

Less Data, More Knowledge: Semantic Communication and Emergent Protocols for Reasoning-Driven Intelligent Systems*

Walid Saad

Department of Electrical and Computer Engineering,
Network intelligEnce, **W**ireless, and **S**ecurity (**NEWS**) Group
Virginia Tech

Email: walids@vt.edu

Group: <http://www.netsciwis.com>

* And a (title -non-conforming 😊) glimpse on how they lead to world models, twins, and physical AI

Semantic Communications



Encoding



Decoding



Receiver

Transmitter

> Stud Hist Philos Sci. 2019 Feb;73:34-43. doi: 10.1016/j.shpsa.2018.06.003. Epub 2018 Jun 19.

No communication without manipulation: A causal-deflationary view of information

Cristian Ariel López ¹, Olimpia Iris Lombardi ²

Affiliations + expand
PMID: 30914122 DOI: 10.1016/j.shpsa.2018.06.003

Analogy to human communication

Abstract Conveying “semantics” or meaning as per Weaver can help us “do better” in communicational contexts, namely, a causal-deflationary one. Our approach draws from Timpon's deflationary view and supplies the field of philosophy of information with new tools that will help to clarify the underlying structure of communication. But what does “semantic” even mean for a communication network? What is the semantic of semantics?

In this paper, we shall discuss how we can understand communication in communicational contexts, namely, a causal-deflationary one. Our approach draws from Timpon's deflationary view and supplies the field of philosophy of information with new tools that will help to clarify the underlying structure of communication. But what does “semantic” even mean for a communication network? What is the semantic of semantics?

In light of our account, communication is not merely the existence of statistical correlations between source and receiver, as usually understood from a purely formal view. Instead, communication is an asymmetric phenomenon involving causal notions: the destination system must be able to be causally manipulated by intervening on the source for successful communication. In a nutshell, we shall support the following lemma: no communication without manipulation.

Keywords: Causation; Communication; Information; Manipulation.



Teacher

Minimally representing the meaning



Apprentice

rationally manipulating the data

What is (is not) Sematic Communication Systems?

Is

Is not

Path to knowledge - driven AI -nativness



By attributing meaning and context (via a representation) to the latent bit-pipeline.

Reasoning -based system



Leverages causality and associational relations in the data to learn a representation and communicate it.

Symmetric Communications



In contrast to a passive receiver that merely reconstructs the conveyed message, an apprentice must be able to generate content from a representation.

Semantic Communication Systems



Data compression (source coding)

Unlike data compression, SC leverages the memory of observations to ultimately learn structure, reuse it, and infer logical decisions.



Application -aware communications


Unlike app-aware communications, SC captures structure, context and attributes a representation to the data at a low-level, that is beyond application level information.



Goal-oriented communications

SC is beyond goal oriented whereby the transmitter and receiver must have cooperative or competing goals with respect to an environment.

Key Characteristics of Semantic Representations and a Semantic Language



- **Minimalism**
 - The capability of characterizing the structure found in the information with the least number of language elements possible
 - Reduction of the number of exchanged messages in the long run as well.
- **Generalizability:**
 - Representing a particular underlying structure (or understanding one at the receiving end) while being invariant to changes in: **a) distribution, b) domain, and c) context.**
 - This mimics the behavior of a natural language to universally use words to describe events.
- **Efficiency:**
 - The ability of the apprentice to re-generate the information with **high fidelity**, in the **least time** possible.
 - ➔ The resolution of the data generated at the apprentice must be equal (or better) to that which could be recovered by a classical receiver.

Reasoning via Causality

- Learn a language by identifying the root-causes of information.
- Generate semantic content and make logical decisions.

Semantic-based KPIs

Novel metrics that can evaluate the performance of semantic networks by characterizing the intrinsic reasoning and level of symmetry of the SC system

Computing resources

The efficiency of SC grows substantially with the ease, dynamicity, and availability of computing resources

Semantic language

A map between a stream of raw data bits and their respective semantic representation.

Minimally sufficient representation

Minimal representations, yet expressive and can capture the semantic structure of the data.




Semantic Communication Backbone

(Some) Benefits of Semantic Communications

**AI-Nativeness
and
Interoperability**

**Robust/Resilient
Channel Control**



**Less Data,
More Knowledge**

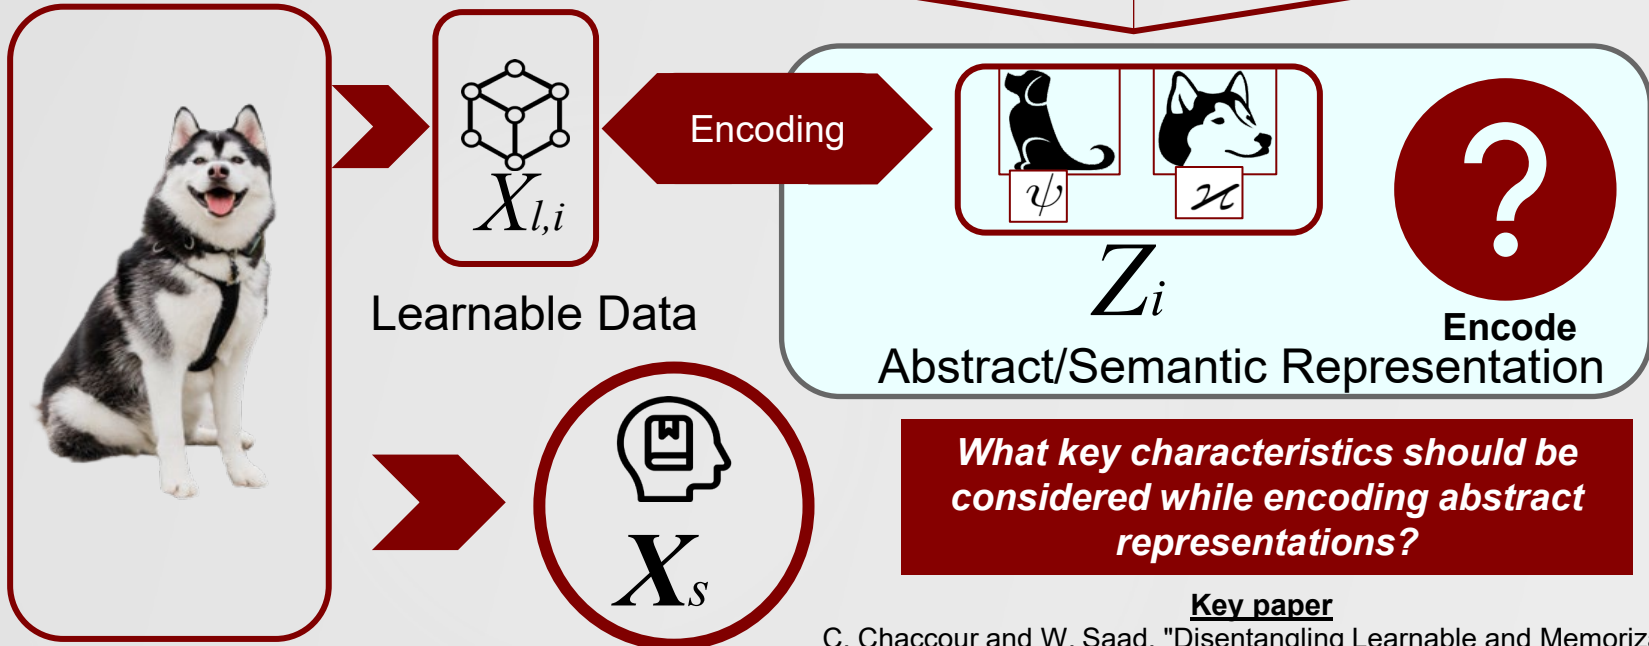
**Intrinsic Contextual
Awareness**

Step 1 - Disentangling Meaning/Semantics

Semantic Content Elements

Y_i

Dog	Siberian Husky
-----	----------------



What key characteristics should be considered while encoding abstract representations?

Key paper

C. Chaccour and W. Saad, "Disentangling Learnable and Memorizable Data via Contrastive Learning for Semantic Communications", in Proc. of the 56th Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA, USA, November 2022.

Encoding

- Semantic communications require generalizable abstract representations
- Key characteristics of **abstract representations** in semantic communications:

Abstraction




- Sufficiently hold their **essential characteristics**

Causality



- Uncover the **causal relations** with other representations

Generalizability




- Maintain a **common generalizable form** for analogical reasoning

Key Result: Disentangling Learnable and Memorizable Data via Contrastive Learning for Semantic Communications

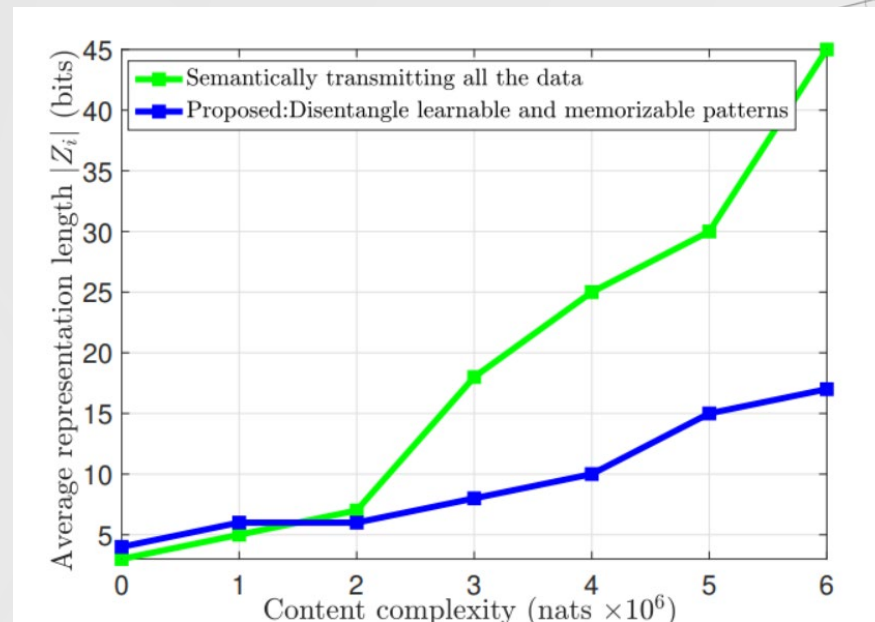
Christina Chaccour and Walid Saad

Published in the Proceedings of the 56th Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, USA




Simulation Results

- This work uses contrastive learning to perform the pre-processing/disentangle process
- **We can see that the average representation length increases with the content complexity:**
 - For a low content complexity, semantically transmitting all the data might result in a smaller representation length. This is because the amount of random information $\rightarrow X_s$ is considerably small.
 - As we \uparrow the content complexity \rightarrow Semantically transmitting all the data is not a feasible approach \rightarrow Representation length steeply increases as we increase the content complexity.
 - Our **representation is minimized by 57.22%** compared to the vanilla semantic approach.
- Now that we know how to disentangle information, let's go deeper into **reasoning and causality**



Step 2 - Causal Reasoning: Why?

- Reasoning and “real” learning can only be performed by asking questions → Queries (counterfactuals and interventions), the emerging framework of causality enable this.
- Reasoning mainly relies on *characterizing causal and associational logic in the data*.
- We cannot rely on state-of-the-art ML frameworks that make assumptions such as:
 - ❖ i.i.d. datasets
 - ❖ Stationarity scenarios
 - ❖ Data has no root-cause



Causal Logic Ladder



Reasoning

Congregating associative, interventional, and counterfactual logic to understand the representations conveyed and generate representations with their proper semantic connotation.



Counterfactual Logic

Learning with retrospection and imagination. The apprentice is attempting to ask the “Why?” questions when it comes to the current representations used by the teacher and their respective semantics. “What is the root cause of a particular representation?”



Interventional Logic

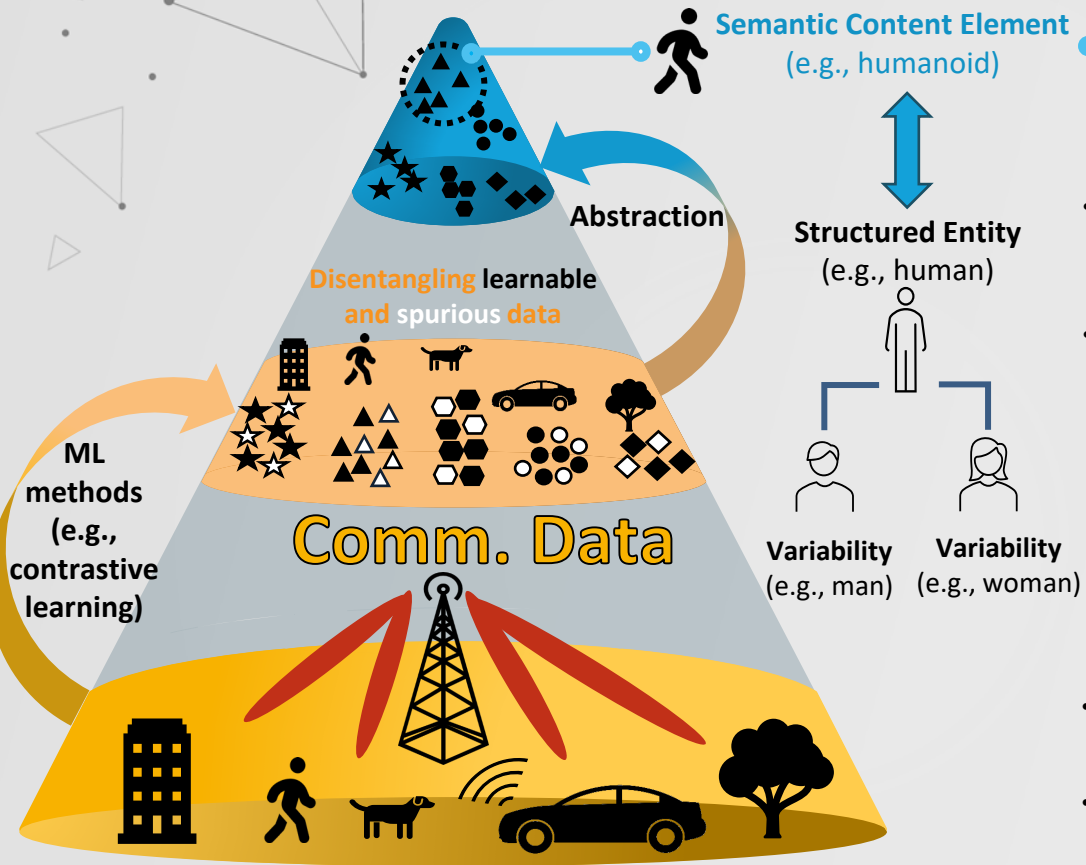
Learning while invoking questions with the do operator. That is, the apprentice is attempting to learn what would happen in case the causes were different. In other words, the apprentice is asking “What if?”. “What would the representation be if the semantics were different?”



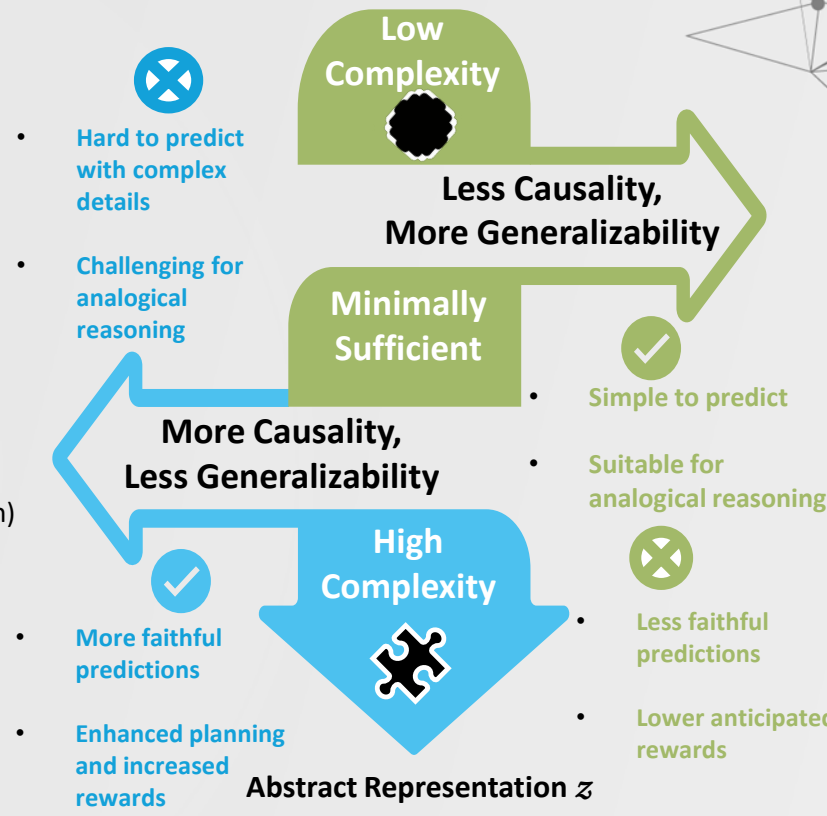
Associative Logic

Learning information based on purely statistical relationships without invoking any causality or semantics within the data. This is a purely observational task on the datastream.

Disentanglement (or Perception!): The Big Picture



Causal vs. Generalizable Representation Learning



KPI 1: Communication Symmetry Index

The communication symmetry index between a teacher b and apprentice d , for a transmission session τ is given by:

$$\eta_{b,d,\tau} = \frac{\zeta_{d,\tau}}{\nu_{b,\tau}} \times \iota_{\tau,Y_i}$$

Number of query packets (e.g., interventions, counterfactuals, etc.) needed to reason

Number of raw data packets accompanying semantics

The semantic impact generated by a semantic representation Z_i during a time τ is the number of packets that would have been needed to be transmitted to regenerate the semantic content element Y_i

Insights:

- Based on the values of $\eta_{b,d,\tau}$ and, ι_{τ,Y_i} , one can determine the level of symmetry between the teacher and the apprentice.
- E.g.: A high $\eta_{b,d,\tau} \rightarrow$ high level of symmetry between teacher and apprentice \rightarrow the apprentice has generative capabilities. (A high $\eta_{b,d,\tau}$, with low ι_{τ,Y_i}) \rightarrow reverse mentorship, teacher's capabilities are also weak.

KPI 2: Reasoning Capacity



The reasoning capacity between a teacher b and an apprentice d is given by:

$$C_R = \Omega \log_2(1 + \eta_{b,d}),$$

Maximum Computing
Resources


$$\rightarrow C_T = C_C + C_R = W \log_2(1 + \gamma) + \Omega \log_2(1 + \eta_{b,d}),$$

The total capacity is no longer limited by Shannon's bound only as a result of the convergence of computing and communications!

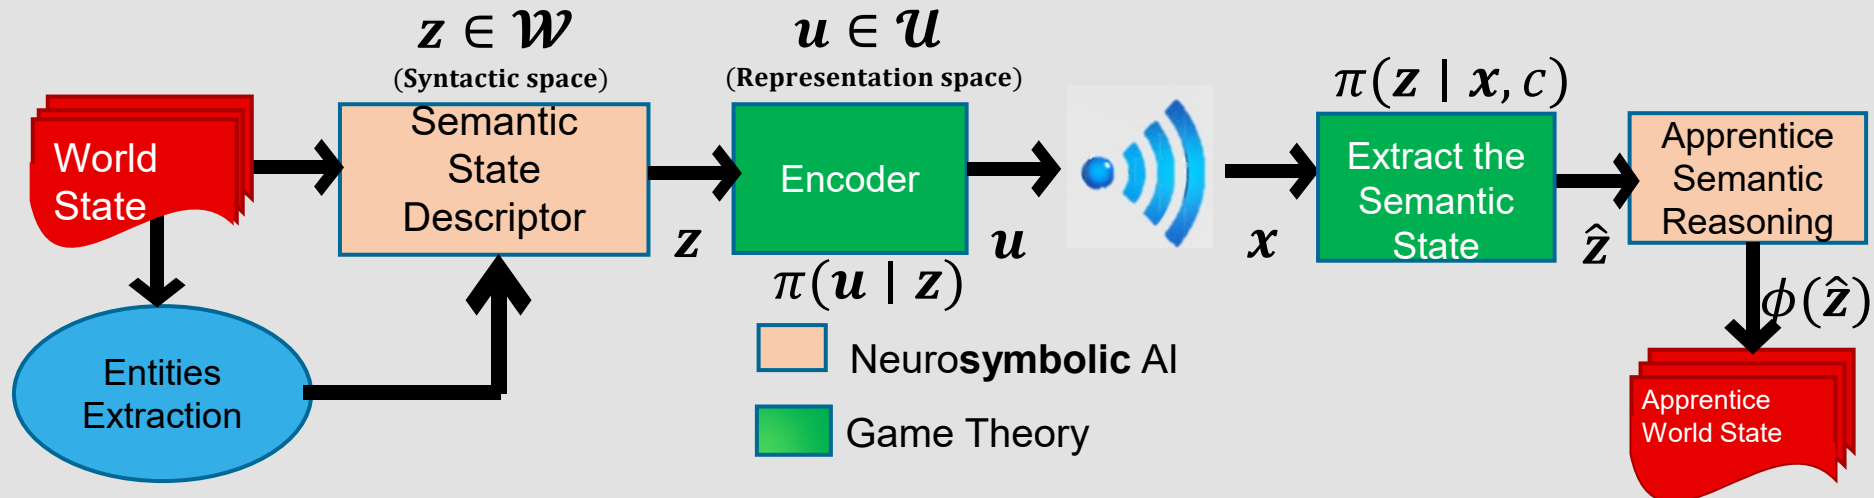
Example Result 1: Neuro-Symbolic Causal Reasoning Meets Signaling Game for Emergent Semantic Communications

Christo Thomas and Walid Saad

IEEE Transactions on Wireless Communications, 2024 :
<https://arxiv.org/abs/2210.12040>




Reasoning Foundations for Emergent Semantic Communication (ESC)



Reasoning Problem


- How to rigorously define and learn the semantics present in the source data ? (**transmitter reasoning problem**)
- How to perform deductive reasoning, by evaluating the logical formulas ($\phi(\hat{z})$)? (**receiver reasoning problem**)

Emergent Language Problem



Compute teacher transmit strategy (**encoder**) $\pi(\mathbf{u} | \mathbf{z})$
 & apprentice inference strategy (**decoder**) $\pi(\mathbf{z} | \mathbf{x}, c)$

Language Problem




Emergent Language (communicating language emerges)

- Compute teacher transmit strategy (encoder) - $\pi(\mathbf{u} | \mathbf{z})$ and apprentice inference strategy (decoder) - $\pi(\mathbf{z} | \mathbf{u}, \mathbf{c})$

Encode based on semantics, benefits

- transmitting semantically similar messages as same signal thus **saving bits/BW**
- removing **redundant semantics**

Causal Reasoning Problem



Semantic state descriptor:

- Infer the hidden relations among the entities (the causal sequence that best explains the event observed)

$$p(s_0, \dots, s_N | e) = \prod_i p(s_i | \text{pa}(s_i)),$$

Parent nodes in the graph

Apprentice semantic reasoning:

$$p(\phi_i = \phi | \hat{\mathbf{z}}), \forall \phi_i \in \Phi,$$

Set of logical formulas

- evaluate the logical formulas

Emergent Language Problem: Two -Player Signaling Game NEWS@VT



Semantic notion of information



Teacher objective

Apprentice wants to

$$\pi_{s,t}^* \in \arg \max_{\pi_{s,t}} -\mathbb{E}_{\mathbf{u}_t} [S_s(\mathbf{z}_t; \mathbf{u}_t | \pi_{s,t}, [\mathbf{z}_{t-1}])],$$

$$\pi_{l,t}^* \in \arg \max_{\pi_{l,t}} \mathbb{E}[S_l(\hat{\mathbf{z}}_t; \mathbf{u}_t | \pi_{l,t}, \pi_{s,t}^*, [\hat{\mathbf{z}}_{t-1}])],$$

s. t. $\mathbb{E}[V(\pi_{s,t}, \pi_{l,t})] \leq D$

01

Transmit less (minimalism)

02

Extract more (under less semantic surprise = high fidelity)

Computing or Communication costs


Apprentice surprise (to minimize)

$$V(\pi_{s,t}, \pi_{l,t}) = c(\mathbf{u}_t) - \log \pi_{l,t}$$

Category theory & semantic information

Information contained in the logical entailments (copesheaves from category theory perspective) that follow from any causal state.

Received semantic information: a fraction (= semantic similarity) of the transmitted information. **Semantic similarity** (≤ 1) can be computed by quantifying the overlap in terms of the copresheaves of \mathbf{z} & $\hat{\mathbf{z}}$



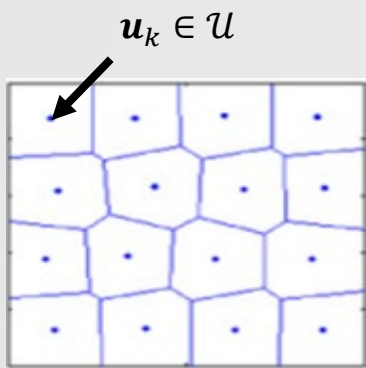
Copresheaf example

Nash Equilibrium Analysis

01	Minimalistic, but low fidelity	Pooling equilibrium $ \mathcal{U} = 1$
	High fidelity, but maximal transmission	Separating equilibrium ($ \mathcal{U} = \mathcal{W} $)
02	Minimalistic, and high fidelity	Partial pooling $ \mathcal{U} < \mathcal{W} $ is realistic. Optimal strategies below:

Teacher – Transmit Signal Partition

Voronoi Tessellation of syntactic space \mathcal{W} s.t. avg. semantic info. extracted at listener is maximum among all possible partitions.



Apprentice Decoding Strategy: Bayesian Estimator

$$\arg \min_{z \in \widehat{\mathcal{W}}} \int_{\widehat{\mathcal{W}}} \|z - \widehat{z}\|_S \pi(d\widehat{z} | u),$$


Not in the **Euclidean space** but in the semantic space!

Key Analytical Results


Semantic Information



Compared to set-theoretic methods [1], using **category theory** helps to represent deductive and logical theorem proving properties



$$F: \mathcal{L} \rightarrow \hat{\mathcal{L}}$$



Copresheaves
(logical conclusions)

Minimalism



Reason or communicate

The **average amount of bits** to represent the state description in an ESC system has **upper and lower bounds less than that of a classical system.**

Semantic Error Probability

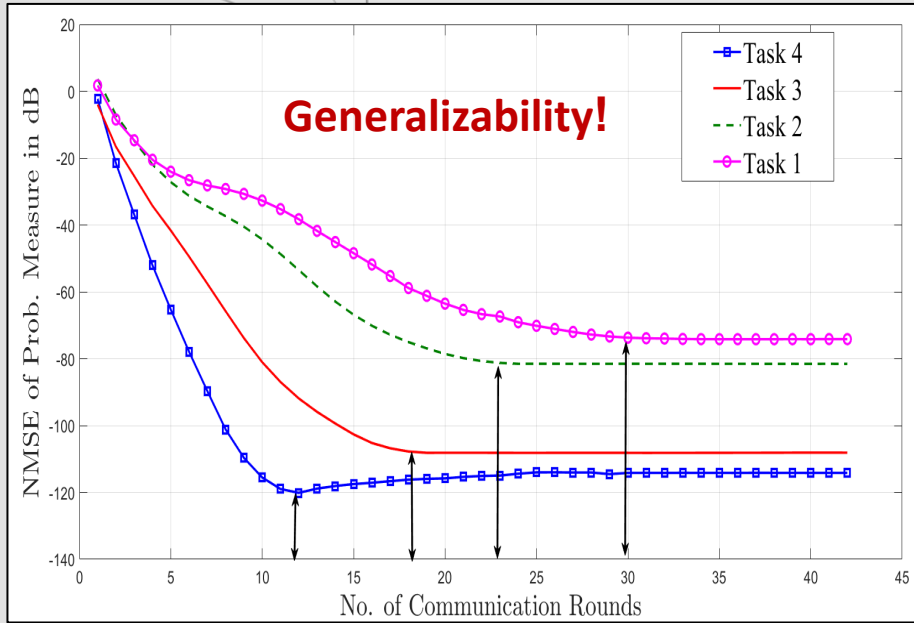


Should we transition to semantic performance measures?

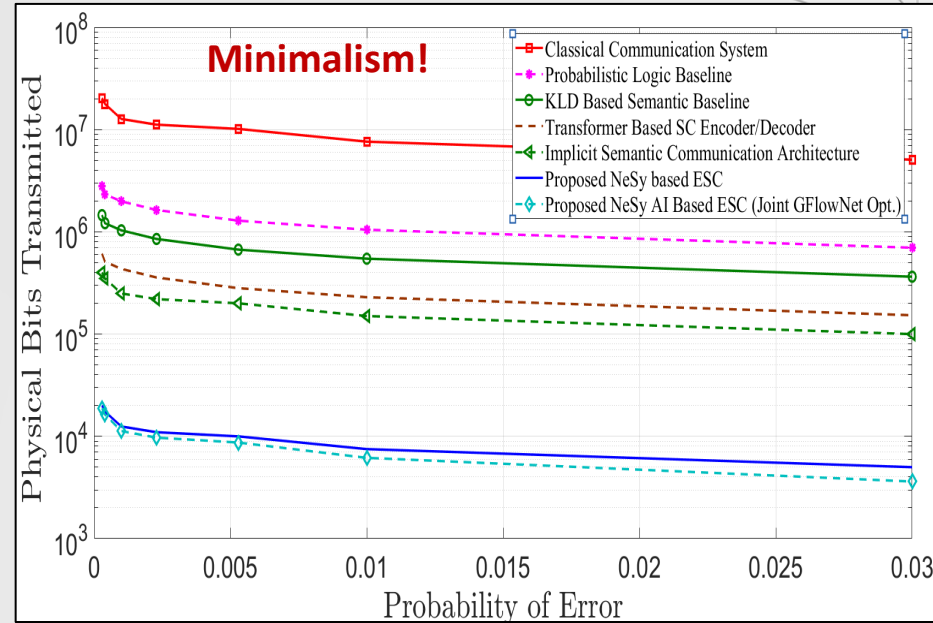
The **lower bound on the semantic error probability** is always **\leq the lower bound on the probability of bit error** measure achieved using classical communication system.

[1] R. Carnap and Y. Bar-Hillel, "An Outline of a Theory of Semantic Information," Technical Report No. 247, Oct. 1952.

Simulation Results



- ❑ Number of training rounds decreases over time which demonstrates how the generalizable aspects of proposed approach help over time.



- ❑ ESC enables efficient communication as the system transmits less in comparison to state of the art while achieving better semantic reliability.

Key Result 2:

Causal Semantic Communication for Digital Twins: A Generalizable Imitation Learning Approach

C. K. Thomas, W. Saad, and Y. Xiao, "Causal Semantic Communication for Digital Twins: A Generalizable Imitation Learning Approach", *IEEE Journal on Selected Areas in Information Theory*, 2023.

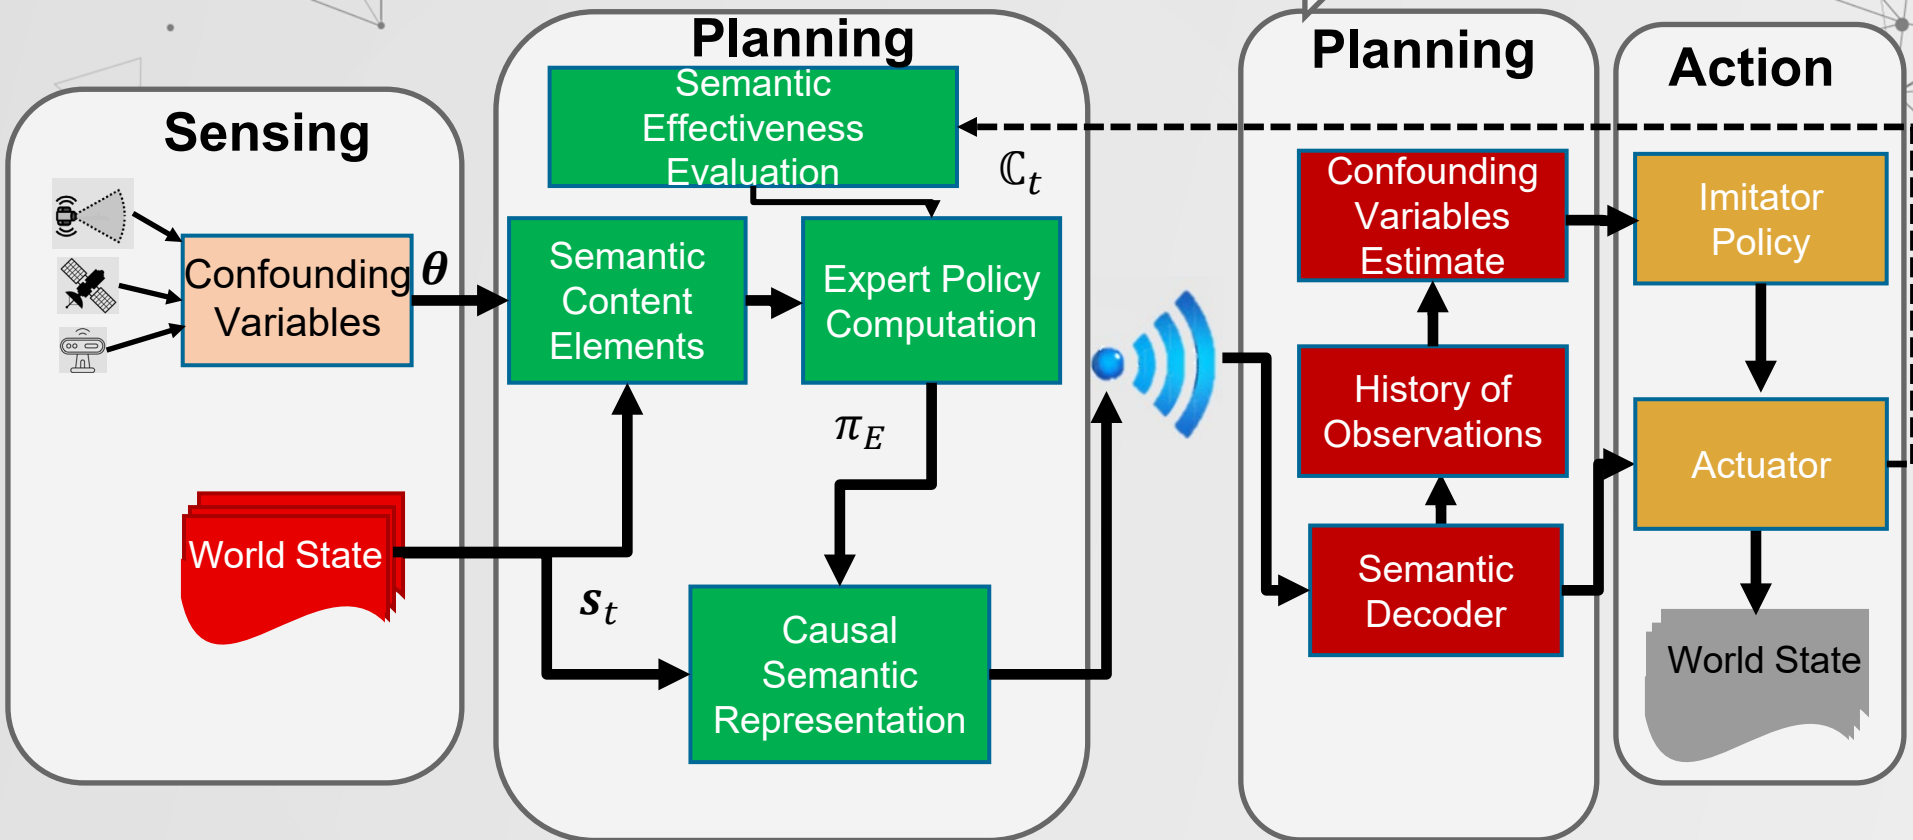


Causal Semantic Communication for DTs

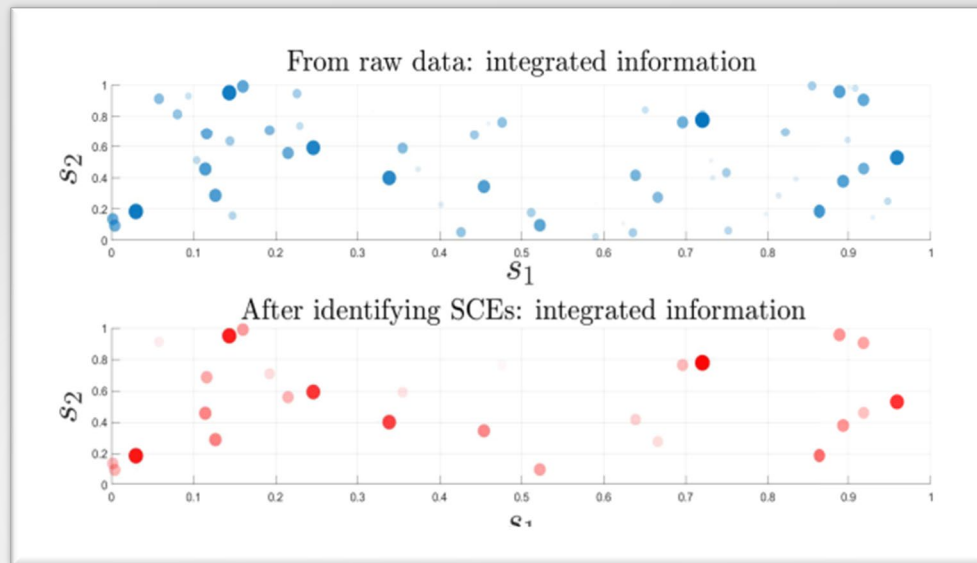

Expert Agent

Imitation Learning

Imitator Agent



Causal Transition Models for Planning



- **Semantic Information:** uses **Integrated information theory (IIT)** from theory of consciousness in neuroscience
- **Intrinsic Information** for State Abstraction
 - Cause and effect information conveyed by any s_i^t (Impact under confounding variables as theoretical result)
- **Information Integration** (via Compositionality, for Identifying **Semantic Content Elements (SCEs)**)
 - Information conveyed by a subset of SCEs, as a whole and beyond sum of information of its parts
- **Semantic concepts**, causal relations among concepts, topological characterization (**abstract cell complex**) as a theoretical result



Key Analytical Results

Confounding variable estimate

The **error in semantic information** learned between expert and imitator nodes can be made arbitrarily small if the learned posterior about the confounding variables has its peak around the true value of confounders.

Causal relations as simplicial complex

The structure of causes and effects in an SCM can be represented using an **abstract simplicial complex**, in which causes and effects are nodes and relations are simplices.

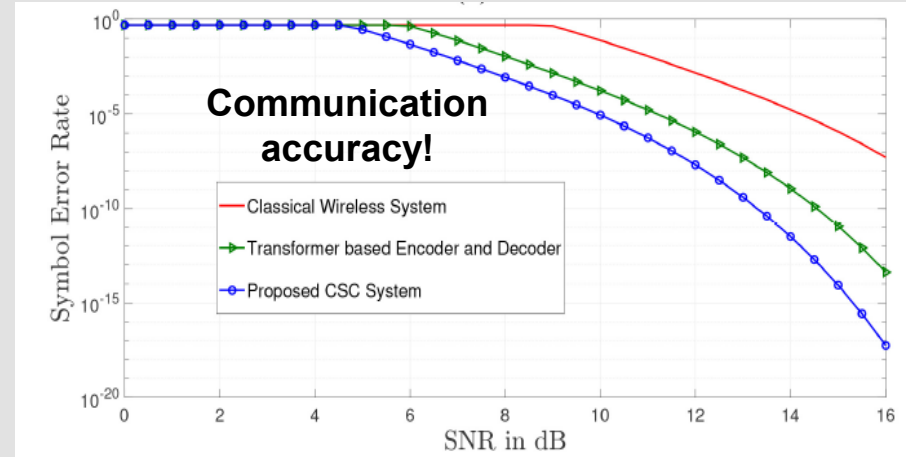
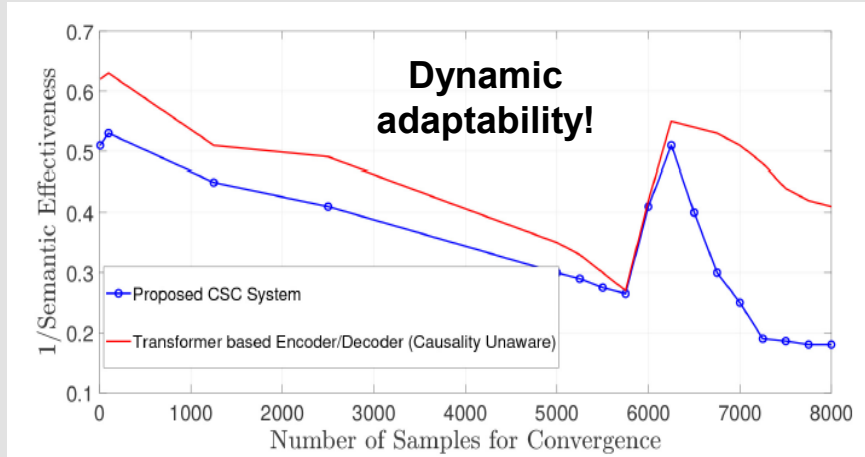
Semantics and Topology

- Topological characterization of semantics as **abstract cell complex** → represents meaning
- Seamless communication across multiple nodes, without dependency on transmit encoders
- Allows a rigorous formulation of **semantic metrics** → reliability, similarity.

Performance characterization

- Discrepancy in transition probability modeling dominates the error in performance (average regret) → accurate physical environment modeling is crucial.
- **Generative AI** can help to close this gap at the receiver as communication progresses.

Simulation Results



- ❑ Proposed CSC system **requires less samples for convergence** as the environment dynamics changes, which demonstrates its dynamic adaptability due to **generalizability**

- ❑ Proposed CSC system **achieves lower symbol error rate** compared to classical wireless systems and state-of-the-art semantic communication systems *that do not incorporate reasoning*



From Semantic Communications to World Models and Digital Twins for Physical AI: **A Glimpse**



Physical AI: The Good, the Bad, and the Ugly

- ❑ There are many recent works that demonstrated some benefits of AI in many fields
- ❑ But that is only the **good** part of the story....There is still the **bad** and **ugly** parts, particularly for **physical AI**, i.e., AI operating under the laws of the real world!

Robot trained to



AI-Native 6G Air Interface



Distribution

Autonomous vehicle



Why is it that today's AI is good enough to win a math Olympiad, but yet it fails in very simple physical world scenarios?

Robot crashes but still crashes



02



01

Millions of driving trials

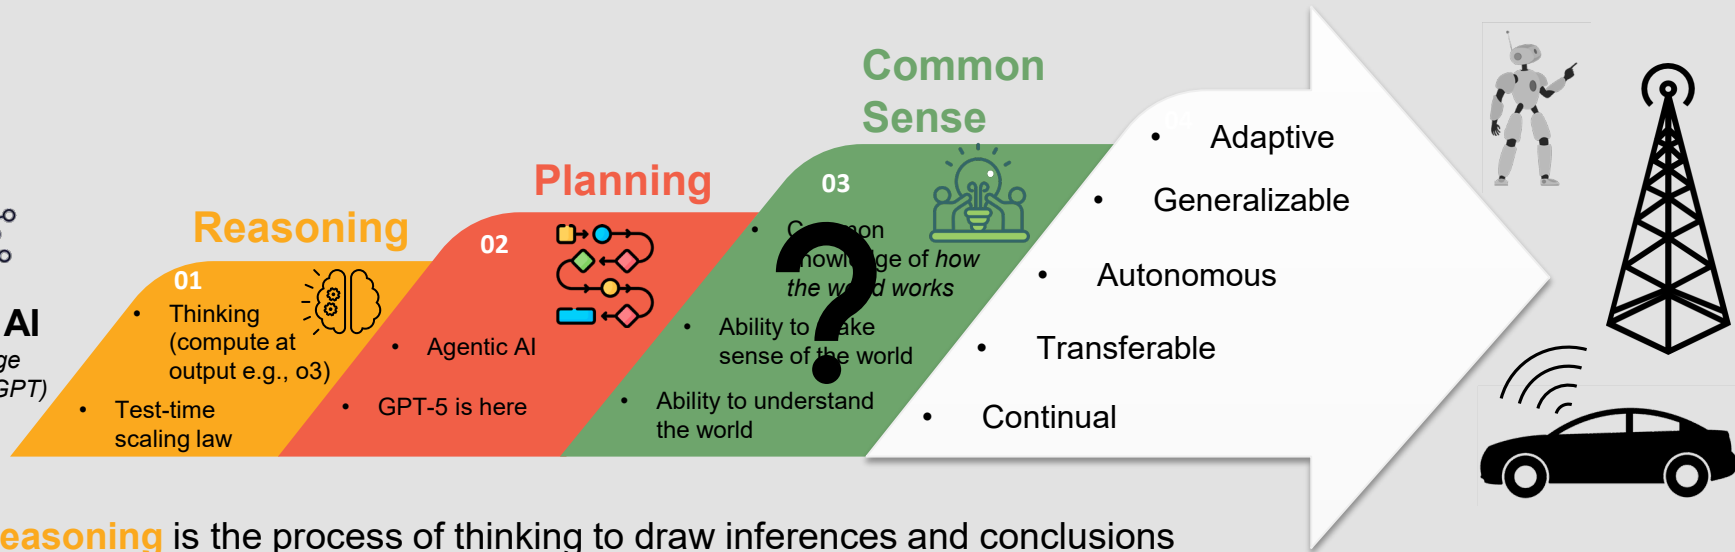


02

AI: The Cognitive Path Forward



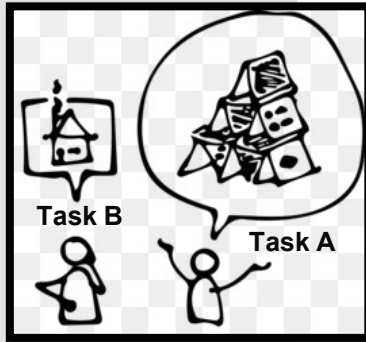
Today's AI
(large language models e.g., GPT)



- ❑ **Reasoning** is the process of thinking to draw inferences and conclusions
- ❑ **Planning** is the capability of breaking down a complex task into a sequence of actions to achieve a goal
- ❑ The current forms of **reasoning/planning** are good for **language models**, *but are not enough for real-world agents*
- ❑ What is the **missing cognitive ability** needed to move forward into the world?

Common Sense: The Missing Link to Next-Generation AI and Wireless Networks

The Four Essential Pillars of Common Sense




Analogical Reasoning

- Learn new skills faster through analogy & limited interactions from the world (via observation)
- Relate elements, situations, and concepts via analogy

Dealing with Unforeseen Scenarios

- Horizontal generalizability
- Leverage common knowledge about the world to deal with out-of-domain, corner cases



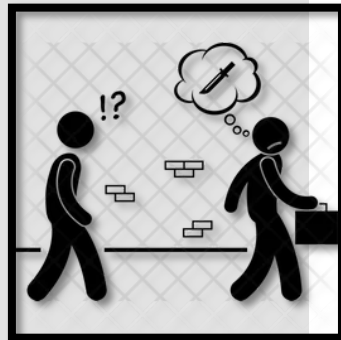
Common Sense

Filling in the Blanks

- Connecting the dots to perform logical reasoning about causes or events
- Insert plausible elements in missing spots as needed


Intuitive Physics

- Background knowledge about the world to infer what is likely to happen next
- Determining future states that are most probable, plausible, or impossible to occur



Common Sense: Between Cognition & AI NEWSVT

- A w
sen
- Wha
➢ A r
○Re
○Ca
○Dy
- The
- Let'
cap
- We
co



SYSTEM 1

Intuition & instinct

95%


Unconscious
Fast
Associative
Automatic pilot

SYSTEM 2

Rational thinking

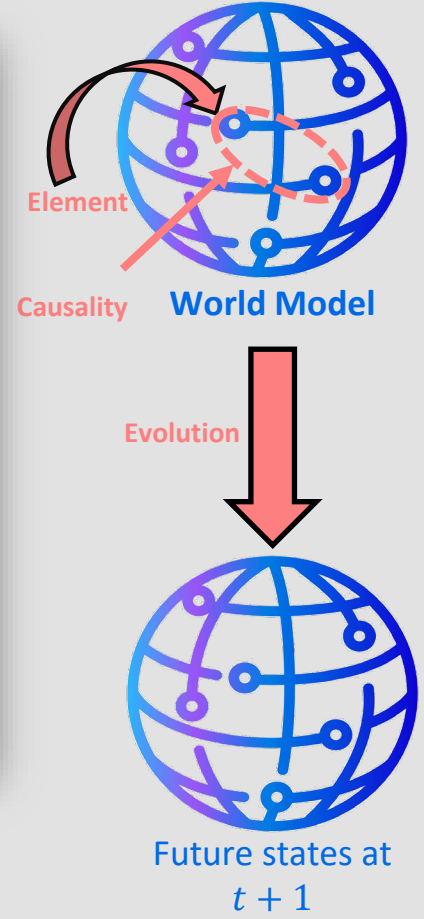
5%

Takes effort
Slow
Logical
Lazy
Indecisive

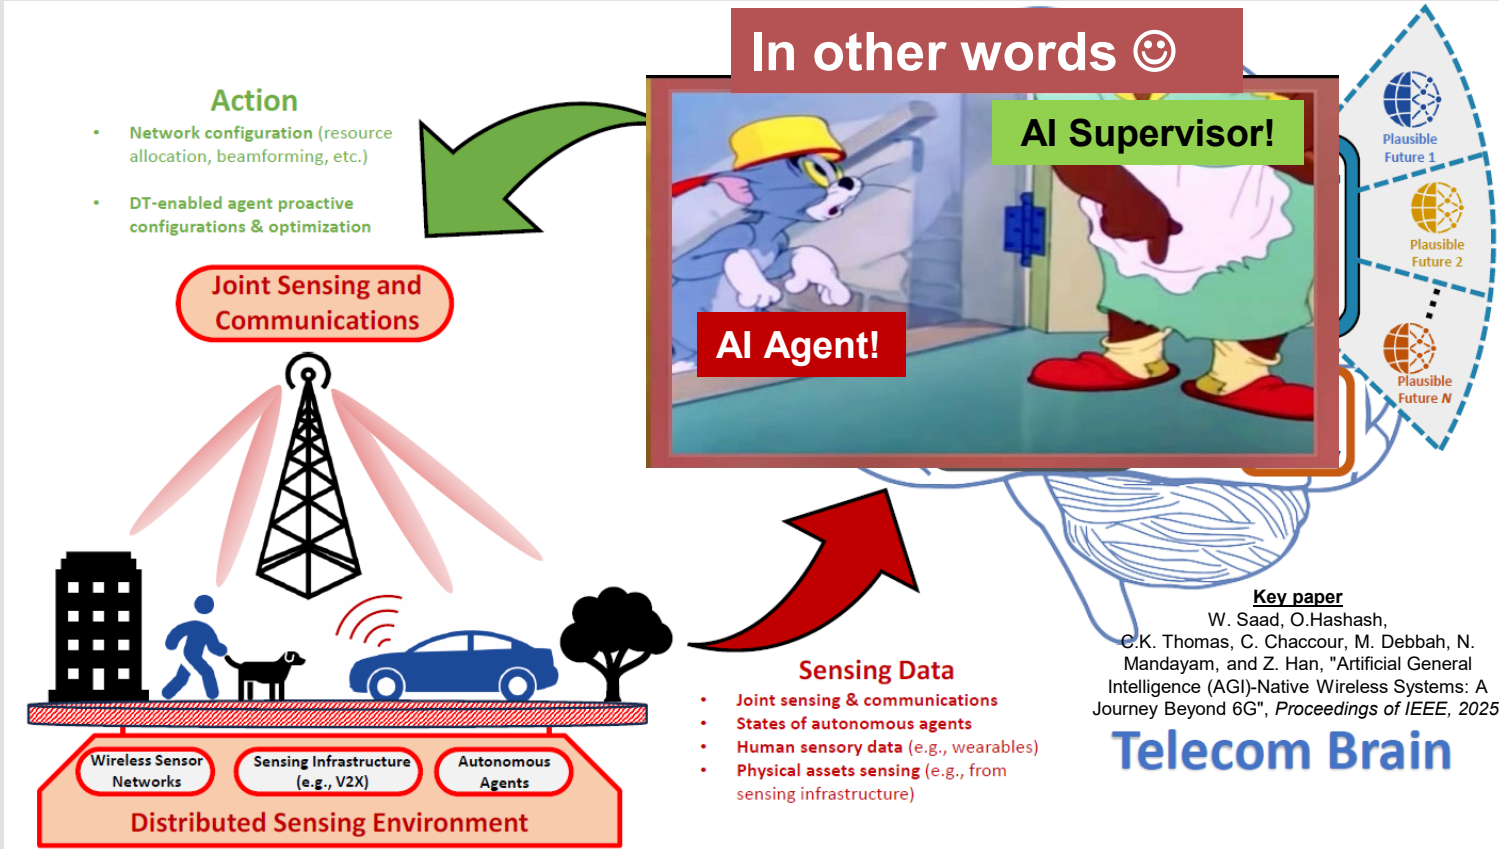


Source: Daniel Kahneman

- We do have “**System 1**” architectures in today’s AI (e.g., LLMs), but are *missing* “**System 2**” architectures
- Towards generalizable intelligence for physical AI



Proposed Vision for Next -Generation Networks




Why AI Still Can't Handle the Physical World

Active Inference as the Test-Time Scaling Law for Physical AI Agents

O. Hashash, C. K. Thomas, W. Saad, M. Debbah, K. Friston, and A. Razi,
“**Active Inference as the Test-Time Scaling Law for Physical AI Agents**”,
(to be submitted) *IEEE Transactions on Pattern Analysis and Machine Intelligence*

Active Inference is a first principle that describes how all living systems survive in the physical world



SYSTEM 1
INTUITION
95%

SYSTEM 2
REASONING
5%

Router

gpt-5-main

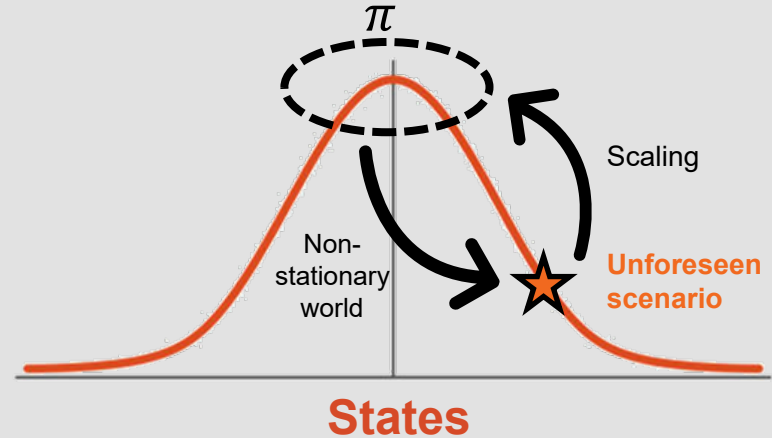
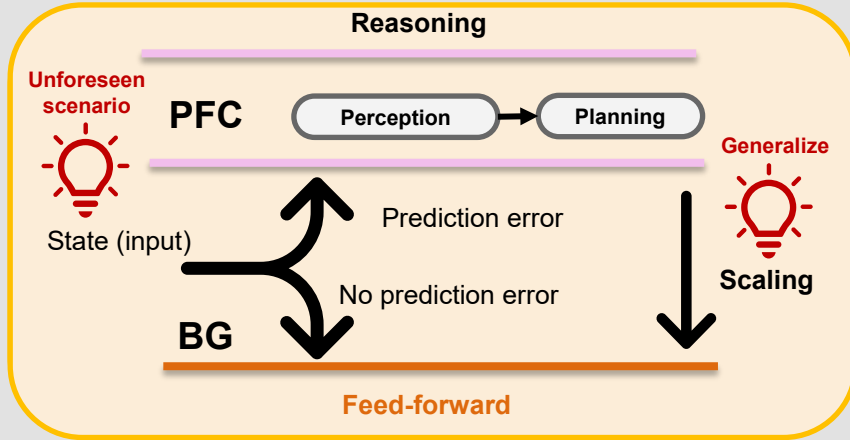
Handles straightforward queries with lightning-fast responses

gpt-5-thinking

Tackles multi-step problems that require careful analysis

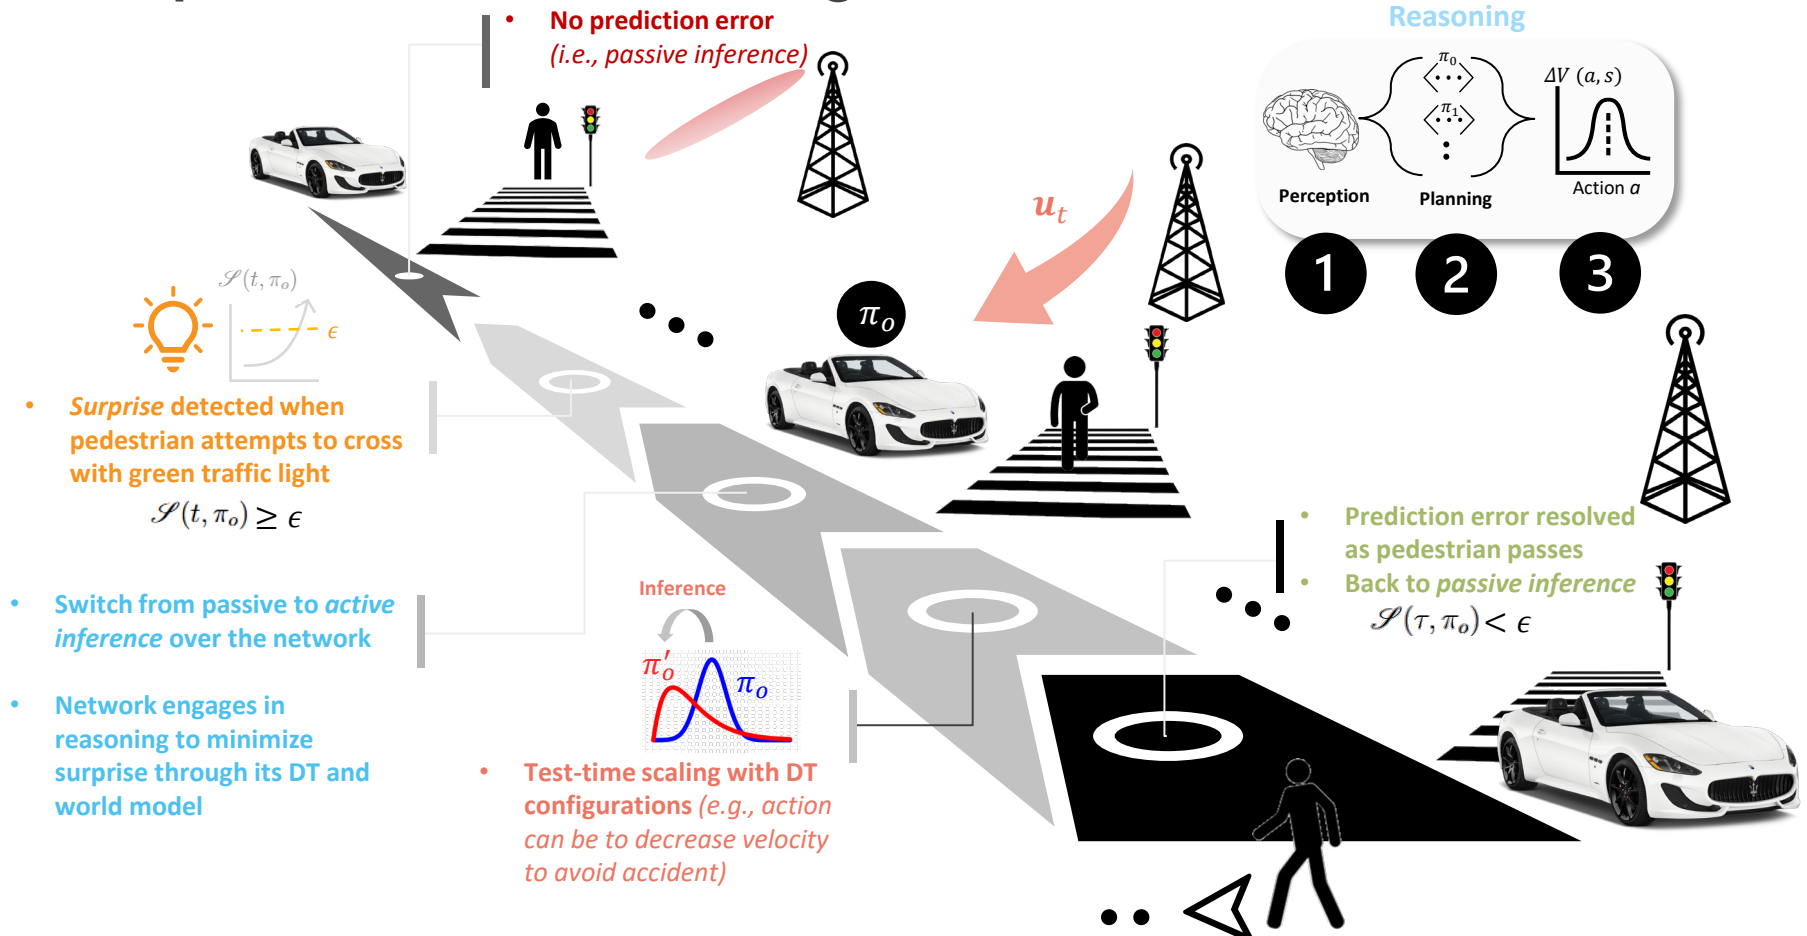
Rewards

Active Inference at Test-time



- ❑ **Prediction error/surprise** = internal predictions vs. received sensory information
- ❑ **Reasoning to reduce surprise** = *minimum amount of reasoning in any living system to survive*
- ❑ Every living system shows a **general (internal) objective** to survive

Example: Test-Time Scaling in an Unforeseen Scenario



Test-Time Scaling Law

- ❑ **Action** is modeled as a soft Bayesian update using *Pearl's method of virtual (counterfactual) evidence* (i.e., from world model):

Theorem 1 (Test-Time Scaling Law). *The test-time scaling law that scales policy $\pi_o(a|s)$ into $\pi'_o(a|s)$ with inference about the actions that are likely to minimize surprise is given by:*

$$\begin{aligned} \underbrace{\pi'_o(a|s)}_{\text{posterior}} &\propto \underbrace{\pi_o(a|s)}_{\text{prior}} \underbrace{p(o_{t+1:T}|\pi)}_{\text{likelihood}} \\ &= \underbrace{[\pi_o(a|s)]^{1-\theta} [\pi_o(a|s) \exp(-\gamma\Delta V(a, s_t))]}_{\text{soft Bayesian update}}^\theta \\ &= \underbrace{\pi_o(a|s)}_{\text{feed-forward}} \underbrace{\exp(-\gamma\theta\Delta V(a, s_t))}_{\text{inference}}. \end{aligned} \tag{5}$$

Action as inference

Normalized surprise

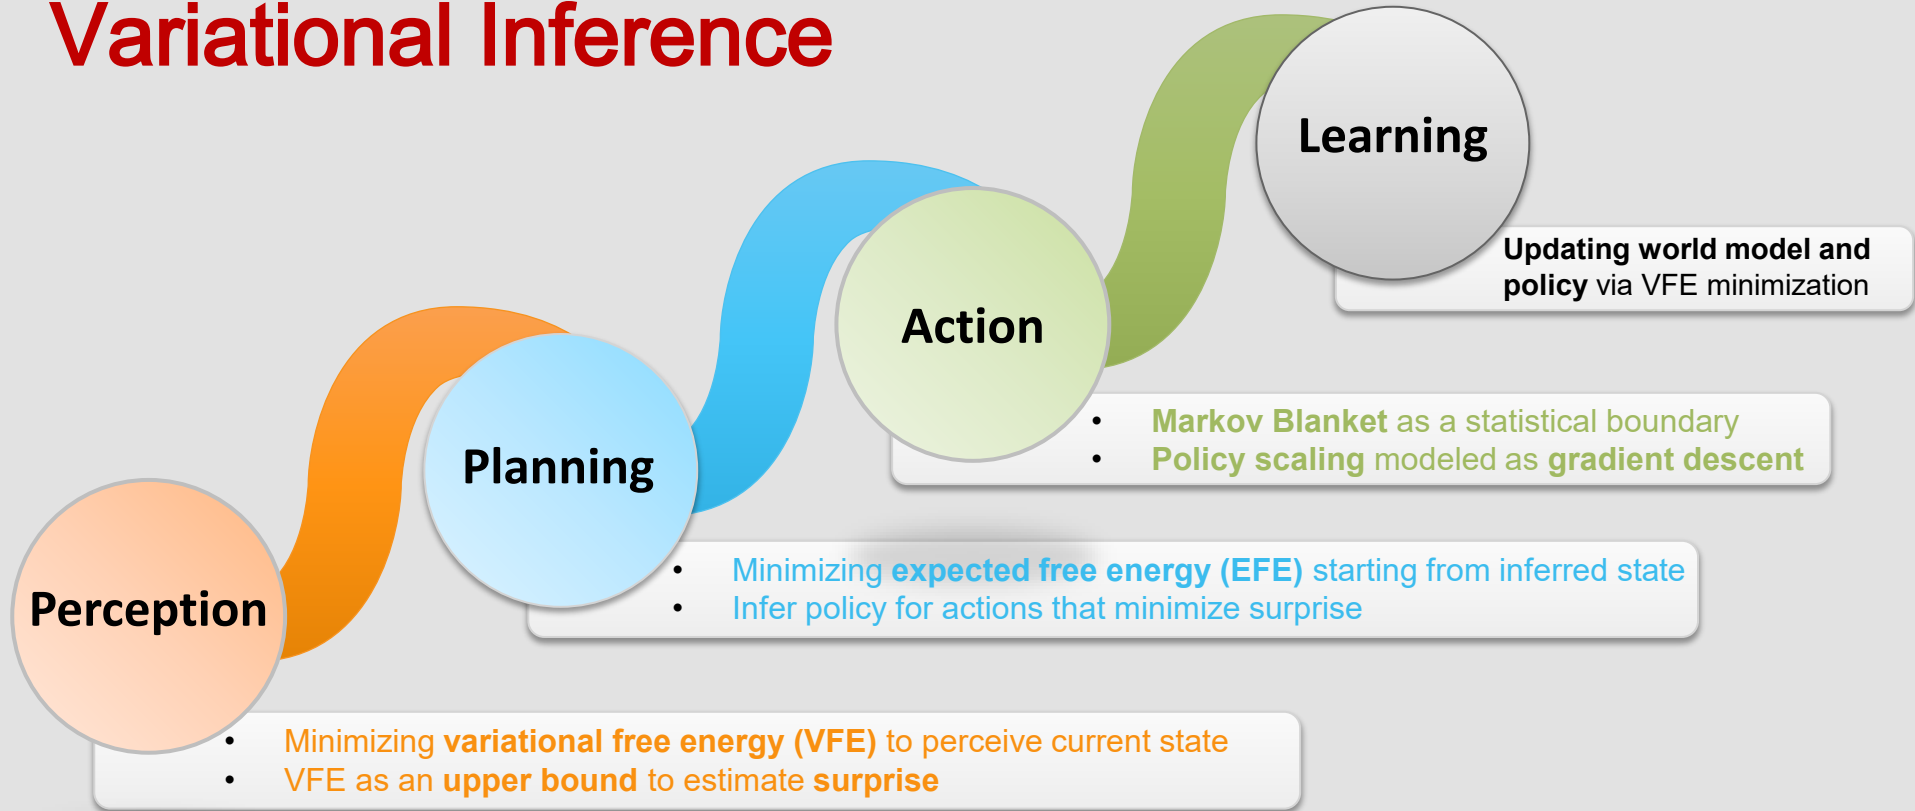
- ❑ When the world is mostly predictable $\Rightarrow \mathcal{S}(t, \pi_o) \approx 0 \Rightarrow \theta \approx 0$.

- ❑ Policies in today's solutions remain limited to the case where *surprise is null due to stationary assumptions (e.g., RL)*, which is clearly a **special case of this equation!**

Inference can be challenging since marginalizing over all the possible states $p(o_{1:t})$ is **computationally intractable** in practice

Solution:
Variational Bayesian Inference


Solution Roadmap: Variational Inference




Simulation Results




Q-Learning: Distance to Intersection




Q-Learning: Velocity




Q-Learning: Actions Taken




Scaled Policy: Distance to Intersection (NO CRASH)



Scaled Policy: Velocity (Maintains v=0 during danger)




Scaled Policy: Actions Taken (Blue=QL, Yellow=AI)



❑ While Q-learning fails in unforeseen scenario, policy scales with *test-time scaling law* to **generalize in an unforeseen scenario**

Simulation Results



- Our solution can detect the unforeseen scenario (i.e., surprise)
- Adapt by scaling its policy to avoid crashing with pedestrian

- Test-time scaling outperforms Q-learning and Bayesian RL in unforeseen scenario:
 - Scaling can sacrifice rewards to ensure success, while efficiently utilizing inference with 36% improvement

Note on Multi-Agent World Models

MetaMind : General and Cognitive World Models in Multi-Agent Systems by Meta - Theory of Mind

Lingyi Wang, Rashed Shelim, Walid Saad, and Naren Ramakrishnan

Submitted, 2026

Key Challenges for World Models in Multi-Agent Systems

- Multi-agent world models must **jointly support** understanding, prediction, and **long-horizon planning**
- Challenges arise due to the agents' **coupled dynamics, hidden goals/beliefs, and heterogeneous policies**
- Existing approaches often rely on **centralized context, explicit communication**, or passive dynamics prediction, which limits decentralized planning and generalization.

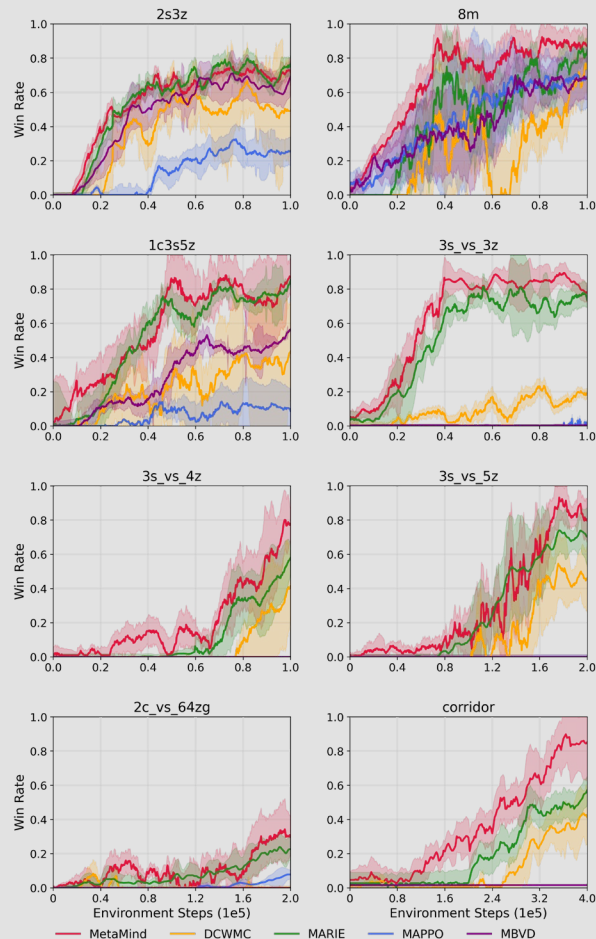


(a) Causal Chain of Goal-Oriented Behavior Principles (b) First-Person Learning for Meta-ToM (c) Third-Person Inference with Meta-ToM

Proposed Framework

- MetaMind introduces a **meta-theory of mind** world model that treats other agents as intentional entities
- Each agent first learns to infer its own goals and beliefs from actions through **self-reflection**
- This metacognitive ability is expanded to others via **analogical inference**

Simulation Results on Robotic Tasks



Red Line: (MetaMind) Our proposed world models

Yellow Line: (MARIE) Centralized world models

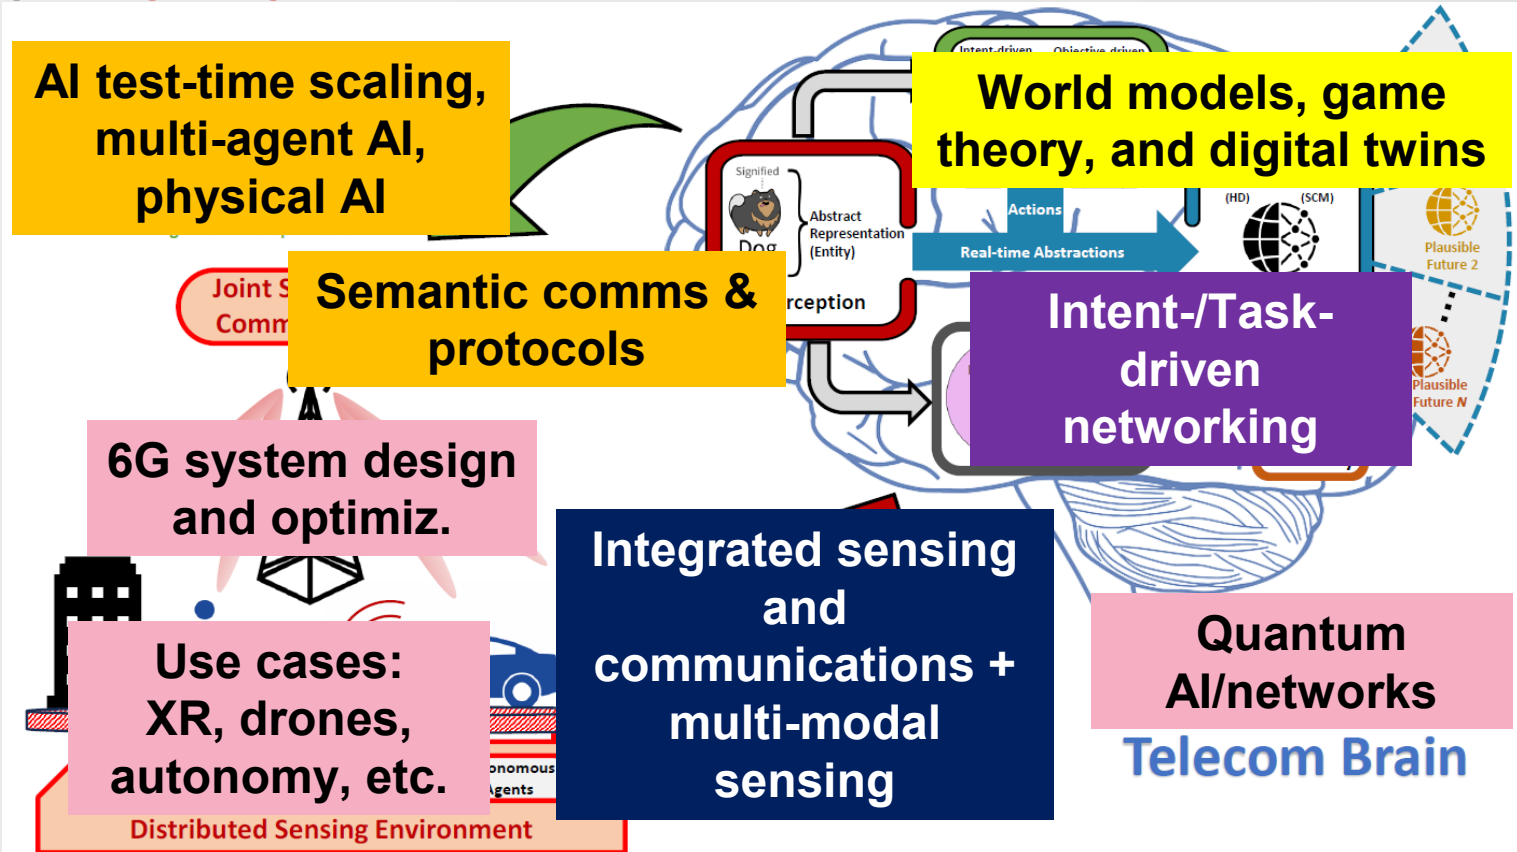
Green Line: (DCWMC) Message-dependent decentralized world model

Blue Line: (MAPPO) Model-free reinforcement learning approach

Purple Line: (MBVD) Model-based reinforcement learning approach

- We conduct the simulation across 8 **Starcraft multi-agent challenge (SMAC) tasks**
- It shows that, compared to existing multi-agent world models and multi-agent reinforcement learning approaches, **MetaMind has a significantly higher data efficiency.**
- MetaMind reaches 80% win rate within 200K steps on {3s_vs_5z}, whereas MARIE and DCWMC can only reach 75% and 52% win rate.
- This is because **MetaMind can provide decentralized, reliable imagination trajectories for policy learning.**

(Some) Ongoing Research Activities



Conclusion and Future Recommendations

Semantic communications may significantly enhance network performance

1

It is not merely a form of minimalism as existing works allude; it can enhance **resilience, reliability,** and overall **capacity** of a network

Advances in AI and computing are necessary

2

More efforts needed on **generalizable, reasoning** and knowledge driven AI as well as judicious computing resources

Mathematical foundations are necessary for multi-agent protocols

3

Autonomous multi-agent communication protocols requires **mathematical foundations** grounded in **category theory, game theory,** and related fields

Active inference is the missing link for general intelligence

4

It allows **test-time scaling**, which has been a key challenge for AI for a while





Thank you
Q&A