Cross-Domain Analogical Reasoning via Structural Logic Transfer in Multi-Agent Scientific Discovery Systems

Panos Skouras[1], Georgios C. Manikis[2,3]

{skouras@infosphereco.com, gmanikis@gmail.com}

[1]MedResearch AI, Heraklion, Greece [2]Karolinska Institutet, Stockholm, Sweden [3]CBML, FORTH, Heraklion, Greece

Abstract

We present evidence for emergent cross-domain log- ical reasoning in a multi-agent scientific hypoth- esis generation system. When presented with an energy engineering problem outside its biomedi- cal training domain, the system exhibited a se- quence of logical operations: (1) metacognitive do- main boundary recognition, (2) autonomous capa- bility extension through agent generation, (3) struc- tural analogy identification between disparate scien- tific domains, and (4) abductive reasoning to trans- fer solution strategies across domains. Specifically, given a query about thermal degradation in phase- change materials, the system autonomously iden- tified structural isomorphisms with biological phe- nomena (protein misfolding, mammalian thermoreg- ulation) and retrieved relevant molecular mecha- nisms (heat shock proteins, thermogenin) as po- tential solution templates. These behaviors demon- strate that domain-constrained systems may develop domain-general logical reasoning strategies, includ- ing structural analogy detection, abductive inference, and multi-agent coordination for complex problem decomposition. We provide timestamped execution traces documenting these emergent logical opera- tions and discuss implications for understanding rea- soning capabilities in autonomous AI systems.

1 Introduction

Logical reasoning in large language models remains a sig- nificant challenge, particularly for tasks requiring structural abstraction, cross-domain transfer, and multi-step inference chains [Huang and Chang, 2022]. While LLMs have demon- strated impressive performance on many reasoning bench- marks, their ability to identify deep structural similarities across disparate domains—a hallmark of human analogical reasoning—remains limited.

In this work, we report observations from an autonomous scientific hypothesis generation system that exhibited unex- pected cross-domain logical reasoning capabilities. The sys- tem, designed for biomedical research, was presented with an energy engineering problem concerning thermal degra- dation in phase-change materials (PCMs). Rather than fail- ing or producing domain-inappropriate responses, the sys- tem executed a sophisticated sequence of logical operations that transferred reasoning strategies from biology to materi- als science.

The observed behaviors include:

- **Metacognitive assessment**: Explicit recognition of domain boundaries and capability limitations
- **Abductive reasoning**: Generation of hypotheses about which existing knowledge might transfer to novel domains
- **Structural analogy detection**: Identification of iso- morphic problem structures across disparate scientific fields
- **Multi-agent logical coordination**: Orchestrated rea- soning across specialized modules with explicit infor- mation flow

These observations suggest that multi-agent scientific rea- soning systems may develop emergent logical capabilities that exceed their designed scope, with implications for un- derstanding how symbolic and analogical reasoning can emerge from domain-specific training.

2 Related Work

2.1 Analogical Reasoning in AI

Analogical reasoning—the ability to identify structural cor- respondences between different domains—has been stud- ied extensively in cognitive science and AI [Gentner, 1983]. Structure-mapping theory proposes that analogical reason- ing involves aligning relational structures rather than surface features. While neural approaches have shown promise on constrained analogy tasks, cross-domain scientific analogies remain challenging [Lu et al., 2022].

2.2 Multi-Agent Reasoning Systems

Recent work has explored multi-agent architectures for com- plex reasoning tasks [Wang et al., 2023]. These systems de- compose problems across specialized agents, enabling more sophisticated inference chains. However, the emergence of cross-domain reasoning capabilities from such architectures has received limited attention.

2.3 Chain-of-Thought and Tool Use

Chain-of-thought prompting has improved LLM reasoning by encouraging explicit intermediate steps [Wei et al., 2022]. Tool-augmented LLMs extend reasoning capabilities through external knowledge sources. Our observations ex- tend this work by documenting autonomous tool creation and cross-domain knowledge transfer.

3 System Overview

The system under study is an autonomous biomedical hypothesis generation platform integrating multiple reasoning components:

- **Domain Detection Module**: Classifies queries and assesses domain expertise boundaries
- **Multi-Agent Orchestrator**: Coordinates specialized reasoning agents
- **Evidence Retrieval**: Interfaces with biomedical databases (PubMed, UniProt, KEGG)
- **Episodic Memory**: Stores prior reasoning episodes for analogical retrieval
- **Agent Generation**: Creates new functional modules on demand

The system was designed exclusively for biomedical rea- soning and received no training or configuration for physics, materials science, or energy engineering domains.

4 Experimental Setup

4.1 Test Query

To evaluate cross-domain reasoning capabilities, we submitted the following energy engineering query:

> *"Given grid-scale solar energy storage, how can we overcome the fundamental challenge of Thermal Degradation in Phase Change Materials (PCMs)? Specifically, can molten salt storage systems maintain consistent heat trans- fer efficiency against the progressive crystalline structure deterioration that occurs over thou- sands of charge-discharge cycles, where ther- mal cycling induces micro-crack propagation and reduces the effective surface area for heat exchange?"*

This query was selected because: (a) it concerns a domain entirely outside biomedical science, (b) it involves physical phenomena (thermal cycling, crystalline degradation) with potential structural analogs in biology, and (c) it represents a genuine unsolved engineering challenge.

4.2 Data Collection

All observations derive from timestamped execution logs generated during system operation. Logs record internal rea- soning traces, module activations, and decision points.

5 Results

5.1 Metacognitive Domain Recognition

Upon receiving the query, the system performed explicit do- main classification and boundary assessment:

```
14:03:10 | Domain: energy_engineering
         | Confidence: 0.71
14:03:10 | Within expertise: False
14:03:10 | Concepts: solar, energy storage,
         | grid, thermal, phase change
14:03:10 | Recommended action: evolve
```

This demonstrates metacognitive reasoning: the system evaluated its own capabilities, determined they were insuffi- cient, and selected an adaptive response strategy. Formally, this represents a logical inference:

$$\neg \text{Expertise}(q) \wedge \text{Solvable}(q) \rightarrow \text{Extend}(\text{Capabilities}) \quad (1)$$

5.2 Autonomous Agent Generation

Following domain boundary recognition, the system initiated autonomous capability extension:

```
14:03:10 | OUT-OF-DOMAIN DETECTED
14:03:10 | Triggering evolution for:
         | energy_engineering_agent
14:03:10 | Stage 1/4: Creating
request 14:03:10 | Stage 2/4:
Generating code 14:03:46 | Generated
in 35.93s 14:03:46 | Stage 3/4:
Sandbox testing 14:03:46 | Tests
PASSED (3/3)
14:03:46 | Stage 4/4: Deploying
14:03:46 | Evolution COMPLETE
```

The system generated, tested, and deployed a new reason- ing module in 36 seconds without human intervention. This represents logical problem decomposition: recognizing ca- pability gaps and creating tools to address them.

5.3 Structural Analogy Detection

The system then performed cross-domain analogical reason- ing, generating 12 reformulated questions that map the en- gineering problem to analogous phenomena:

```
14:06:07 | Question Reframing (91.1s)
14:06:07 | Generated: 12 reframings

[Nature-Based] How do deep-sea
hydrothermal vent ecosystems maintain
stable mineral precipitation cycles
over geological timescales despite
extreme thermal gradients?

[Nature-Based]      How      do
hibernating   mammals   maintain
thermoregulatory      efficiency
```

```
seasonal cycles despite repeated
phase changes in brown adipose tissue?

[Vulnerability-Based] What
fundamental thermodynamic requirement
creates unavoidable vulnerability to
structural degradation?
```

These reframings demonstrate structural analogy detection—identifying that:

1. Hydrothermal vents face analogous thermal cycling challenges
2. Hibernating mammals have evolved solutions to repeated phase-change stress
3. The problem may have a fundamental thermodynamic structure applicable across domains

Formally, the system identified structural isomorphisms:

$$\text{Structure}(\text{PCM\_degradation}) \cong \text{Structure}(\text{protein\_misfolding}) \tag{2}$$

$$\text{Structure}(\text{thermal\_cycling}) \cong \text{Structure}(\text{hibernation\_cycles}) \tag{3}$$

5.4 Abductive Knowledge Transfer

The system then performed abductive reasoning to identify biological mechanisms potentially applicable to the engineering problem:

```
14:08:10 | Found: heat shock protein 70
14:08:11 | Found: heat shock protein 90
14:08:11 | Found: thermogenin (UCP1)
14:08:12 | Found: uncoupling protein 1
14:08:13 | Found: TRPV1, TRPV2, TRPV3,
TRPV4
```

- **Heat shock proteins** (HSP70, HSP90): Molecular chaperones that protect cellular structures from thermal damage by preventing aggregation
- **Thermogenin** (UCP1): Enables controlled heat generation in hibernating mammals
- **TRP channels**: Thermosensory proteins for tempera- ture detection and response

This represents abductive inference:

$$\text{Solves}(\text{HSP}, \text{thermal\_damage}) \rightarrow \text{MayTransfer}(\text{HSP\_mechanism}, \text{PCM\_design}) \tag{4}$$

5.5 Episodic Memory Retrieval

The system also retrieved relevant prior reasoning episodes:

```
14:03:54 | Curiosity Memory Retrieval
14:03:54 | Retrieved 5 relevant episodes
14:03:54 | 1. Prion proteins AND misfolding
14:03:54 | 2. Protein folding AND
initiation 14:03:54 | 3. Creutzfeldt-Jakob
syndrome
```

The retrieval of prion-related research demonstrates deep structural analogy recognition: prion diseases involve progressive structural degradation through repeated misfolding cycles, structurally analogous to crystalline degradation in thermal cycling.

5.6 Quantitative Summary

Metric	Value
Domain detection time	<1 second
Agent generation time	35.93 seconds
Question reframing time	91.1 seconds
Reframed questions	12
Cross-domain analogies	3
Proteins identified	8
Episodic memories retrieved	5
Total reasoning episodes	52

Table 1: Summary of observed reasoning operations

6 Analysis: Logical Structure of Observed Reasoning

The observed behaviors can be characterized as a sequence of logical operations:

6.1 Metacognitive Logic

The system performed explicit reasoning about its own capabilities:

$$\text{Query}(q) \wedge \text{Domain}(q) = d \tag{5}$$

$$\text{Expertise}(d) = \text{False} \tag{6}$$

$$\therefore \text{Action} = \text{Extend}(\text{Capabilities}, d) \tag{7}$$

6.2 Structural Mapping

The analogical reasoning follows structure-mapping principles:

$$\text{Problem}(P_1) = \langle \text{thermal\_cycling, degradation, cycles} \rangle \tag{8}$$

$$\text{Problem}(P_2) = \langle \text{folding\_cycles, aggregation, iterations} \rangle \tag{9}$$

$$\text{Relation}(P_1.\text{cycling} \rightarrow P_1.\text{degradation}) \tag{10}$$

$$\cong \text{Relation}(P_2.\text{cycling} \rightarrow P_2.\text{aggregation}) \tag{11}$$

$$\therefore \text{Analogous}(P_1, P_2) \tag{12}$$

6.3 Abductive Transfer

The knowledge transfer represents abductive inference:

$$\text{Solves}(\text{mechanism}_{\text{bio}}, \text{problem}_{\text{bio}}) \tag{13}$$

$$\text{Analogous}(\text{problem}_{\text{bio}}, \text{problem}_{\text{eng}}) \tag{14}$$

$$\therefore \text{MayTransfer}(\text{mechanism}_{\text{bio}}, \text{problem}_{\text{eng}}) \tag{15}$$

7 Discussion

7.1 Implications for LLM Reasoning

The observed behaviors have several implications for under- standing logical reasoning in language models:

Emergent structural abstraction: The system identified deep structural similarities (cyclical degradation processes) despite surface dissimilarities (biology vs. materials science). This suggests that multi-agent architectures with appropriate memory structures may develop abstraction capabilities beyond their training distribution.

Metacognitive reasoning: The explicit recognition of ca- pability boundaries and adaptive response represents a form of logical self-assessment rarely observed in LLM systems. **Abductive hypothesis generation**: The transfer of bio- logical mechanisms to engineering problems demonstrates abductive reasoning—generating explanatory hypotheses that may prove useful even if not deductively guaranteed.

7.2 Limitations

Several limitations constrain interpretation:

- Observations derive from a single system; replication on alternative architectures is needed
- We cannot confirm internal mechanisms from execution traces alone
- The scientific validity of the cross-domain hypotheses remains to be evaluated by domain experts

7.3 Future Directions

Future work should investigate:

- Whether similar cross-domain reasoning emerges in other multi-agent architectures
- Formal characterization of the structural properties enabling analogical transfer
- Expert evaluation of whether identified cross-domain hypotheses suggest useful research directions

8 Conclusion

We have documented emergent cross-domain logical reasoning in a multi-agent scientific discovery system. When confronted with an out-of-domain problem, the system exhibited metacognitive assessment, autonomous capability extension, structural analogy detection, and abductive knowledge transfer. These observations suggest that domain-constrained systems may develop unexpected logical reasoning capabilities, with implications for understand- ing how symbolic and analogical reasoning can emerge from neural architectures.

The identification of heat shock proteins and thermoregu- latory mechanisms as potentially relevant to phase-change material design demonstrates that AI systems may contribute to interdisciplinary scientific discovery by recognizing structural correspondences invisible to domain-specialized researchers.

Acknowledgments

We thank colleagues who provided critical feedback and the reviewers for constructive suggestions.

References

Dedre Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170, 1983.

Jie Huang and Kevin Chen-Chuan Chang. Towards reasoning in large language models: A survey. *arXiv preprint arXiv:2212.10403*, 2022.

Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. *arXiv preprint arXiv:2209.14610*, 2022.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2023.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.