Some Notes on Our Ontologies and AI Track

Gary Berg-Cross, Board Member Ontolog Forum
(Retired Knowledge Engineer)

Intelligence = <u>Learned Knowledge</u> from Data & Experience + <u>Reasoning</u>???

Ideas advancing rapidly

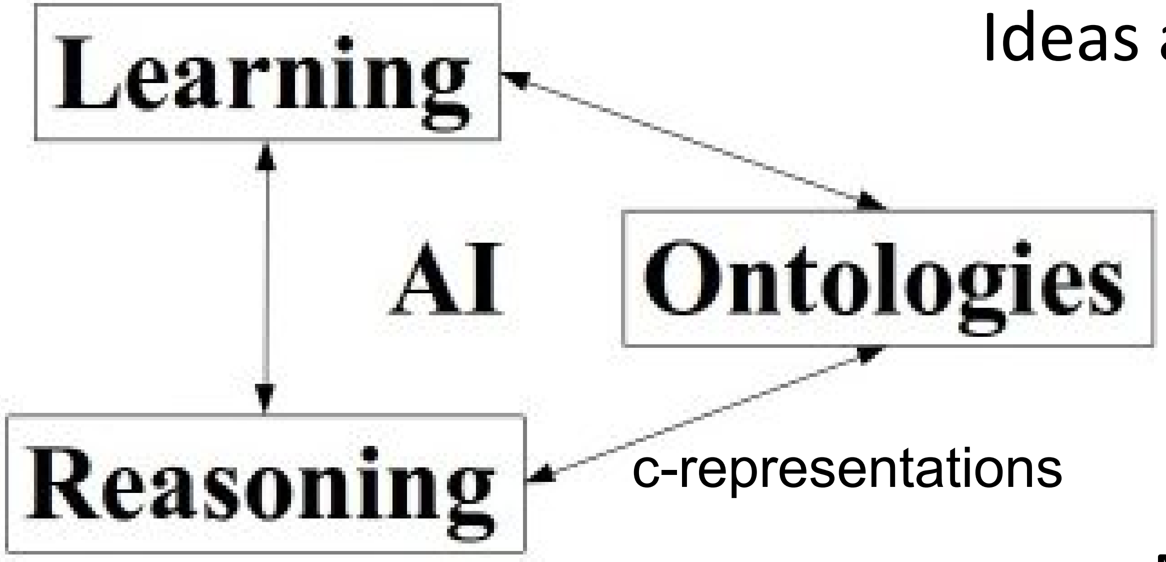

(Knowledge representation Stand in for Semantics...)

Many such connections...

“ ..knowledge representation, which is not a fashionable concept in modern neural net AI.” Ben Goertzel

Talk Outline

- Past: Some retrospectives 2017 -2025 Summits
  - Critiques of LLMs from 2024

    Present (including risk)
  - Integrating GenAI into KE/KG tasks
  - How do GenAIs assist? What limits?
- Many views of the Future
  - Architectures -Kinds of hybrids including how to use ontologies for GenAI Robustness
  - Hybrid robot types as examples
  - How should common-sense knowledge be learned and reasoned about?
- Another Look at Track Mission, Track plan & Speakers
- Is there a role for meaningful explanation? Etc...

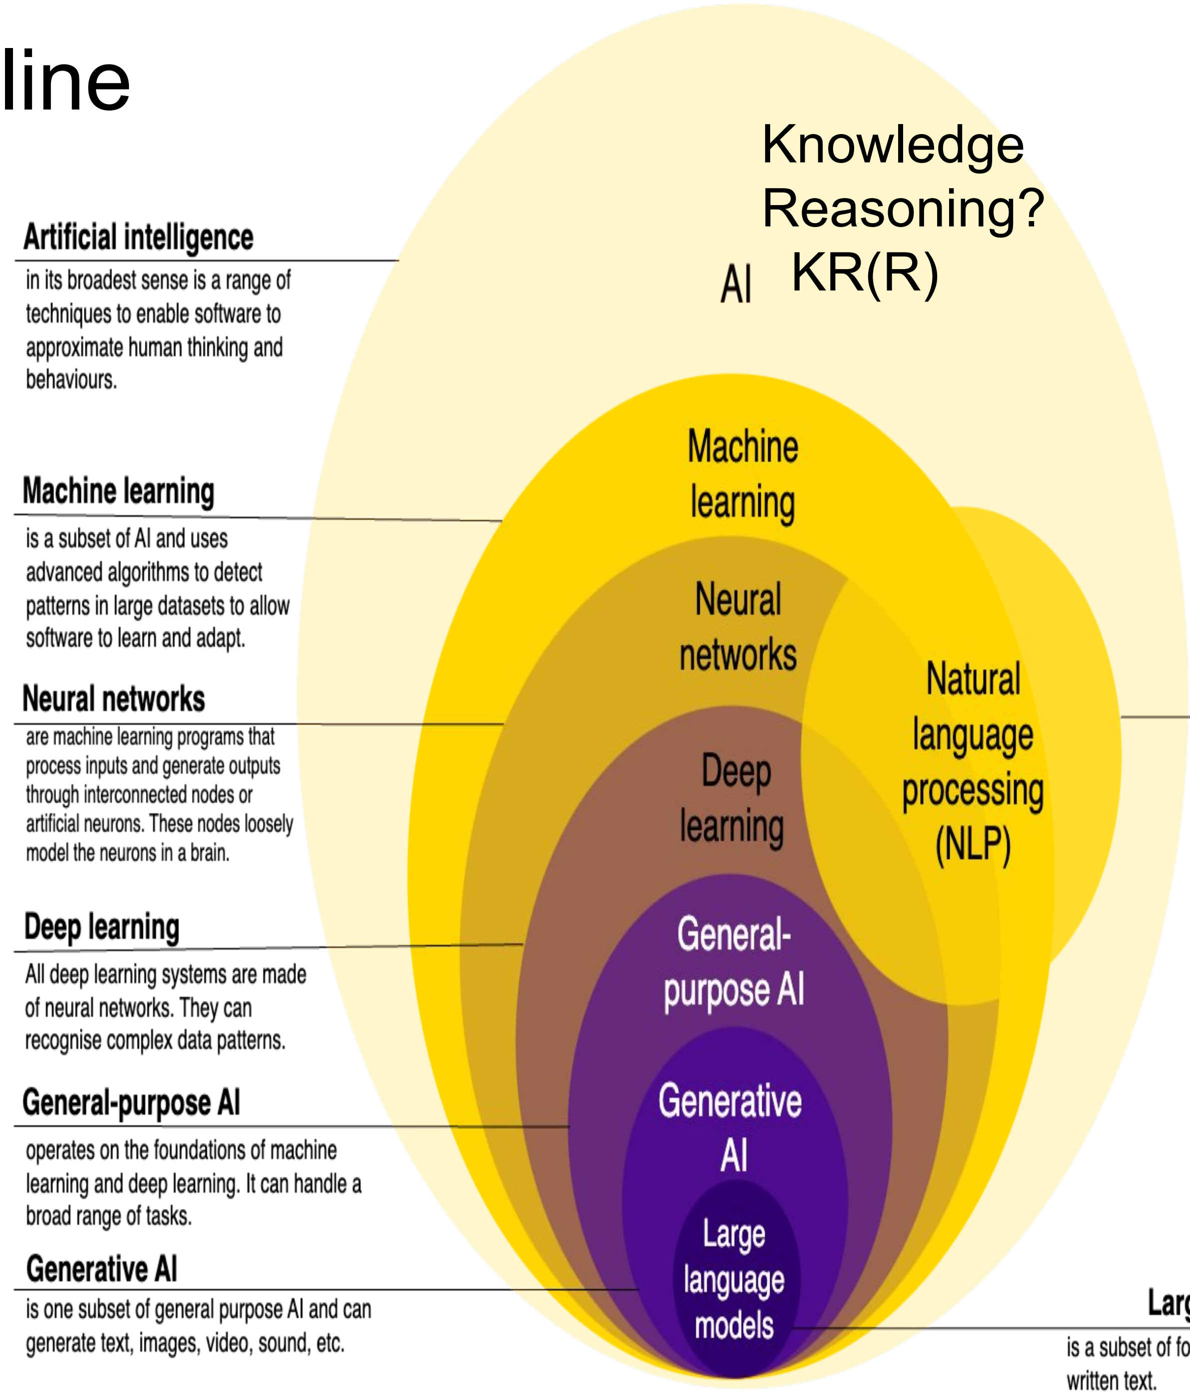

PAST: 9 Years of Ontological Summits

OntologySummit2025 - **"Conceptualization, Analysis and Formalization"** (Foundations)

OntologySummit2024 - **"Neuro-Symbolic Techniques for and with Ontologies and Knowledge Graphs"**(pros & cons & Hybrid views & more cognitive)

OntologySummit2023 - **"Helping scientific researchers make better use of ontologies"**

OntologySummit2022 - **"Dealing with Disasters"**

OntologySummit2021 - **"Ontology Generation and Harmonization" (AI helps)**

OntologySummit2020 - **"Knowledge Graphs" (**the accumulation of large factual knowledge (facts or instances, typically stored as triples) can be usefully structured with the help of ontologies...but this is a light use)

OntologySummit2019 - **"Explanations"** (including the role/issue of commonsense)

OntologySummit2018 - **"Contexts in Context"** (need to take the intentional context of cognitive agents into account – gets cognitive)

OntologySummit2017 - **"AI, Learning, Reasoning, and Ontologies"**

Topics: "Using Automation and ML to Extract Knowledge and Improve Ontologies" Alessandro Oltramari (Research Scientist at Bosch)

"*From machines that learn to machines that know: the role of ontologies in machine intelligence*"

The last few years have had more on the intersection and harmonization that supports AI/ML.

1. The 2021 Ontology Summit examined the overall landscape of ontologies, the many kinds of ontology generation & harmonization, as well as the sustainability of ontologies.

2. Advances in machine learning and the development of very large KGs.

3. These ontologies are commonly developed independently, and as a result, it can be difficult to communicate about and between them

Need agree on how their respective terminologies and formalizations relate to each other via process called "harmonization."

One impediment to harmonization is the relatively poor quality of natural language definitions – misleads axiomitization.

Turns out that GenAI can help with that.

Present: Hybrid approaches (e.g. AlphaGo) demonstrated success in some domains

Scientific discovery: Combining neural pattern recognition with symbolic hypothesis generation in drug discovery -EG. SARA Scientific Autonomous Reasoning agent (Yolanda Gil might talk about such systems).

Robotics: Integrating neural perception with symbolic task planning for complex manipulation tasks (more later).

Natural language processing: Augmenting large language models with symbolic knowledge bases for more reliable reasoning (generalization, extrapolation, abductive reasoning etc.) (More during John Sowa's track?).

Verification and validation: Using symbolic methods to formally verify properties/outputs/states of neural networks. (Perhaps Pascal Hitzler talk)

BUT…..Critiques of LLMs as a direct path of AGI from 2024 Summit

Gary Marcus No AGI (and no Trustworthy AI) w/o Neurosymbolic &

John Sowa w/o Ontology, LLMs are clueless

Foundations Needed (for **AGI**)

1 Rich **cognitive** models that describe mental processes in detail that keep track of dynamic environments

2 Extensive real **world knowledge** > text

3 A representation for relationships between entities e.g. understanding causal relationships and being able to manipulate the world to achieve specific goals

4 Composability Of wholes and parts….

5 Commonsense knowledge developed over time through embodied experience

6 Sophisticated reasoning explicitly using symbols, logic, and rules that provide a symbolic foundation

7 Human values – e.g. important for medical apps

Some Limitations of LLMs (and AGI needs)

1 No fixed set of meanings can adequately describe a continuous, dynamically changing world.

2 Written language is isolated from perception, feelings, actions, and reactions of people in a dynamically changing world.

3 Mental models are more fundamental than language or logic and are needed.

4 Much of human intelligence & underlying mental models are probably lost in a mapping to LLMs?

5 A linear language or notation is not ideal for thinking or communicating complex spatial patterns.

Happening Now – A shift to world models

A discernible shift is occurring within AI research, moving from generative models for language and images toward the development of world models with spatial intelligence.

World models help simulate spatial relations and reason about the environment before interacting with it, essential for applications like robotics and self-driving cars that operate in real world space.

- Google is testing its Genie models:
    - Given a text prompt, Genie 3 can generate dynamic worlds that you can navigate in real time at 24 frames per second, retaining consistency for a few minutes at a resolution of 720p.
- Nvidia is developing its Omniverse and Cosmos platforms for physical AI.

One View of an AI future Dario Amodei, co-founder and CEO of Anthropic, a safety-focused CEO who identified five major categories of AI risk.

Imagine it's 2027 or 28.
A new country appears overnight.

50 million citizens, every one smarter than any Nobel Prize winner who has ever lived.
They think 10 to 100 times faster than any human.

They never sleep.
They can use the internet, control robots, direct experiments, and operate anything with a digital interface.

What would a national security advisor say? ….. the answer is obvious:

"the single most serious national security threat we've faced in a century, possibly ever."

Explanation, Semantics, and Ontology Giancarlo Guizzardi (2025)

Take Away Messages

1. All science is data science
2. **Interoperability** is fundamental for science
3. No interoperability without **understanding**
4. "No understanding without **explanation**" (Strevens)
5. "No explanation without **semantics**" (Brown and Swift)
6. No semantics without **ontology** (Quine)
7. "No ontology without **Ontology**" (Varzi)

There are still many ingredients to a quality KG

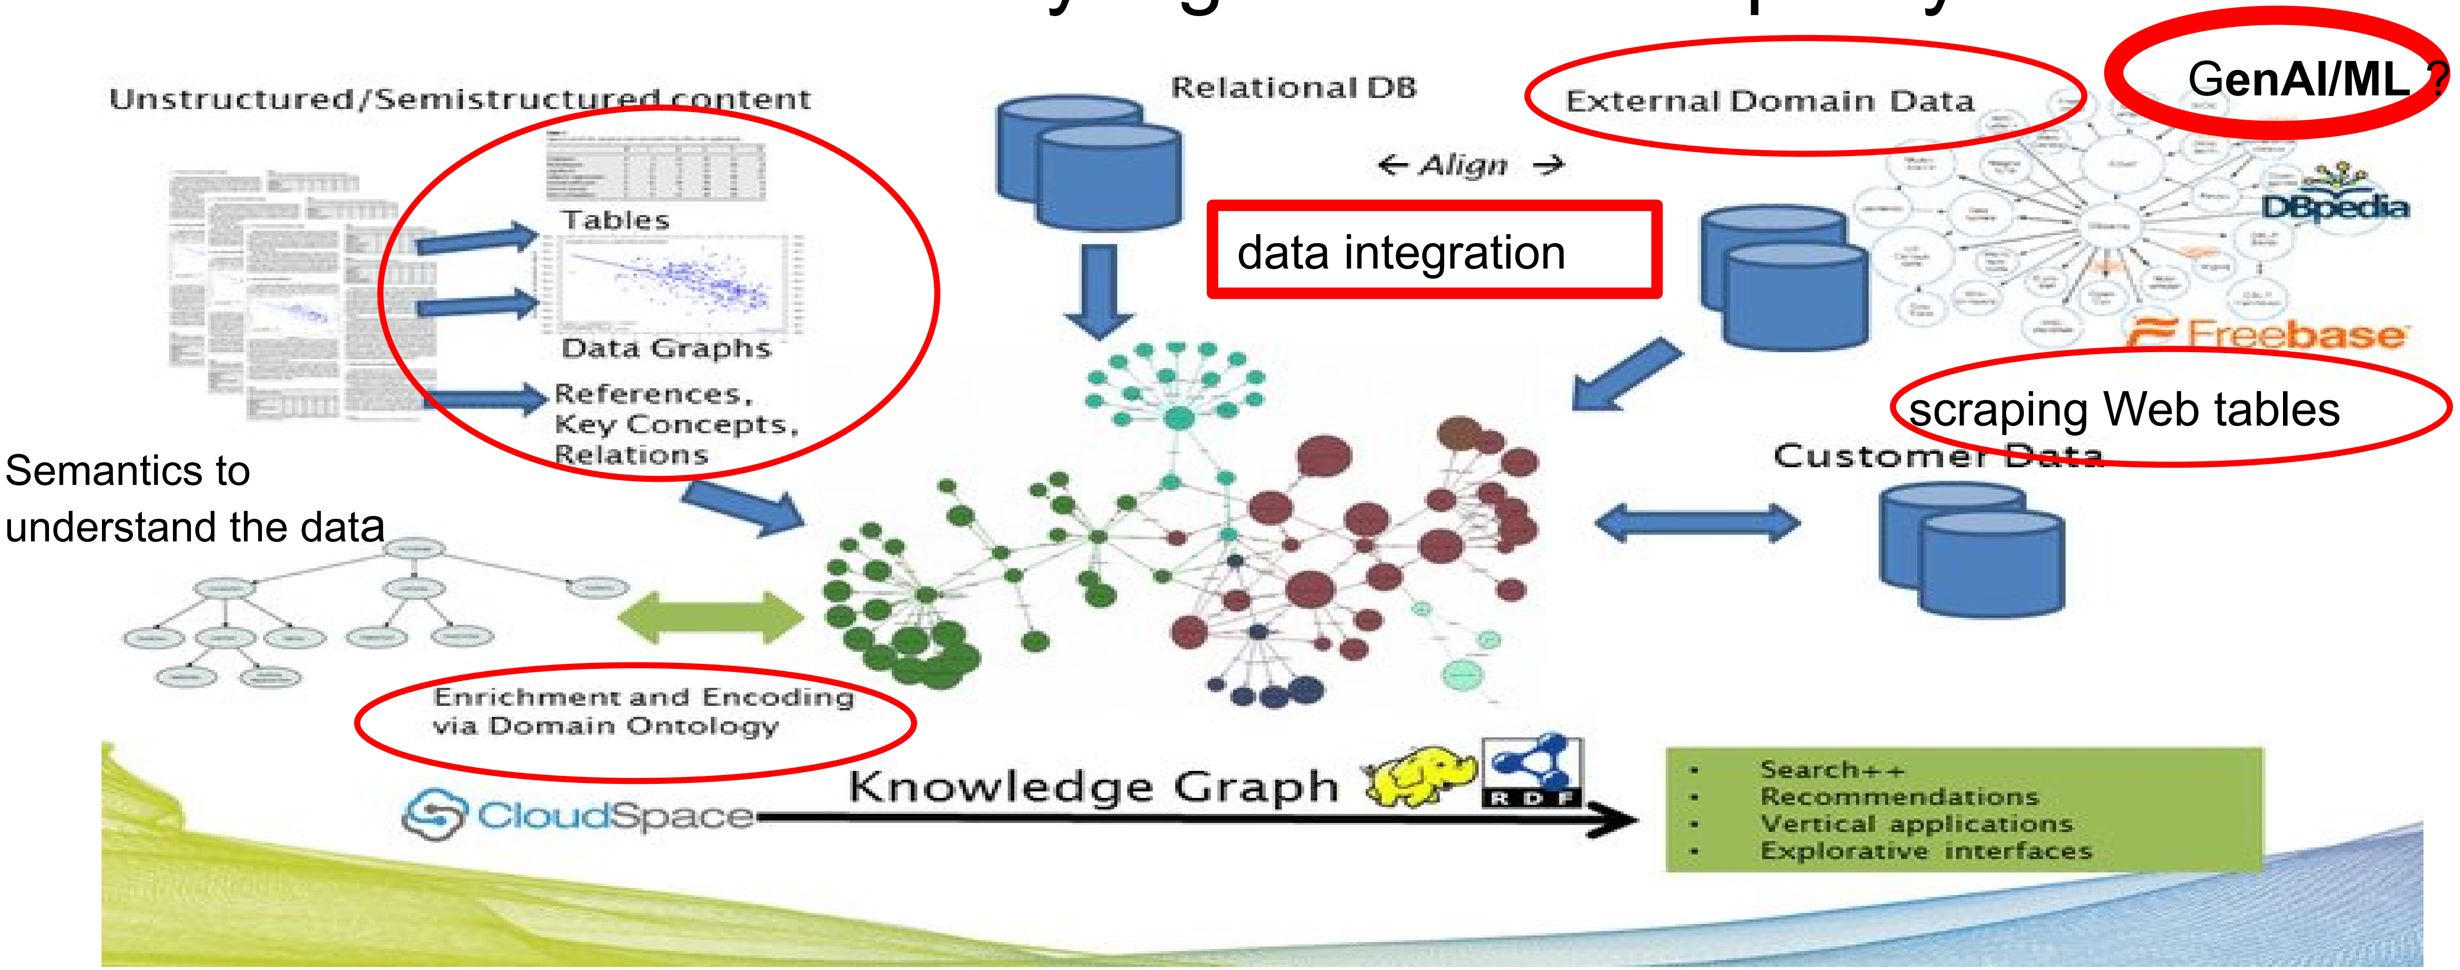

After https://www.researchgate.net/post/What_is_Knowledge_Graphs Ajit kumar Roy’s What is Knowledge Graphs?

Efforts for communities related to GenAI, Knowledge Graphs, and NLP join their forces in order to develop more effective algorithms and applications.

There are several ways of “Using” GenAIs to Assist Phases of the Knowledge Engineering Process.

- GenAIs leveraging LLM models can help add text for KGs & Ontologies.
- It can identify & define major concepts, entities, & relationships within the domain & represent them in a logical language.
- But w/o help they have problems organizing the concepts/relationships into a meaningful structure.
- Needs help with expert review & iterative refinement.

Early work showed that while many entities are extracted, many are isolated and among connected ones the number of unique word based relations is far too large for a GenAI to produce a typical ontology (Trajanoska et al, 2023) Without guidance on formal relations a KG suffers.

“Using” GenAIs CAN help Enhance Aspects of Ontological Engineering too (or Symbolic AI will be advanced by GenAI)

Findings

(1) ChatGPT-4 achieves intermediate-to-expert scores on an ontology modeling qualification test;

(2) the model performs ontology restriction verification with accuracy of 92.22%;

(3) combining model answers on the same ontology axiom represented in different formalisms improves the accuracy to 96.67% (from 92.22%); and

(4) higher accuracy is observed in identifying defects related to the incompleteness of ontology axioms compared to errors due to restrictions misuse.

From Tsaneva, Stefani, Stefan Vasic, and Marta Sabou. "Llm-driven ontology evaluation: Verifying ontology restrictions with chatgpt."The semantic web: ESWC satellite events 2024.

Some **Best Practices** have been suggested (Trajanoska et al, 2023*) (Remember that semantic spectrum?)

1) Triples should be concise.
2) Contextual information of entities should be captured.
3) Knowledge graph should not contain redundant triples.
4) Knowledge graph can be updated dynamically
5) Entities should be <u>densely</u> connected.
6) Relations among different types of entities should be included.....

9) Synonyms should be mapped, and ambiguities should be eliminated to ensure reconcilable expressions.

10) Knowledge graph should be organized in <u>structured</u> triples for easy machine processing...

12) The attributes of the entities should be included.
13) Knowledge graph should be publicly available ...

15) Knowledge graph should be concentrated.
16) The triples should not contradict each other.

* Trajanoska, Milena, Riste Stojanov, and Dimitar Trajanov. "Enhancing knowledge graph construction using large language models." arXiv preprint arXiv:2305.04676 (2023).

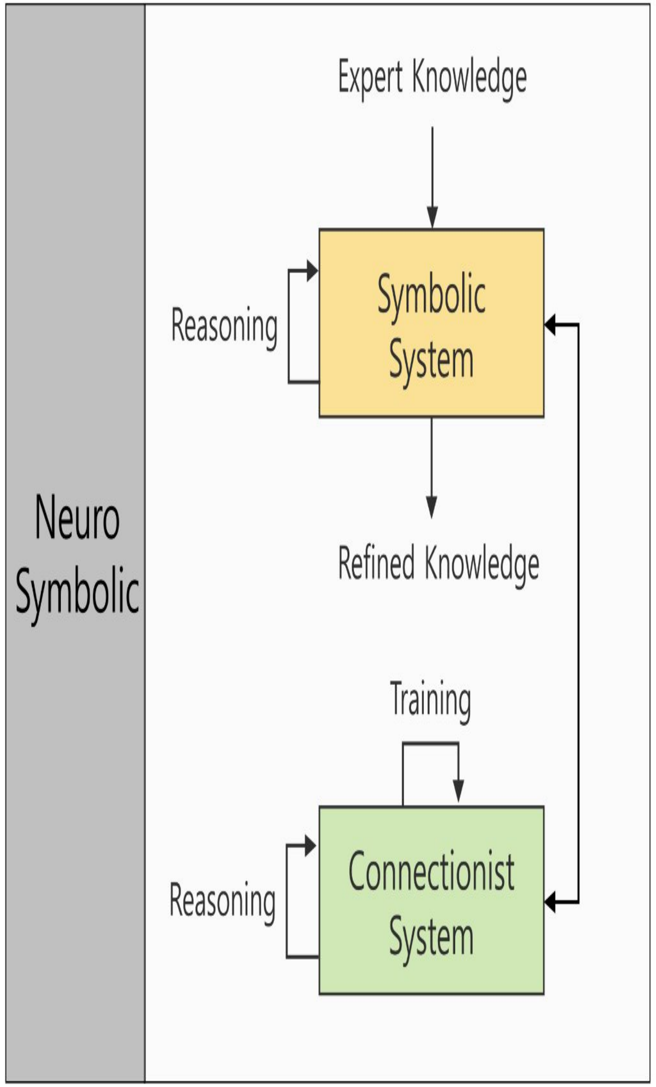

Plugin Framework-Based Neuro-Symbolic Grounded Task Planning for Multi-Agent System by Jiyoun Moon

What About Rich Hybrid Cognitive Architectures?

- Robust AI system will include a sound reasoning layer in combination with deep learning.
- Many proposed types to operate (Kautz).
- The integration of these models remains a challenge but…
  - Hybrid should be useful for addressing complex AI problems that cannot be solved by purely symbolic or neural means so there are <u>benefits on integrating</u> the accomplishments of both paradigms.(Inter-operation)
- Coupling may be through different methods, including:
  - the calling of deep learning systems <u>within</u> a symbolic algorithm, or
  - the <u>acquisition of symbolic rules during training</u>.

Combined strengths of (Various) Hybrids:

Neurosymbolic (**NeSy**) taxonomy of Hybrid Systems by Henry Kautz

1. Symbolic Neuro symbolic
2. Symbolic [Neuro]
3. Neuro | Symbolic   Intertwined?
4. Neuro:Symbolic → Neuro
5. Neuro_{Symbolic}
6. Neuro [Symbolic]

1. Symbolic Neural symbolic is the current approach of many neural models in natural language processing,
 2. Symbolic[Neural] is exemplified by AlphaGo, where symbolic techniques are used to <u>invoke</u> neural techniques (sub-routine).

<u>5. NeuralSymbolic uses a neural net that is generated from symbolic rules. EG the Neural Theorem Prover, which constructs a neural network from an AND-OR proof tree generated from KB rules/ terms. Logic Tensor Networks also fall into this category.</u>

Neural networks handle pattern identification and <u>learning</u>, while integrated symbolic reasoning might ensure organized decision-making and explicability – proper symbolic reasoning can facilitate post-hoc <u>explanations</u> which are essential in sensitive domains such as healthcare and legal reasoning. They might drive adoption.(See also Van Bekkkum et al)

The Case for Some Sort of "Hybrid" "Adding" Ontologies can help GenAI Robustness

Neuro-Symbolic AI hybrids bridges the gap between isolated neuro-inspired and symbolic systems and may overcome the inherent limitations of conventional AI architectures by:

- Integrating neural networks and symbolic representations: It allows symbolic knowledge to support structured reasoning & guide the learning process of neural networks, making them more interpretable (understandable??) and efficient.
- This approach inherits the efficient, data-driven learning power of neural networks and the reasoning capabilities of symbolic AI, aiming to create more robust, somewhat cognitive and flexible AI systems

But as Gary Marcus argued:
"We cannot construct rich cognitive models in an adequate, automated way without the triumvirate of hybrid architecture, rich prior knowledge, and sophisticated techniques for reasoning.

An Example: Robotic learning in a hybrid environment: it might mix deliberative and reactive behavior taking human and agent intentions and goals to account (Luc Steel)

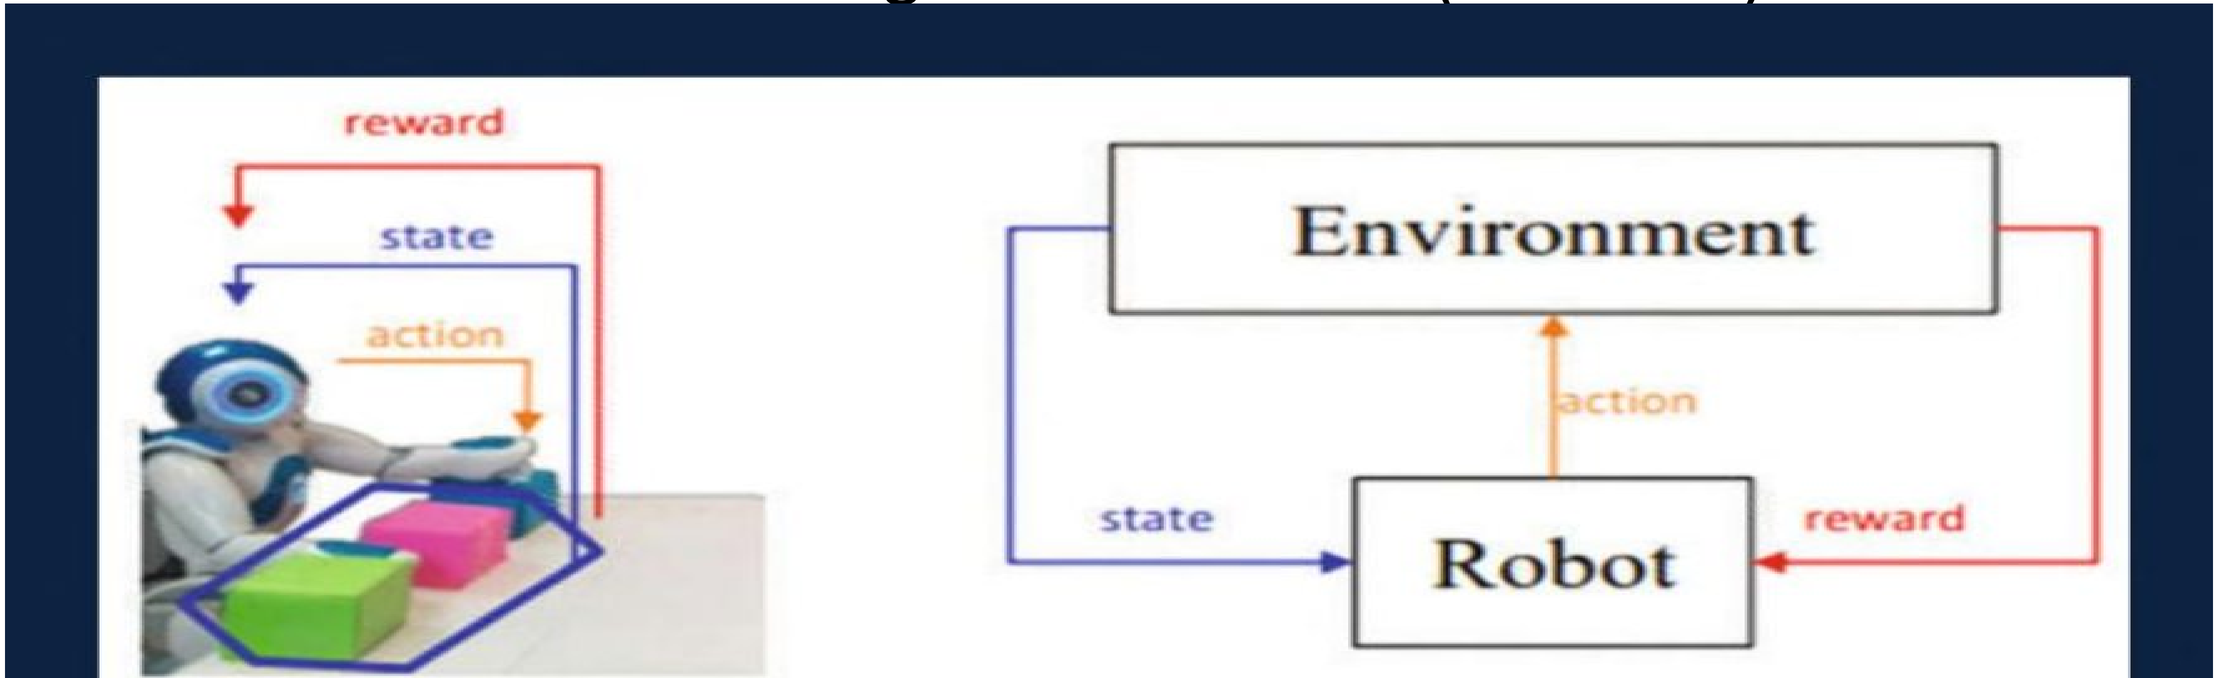

Theorizing, mental models….

From *Human-Centered Artificial Intelligence: Advanced Lectures* 13500 (2023)
See also Gary Berg-Cross *Applying developmental-inspired principles to the field of developmental robotics, PerMIS '08: Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems*

Types of Neuro-Symbolic Robotic Hybrids (Neuro-Symbolic Robotics: Capabilities, Limitations, and Challenges)

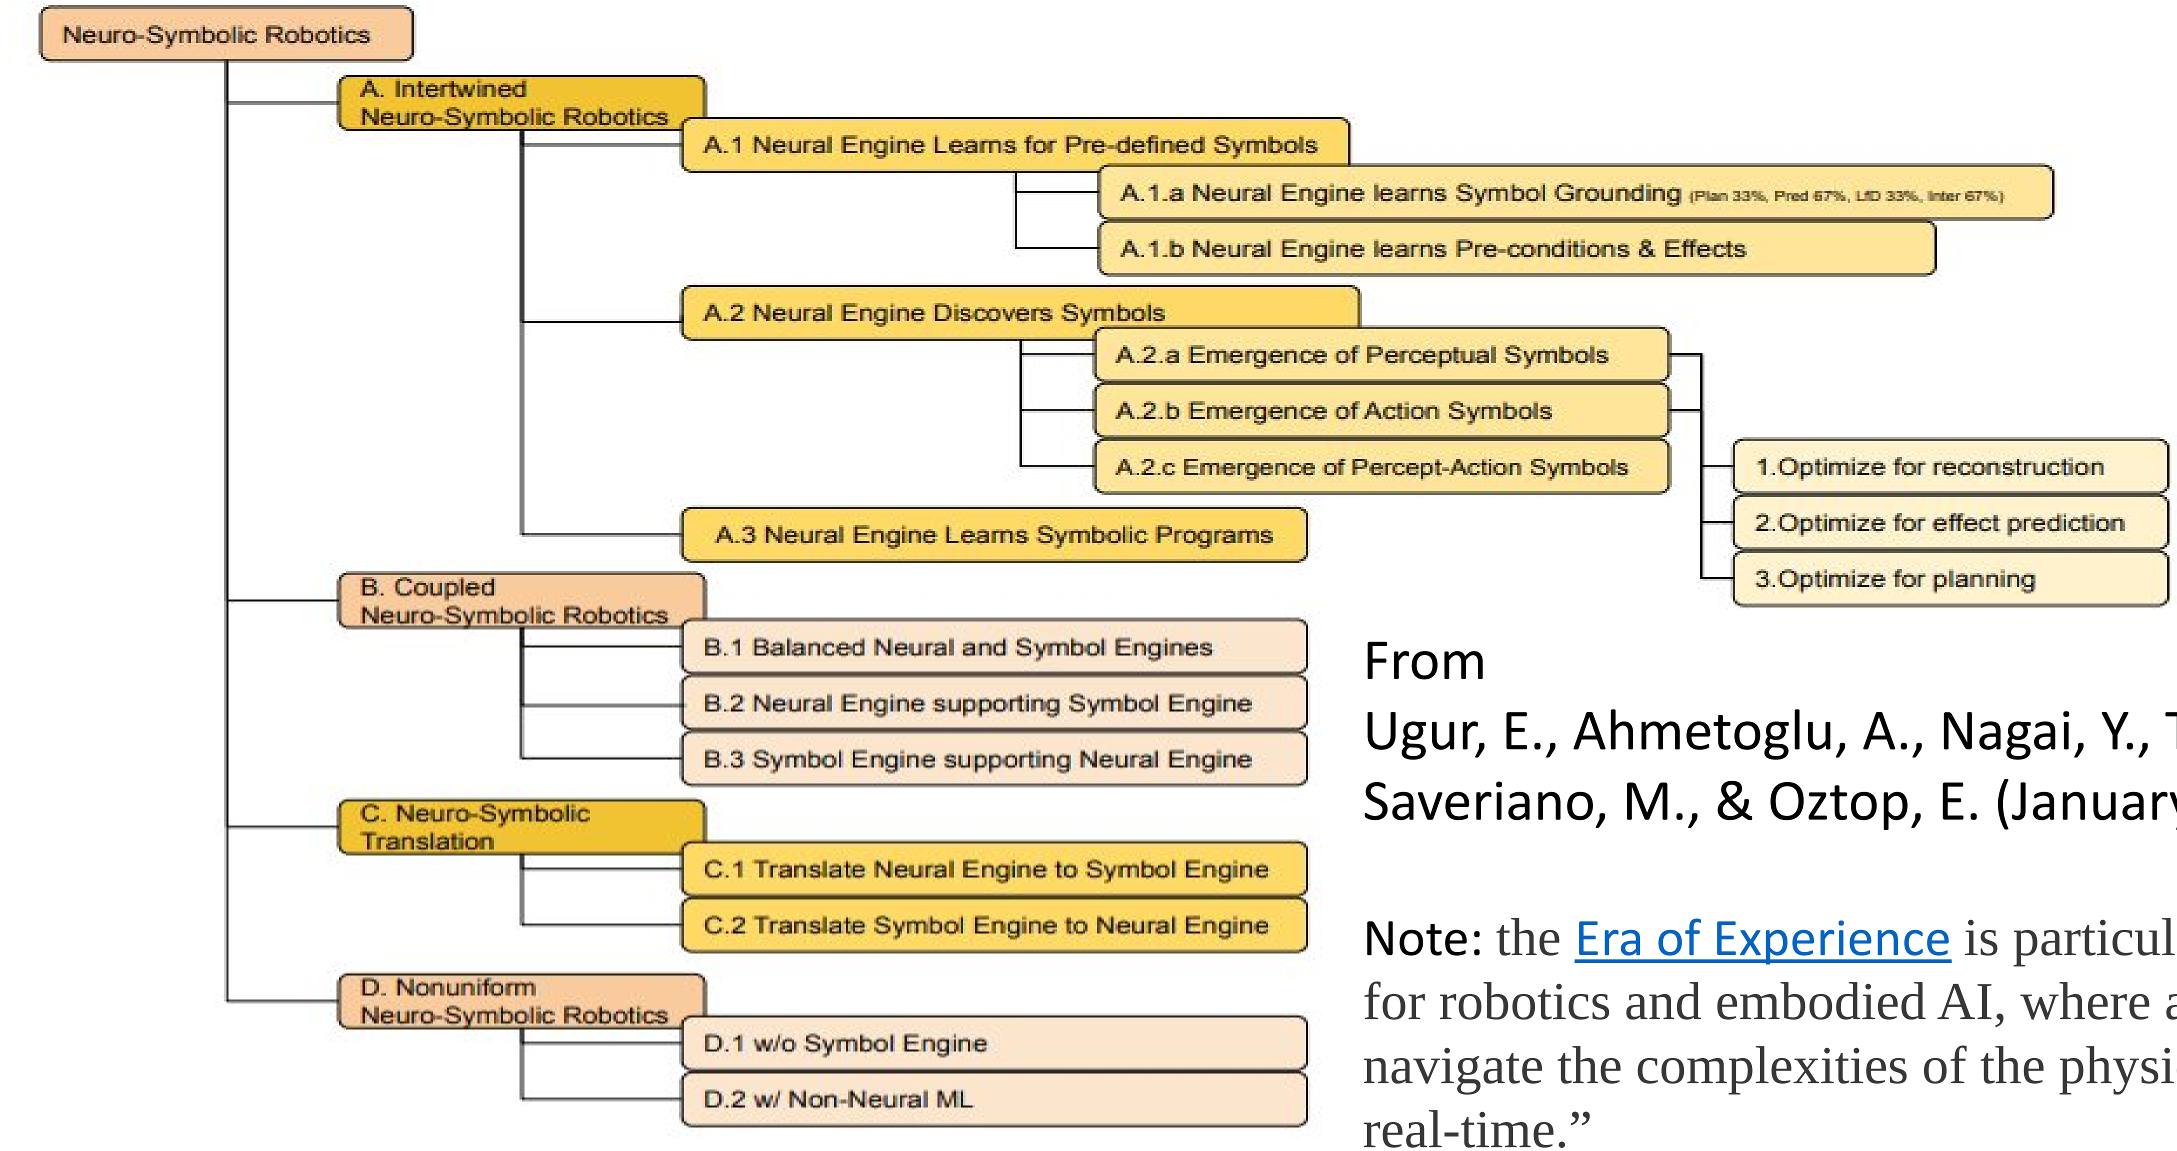

Fig. 2. A taxonomy of Neuro-Symbolic Robotics

From
Ugur, E., Ahmetoglu, A., Nagai, Y., Taniguchi, T., Saveriano, M., & Oztop, E. (January 2025)

Note: the Era of Experience is particularly “relevant for robotics and embodied AI, where agents must navigate the complexities of the physical world in real-time.”

Next Talks Planned for the AI Track and Cover: Yesterday, Today and Tomorrow

1 Randy Goebel 3/11 (Alberta Machine Intelligence Institute) interested in the logic of machines, Explainable AI, and Debugging Foundational Models as an eventual path to AGI via a stackable neurosymbolic framework.

2 Pascal Hitzler 3/18 University Distinguished Professor, Lloyd T. Smith Chair, Kansas State U. "Ontologies in the Era of Large Language Models – a perspective"

3 Manas Gaur March 25, Neurosymbolic AI for High Stakes Applications in the LLM Era

* Note, we hope that some of this discussion will help prepare for later tracks such as on foundational issues and proper ways of thinking about the future development of knowledgeable/intelligent systems that blend many capacities and representations.

Randy Goebel on a Scientific Basis for Debugging Foundational Models

(Homogenization provides powerful leverage but demands caution)

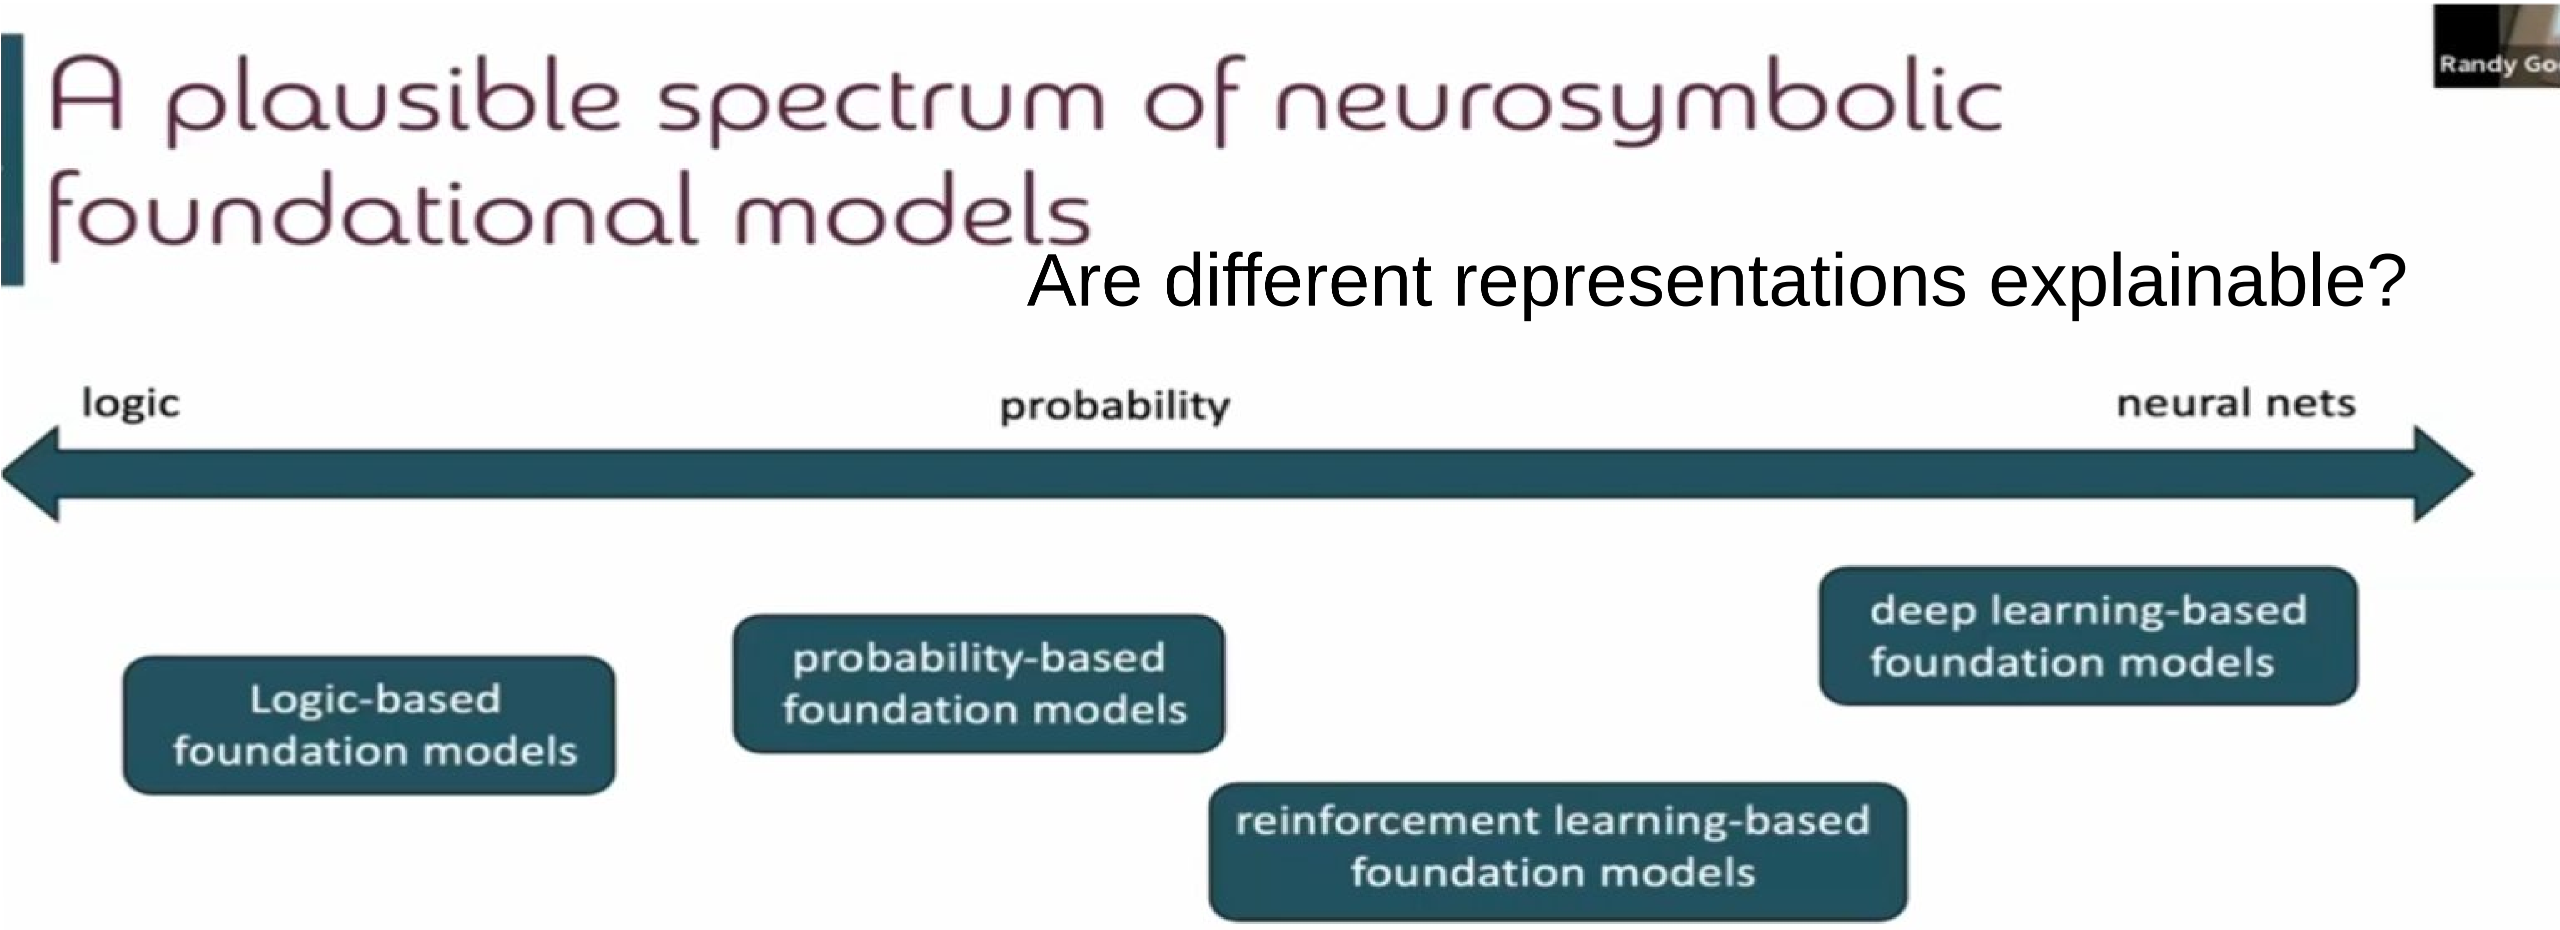

Are different representations explainable?

Can they be blended to cooperate dynamically choosing the best representation for processing?

How important are meaningful explanations?....

Backup Slides

The New Field of Neuro-Symbolic AI

Neuro-Symbolic AI bridges the gap between neuro-inspired and symbolic systems by:

Integrating neural networks and symbolic representations: It allows symbolic knowledge to support structured reasoning & <u>guide</u> the learning process of neural networks, making them more interpretable and efficient.

Combines strengths: This approach inherits the learning power of neural networks and the reasoning capabilities of symbolic AI, aiming to create more robust and flexible AI systems.

Sources:

Hitzler, Pascal, et al. "Neuro-symbolic approaches in artificial intelligence." National Science Review 9.6 (2022): nwac035.

Bhuyan, Bikram Pratim, et al. "Neuro-symbolic artificial intelligence: a survey." Neural Computing and Applications 36.21 (2024): 12809-12844.

Neurosymbolic AI, approach aiming to combine the strengths of Neural Networks and Symbolic AI

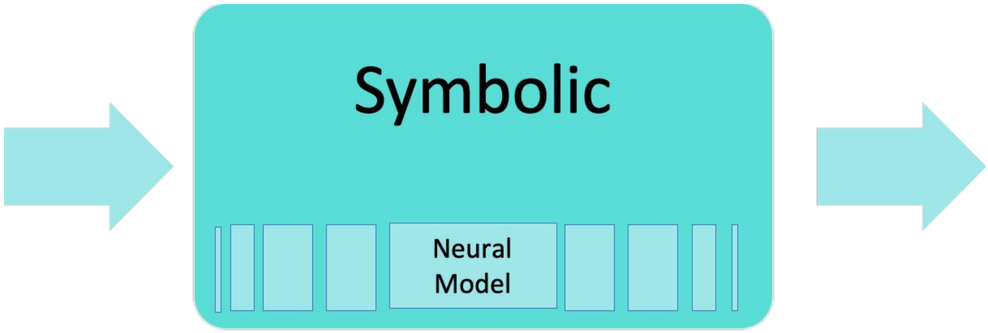

Symbolic[Neuro] Inverse is
Graph Neural Network

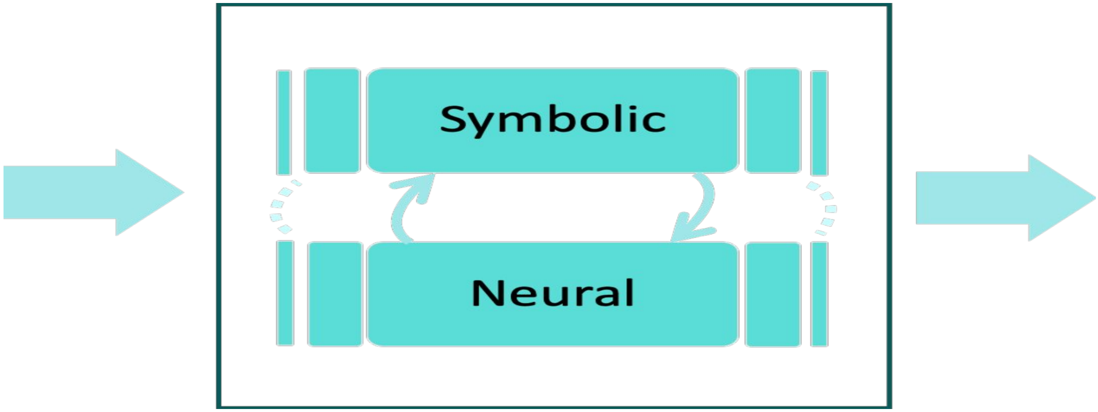

Neuro;Symbolic communicate

One example - RoboTurk combines symbolic reasoning and reinforcement learning to enable robots to manipulate objects in complex environments.
The symbolic component helps the robot understand task goals and constraints, while the neural network learns movement strategies through trial and error.

From "Types of Neuro-Symbolic Systems" by Harsha Kokel but based on Henry Kautz
https://harshakokel.com/posts/neurosymbolic-systems/

taxonomy of Hybrid Systems by Henry Kautz

1. Symbolic Neural symbolic is the current approach of many neural models in natural language processing, where words or subword tokens are the ultimate input and output of large language models.
Examples include BERT, RoBERTa, and GPT-4.

2. Symbolic[Neural] is exemplified by AlphaGo, where symbolic techniques are used to <u>invoke</u> neural techniques. In this case, the symbolic approach is Monte Carlo tree search and the neural techniques learn how to evaluate game positions.

3. Neural | Symbolic uses a neural architecture to <u>interpret perceptual data as symbols</u> and relationships that are reasoned about symbolically. Neural-Concept Learner is an example.

4. Neural: Symbolic → Neural relies on symbolic reasoning to generate or label training data that is subsequently learned by a deep learning model, e.g., to train a neural model for symbolic computation by using a Macsyma-like symbolic mathematics system to create or label examples.

5. NeuralSymbolic uses <u>a neural net that is generated from symbolic rules.</u> An example is the Neural Theorem Prover, which constructs a neural network from an AND-OR proof tree generated from knowledge base rules and terms. Logic Tensor Networks also fall into this category.

Neural[Symbolic] according to Kautz, this approach <u>embeds true symbolic reasoning inside a neural network</u>. These are tightly-coupled neural-symbolic systems, in which the logical inference rules are internal to the neural network. This way, the neural network internally computes the inference from the premises and learns to reason based on logical inference systems. Early work on connectionist modal and temporal logics by Garcez, Lamb, and Gabbay is aligned with this approach.

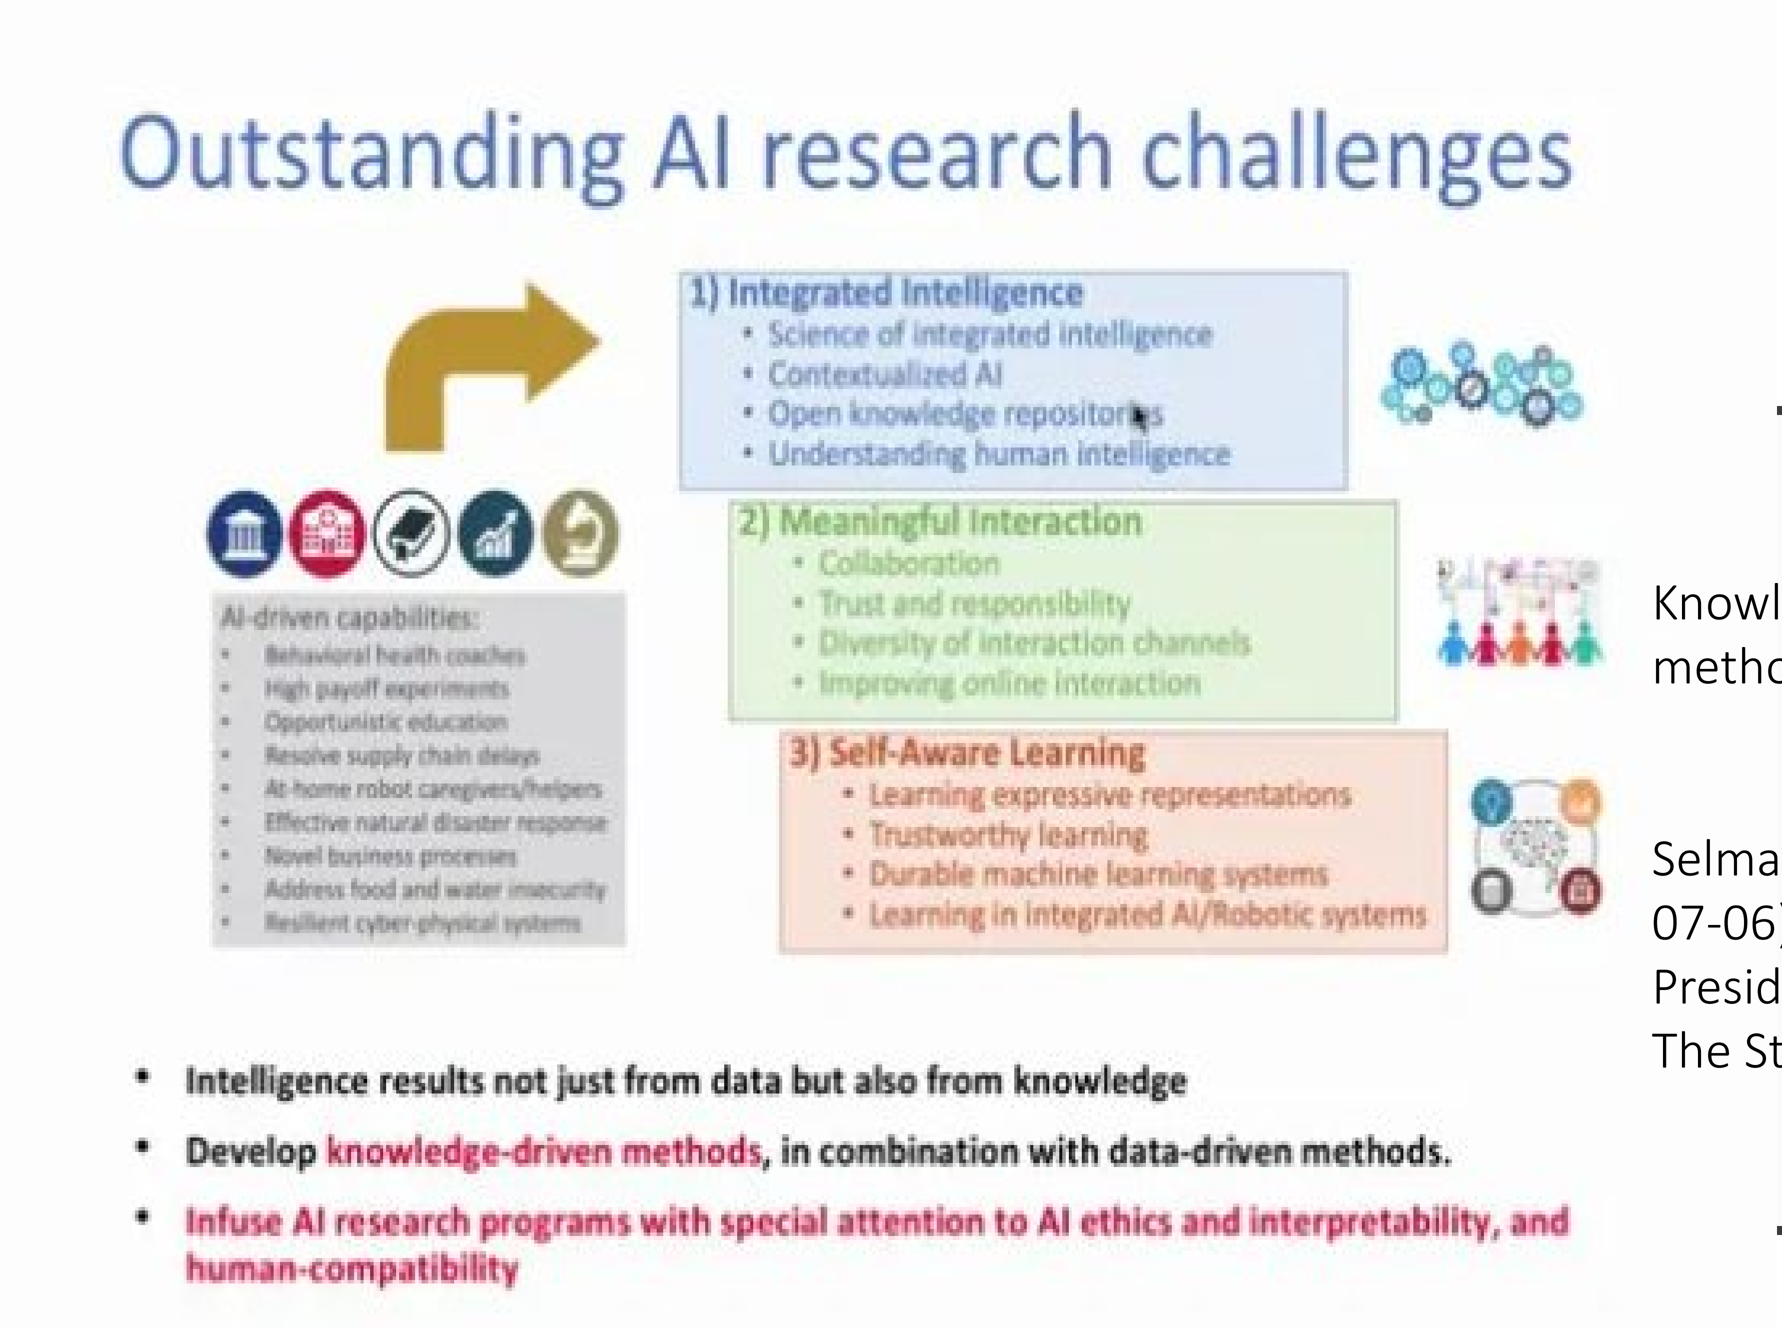

Knowledge driven methods

Selman, Bart (2022-07-06). "AAAI2022: Presidential Address: The State of AI

“Using” GenAIs can help enhance KGs….but

GenAIs using LLMs, have sparked interest in their application to knowledge engineering (KE) tasks, a topic explored in our 2024 Summit.

Knowledge Graph construction & structuring is complex using in part, raw texts & GenAIs leveraging LLM models can help add material, as discussed at our 2020 Summit. Richer forms of knowledge may also be involved.

Early work however, showed that while many entities are extracted many are isolated and among connected ones the number of unique word based relations is far too large for a GenAI to produce a typical ontology (Trajanoska et al, 2023)
GenAIs may also help leverage knowledge higher up on the semantic spectrum. But there are warnings:

“**Integrating generative AI into KE tasks needs to be done with awareness of potential risks and harms.” from** “Knowledge Prompting: How Knowledge Engineers Use Generative AI” Elisavet Koutsiana et al, 2026, Journal of Web Semantics
Without guidance on formal relations a KG suffers.

Selective NeSy Sources & Issues

Sources:

Hitzler, Pascal, et al. "Neuro-symbolic approaches in artificial intelligence." National Science Review 9.6 (2022): nwac035.

Bhuyan, Bikram Pratim, et al. "Neuro-symbolic artificial intelligence: a survey." Neural Computing and Applications 36.21 (2024): 12809-12844.

Some Issues:

What's the best representation, learning, reasoning, architecture & integration process for a hybrid system?

Are there many competing versions of these perhaps by task or domain?