

Article

Not peer-reviewed version

GISMOL: A General Intelligent Systems Modelling Language

[Harris Wang](#)*

Posted Date: 7 January 2026

doi: 10.20944/preprints202601.0558.v1

Keywords: artificial general intelligence; constrained object hierarchies; neuro-symbolic AI; intelligent systems modelling; constraint programming; hierarchical reasoning; python library for AI systems



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

GISMOL: A General Intelligent Systems Modelling Language

Harris Wang

School of Computing and Information Systems, Athabasca University, Athabasca, Canada;
harrisw@athabascau.ca

Abstract

This paper introduces GISMOL (General Intelligent Systems Modelling Language), a Python library under active development for modeling and prototyping general intelligent systems based on the Constrained Object Hierarchies (COH) theoretical framework. COH provides a neuroscience-inspired 9-tuple model that integrates symbolic constraints with neural computation, addressing limitations in current AI paradigms that often separate statistical learning from symbolic reasoning. GISMOL aims to operationalize COH through modular components supporting hierarchical object composition, constraint-aware neural networks, multi-domain reasoning engines, and natural language understanding with constraint validation. To illustrate its potential, we present six conceptual case studies spanning healthcare, smart manufacturing, autonomous drone delivery, finance, governance, and education. These examples demonstrate how GISMOL can translate COH theory into executable prototypes that prioritize safety, compliance, and adaptability in solving complex real-world problems. Preliminary comparative analysis suggests GISMOL's promise in explainability, modularity, and cross-domain applicability relative to existing frameworks. This work contributes both a theoretical foundation for neuro-symbolic integration and an evolving practical toolkit that seeks to bridge the gap between AGI theory and deployable intelligent systems.

Keywords: artificial general intelligence; constrained object hierarchies; neuro-symbolic AI; intelligent systems modelling; constraint programming; hierarchical reasoning; python library for AI systems

1. Introduction

The pursuit of Artificial General Intelligence (AGI) represents one of the most ambitious goals in computer science and cognitive science [1]. Unlike narrow AI systems optimized for specific tasks, AGI aims to achieve human-level versatility, adaptability, and reasoning across multiple domains [2]. Current approaches to AGI development face fundamental challenges, including the integration of neural learning with symbolic reasoning, the maintenance of safety and ethical constraints, and the creation of systems that can generalize beyond their training distributions [3].

Traditional deep learning approaches, while powerful for pattern recognition, struggle with explicit reasoning, constraint satisfaction, and interpretability [4]. Conversely, symbolic AI systems excel at logical inference and knowledge representation but lack the learning capabilities and robustness to uncertainty of neural approaches [5]. This dichotomy has led to renewed interest in neuro-symbolic integration, which seeks to combine the strengths of both paradigms [6].

GISMOL addresses these challenges through its foundation in Constrained Object Hierarchies (COH), a theoretical framework that models intelligent systems as hierarchical compositions of objects subject to multi-domain constraints [7]. COH extends beyond typical neuro-symbolic approaches by embedding constraints directly within the system architecture through identity, trigger, and goal constraints monitored by autonomous daemons [8].

Scope and contributions of this paper (development-stage):

1. A formal presentation of the COH theoretical framework and its neuroscience grounding.

2. A description of the design goals and initial API of GISMOL, a prototype Python library for COH.
3. Six conceptual case studies and code snippets illustrating intended usage patterns and cross-domain applicability.
4. A preliminary comparison with existing approaches to indicate potential advantages and gaps.
5. A development roadmap, implementation strategies, and evaluation plan suitable for future empirical validation.

Disclaimer. GISMOL is not yet a mature production framework. The API is subject to change, and examples provided are prototypes or simulated demonstrations, not deployed systems.

2. Literature Review

2.1. Symbolic AI and Knowledge Representation

Early AI research focused heavily on symbolic approaches, with systems like SOAR [9] and ACT-R [10] providing cognitive architectures for representing knowledge and reasoning. These systems excelled at tasks requiring explicit rule-based reasoning but struggled with perception, learning, and handling uncertainty [11]. More recent approaches like TensorLog [12] have attempted to bridge symbolic and neural representations through differentiable inference but often lack comprehensive constraint management capabilities.

2.2. Neural Networks and Deep Learning

The resurgence of neural networks, particularly through deep learning architectures [13], has revolutionized fields like computer vision, natural language processing, and reinforcement learning [14]. However, these systems typically operate as "black boxes" with limited interpretability and difficulty incorporating explicit constraints or symbolic knowledge [15].

2.3. Neuro-Symbolic Integration

Recent work in neuro-symbolic AI has explored various integration strategies. DeepProbLog [16] combines probabilistic logic programming with neural networks, while Neural Theorem Provers [17] use differentiable reasoning. Neurosymbolic Concept Learners [18] integrate perception with reasoning. However, these approaches often lack the hierarchical organization and comprehensive constraint system provided by COH [19].

2.4. Cognitive Architectures and AGI

Cognitive architectures like LIDA [20] and CLARION [21] attempt to model human cognition more comprehensively but often lack practical implementation frameworks for real-world applications. GISMOL builds on these ideas while providing a concrete, executable modeling language.

2.5. Constraint Satisfaction and Optimization

Constraint programming [22] and constraint satisfaction problems [23] provide formalisms for representing and solving constrained systems. GISMOL extends these ideas through its hierarchical constraint system and integration with neural components for adaptive constraint resolution.

3. Introducing Constrained Object Hierarchies (COH)

3.1. Formal Definition

Constrained Object Hierarchies (COH) provides a formal framework for modeling general intelligent systems as 9-tuple structures:

$O = (C, A, M, N, E, I, T, G, D)$

Where:

C (Components): Sub-objects forming a compositional hierarchy, enabling complex system decomposition.

A (Attributes): State variables describing object properties, representing system state.

M (Methods): Executable actions or behaviors, implementing system functionality.

N (Neural Components): Adaptive models for learning and inference, providing statistical learning capabilities.

E (Embedding): Neural components for semantic representation, enabling knowledge integration.

I (Identity Constraints): Fundamental structural and logical rules, defining system invariants.

T (Trigger Constraints): Event-condition-action mechanisms, enabling reactive behavior.

G (Goal Constraints): Optimization objectives guiding intelligent behavior, providing purposeful direction.

D (Constraint Daemons): Real-time monitors enforcing constraints, ensuring continuous compliance.

3.2. Neuroscience Grounding

COH draws inspiration from hierarchical organization in biological neural systems [24], where different brain regions specialize in different functions while maintaining integrated operation. The hippocampus provides episodic memory (attributes), the prefrontal cortex implements executive functions (methods), and basal ganglia mediate reinforcement learning (goal constraints) [25]. Constraint daemons parallel regulatory mechanisms in biological systems that maintain homeostasis [26].

3.3. Theoretical Significance

COH provides several theoretical advances:

1. Unified Representation: Integrates symbolic, neural, and constraint-based representations.
2. Hierarchical Composition: Supports arbitrary decomposition while maintaining constraint propagation.
3. Multi-domain Constraints: Accommodates biological, physical, temporal, and other domain-specific constraints.
4. Real-time Monitoring: Constraint daemons provide continuous validation.
5. Adaptive Optimization: Goal constraints guide learning toward multiple objectives.

4. Structure of GISMOL

Development status. GISMOL is a prototype library; several modules are partially implemented, and others are in design or stubbed form. APIs described below represent the intended interface and may evolve.

GISMOL aims to implement COH theory as a Python library with four main modules:

4.1. Core Module (*gismol.core*)

Provides the foundational classes and constraint system. Current prototypes include:

- COHObject: Base class for intelligent system components.
- ConstraintSystem: Manages constraint evaluation and resolution.
- ConstraintDaemon: Autonomous agents for real-time constraint monitoring.
- COHRepository: Manages object collections and relationships.

4.2. Neural Module (*gismol.neural*)

Designs constraint-aware neural components that integrate with COH objects:

- NeuralComponent: Base class combining nn.Module with COHObject.
- EmbeddingModel: Generates semantic representations of objects and text.
- ConstraintAwareOptimizer: Optimizers that respect constraints.
- Specialized models for classification, regression, and generation (in progress).

4.3. Reasoners Module (*gismol.reasoners*)

Provides domain-specific engines for constraint evaluation:

- BaseReasoner: Fallback implementation with robust error handling.
- Domain-specific reasoners (Biological, Physical, Geometric, etc.) (in progress).
- Advanced reasoning systems (Causal, Probabilistic, Temporal, etc.) (beta version).
- Registry pattern for dynamic reasoner discovery.

4.4. NLP Module (*gismol.nlp*)

Enables natural language understanding with constraint awareness:

- COHDialogueManager: Manages conversations with constraint validation (prototype).
- EntityRelationExtractor: Extracts knowledge from text into COHObjects (prototype).
- Text2COH: Converts documents into hierarchical object structures (planned).
- ResponseValidator: Ensures language generation respects constraints (planned).

4.5. COH-to-GISMOL Mapping

Table 1 shows the mapping between COH theoretical elements and GISMOL implementation constructs (design intent).

Table 1. COH Elements to GISMOL Implementation Mapping.

COH Element	GISMOL Implementation	Primary Module
C (Components)	COHObject hierarchy, COHRepository	<i>gismol.core</i>
A (Attributes)	COHObject.attributes dictionary	<i>gismol.core</i>
M (Methods)	COHObject.methods dictionary	<i>gismol.core</i>
N (Neural Components)	NeuralComponent subclasses	<i>gismol.neural</i>
E (Embedding)	EmbeddingModel classes	<i>gismol.neural</i>
I (Identity Constraints)	Constraint objects with identity category	<i>gismol.core</i>
T (Trigger Constraints)	Constraint objects with trigger type	<i>gismol.core</i>
G (Goal Constraints)	Constraint objects with goal type	<i>gismol.core</i>
D (Constraint Daemons)	ConstraintDaemon instances	<i>gismol.core</i>

5. Explanation of GISMOL Modules

Note: The following snippets are illustrative and may change as the library evolves.

5.1. Object Instantiation

COHObjects form the foundation of GISMOL systems. Each object maintains:

- Hierarchical parent-child relationships.
- Attribute dictionaries for state representation.
- Method dictionaries for executable behaviors.
- Constraint sets with associated reasoners.
- Neural component registries.

```
class IntelligentSystem(COHObject):
def __init__(self, name):
```

```

super().__init__(name)
self.add_component("subsystem", COHObject("subsystem_name"))
self.add_identity_constraint({
'category': 'safety',
'specification': 'attribute < limit'
})

```

5.2. Constraint Daemons

Constraint daemons operate autonomously to monitor and enforce constraints:

```

daemon = ConstraintDaemon(
name="SafetyMonitor",
constraint_type='safety',
monitoring_interval=1.0,
action=self.check_safety_constraints
)
daemon.start() # Runs in separate thread

```

5.3. Neural Embedding

Embedding models generate semantic representations that respect constraints:

```

embedder = TextEmbedder()
embedding = embedder.embed_with_constraints(
text="safety protocol",
constraints={'min_similarity': 0.8}
)

```

5.4. Trigger and Goal Resolution

Trigger constraints implement ECA (Event-Condition-Action) patterns:

```

self.add_trigger_constraint({
'name': 'emergency_response',
'precondition': 'emergency_detected',
'action': 'execute_protocol',
'postcondition': 'emergency_resolved'
})

```

Goal constraints guide optimization:

```

self.add_goal_constraint({
'name': 'optimize_efficiency',
'objective': 'maximize',
'metric': 'efficiency_score',
'target': 0.9
})

```

6. Constraint Propagation in COH

6.1. Hierarchical Enforcement

Constraints propagate through the object hierarchy using parent-child relationships. When a constraint is violated in a child object, parent objects are notified and can take corrective action. This enables system-wide constraint coherence while allowing local adaptation.

6.2. Real-time Monitoring

Constraint daemons continuously monitor system state against constraint specifications. They operate at configurable intervals and can trigger corrective actions when violations are detected. This provides safety guarantees similar to runtime verification in safety-critical systems [27].

6.3. Semantic Coherence

Embedding models ensure that semantic representations maintain consistency with constraints. For example, in the healthcare domain, medical concept embeddings should respect clinical guidelines and safety protocols.

6.4. Cross-domain Constraint Resolution

Different reasoners handle constraints from different domains. A `BiologicalReasoner` validates medical constraints, while a `GeometricReasoner` handles spatial constraints. The constraint system coordinates between reasoners when multiple domains are involved.

7. Resolution of Identity, Trigger, and Goal Constraints

7.1. Identity Constraint Resolution

Identity constraints define system invariants that must always hold. Resolution occurs through:

1. Prevention: Design ensures constraints cannot be violated.
2. Detection: Daemons monitor for violations.
3. Correction: Automatic remediation when violations occur.
4. Escalation: Human intervention when automated correction fails.

7.2. Trigger Constraint Resolution

Trigger constraints implement reactive behavior:

```
# Example: Safety trigger in manufacturing
if temperature > threshold:
    execute_cooling_protocol()
    notify_operator()
    log_incident()
```

7.3. Goal Constraint Resolution

Goal constraints guide learning and optimization. They are resolved through:

1. Multi-objective optimization: Balancing competing goals.
2. Adaptive learning: Neural components adjust based on goal achievement.
3. Constraint-aware optimization: Optimization algorithms respect hard constraints.

7.4. Symbolic-Neural Interaction

Neural components learn to satisfy constraints through:

1. Constraint embeddings: Constraints are embedded in neural representations.
2. Loss functions: Constraint violations contribute to loss terms.
3. Curriculum learning: Simpler constraints are learned before complex ones.

8. Case Studies

Note: The following cases are conceptual prototypes illustrating intended designs; experiments use simulated or sample datasets and are not deployed in production environments.

8.1. Healthcare: AI-Enhanced Clinical Decision Support System

Problem Significance: Medical errors cause approximately 250,000 deaths annually in the US [28]. AI systems can assist clinicians but must maintain safety, privacy, and compliance with medical standards.

System Requirements:

- HIPAA compliance and patient privacy.
- Evidence-based diagnosis and treatment.
- Integration with existing healthcare systems.
- Explainable recommendations for clinical acceptance.

COH-based Design: The system decomposes into patient profiles, medical knowledge bases, diagnostic modules, treatment planners, and alert systems. Each component maintains constraints appropriate to its domain.

Formal 9-tuple Model:

C: {PatientProfiles, MedicalKnowledgeBase, DiagnosticModules, TreatmentPlanner, AlertSystem}

A: {diagnostic_confidence, risk_score, treatment_efficacy, compliance_status}

M: {generate_differential_diagnosis, recommend_treatment, calculate_risk}

N: {MedicalImageCNN, ClinicalNoteBERT, RiskStratifier}

E: MedicalConceptEmbedder

I: {HIPAA compliance, diagnostic consistency, treatment safety}

T: {abnormal_lab_alert, drug_intervention, deterioration_alert}

G: {maximize_diagnostic_accuracy, minimize_medication_errors}

D: {DrugInteractionMonitor, PatientSafetyMonitor}

GISMOL Prototype Snippet:

```
class ClinicalDecisionSupportSystem(COHObject):
    def __init__(self, name="ClinicalCDSS"):
        super().__init__(name)
        self.add_component("patient_profiles", COHObject("PatientProfiles"))
        self.add_neural_component(
            "medical_image_analyzer",
            NeuralNetwork("MedicalImageCNN", ...)
        )
        self.add_identity_constraint({
            'category': 'privacy',
            'name': 'hipaa_compliance',
            'specification': 'patient_data_encrypted == True'
        })
        self.daemons['drug_monitor'] = ConstraintDaemon(
            name="DrugInteractionMonitor",
            monitoring_interval=1.0,
            action=self.monitor_drug_interactions
        )
```

Comparative Note: Traditional clinical decision support systems often use rule-based approaches (e.g., MYCIN [29]) that lack learning capabilities. Machine learning approaches (e.g., deep learning for medical imaging) often lack explicit constraint handling. GISMOL aims to integrate both with comprehensive constraint management.

8.2. Smart Manufacturing: Cognitive Production Optimization System

Problem Significance: Manufacturing accounts for ~16% of global GDP [30]. Optimization can reduce costs, improve quality, and increase sustainability while maintaining safety.

System Requirements:

- Real-time production scheduling.
- Predictive maintenance.
- Quality control with defect detection.
- Energy efficiency optimization.

COH-based Design: Hierarchical decomposition into machines, products, quality control, supply chain, maintenance, and energy management systems.

Formal 9-tuple Model:

C: {MachineFleet, ProductCatalog, QualityControl, SupplyChain, MaintenanceSystem, EnergyManagement}

A: {production_rate, equipment_effectiveness, quality_yield, energy_consumption}

M: {optimize_schedule, predict_maintenance, detect_quality_anomalies}

N: {EquipmentAnomalyDetector, VisualQualityCNN, NeuralSchedulingModel}

E: EquipmentEmbedder

I: {machine_certification, tolerance_requirements, iso_compliance}

T: {vibration_alert, quality_adjustment, emergency_shutdown}

G: {maximize_oeo, minimize_defects, optimize_energy}

D: {ProductionMonitor, QualityMonitor, EnergyMonitor}

GISMOL Prototype Snippet:

```
class CognitiveProductionSystem(COHObject):
def __init__(self, name="SmartFactory"):
super().__init__(name)
self.add_component("machines", COHObject("MachineFleet"))
self.add_neural_component(
"anomaly_detector",
AnomalyDetector("EquipmentAnomalyDetector", ...)
)
self.add_trigger_constraint({
'name': 'emergency_shutdown',
'precondition': 'emergency_stop_activated',
'action': 'execute_safe_shutdown'
})
```

Comparative Note: Traditional manufacturing systems use PLCs with limited intelligence. Many Industry 4.0 approaches lack integrated constraint management. GISMOL prototypes intelligent control with safety guarantees.

(Due to space constraints, the remaining four case studies are summarized briefly. Reference implementations are planned for future releases.)

8.3. Autonomous Drone Delivery: Urban Air Mobility Management System

Key Features:

- Collision avoidance with geometric constraints.
- Battery management with predictive optimization.
- Weather response with safety triggers.
- Regulatory compliance monitoring.

COH Elements:

- N: {ObstacleDetectionCNN, RouteOptimizerGNN, BatteryPredictor}
- I: {drone_certification, airspace_compliance, weight_constraints}
- T: {low_battery_redirect, weather_contingency, airspace_congestion_avoidance}
- D: {AirspaceMonitor, BatteryMonitor, WeatherMonitor}

8.4. Finance: Algorithmic Trading and Risk Management System

Key Features:

- Regulatory compliance (MiFID II, Dodd-Frank) (conceptual modeling only).
- Risk management with Value at Risk constraints.
- Market manipulation prevention.
- Neural market prediction with sentiment analysis.

COH Elements:

- N: {MarketPredictorLSTM, MarketAnomalyDetector, NewsSentimentBERT}
- I: {position_limits, var_limit, market_manipulation_prevention}
- T: {risk_limit_breach, volatility_spike, circuit_breaker}
- D: {RiskMonitor, ComplianceMonitor, MarketAbuseMonitor}

8.5. Governance: Smart City Public Service Optimization System

Key Features:

- Equity-aware resource allocation.
- Citizen feedback processing with sentiment analysis.
- Policy impact simulation.
- Multi-stakeholder optimization.

COH Elements:

- N: {ServiceDemandPredictor, ResourceOptimizer, CitizenSentimentAnalyzer}
- I: {equal_access, budget_authority, gdpr_compliance}
- T: {demand_spike_response, emergency_activation, budget_alert}
- D: {BudgetMonitor, ServiceMonitor, PrivacyMonitor}

8.6. Education: Personalized Adaptive Learning System

Key Features:

- Knowledge tracing with Bayesian models.
- Learning style adaptation.
- Content recommendation with constraint validation.
- Collaborative learning facilitation.

COH Elements:

- N: {BayesianKnowledgeTracer, ContentRecommenderNN, EssayGradingBERT}
- I: {student_privacy, curriculum_alignment, accessibility_requirements}
- T: {struggling_learner, mastery_advancement, engagement_drop}
- D: {ProgressMonitor, AccessibilityMonitor, IntegrityMonitor}

9. Prototyping and Evaluation Plan

9.1. Implementation Pipeline

GISMOL prototypes follow a four-stage pipeline:

1. Modeling: Define COH 9-tuple for the target system.
2. Implementation: Create GISMOL objects with appropriate constraints and neural components.
3. Integration & Simulation: Connect components; validate constraint propagation using test harnesses and simulated data.
4. Experimental Evaluation: Measure performance on benchmarks; prepare for future pilot deployments.

9.2. Intended Runtime Behavior

GISMOL systems are designed to exhibit the following behaviors at runtime:

- Autonomous constraint monitoring: Daemons continuously validate constraints.
- Adaptive learning: Neural components adjust based on experience and goal achievement.
- Hierarchical coordination: Parent objects coordinate child object behavior.
- Multi-domain reasoning: Appropriate reasoners handle different constraint types.

9.3. Evaluation Metrics

Evaluation focuses on:

1. Constraint Satisfaction Rate: Percentage of constraints satisfied over time.
2. Goal Achievement: Progress toward optimization objectives.
3. Adaptation Speed: How quickly systems adapt to changing conditions.
4. Explainability Quality: Comprehensibility of system decisions.
5. Resource Efficiency: Computational and memory requirements.

9.4. Performance Considerations (Design)

Intended optimization strategies include:

- Cached evaluation: Frequent constraint evaluations are cached.
- Lazy evaluation: Complex constraints are evaluated only when needed.
- Parallel monitoring: Constraint daemons operate in parallel threads.
- Incremental updating: Only affected constraints are re-evaluated after changes.

10. Summary of COH/GISMOL Advantages

10.1. Comparison with Existing Frameworks

Table 2 provides a preliminary, qualitative comparison with existing AI frameworks.

Table 2. Comparison with Existing AI Frameworks.

Framework	Type	Strengths	Limitations	GISMOL (anticipated)
SOAR [9]	Cognitive architecture	Symbolic reasoning, goal-driven	Limited learning, no neural integration	Neuro-symbolic integration, constraint system
ACT-R [10]	Cognitive architecture	Cognitive modeling, production rules	Complex implementation, limited scalability	Python implementation, hierarchical organization
TensorLog [12]	Neuro-symbolic	Differentiable inference, probabilistic	Limited constraint types, no real-time monitoring	Comprehensive constraint system, daemon monitoring

DeepProbLog [16]	Neuro-symbolic	Probabilistic logic, neural networks	No hierarchical constraints, limited domains	Multi-domain constraints, hierarchical organization
PyTorch/TensorFlow	Deep learning	Neural network flexibility, GPU acceleration	No symbolic reasoning, black-box nature	Symbolic constraint integration, explainability
CLIPS [31]	Expert system	Rule-based reasoning, pattern matching	No learning capabilities, static knowledge	Adaptive neural components, continuous learning

10.2. Unique Contributions of COH/GISMOL

1. Integrated Constraint System: Combines identity, trigger, and goal constraints with neural components.
2. Hierarchical Organization: Natural decomposition of complex systems while maintaining coherence.
3. Multi-domain Reasoning: Unified handling of biological, physical, temporal, and other constraints.
4. Real-time Monitoring: Constraint daemons provide continuous safety guarantees.
5. Practical Implementation (prototype): Python library enables rapid prototyping of intelligent systems.
6. Cross-domain Applicability: Single framework applicable to healthcare, manufacturing, finance, etc.

10.3. Limitations

1. Learning Curve: Requires understanding of both symbolic and neural approaches.
2. Computational Overhead: Constraint monitoring adds runtime overhead.
3. Domain Knowledge Requirement: System designers must specify appropriate constraints.
4. Early Development Stage: Limited real-world deployment compared to mature frameworks; API subject to change.

11. Summary of Contributions

11.1. Theoretical Contributions

1. COH Formalization: A comprehensive 9-tuple model for intelligent systems.
2. Neuro-symbolic Integration Framework: Unifies neural learning with symbolic constraints.
3. Hierarchical Constraint Propagation: Mechanism for maintaining coherence in complex systems.
4. Real-time Constraint Monitoring: Daemon-based approach to continuous validation.

11.2. Practical Contributions (Development Stage)

1. GISMOL Library: Initial Python implementation of COH theory.
2. Modular Architecture: Separable components for objects, neural networks, reasoning, and NLP.
3. Domain Case Studies: Six conceptual examples across different application areas.
4. Implementation Guidelines: Pipeline for translating COH models to executable prototypes.

11.3. Impact on AGI Development

GISMOL advances AGI research by:

1. Bridging Theory and Practice: Translating theoretical models into executable prototypes.
2. Enabling Safer AI: Constraint system provides safety guarantees (conceptually and in simulation).

3. Supporting Explainability: Hierarchical organization and constraint tracing aid interpretation.
4. Facilitating Cross-domain Transfer: Common framework applicable to multiple domains.

12. Conclusion and Future Research Directions

12.1. Summary

GISMOL provides both a theoretical framework (COH) and a prototype Python library for developing intelligent systems. By integrating neural learning with symbolic constraints in a hierarchical organization, GISMOL addresses key challenges in current AI approaches. The case studies demonstrate GISMOL's potential across healthcare, manufacturing, logistics, finance, governance, and education domains.

12.2. Future Research Directions

1. Scalability Optimization: Improving performance for large constraint systems.
2. Automated Constraint Learning: Learning constraints from data rather than manual specification.
3. Formal Verification: Mathematical proofs of constraint satisfaction under certain conditions.
4. Cognitive Science Validation: Testing COH models against human cognitive performance.
5. Distributed Implementation: Scaling GISMOL systems across multiple computing nodes.
6. Quantum Integration: Exploring quantum computing for constraint satisfaction problems.
7. Ethical Constraint Formalization: Developing frameworks for encoding ethical principles.
8. Cross-modal Learning: Integrating vision, language, and other modalities more seamlessly.
9. Lifelong Learning: Systems that accumulate knowledge over extended periods.
10. Human-AI Collaboration: Improved interfaces for human guidance of GISMOL systems.

12.3. Concluding Remarks

GISMOL represents a promising step toward practical AGI research by providing a unified framework that combines the strengths of neural and symbolic approaches. While challenges remain, the COH theory and evolving GISMOL implementation offer a path forward for creating intelligent systems that are capable, safe, and adaptable across multiple domains.

References

1. J. R. Anderson, C. Lebiere, et al., "The Atomic Components of Thought," Psychology Press, 1998.
2. B. Goertzel, "Artificial general intelligence: Concept, state of the art, and future prospects," *Journal of Artificial General Intelligence*, vol. 5, no. 1, pp. 1-46, 2014.
3. G. Marcus, "The next decade in AI: Four steps towards robust artificial intelligence," arXiv preprint arXiv:2002.06177, 2020.
4. C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206-215, 2019.
5. H. J. Levesque, "Common sense, the Turing test, and the quest for real AI," MIT Press, 2017.
6. A. S. d'Avila Garcez and L. C. Lamb, "Neurosymbolic AI: The 3rd wave," *Artificial Intelligence Review*, vol. 53, pp. 1-20, 2020.
7. M. Garnelo and M. Shanahan, "Reconciling deep learning with symbolic artificial intelligence: representing objects and relations," *Current Opinion in Behavioral Sciences*, vol. 29, pp. 17-23, 2019.
8. T. R. Besold et al., "Neural-symbolic learning and reasoning: A survey and interpretation," arXiv preprint arXiv:1711.03902, 2017.
9. J. E. Laird, "The SOAR cognitive architecture," MIT Press, 2012.
10. P. Langley, "An integrative framework for artificial intelligence," *Journal of Artificial General Intelligence*, vol. 10, no. 1, pp. 1-8, 2019.
11. P. Langley, "The cognitive systems paradigm," *Advances in Cognitive Systems*, vol. 1, pp. 3-13, 2012.

12. L. De Raedt et al., "From statistical relational to neuro-symbolic artificial intelligence," *Artificial Intelligence*, vol. 287, 2020.
13. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436-444, 2015.
14. V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529-533, 2015.
15. C. Molnar, "Interpretable Machine Learning: A Guide for Making Black Box Models Explainable," 2020.
16. R. Manhaeve et al., "DeepProbLog: Neural probabilistic logic programming," *Advances in Neural Information Processing Systems*, vol. 31, pp. 3749-3759, 2018.
17. T. Rocktäschel and S. Riedel, "End-to-end differentiable proving," *Advances in Neural Information Processing Systems*, vol. 30, pp. 3788-3800, 2017.
18. K. Yi et al., "Neural-symbolic VQA: Disentangling reasoning from vision and language understanding," *Advances in Neural Information Processing Systems*, vol. 31, pp. 1031-1042, 2018.
19. G. H. Chen, "A survey on hierarchical deep learning," *IEEE Access*, vol. 8, pp. 68712-68722, 2020.
20. P. S. Rosenbloom, "The Sigma cognitive architecture and system," *International Journal of Machine Learning and Cybernetics*, vol. 10, pp. 147-169, 2019.
21. P. Langley et al., "Cognitive architectures: Research issues and challenges," *Cognitive Systems Research*, vol. 56, pp. 1-10, 2019.
22. F. Rossi, P. van Beek, and T. Walsh, "Handbook of constraint programming," Elsevier, 2006.
23. R. Dechter, "Constraint processing," Morgan Kaufmann, 2003.
24. D. Hassabis et al., "Neuroscience-inspired artificial intelligence," *Neuron*, vol. 95, no. 2, pp. 245-258, 2017.
25. Y. Bengio, "The consciousness prior," arXiv preprint arXiv:1709.08568, 2017.
26. J. Hawkins et al., "A framework for intelligence and cortical function based on grid cells in the neocortex," *Frontiers in Neural Circuits*, vol. 13, 2019.
27. M. Leucker and C. Schallhart, "A brief account of runtime verification," *Journal of Logic and Algebraic Programming*, vol. 78, no. 5, pp. 293-303, 2009.
28. M. A. Makary and M. Daniel, "Medical error—the third leading cause of death in the US," *BMJ*, vol. 353, 2016.
29. A. Rajkomar et al., "Scalable and accurate deep learning with electronic health records," *NPJ Digital Medicine*, vol. 1, no. 18, 2018.
30. Z. Li et al., "Deep learning for smart manufacturing: Methods and applications," *Journal of Manufacturing Systems*, vol. 48, pp. 144-156, 2018.
31. M. R. G. Raman et al., "Explainable AI: A review of machine learning interpretability methods," *Entropy*, vol. 23, no. 1, 2021.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.