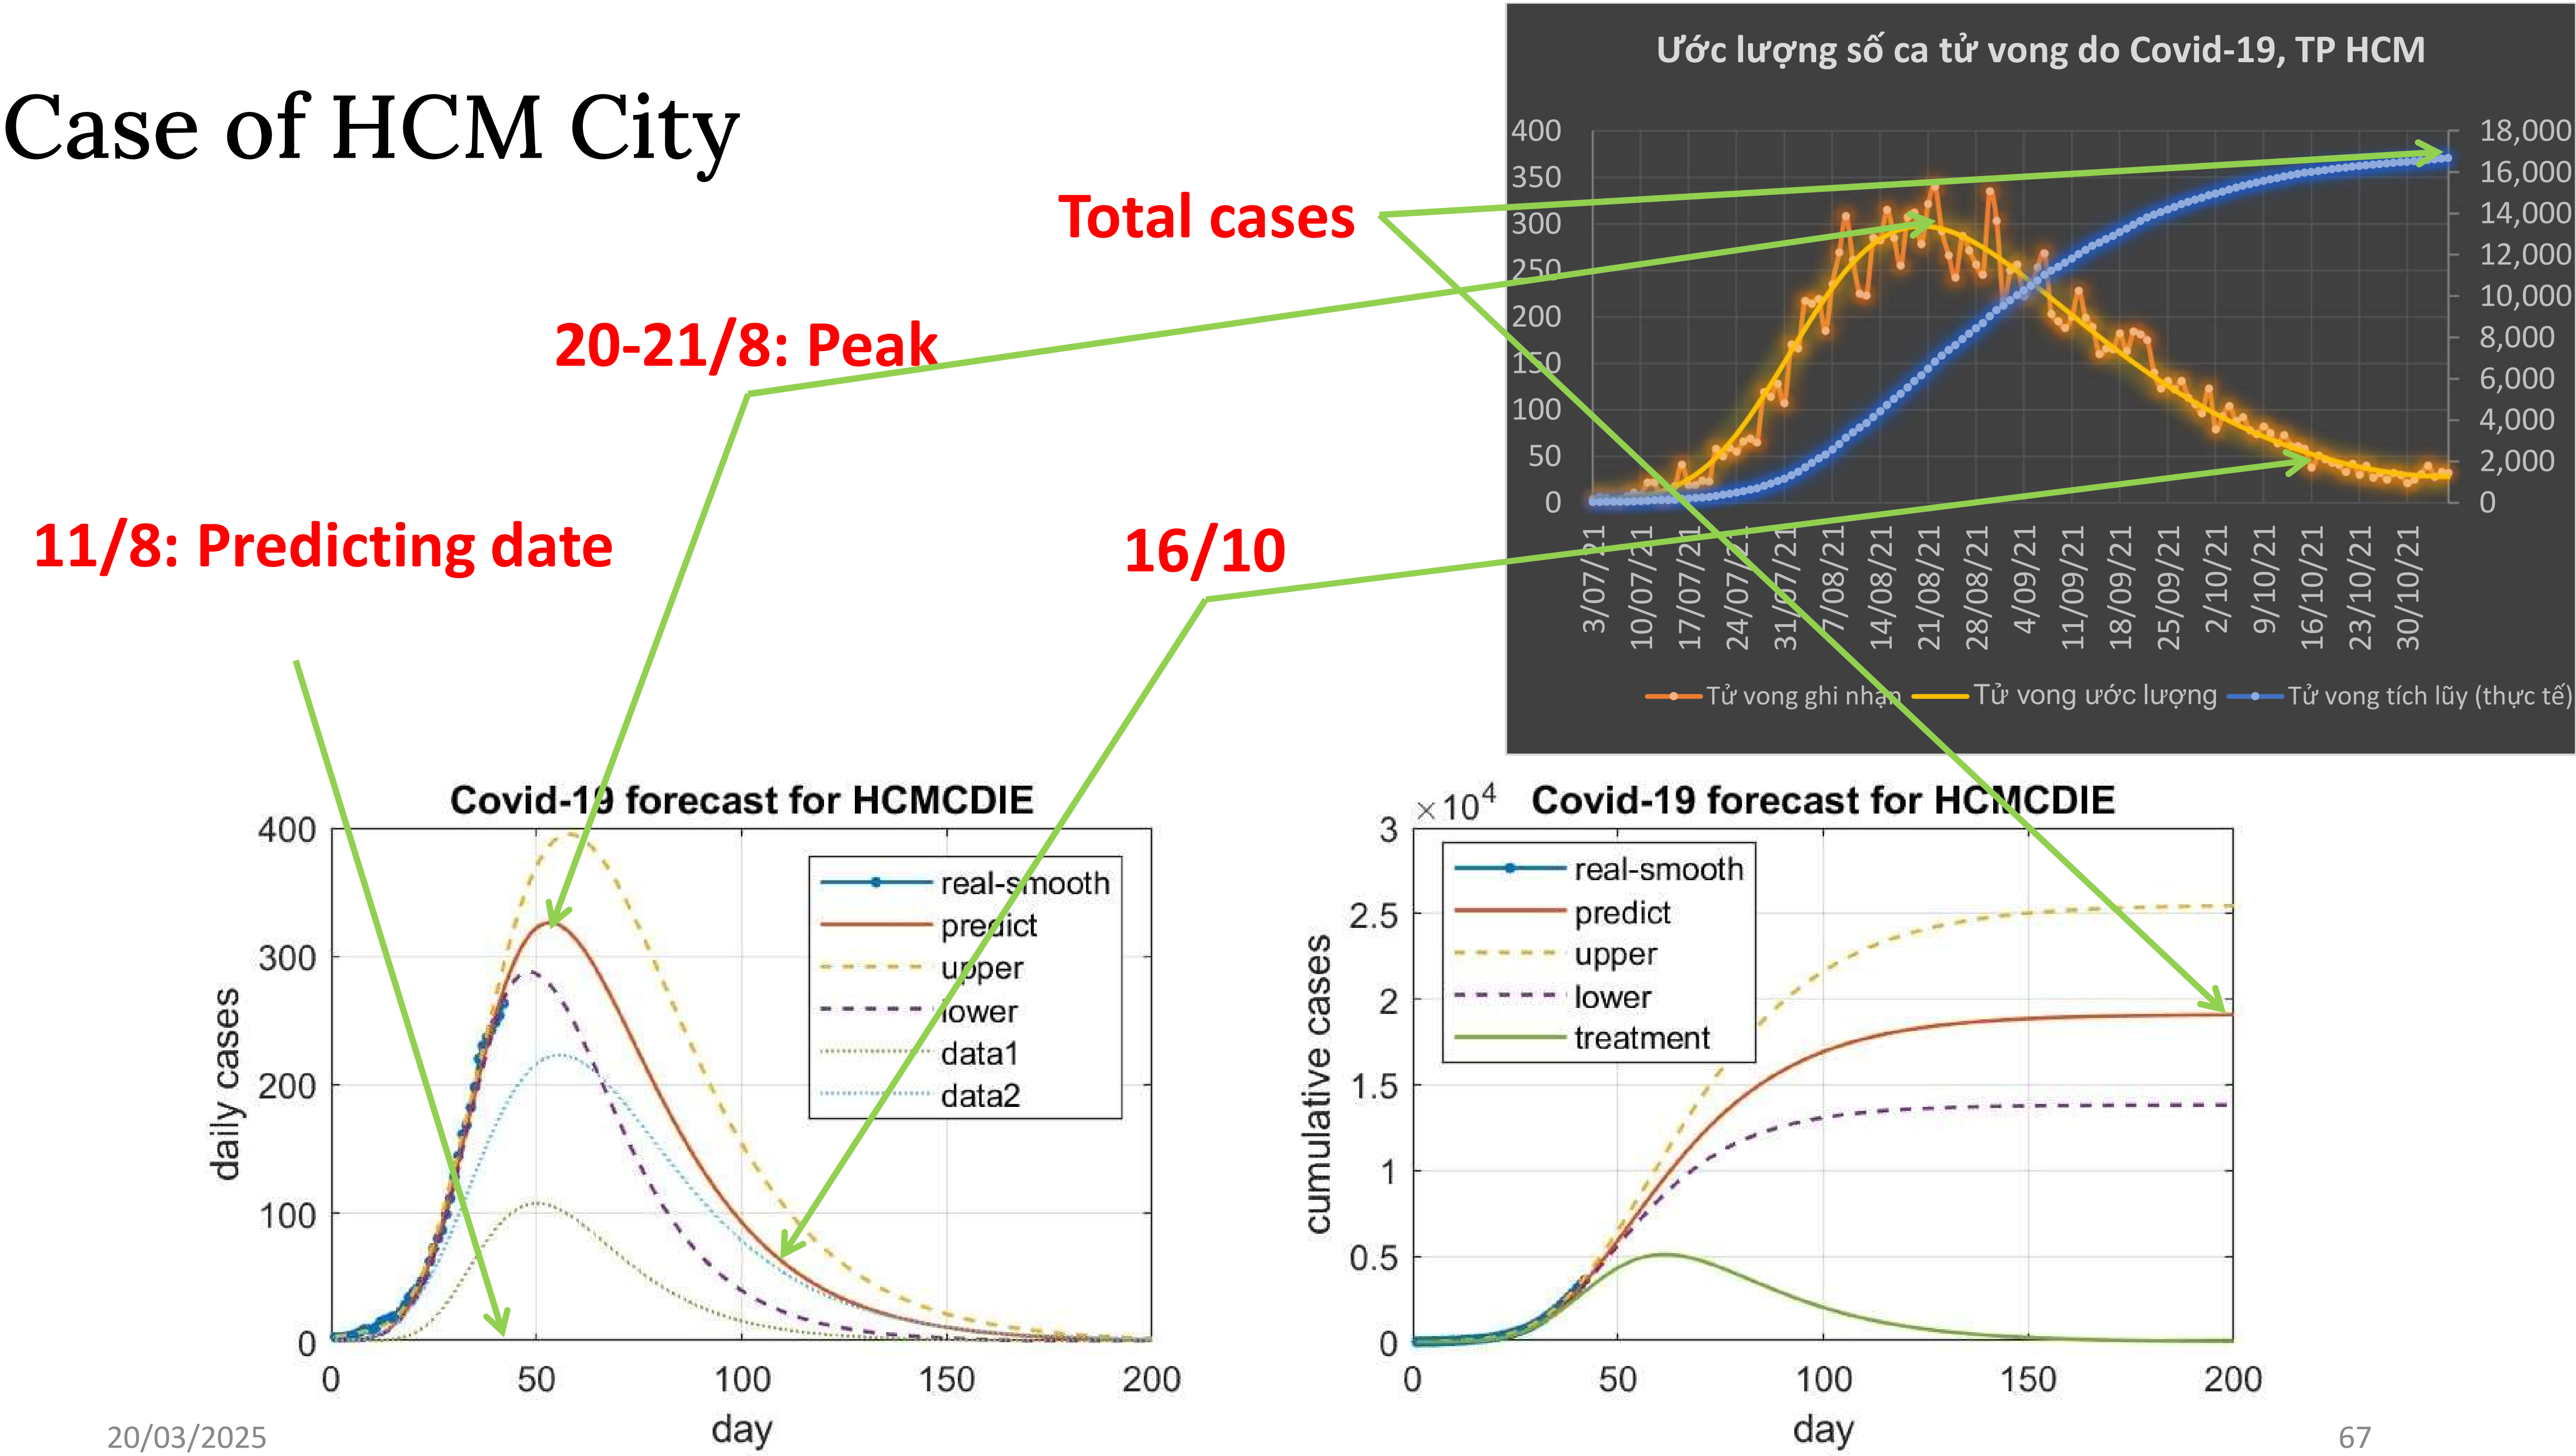

DFT-informed GNN: Embedding material graphs using the electron-ion potential

- Integrate the external (electron-atom) potential into GNN!

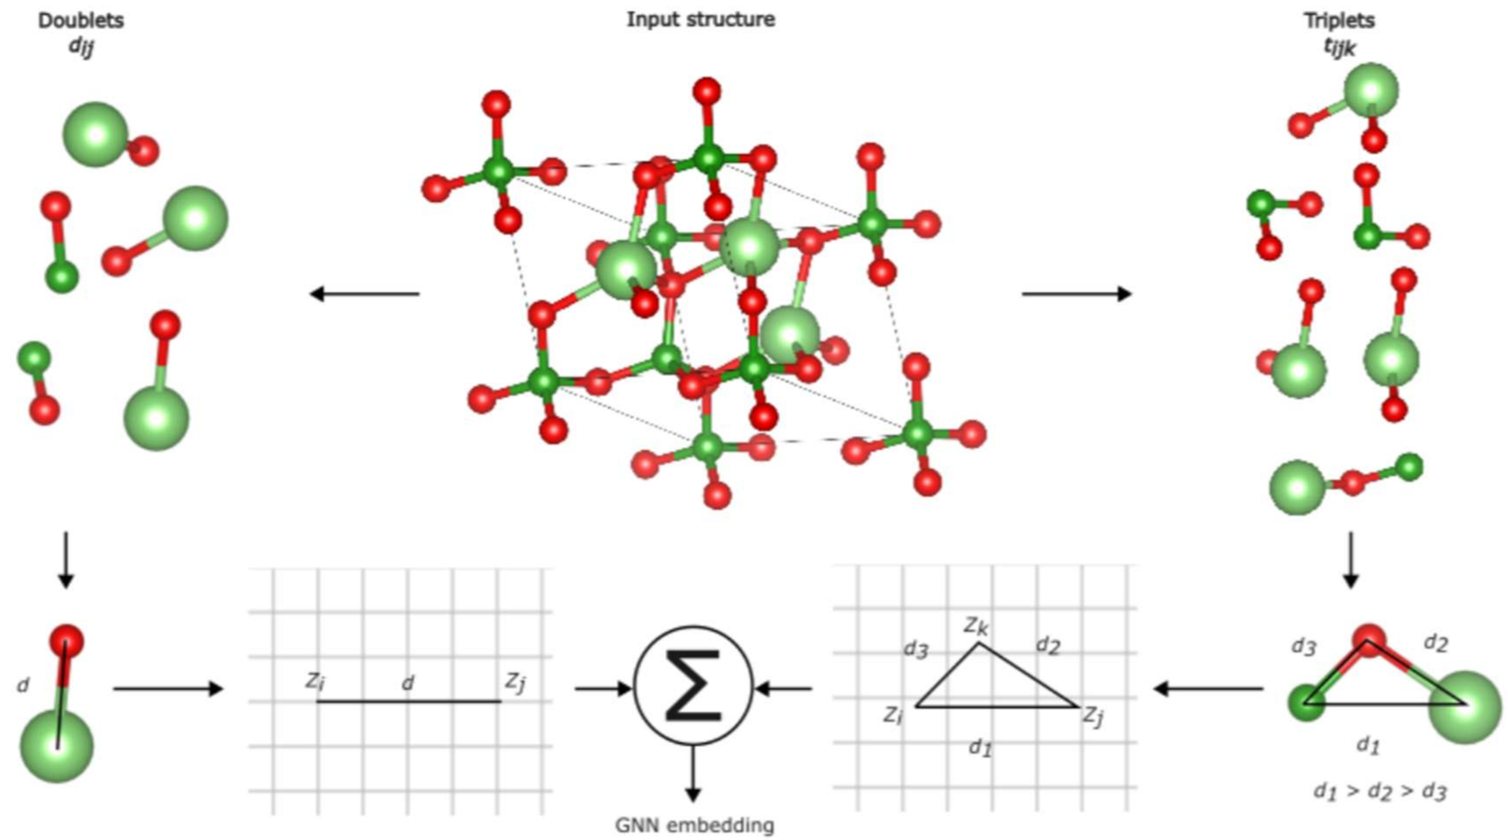

Tawfik, Sherif Abdulkader, Tri Minh Nguyen, Salvy P. Russo, Truyen Tran, Sunil Gupta, and Svetha Venkatesh. "Embedding material graphs using the electron-ion potential: application to material fracture." *Digital Discovery* (2024).

Topics

- ~~Why AI for Science?~~
- ~~Representation~~
- ~~Prediction~~
- ~~Optimization & Generalization~~
- ~~Explanation~~
- ~~Physics-informed ML~~
- AI Co-scientist
- Future & Risks

Social AI scientists

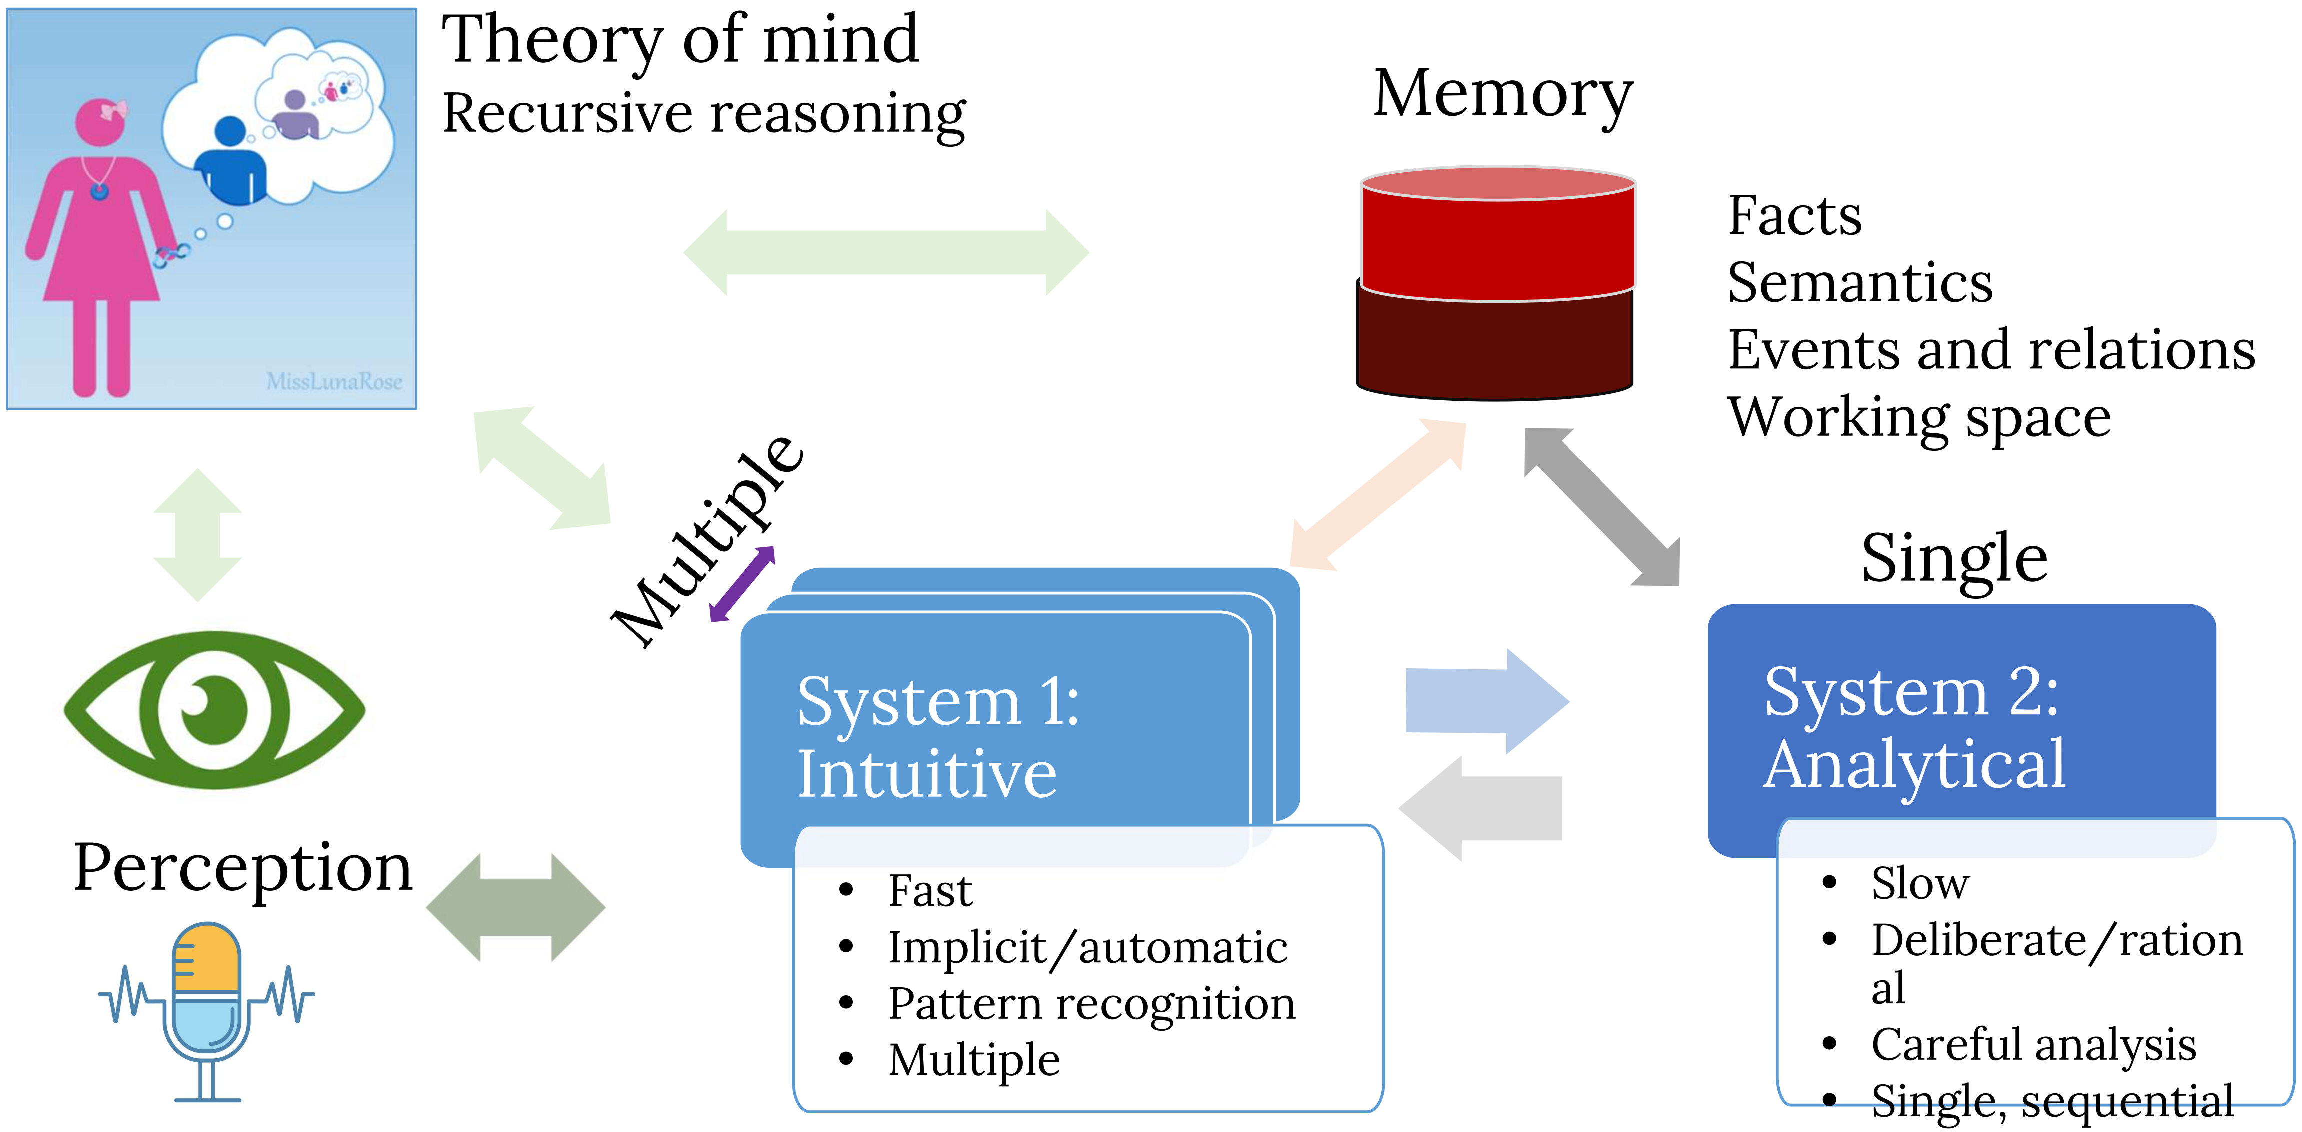

Image credit: VectorStock | Wikimedia

AI alignment

1. The machine's only objective is to maximize the realization of human preferences.
2. The machine is initially uncertain about what those preferences are.
3. The ultimate source of information about human preferences is human behavior.

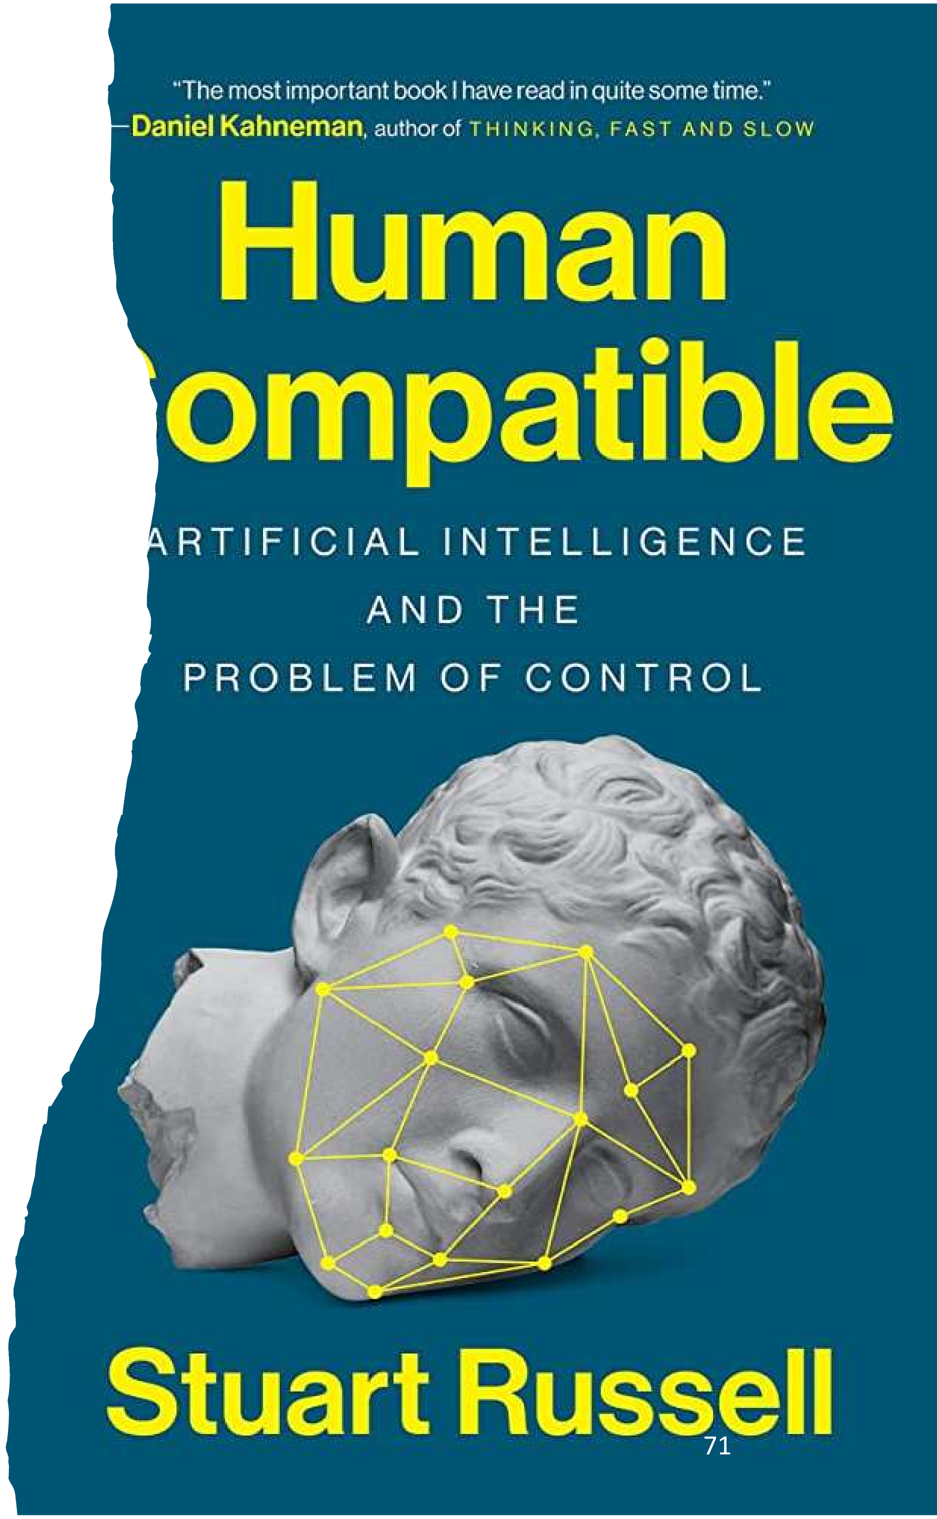

Generative AI => Agentic workflow

GenAIs are compression engine

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

GenAIs are approximate program database

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

GenAIs are World Model

We can live entirely in simulation!

Agents for science – tool use

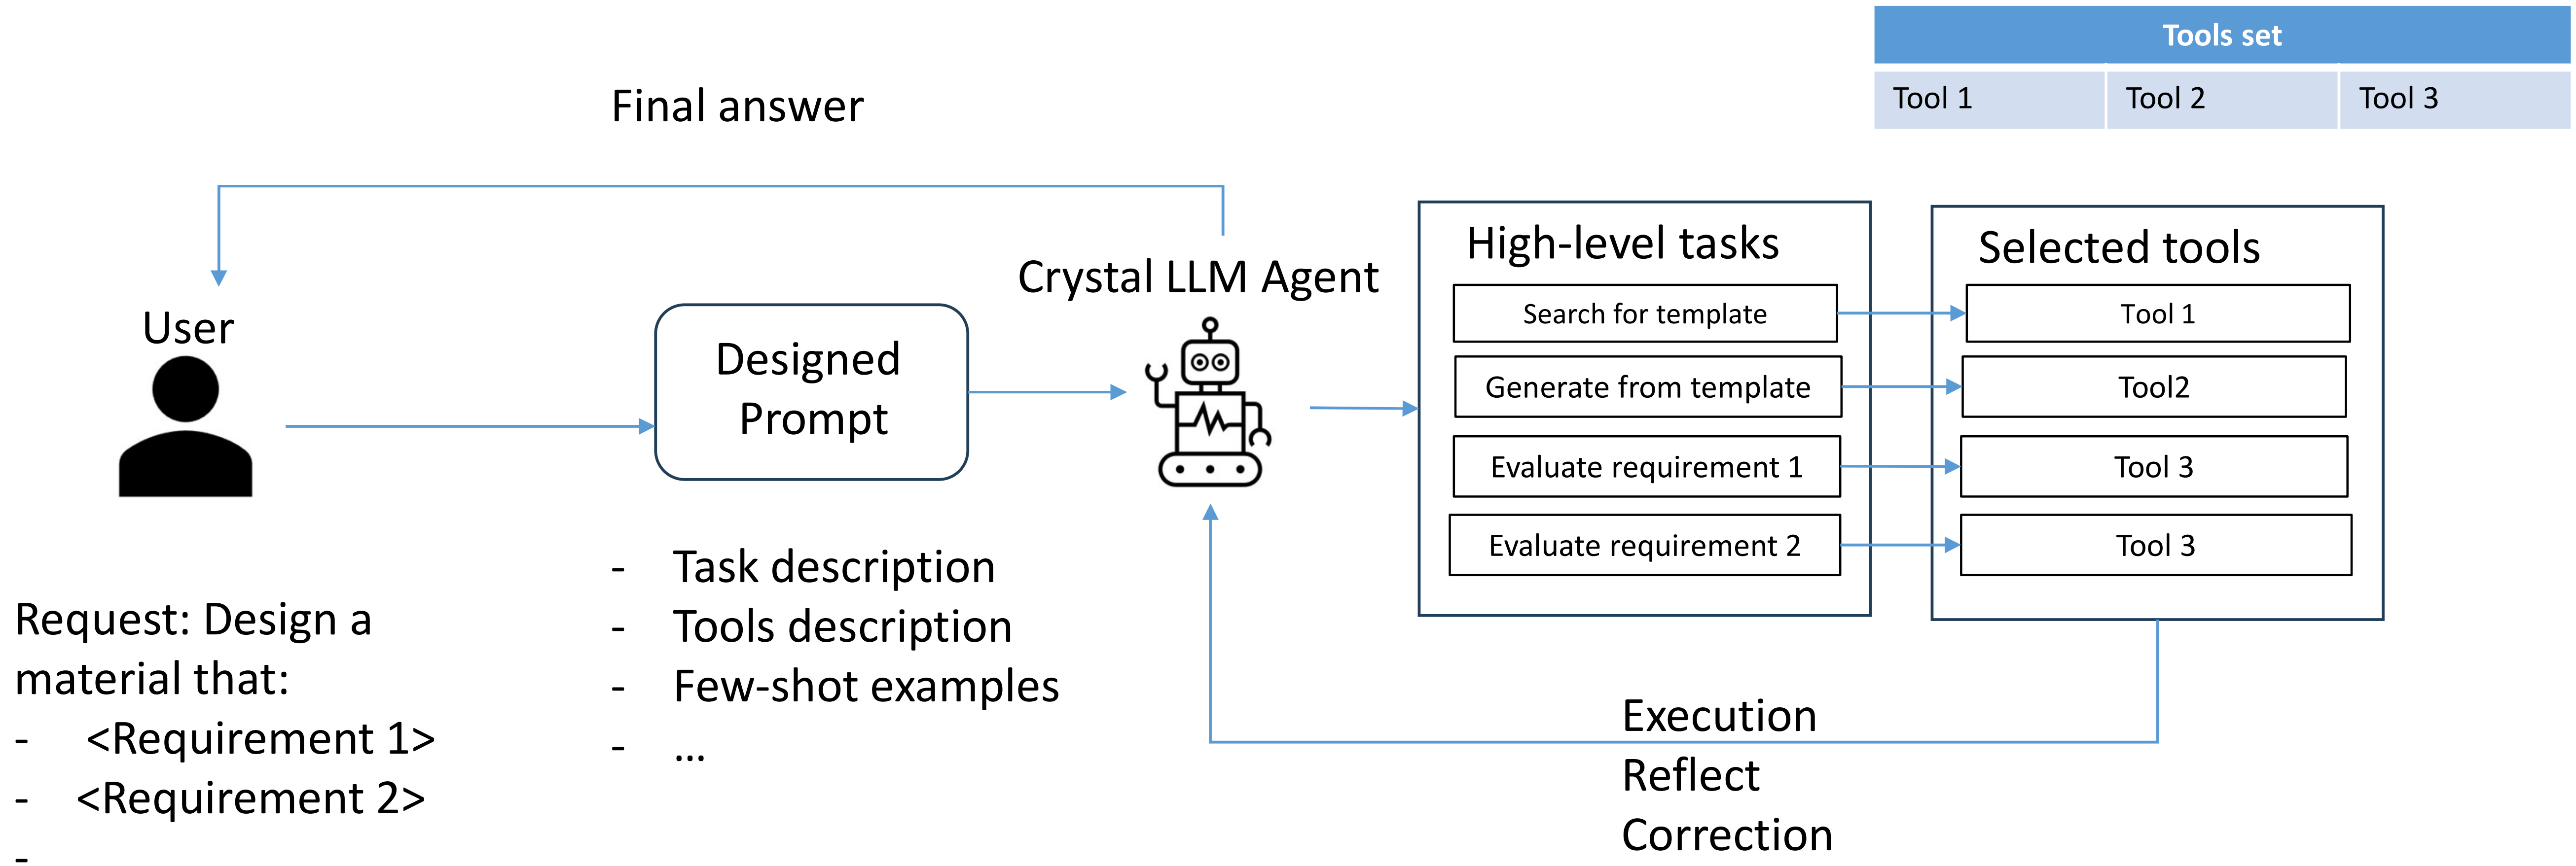

LLM for inverse design

- Input prompt: “You are machine learning crystal structure prediction model. Predict crystal structure in Crystallographic Information File format to representing crystallographic information. You are given formula, formation energy, band gap. Formula: {formula}. Formation energy: {forme}. Band gap: {bandgap}.”
- Output: Structure in CIF format
- Evaluation: Predict the formation energy of output structure using matgl and compare with the {forme} in the promp.

Results – formation energy inverse design

Codegemma – finetuned on mp-20

	RMSE	MSE	Pearson	Spearman
Train (temp. = 0.7)	0.764	0.583	0.792	0.722
Train (temp. = 1.2)	0.962	0.926	0.695	0.657
Val (temp. = 0.7)	0.866	0.750	0.743	0.723
Val (temp. = 1.2)	0.834	0.695	0.774	0.633

Llama3-8B-instruct – no finetune

	RMSE	MSE	Pearson	Spearman
Val (temp. = 0.7)	1.554	2.417	0.406	0.381

Llama-3-70B-Instruct-Q4_K_M.gguf – no finetune

	RMSE	MSE	Pearson	Spearman
Val temp 0.8	1.337	1.787	0.611	0.573

Results – band gap inverse design

	RMSE	MSE	Pearson	Spearman
Train (temp. = 0.7)	0.998	0.996	0.682	0.492
Train (temp. = 1.2)	1.039	1.080	0.676	0.516
Val (temp. = 0.7)	0.993	0.987	0.693	0.470
Val (temp. = 1.2)	0.909	0.826	0.615	0.442

Prediction properties task

- Dataset: Material project dataset (55k samples) and mp-20 dataset.
- Model: code-gemma 7b tuned on mp-20

Method	Dataset	RMSE	MSE	Pearson	Spearman	MAE
code-gemma 7b tuned on mp-20	Material project	0.654	0.428	0.845	0.852	0.361
	Mp-20 (train set)	0.187	0.034	0.983	0.977	0.124
	Mp-20 (val set)	0.192	0.037	0.982	0.976	0.129
Matformer						0.021

AI Co-Scientist, Google DeepMind

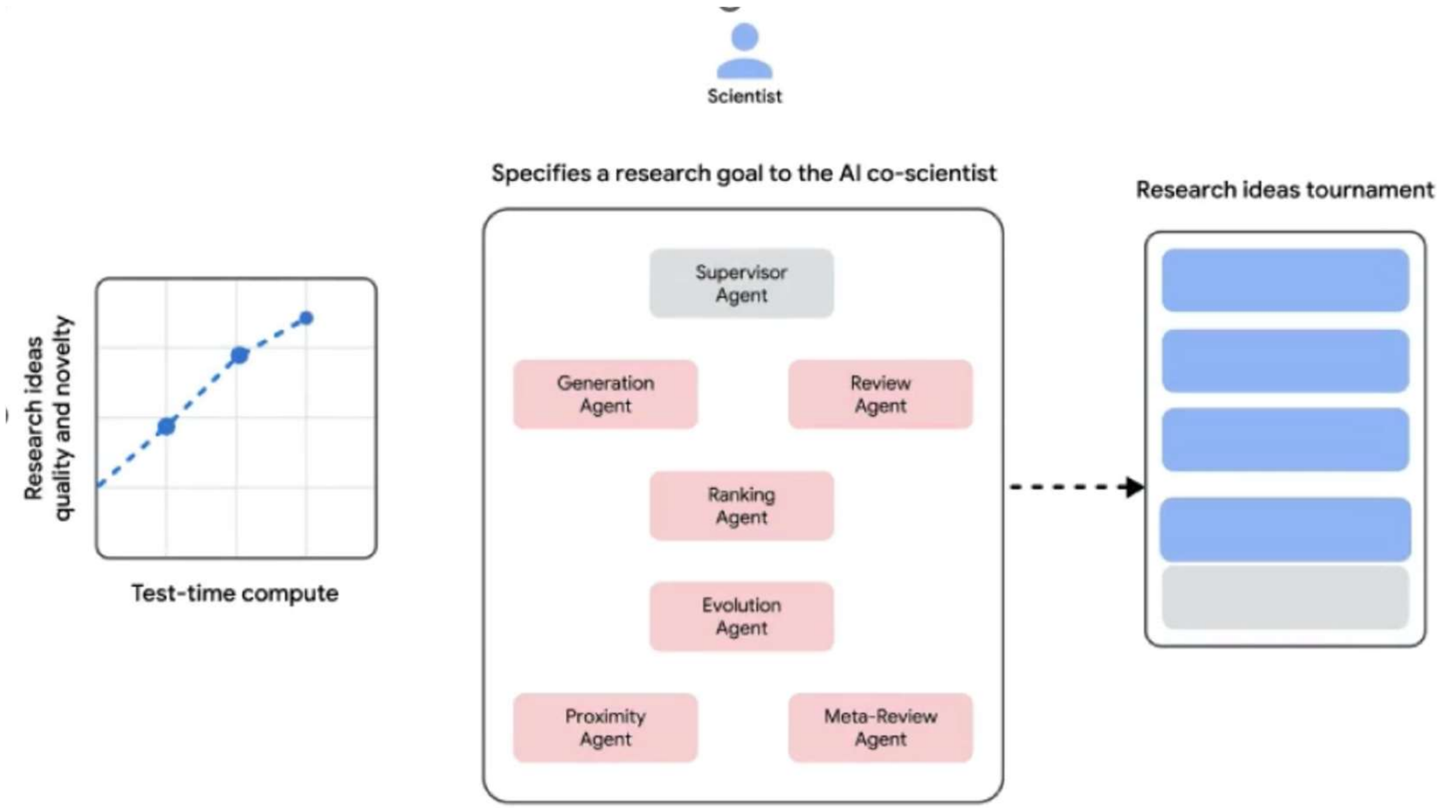

Topics

- ~~Why AI for Science?~~
- ~~Representation~~
- ~~Prediction~~
- ~~Optimization & Generalization~~
- ~~Explanation~~
- ~~Physics-informed ML~~
- ~~AI Co-scientist~~
- Future & Risks

AI's technical success formula

DATA

KNOWLEDGE

COMPUTE

SCALABLE ALGORITHMS

TALENT

Example: AI to tackle COVID-19

Action areas

- Robot deployment
- Rapid treatment
- Vaccine development
- Early warning
- Social distancing
- Home isolation
- Quarantine
- Mental health
- Information
- Personal actions
- Collective decisions
- Education
- Finance

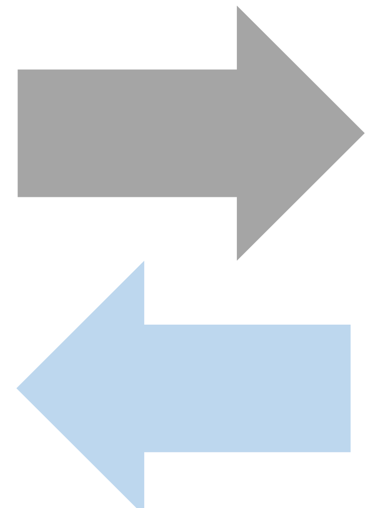

AI areas

- Computer vision
- NLP
- Reinforcement learning
- Time-series
- Interpretable learning
- Reasoning, causality
- Transfer, continual and lifelong learning
- Uncertainty quantification
- Unsupervised learning
- Structured data
- Knowledge-driven ML

Inspired by Rolnick, David, et al. "Tackling Climate Change with Machine Learning." arXiv preprint arXiv:1906.05433 (2019).

Looking into the future

Giorgio Parisi, 2021 Nobel in Physics for complex systems

AI as a discovery tool: e.g., quantum mechanics, materials science, and complex systems.

Interdisciplinary: Physics + AI + cognitive sciences for study of universe and human cognition.

Philosophical implications: Informational fabric => the nature of consciousness, intelligence, and the universe's computational structure.

Future-ready scientists: Technology + science.

Credit: nobelprize.org

BUSINESS • TECHNOLOGY

Anthropic CEO Dario Amodei Believes A.I. Could Double Human Lifespans in 5 Years

Because "we can make 100 years of progress in areas like biology in five to ten years," the former OpenAI executive said at Davos this week.

By Alexandra Tremayne-Pengelly • 01/24/25 2:22pm

AI as physics

> Intelligence as self-organizing phenomena: reducing ignorance/entropy

> Neural networks as a statistical mechanical system

> Learning as variational optimization

> Reasoning free-energy minimization

> Phase transition may occur in AI systems

> Ultimate AI must solve the **consciousness problem**, which may require new physics.

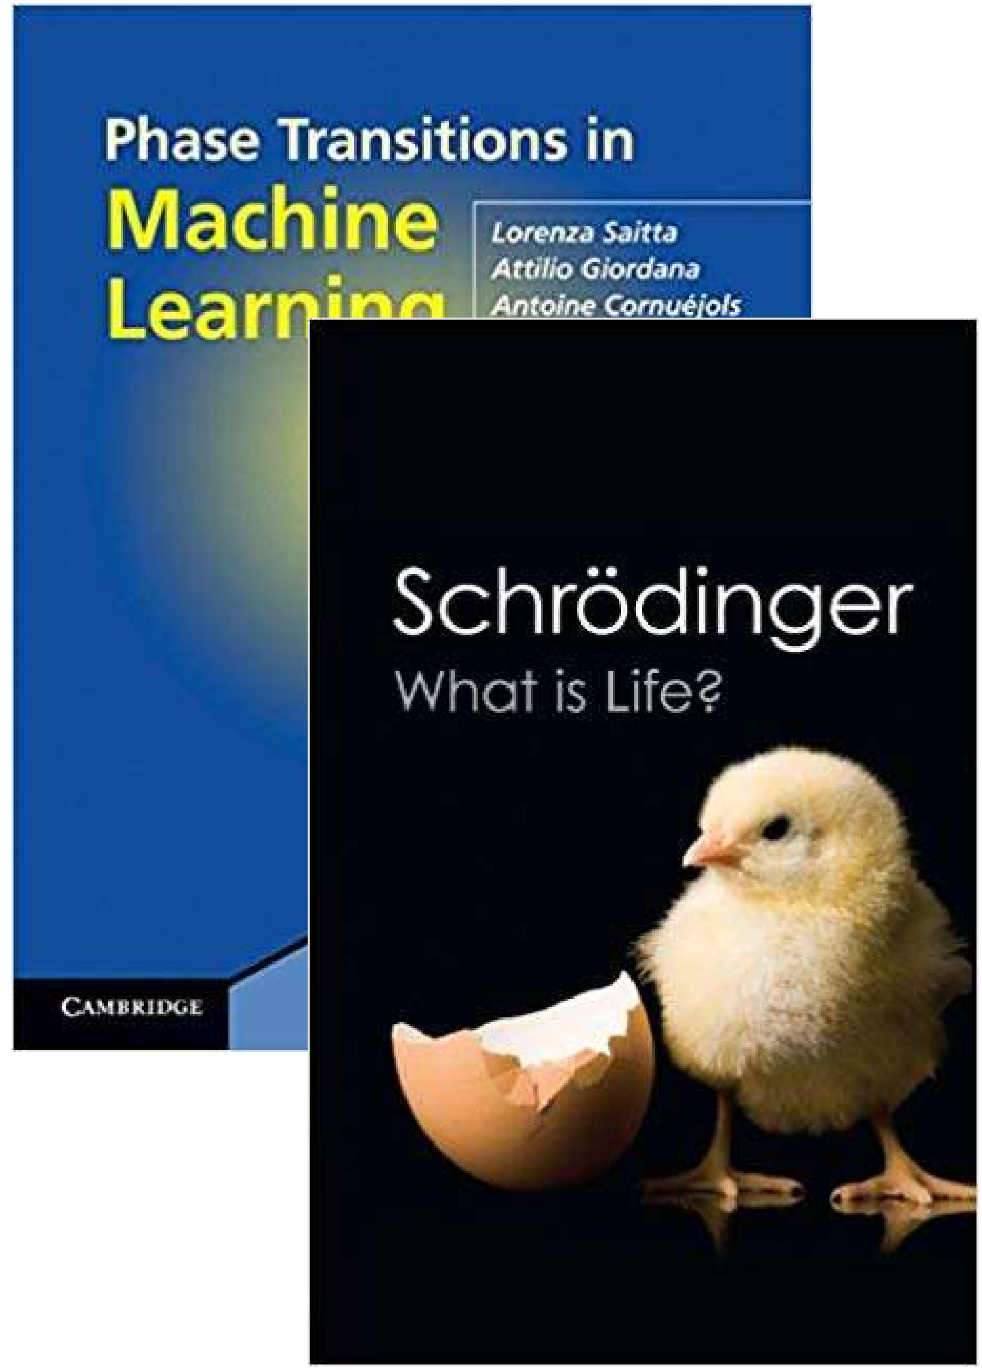

New dynamics of knowledge

- Printing technology (Gutenberg Bible in 1455): Spreading human thought
- AI: Distillation and elaboration of thought

=> new concepts of human thought and interaction with machines

⇒new challenges on a scale not experienced since the beginning of the Enlightenment.

⇒More here

Photo by Chip Somodevilla/Getty Images

WSJ | OPINION

OPINION | COMMENTARY

ChatGPT Heralds an Intellectual Revolution

Generative artificial intelligence presents a philosophical and practical challenge on a scale not experienced since the start of the Enlightenment.

By Henry Kissinger, Eric Schmidt and Daniel Huttenlocher

Feb. 24, 2023 2:17 pm ET

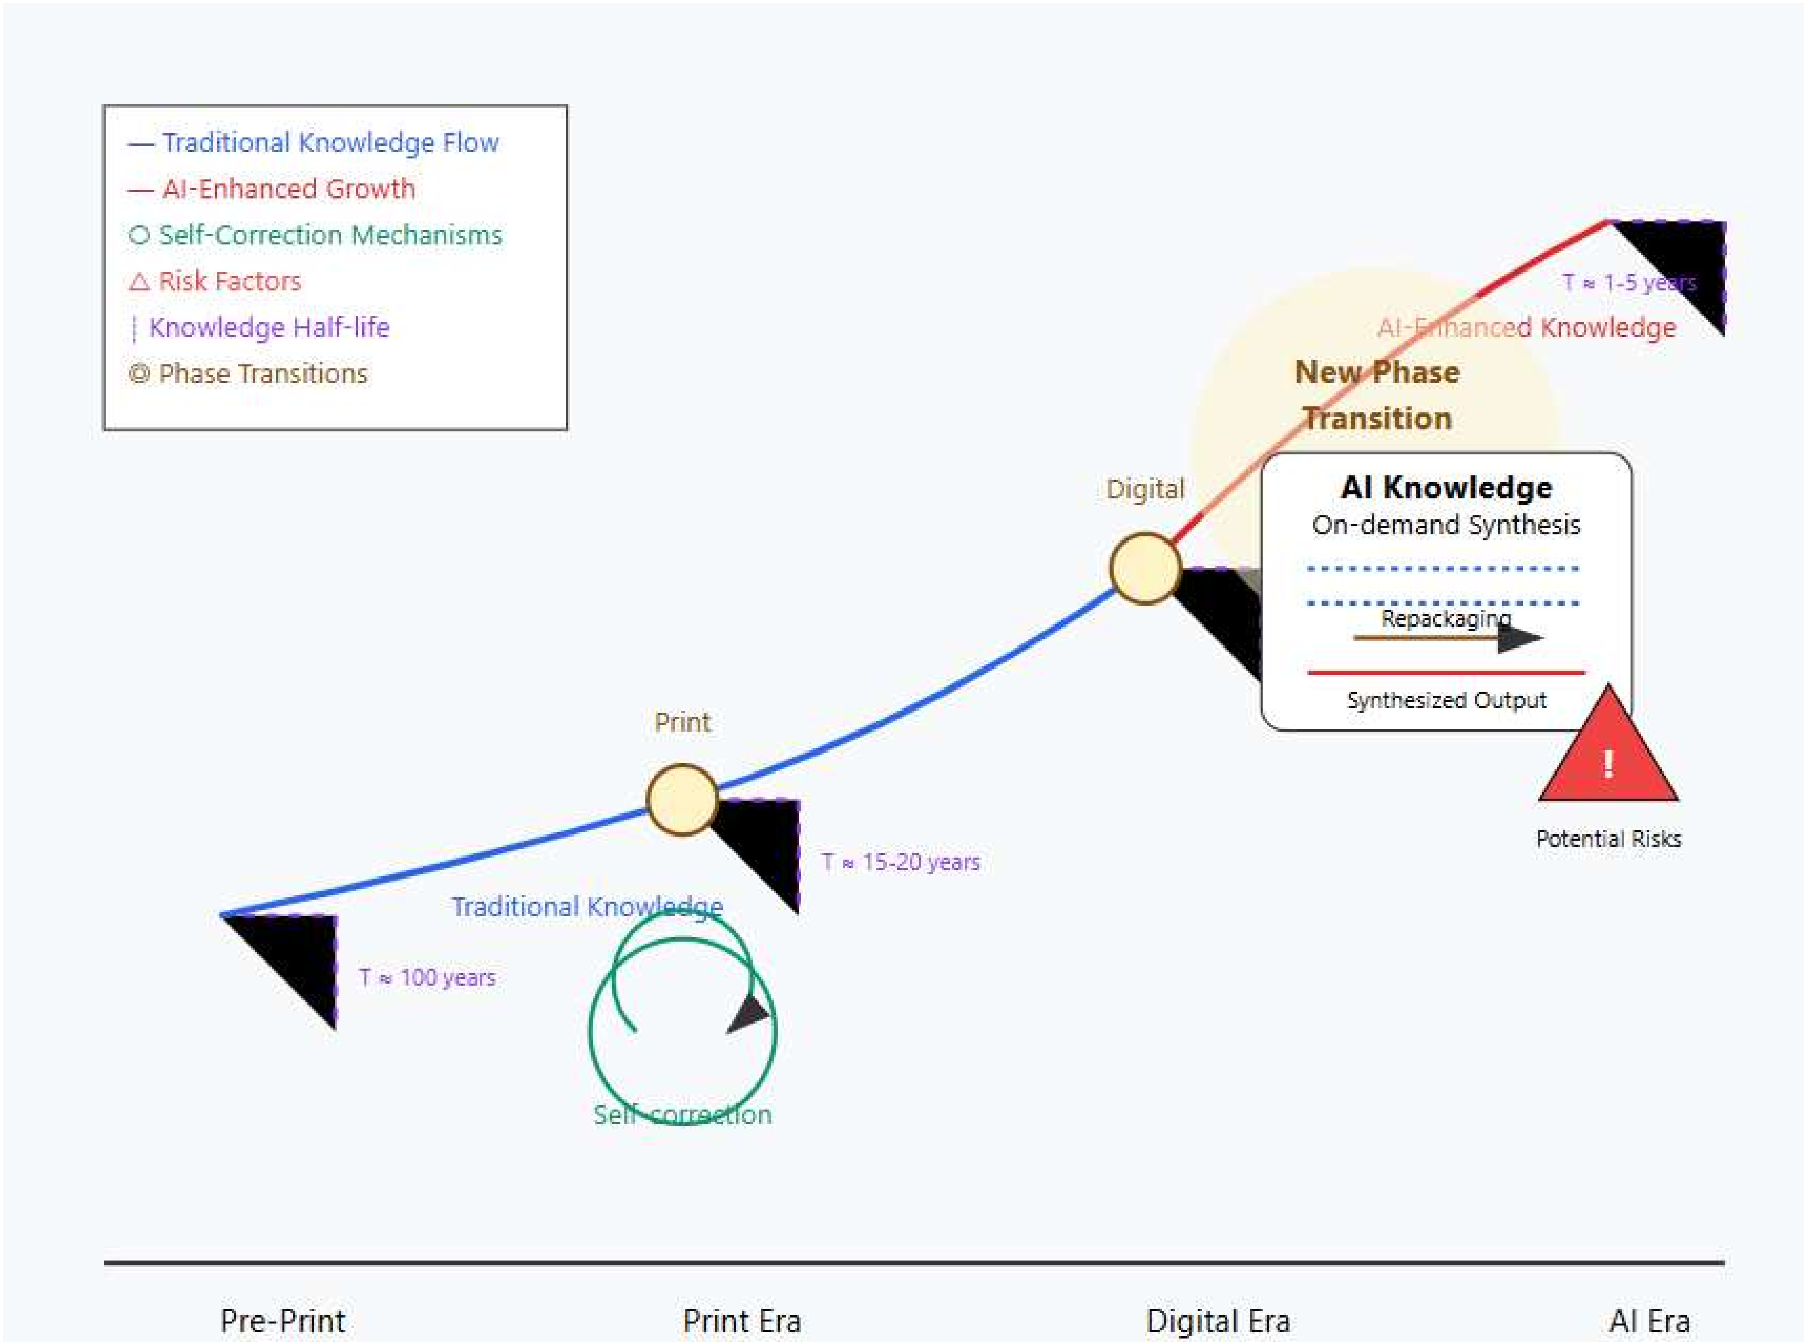

— Traditional Knowledge Flow
— AI-Enhanced Growth
○ Self-Correction Mechanisms
△ Risk Factors
┆ Knowledge Half-life
◎ Phase Transitions
T ≈ 1-5 years
AI-Enhanced Knowledge
New Phase
Transition
Digital
AI Knowledge
On-demand Synthesis
Repackaging
Synthesized Output
!
Potential Risks
Print
T ≈ 15-20 years
Traditional Knowledge
T ≈ 100 years
Self-correction
Pre-Print
Print Era
Digital Era
AI Era

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?

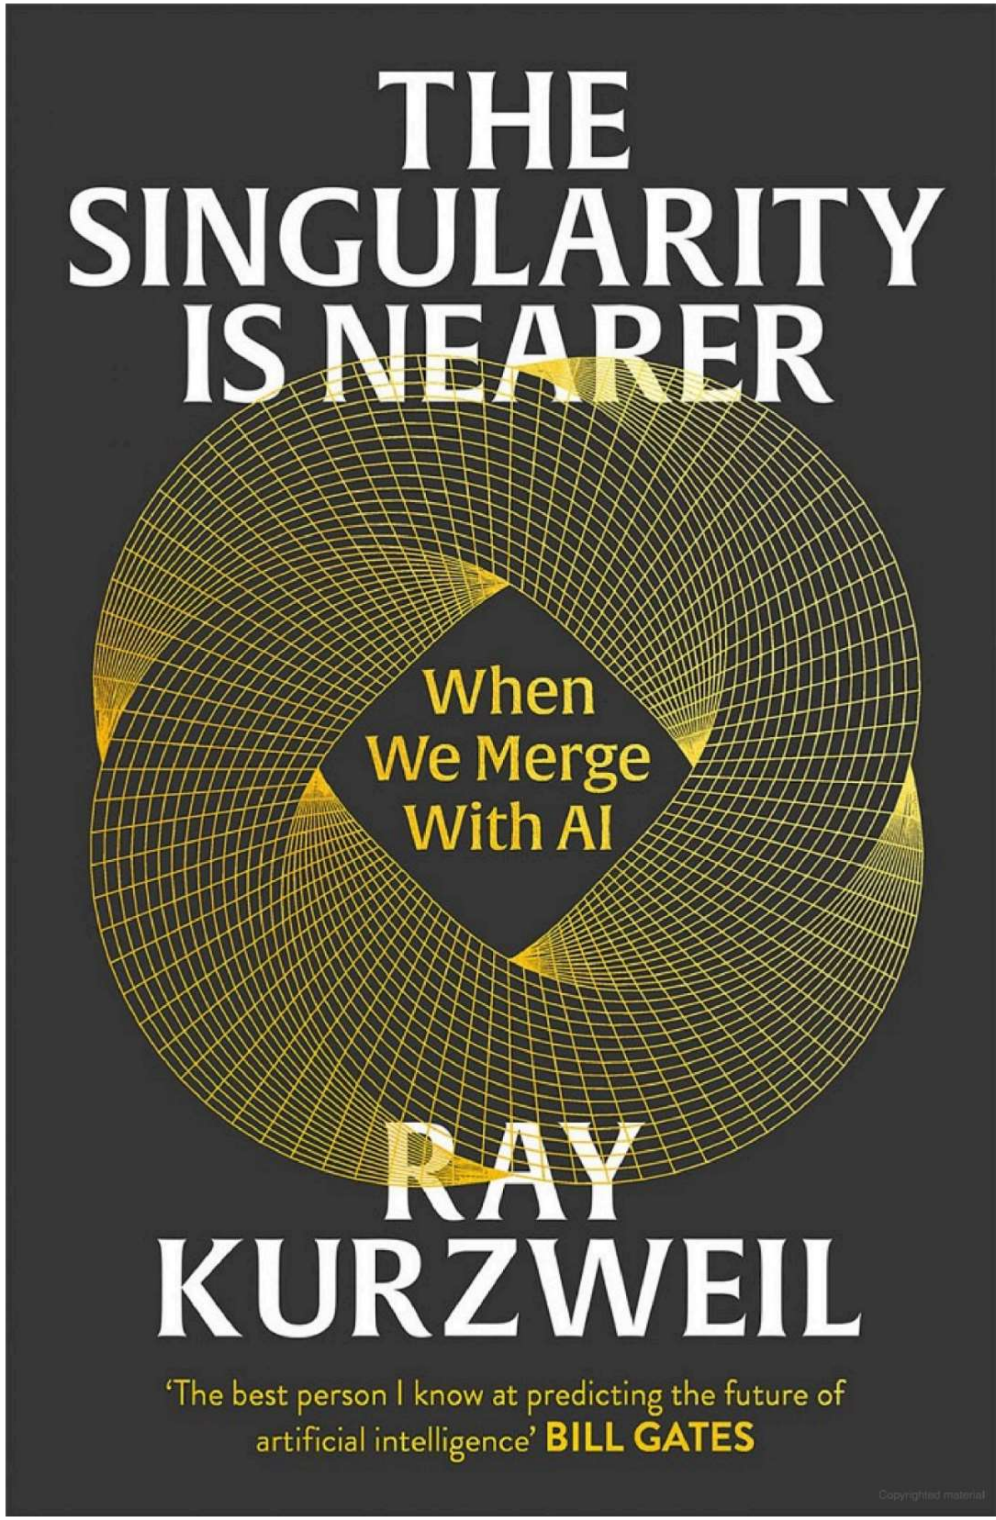

Prediction versus understanding

- We can predict well without understanding (e.g., planet/star motion Newton).
- Guessing the God's many complex behaviours versus knowing his few universal laws.

Warning: Illusion of progress

- AI can learn to mimic human from texts → Chinese room argument (by John Searle)
- Can pass Turing test without any understanding.

Credit: Alan Tan

To sum up

AI is a General-Purpose Technology (GPT)

- Just like electricity

Why AI for Science?

- Automation, scalability, knowledge and data integration
- Assisting in decision making
- Rational in an irrational world of politics.

Can AI fail?

- Yes. We are still learning.
- It is subject to misuse.
- It can be wrongly aligned with human scientist values.

Source: opgal

A church in my home village in Vietnam in style of Buddhist temple

Fusing traditions can be very beautiful