



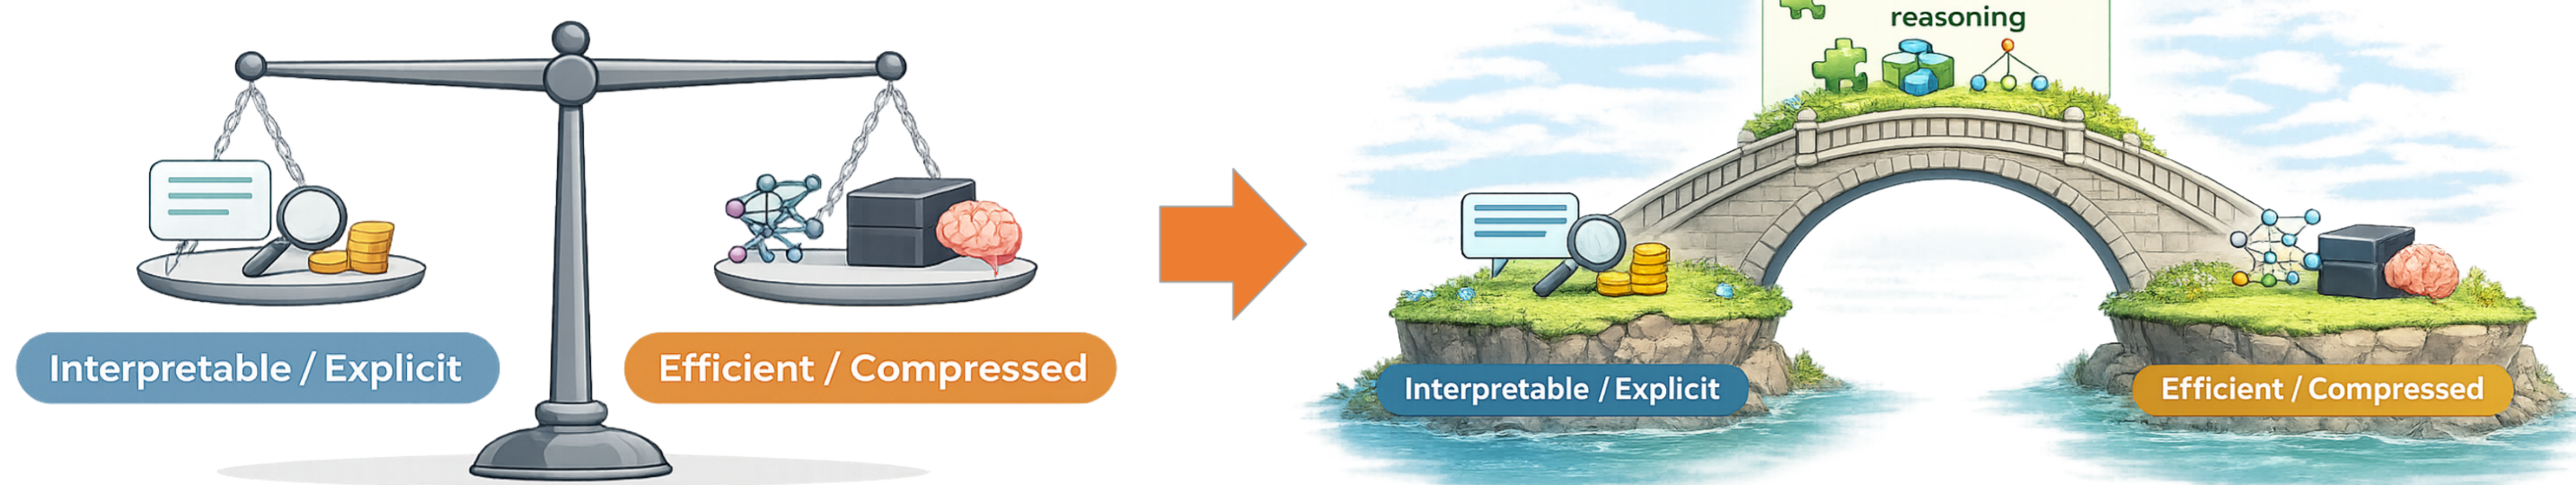
Core question: Can we obtain efficient reasoning without giving up interpretability?

1. PhD Problem and Hypothesis

Move Beyond CoT

- It is **easy to read/debug** and **supervise**
- But also **slow, token-costly**, and **not always faithful**
- Many tasks require structured, or iterative forms of computation for planning and exploration

Main hypothesis



- Reasoning over **structured concepts** or **symbolic abstractions** may improve both efficiency and interpretability.
- Structured latent reasoning as a bridge between neural and symbolic reasoning.**

PhD Goal

Develop efficient neuro-symbolic reasoning models

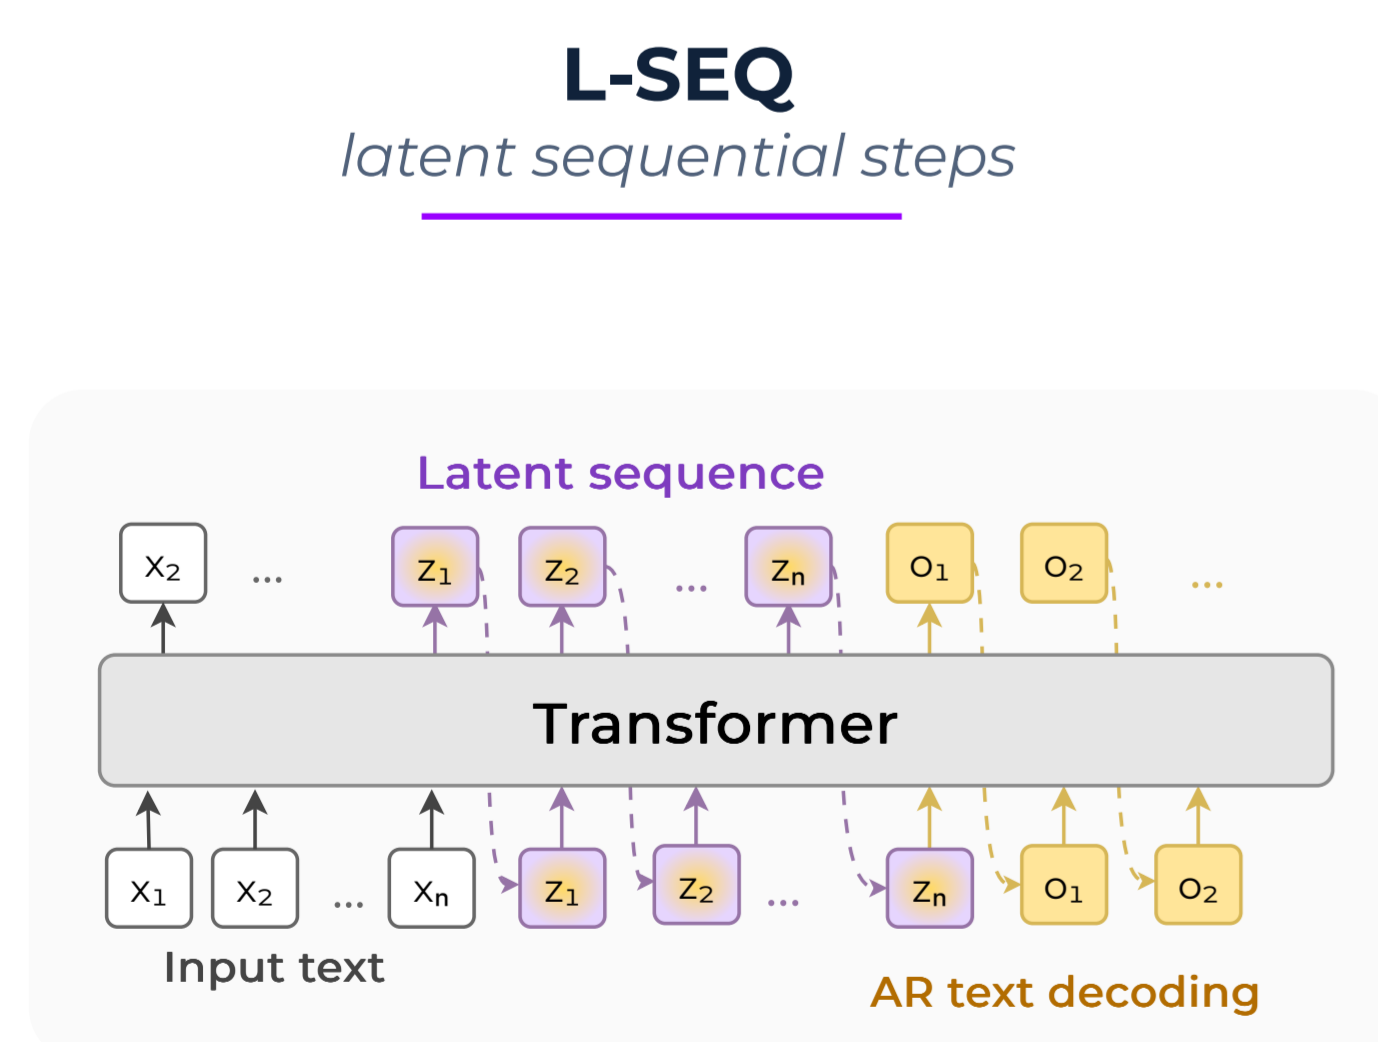
2. Latent Reasoning Overview

Explicit CoT vs. Latent Reasoning

| Dimension | Explicit CoT | Latent Reasoning |
|-------------------------------|---|--|
| Efficiency | Verbose, high cost and latency | Compact, faster, resource-efficient |
| Interpretability | Directly observable but may be unfaithful | Indirect |
| Control | Easier to guide, inspect, and correct | No built-in mechanism to steer or verify |
| Diversity | One trajectory | Multiple alternatives |
| Supervision Granularity | Explicit, step-aware supervision | Guided by latent objectives |
| Alignment with Human Thinking | Explains thoughts aloud | Thinking silently |

Why?

- Removes the language bottleneck and the token-level serialization cost
- Enables richer, **higher-bandwidth** internal communication than text
- May support **planning, exploration, compositional reasoning, out-of-distribution generalization**, and **agentic reasoning** more naturally
- Possible modular support for **Multilinguality** and **Multimodality**



- Compressed** step-by-step reasoning
- Keeps **explicit** latent **intermediate states**
- A natural notion of latent "steps"
- Promising for **planning** and **exploration**

L-SEQ and L-LOOP may be **complementary** rather than competing paradigms.

Main Limitations and Open Challenges

- Supervision is harder:** less directly supervisable than explicit CoT.
- Interpretability is indirect:** latent states are not directly readable.
- Field **lacks robust benchmarks** on efficiency, effectiveness, and faithfulness
- Emerging planning benefits remain unproven:** does it truly enable better planning, or mainly compress token-level reasoning?

3. Research Questions

When is latent reasoning useful?

Tasks, settings, and regimes where latent reasoning outperforms explicit CoT

How efficient is it?

Scaling with training data and test-time reasoning budget

How interpretable is it?

What is lost and what can still be recovered when reasoning moves into latent space

Structure to bridge the gap?

Can modularity, hierarchy, and interventions improve both efficiency & interpretability

4. Path to Efficient Neuro-Symbolic Reasoning

Long-term Roadmap

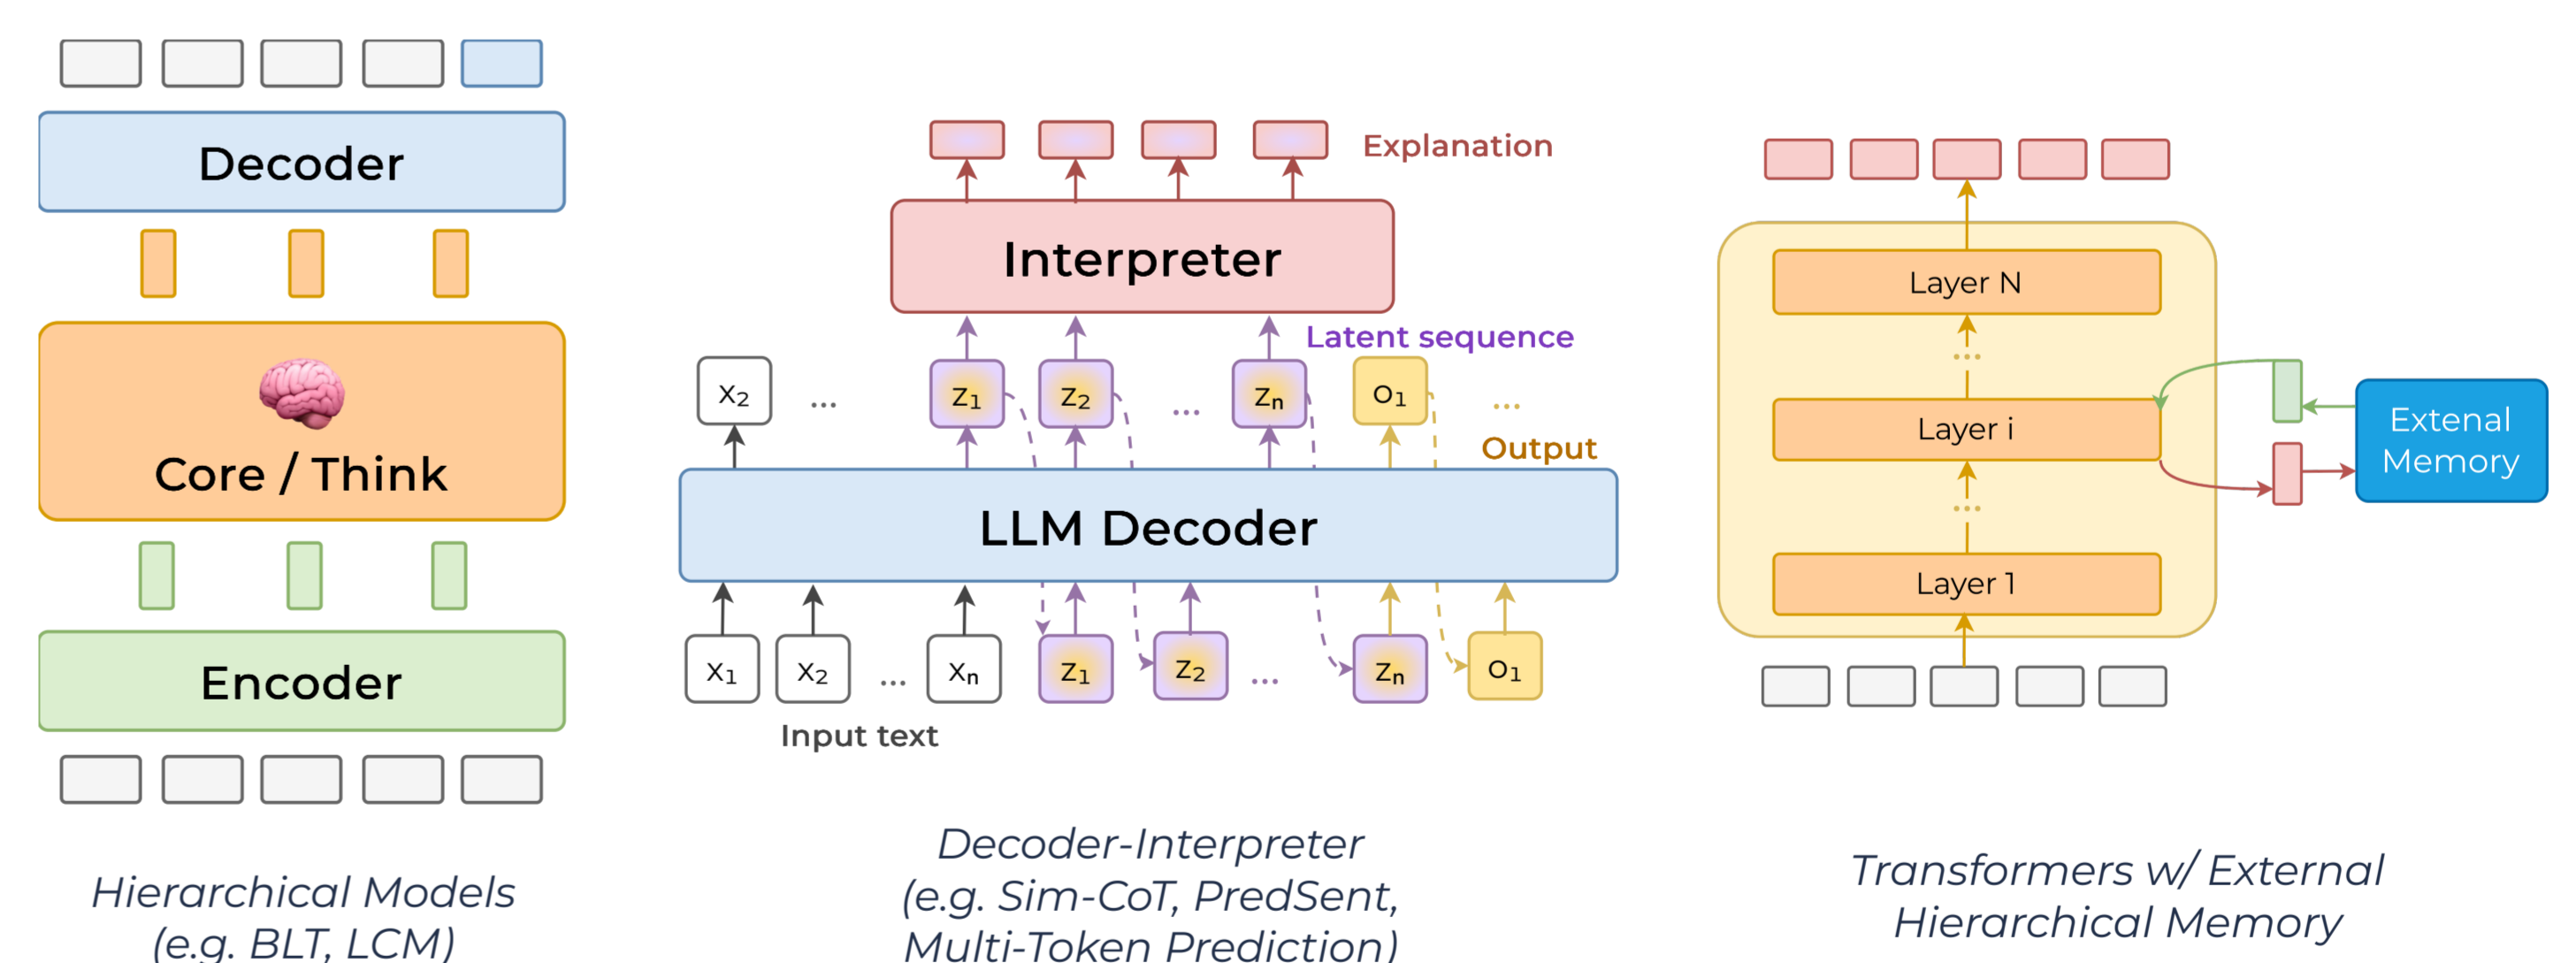


Target properties

- reason in continuous space** rather than only in language
- expose **meaningful internal structure**
- support **interventions** and control
- operate over reusable concepts, subgoals, or symbolic abstractions**

Possible directions (next 1-year)

- Benchmark Explicit CoT, L-SEQ, and L-LOOP on **performance, efficiency, generalization, and interpretability.**
- Study differences and complementarities between L-SEQ and L-LOOP.
- Probe latent steps for meaningful reasoning phases
 - intermediate variables, subgoals, etc.
- Explore **modularity, hierarchy, memory, and supervision** as routes toward structured latent reasoning.



References

- Chen, X., et al. "Reasoning beyond language: A comprehensive survey on latent chain-of-thought reasoning." arXiv:2505.16782 (2025).
- Li, J., et al. "Implicit reasoning in large language models: A comprehensive survey." arXiv:2509.02350 (2025).
- Pan, L., et al. "Opening the Black Box: A Survey on the Mechanisms of Multi-Step Reasoning in Large Language Models." arXiv:2601.14270 (2026).
- Zhu, L., et al. "A Survey on Latent Reasoning." arXiv:2507.06203 (2025).
- Hao, S., et al. "Training Large Language Models to Reason in a Continuous Latent Space." COLM 2025.
- Hwang, H., et al. "Latent Reasoning via Sentence Embedding Prediction." arXiv:2505.22202 (2025).
- Wei, X., et al. "SIM-CoT: Supervised Implicit Chain-of-Thought." arXiv:2509.20317 (2025).
- Pagnoni, A., et al. "Byte latent transformer: Patches scale better than tokens." ACL 2025.
- Barraut, Loïc, et al. "Large concept models: Language modeling in a sentence representation space." arXiv:2412.08821 (2024).
- Koishikenov, Y., Lipani, A. & Cancedda, N. "Encode, Think, Decode: Scaling test-time reasoning with recursive latent thoughts." arXiv:2510.07358 (2025).
- Zhu, R., et al. "Scaling latent reasoning via looped language models." arXiv:2510.25741 (2025).
- Pouransari, H., et al. "Pretraining with hierarchical memories: separating long-tail and common knowledge." arXiv:2510.02375 (2025).

Four Ways to Express Reasoning

Problem. A team starts with 3 members. They recruit 5 new members. Then each current member recruits 2 additional people. How many people are there now on the team?

answer

24

LLM Textual Chain-of-Thought

Okay, let's solve this math problem step by step.

So initially, there are **3** team members.

The team recruits 5 new members, so the total becomes **$3 + 5 = 8$** .

Now, each of these 8 current members recruits 2 additional people. That means the number of additional recruits is **$8 \times 2 = 16$** .

Finally, we add these 16 new recruits to the existing 8 members: **$8 + 16 = 24$** .

So the final answer is 24.

Human Compressed Mental Reasoning

Start with 3
 $3 + 5 = 8$

Now 8 people each recruit 2:
 $8 \times 2 = 16$

Total:
 $8 + 16 = 24$

Symbolic State Updates

```
# initial state  
members = 3
```

```
# recruit 5  
members += 5
```

```
# each current member recruits 2  
members += 2*members
```

```
# result  
print(members)
```

Computation Tree

