

Article

Not peer-reviewed version

---

# Constrained Object Hierarchies as a Unified Framework for Artificial General Intelligence

---

[Harris Wang](#) \*

Posted Date: 22 January 2026

doi: 10.20944/preprints202601.1685.v1

Keywords: artificial general intelligence; Constrained Object Hierarchies; neural-symbolic integration; constraint satisfaction; hierarchical systems; multi-domain modeling; cognitive architecture; universal world model; agentic systems; jagged intelligence



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

# Constrained Object Hierarchies as a Unified Framework for Artificial General Intelligence

Harris Wang

School of Computing and Information Systems, Athabasca University, Harrisw AT Athabascau DOT Ca;  
harrisw@athabascau.ca

## Abstract

Constrained Object Hierarchies (COH) presents a neuroscience-grounded theoretical framework for modeling artificial general intelligence (AGI) systems across diverse domains. This paper introduces COH as a 9-tuple formalism that integrates compositional structure, adaptive neural components, and multi-type constraints into a unified representation as a unified framework for AGI. We present GISMOL (General Intelligent System Modelling Language) as a practical implementation of COH, providing a toolkit for developing constraint-aware intelligent systems. Through six case studies across healthcare, manufacturing, social science, finance, biology, and astronomy, we demonstrate COH's ability to model complex adaptive systems with explicit constraint management. The framework addresses the "jagged intelligence" problem by ensuring consistent constraint enforcement across abstraction levels, while providing mechanisms for universal world modeling and agentic system implementation. Our results show that COH/GISMOL enables the development of explainable, safety-critical AGI systems with verifiable constraint satisfaction across multiple domains simultaneously.

**Keywords:** artificial general intelligence; Constrained Object Hierarchies; neural-symbolic integration; constraint satisfaction; hierarchical systems; multi-domain modeling; cognitive architecture; universal world model; agentic systems; jagged intelligence

---

## 1. Introduction

The pursuit of Artificial General Intelligence (AGI) has long been hindered by the challenge of creating systems that can operate effectively across diverse domains while maintaining consistent reasoning capabilities and safety constraints. Current AI approaches often excel in narrow domains but fail to generalize or integrate knowledge across different contexts—a phenomenon sometimes called "jagged intelligence" [1]. This paper introduces Constrained Object Hierarchies (COH) as a unified theoretical framework for AGI that addresses these limitations through a neuroscience-grounded approach to system modeling.

Constrained Object Hierarchies formalizes intelligent systems as hierarchical structures of objects subject to multiple constraint types, expressed as a 9-tuple:  $O = (C, A, M, N, E, I, T, G, D)$ . This representation captures the essential elements of intelligent systems: compositional structure (C), state variables (A), executable actions (M), adaptive neural components (N), semantic embeddings (E), and three constraint types—identity (I), trigger (T), and goal (G)—monitored by real-time daemons (D). The framework is grounded in principles from cognitive neuroscience, particularly the hierarchical organization of cortical processing and the constraint-satisfaction nature of neural computation [2].

We present GISMOL (General Intelligent System Modelling Language) as a Python implementation of the COH framework, providing developers with practical tools for building constraint-aware intelligent systems. GISMOL includes core modules for object representation (gismol.core), neural integration (gismol.neural), multi-domain reasoning (gismol.reasoners), and natural language processing (gismol.nlp). Together, COH theory and GISMOL implementation offer

a comprehensive solution for AGI development that addresses both theoretical foundations and practical implementation challenges.

This paper makes three primary contributions: (1) formalization of the COH framework with detailed 9-tuple semantics, (2) demonstration through six complex system case studies across diverse domains, and (3) implementation guidelines showing how COH models translate to executable GISMOL systems. We argue that COH provides a path toward AGI by offering a unified representation that scales across domains while maintaining verifiable constraint satisfaction.

## 2. Literature Review

Research in AGI has followed multiple paths, including symbolic AI, connectionist approaches, and hybrid systems. Early symbolic systems [3] demonstrated robust reasoning but lacked learning capabilities, while modern deep learning [4] excels at pattern recognition but struggles with explicit reasoning and constraint enforcement. Recent approaches have explored neural-symbolic integration [5], cognitive architectures [6], and world modeling [7], yet none have fully addressed the challenge of multi-domain constraint management.

Hierarchical representations have long been recognized as fundamental to intelligence, from Miller's chunking theory [8] to Hawkins' hierarchical temporal memory [9]. Constraints have been used in various AI contexts, from constraint satisfaction problems [10] to constraint programming [11], but typically as isolated mechanisms rather than integrated frameworks. Neuroscience research supports the hierarchical organization of cortical processing [12] and the constraint-based nature of neural computation [13], providing biological grounding for the COH approach.

Multi-domain reasoning has been explored in cognitive architectures like SOAR [14] and ACT-R [15], but these often lack seamless neural integration. Recent work on foundation models [16] demonstrates impressive generalization but suffers from inconsistent behavior and lack of explicit constraint enforcement—the "jagged intelligence" problem [17]. Safety-critical systems research [18] emphasizes constraint management but typically focuses on specific domains rather than general intelligence.

World modeling approaches [19,20] attempt to create comprehensive environment representations but often lack structured constraint mechanisms. Agent-based systems [21] provide autonomy but struggle with coordinated constraint management across hierarchies. The COH framework uniquely integrates these elements: hierarchical representation, multi-type constraints, neural adaptation, and domain generality, offering a comprehensive solution to AGI development.

## 3. Workflow of Modelling and Implementing Complex Systems

The COH/GISMOL framework provides a systematic workflow for modeling and implementing complex intelligent systems across domains. This workflow consists of four main phases: analysis, design, modeling, and implementation, each supported by specific GISMOL components.

The analysis phase begins with domain decomposition, where complex systems are broken down into hierarchical components. GISMOL's COHObject class provides the foundation, allowing developers to represent entities at multiple abstraction levels. During this phase, system requirements are translated into constraint specifications, identifying identity constraints (fundamental properties), trigger constraints (event-driven behaviors), and goal constraints (optimization objectives). The COHRepository helps manage component relationships, while the EntityRelationExtractor from gismol.nlp can assist in extracting structured knowledge from domain documentation.

The design phase focuses on constraint specification and neural component integration. Developers use the ConstraintSystem from gismol.core to define constraints with appropriate categories (safety, temporal, resource, etc.) and severity levels. Neural components are selected based on adaptive requirements: Classifier and Regressor for prediction tasks, NeuralSchedulingModel for

planning, and EmbeddingModel for semantic operations. The design must ensure that neural components respect system constraints through the ConstraintAwareOptimizer.

The modeling phase formalizes the system as a COH 9-tuple, specifying all elements: components (C), attributes (A), methods (M), neural components (N), embedding model (E), identity constraints (I), trigger constraints (T), goal constraints (G), and daemon configurations (D). This formal model serves as both documentation and executable specification, with each element mapping directly to GISMOL implementation constructs.

The implementation phase translates the COH model into executable GISMOL code. The COHObject hierarchy is instantiated with specific constraints, neural components are trained with constraint-aware optimization, and daemons are configured for real-time monitoring. Integration testing ensures that constraints are properly enforced across the hierarchy, with the SafetyReasoner providing redundant validation for critical systems. The complete system can then operate autonomously while maintaining all specified constraints.

This workflow enables consistent development of intelligent systems across domains, with explicit constraint management ensuring safety and reliability. The following sections demonstrate this workflow through six detailed case studies.

#### 4. Analysis, Design, Modeling, and Implementation of Complex Systems - Adaptive Pandemic Response System

Due to limitation of space, we will only include the Analysis, Design, Modeling, and Implementation of Adaptive Pandemic Response System in the paper. Other case studies can be found in supplemental material A.

**System Description and Significance:** Modern pandemics require coordinated responses across healthcare infrastructure, population management, and resource allocation. Traditional approaches struggle with dynamic optimization under uncertainty, leading to suboptimal resource distribution and delayed interventions. An intelligent pandemic response system must balance multiple objectives: minimizing mortality, maintaining healthcare capacity, and preserving socioeconomic function.

**Analysis and Requirements:** The system requires multi-scale modeling from individual patients to regional healthcare networks. Key requirements include: (1) predictive modeling of disease spread, (2) dynamic resource allocation, (3) policy impact simulation, (4) equity constraints in resource distribution, and (5) real-time adaptation to new information. Constraints must handle uncertainty while ensuring critical thresholds are never violated.

**COH Solution Design:** The system is designed as a hierarchical network of regional healthcare clusters, population segments, and resource distribution networks. Each component has adaptive neural models for prediction and optimization, with constraints ensuring ethical and operational requirements. Trigger constraints implement responsive policy adjustments, while goal constraints optimize multi-objective outcomes.

##### 9-Tuple Formalization:

```

C: {RegionalHealthcareClusters, PopulationSegments, ResourceNetworks,
    DataStreams}

A: {R0: float, BedOccupancy: float, VaccineCoverage: Dict, SupplyLeadTime:
    float}

M: {reallocate_resources(), adjust_policy(), update_models(),
    trace_contacts()}

N: {SpreadPredictor: Transformer, InterventionOptimizer: RL,
    SupplyChainOptimizer: GNN}

```

**E:** MultiModalEmbedder connecting epidemiological, operational, socioeconomic concepts

**I:** {Capacity  $\leq$  120%, Distribution  $\geq$  WHO\_equity, Privacy  $\equiv$  preserved}

**T:** {ICU > 90%  $\rightarrow$  emergency\_protocol, Variant > 50%  $\rightarrow$  update\_models}

**G:** {Minimize: Mortality + EconomicDisruption, Maximize: Resilience + Equity}

**D:** {CapacityMonitor, EquityEnforcer, ComplianceTracker}

**Modeling Justification:** The hierarchical decomposition allows localized adaptation while maintaining global coordination. Neural components provide predictive capability under uncertainty, while constraints ensure ethical and operational boundaries. Embedding integration enables semantic understanding across different data modalities.

#### GISMOL Implementation code snippets:

```
class PandemicResponseSystem(COHOject):
    def __init__(self):
        super().__init__(name="PandemicResponseSystem")

        # Components
        self.add_component(COHOject("HealthcareClusters"))
        self.add_component(COHOject("PopulationSegments"))

        # Neural components
        self.add_neural_component(
            "spread_predictor",
            TransformerPredictor(input_dim=256, output_dim=7)
        )
        self.add_neural_component(
            "resource_optimizer",
            GraphNeuralNetwork(node_dim=32, edge_dim=16)
        )

        # Identity constraints
        self.add_identity_constraint({
            'name': 'capacity_limit',
            'specification': 'hospital_occupancy <= 1.2',
            'category': 'safety',
            'severity': 10
        })

        # Trigger constraints
        self.add_trigger_constraint({
            'name': 'emergency_activation',
            'precondition': 'icu_occupancy > 0.9',
            'action': 'activate_emergency_protocol',
            'postcondition': 'emergency_active == true'
```

```

    })

    # Goal constraints
    self.add_goal_constraint({
        'name': 'multi_objective',
        'specification': 'minimize mortality + economic_disruption',
        'weights': {'mortality': 0.6, 'economic': 0.4}
    })

    # Start monitoring daemons
    self.start_daemons()

```

**Benefits of Hierarchical Decomposition:** Regional clusters operate semi-autonomously while respecting global constraints, enabling localized optimization. Neural components at different levels specialize in their domains while sharing embeddings for coordination. Constraints propagate both upward (aggregate metrics) and downward (resource allocations), ensuring consistency.

**Neural-Constraint Interaction:** Spread predictors update based on constraint violations (e.g., when distancing compliance drops), creating feedback loops. The resource optimizer respects equity constraints through penalty terms in its loss function, implemented via ConstraintAwareOptimizer.

**Constraint Propagation:** Capacity constraints at hospital level propagate upward to regional constraints, triggering resource redistribution when violated. Equity constraints flow downward to ensure fair distribution across demographic groups.

**Constraint Roles:** Identity constraints define operational boundaries (capacity limits), trigger constraints implement responsive actions (emergency protocols), and goal constraints guide multi-objective optimization (balancing health and economic outcomes).

**Critical Comparison:** Traditional compartmental models (SIR/SEIR) [22] lack resource constraints and optimization capabilities. Agent-based simulations [23] capture heterogeneity but scale poorly. COH integrates predictive modeling, constraint optimization, and multi-scale coordination in a unified framework with explicit ethical constraints.

## 5. Implementing COH Models into Executable Live Applications

The transition from COH theoretical models to executable GISMOL systems follows a systematic implementation pattern that preserves constraint integrity while enabling adaptive behavior. This section demonstrates the implementation process through comprehensive examples.

**Implementation Architecture:** GISMOL implements COH models through four core modules: gismol.core for object representation, gismol.neural for adaptive components, gismol.reasoners for constraint evaluation, and gismol.nlp for natural language interaction. The implementation maintains the 9-tuple structure while providing practical programming interfaces. Sample implementation of the **Pandemic Response System** can be found in the supplemental materials of the paper. In this section we will explain the key techniques involved in the implementation.

**Neural Component Integration:** Neural components in GISMOL are first-class COHObjects that participate in constraint validation. During training, ConstraintAwareOptimizer incorporates constraint violations as penalty terms. During inference, neural outputs are validated against relevant constraints before being accepted.

```

# Example: Constraint-aware neural training
class ConstraintAwarePredictor(NeuralComponent):
    def train(self, data, constraints=None):
        # Convert constraints to loss penalties

```

```

    if constraints:
        penalty_terms = []
        for constraint in constraints:
            if constraint['type'] == 'inequality':
                penalty = torch.relu(self._evaluate_constraint(constraint) -
constraint['threshold'])
                penalty_terms.append(penalty)

        constraint_loss = sum(penalty_terms) * self.penalty_weight

    # Combine with primary loss
    total_loss = primary_loss + constraint_loss

    # Optimization step
    self.optimizer.zero_grad()
    total_loss.backward()
    self.optimizer.step()

```

**Constraint Propagation Mechanism:** Constraints in hierarchical systems propagate through parent-child relationships. When a child object violates a constraint, it can trigger re-evaluation at parent level, potentially leading to reallocation or adaptation.

```

# Example: Hierarchical constraint propagation
def propagate_constraint_violation(self, child_object, violation):
    """Propagate constraint violation up the hierarchy"""
    # Log violation
    self.logger.warning(f"Constraint violation in {child_object.name}:
{violation}")

    # Check if local resolution is possible
    if child_object.attempt_resolution(violation):
        return True

    # Propagate to parent
    parent = self.get_parent(child_object)
    if parent:
        # Parent may reallocate resources or adjust parameters
        resolution = parent.handle_child_violation(child_object, violation)

        # Continue propagation if needed
        if not resolution:
            return self.propagate_constraint_violation(parent, violation)

    return False

```

**Real-time Monitoring with Daemons:** Constraint daemons provide continuous monitoring with configurable sensitivity. They can operate at different frequencies based on constraint criticality and system state.

```
# Example: Safety constraint daemon
class SafetyConstraintDaemon(ConstraintDaemon):
    def __init__(self, constraints, monitoring_interval=1.0):
        super().__init__(constraints)
        self.monitoring_interval = monitoring_interval
        self.violation_history = []

    def monitor(self):
        while self.active:
            # Check all safety constraints
            for constraint in self.constraints:
                result = self.evaluate_constraint(constraint)

                if not result['satisfied']:
                    # Handle violation
                    self.handle_violation(constraint, result)

                    # Log for analysis
                    self.violation_history.append({
                        'constraint': constraint.name,
                        'time': time.time(),
                        'severity': constraint.severity,
                        'context': result['context']
                    })

            # Adaptive monitoring interval
            if len(self.violation_history) > 0:
                # Increase frequency if recent violations
                recent_violations = [v for v in self.violation_history
                                     if time.time() - v['time'] < 300]
                if len(recent_violations) > 3:
                    self.monitoring_interval = 0.1 # More frequent
                else:
                    self.monitoring_interval = 1.0 # Normal

            time.sleep(self.monitoring_interval)

    def handle_violation(self, constraint, result):
        """Handle constraint violation"""
        # Attempt automatic resolution
        if constraint.auto_resolve:
```

```

        resolution = self.attempt_resolution(constraint, result)
        if resolution:
            return

    # Escalate if critical
    if constraint.severity >= 9:
        self.escalate_to_supervisor(constraint, result)

    # Execute emergency action if defined
    if hasattr(constraint, 'emergency_action'):
        constraint.emergency_action(result['context'])

```

This implementation approach ensures that COH models translate into executable systems that maintain theoretical properties while providing practical functionality. The hierarchical structure enables scalable deployment, while constraint mechanisms ensure safety and reliability.

## 6. Avoiding "Jagged Intelligence" with COH/GISMOL

"Jagged intelligence" refers to the inconsistent performance of AI systems across different domains or task types, where a system may excel in one area but fail catastrophically in another [17]. This problem arises from several factors: lack of unified representation, inconsistent constraint enforcement, and domain-specific optimization that doesn't generalize. COH/GISMOL addresses these issues through several key mechanisms.

**Unified Representation Across Domains:** COH provides a consistent 9-tuple representation that applies equally to all domains. This uniformity ensures that systems developed for different applications share common structure and semantics, enabling knowledge transfer and consistent behavior.

```

# Example: Unified representation enables cross-domain consistency
class DomainAgnosticSystem(COHObject):
    def __init__(self):
        # Same structure regardless of domain
        self.components = [] # C
        self.attributes = {} # A
        self.methods = {} # M
        self.neural_components = {} # N
        self.embedding_model = None # E
        self.identity_constraints = [] # I
        self.trigger_constraints = [] # T
        self.goal_constraints = [] # G
        self.daemons = {} # D

    # Methods work consistently across domains
    def evaluate_constraints(self, context):
        # Uses same evaluation logic for healthcare, finance, etc.
        return self.constraint_system.evaluate_all(context)

```

**Consistent Constraint Enforcement:** The constraint system in GISMOL ensures that rules are applied uniformly regardless of context. Constraints are evaluated using specialized reasoners that understand domain semantics but follow consistent evaluation protocols.

```
# Example: Consistent constraint evaluation
def evaluate_across_domains(self, domains):
    """Evaluate constraints consistently across multiple domains"""
    results = {}

    for domain, context in domains.items():
        # Get appropriate reasoner for domain
        reasoner = self.get_reasoner_for_domain(domain)

        # Evaluate all constraints with same interface
        domain_results = reasoner.evaluate_all(
            self.constraints,
            context,
            evaluation_mode='strict'
        )

        # Apply consistent validation criteria
        valid = all(r['satisfied'] for r in domain_results)
        results[domain] = {
            'valid': valid,
            'violations': [r for r in domain_results if not r['satisfied']],
            'confidence': self.calculate_confidence(domain_results)
        }

    # Ensure consistency: system is only valid if ALL domains are valid
    system_valid = all(r['valid'] for r in results.values())

    return system_valid, results
```

**Cross-Domain Knowledge Transfer:** The embedding component (E) creates a unified semantic space that enables knowledge transfer between domains. Similarities in constraint structures or component relationships can be recognized and leveraged.

```
# Example: Cross-domain knowledge transfer via embeddings
def transfer_knowledge(self, source_domain, target_domain):
    """Transfer constraint knowledge between domains"""
    # Get embeddings of constraints from source domain
    source_constraints = self.get_domain_constraints(source_domain)
    source_embeddings = self.embedding_model.embed_constraints(source_constraints)

    # Find similar constraints in target domain
    target_constraints = self.get_domain_constraints(target_domain)
```

```

target_embeddings = self.embedding_model.embed_constraints(target_constraints)

# Calculate similarities
similarities = cosine_similarity(source_embeddings, target_embeddings)

# Transfer knowledge for highly similar constraints
for i, source_constraint in enumerate(source_constraints):
    max_similarity = np.max(similarities[i])
    if max_similarity > 0.8: # Threshold for transfer
        j = np.argmax(similarities[i])
        self.transfer_constraint_knowledge(
            source_constraint,
            target_constraints[j]
        )

```

**Hierarchical Generalization:** The compositional hierarchy enables generalization from specific instances to broader categories. Constraints defined at higher levels apply to all sub-components, ensuring consistent behavior across the hierarchy.

```

# Example: Hierarchical constraint generalization
class HierarchicalGeneralization:
    def apply_parent_constraints(self, child_object):
        """Apply parent constraints to child objects"""
        parent = self.get_parent(child_object)

        if parent:
            # Inherit identity constraints
            for constraint in parent.identity_constraints:
                if constraint.inheritable:
                    child_object.add_identity_constraint(constraint.clone())

            # Check for conflicts
            self.resolve_constraint_conflicts(child_object)

    def resolve_constraint_conflicts(self, object):
        """Resolve conflicts between inherited and local constraints"""
        conflicts = []

        for i, constraint1 in enumerate(object.identity_constraints):
            for j, constraint2 in enumerate(object.identity_constraints[i+1:]):
                if self.constraints_conflict(constraint1, constraint2):
                    conflicts.append((constraint1, constraint2))

        # Resolve based on priority rules
        for constraint1, constraint2 in conflicts:

```

```

    if constraint1.priority > constraint2.priority:
        object.remove_constraint(constraint2)
    else:
        object.remove_constraint(constraint1)

```

**Adaptive Constraint Relaxation:** In situations where strict constraint enforcement would cause failure, COH/GISMOL provides mechanisms for controlled constraint relaxation with compensatory measures, avoiding the brittle behavior typical of "jagged" systems.

```

# Example: Adaptive constraint relaxation
class AdaptiveConstraintManager:
    def handle_impossible_constraints(self, context):
        """Handle situations where constraints cannot all be satisfied"""
        # Identify which constraints are violatable
        violatable = [c for c in self.constraints if not c.unviolatable]
        unviolatable = [c for c in self.constraints if c.unviolatable]

        # First ensure unviolatable constraints are satisfied
        for constraint in unviolatable:
            if not self.evaluate_constraint(constraint, context):
                raise CriticalConstraintViolation(f"Cannot violate:
{constraint.name}")

        # For violatable constraints, find optimal relaxation
        relaxation_plan = self.optimize_constraint_relaxation(violatable, context)

        # Apply relaxation with compensatory actions
        self.apply_relaxation_plan(relaxation_plan)

        # Monitor for when constraints can be restored
        self.schedule_constraint_restoration(relaxation_plan)

    def optimize_constraint_relaxation(self, constraints, context):
        """Find optimal constraint relaxation plan"""
        # Use goal reasoner to minimize impact
        goal_reasoner = GoalReasoner()

        relaxation_plan = goal_reasoner.optimize_relaxation(
            constraints=constraints,
            context=context,
            objective=self.relaxation_objective_function
        )

        return relaxation_plan

```

```

def relaxation_objective_function(self, relaxation):
    """Objective: minimize harm from constraint relaxation"""
    harm_score = 0

    for constraint, degree in relaxation.items():
        harm_score += constraint.severity * degree * constraint.importance

    return harm_score

```

**Continuous Consistency Validation:** GISMOL includes mechanisms for continuously validating consistency across the system, detecting and correcting "jagged" behavior before it causes problems.

```

# Example: Continuous consistency validation
class ConsistencyValidator:
    def __init__(self):
        self.consistency_metrics = {}
        self.validation_schedule = {}

    def validate_cross_domain_consistency(self):
        """Validate consistency across all domains"""
        inconsistencies = []

        # Check constraint consistency
        for domain1, domain2 in combinations(self.domains, 2):
            if not self.constraints_consistent(domain1, domain2):
                inconsistencies.append({
                    'type': 'constraint_inconsistency',
                    'domains': (domain1, domain2),
                    'description': 'Conflicting constraints between domains'
                })

        # Check behavioral consistency
        behavioral_tests = self.generate_cross_domain_tests()
        for test in behavioral_tests:
            result = self.execute_behavioral_test(test)
            if not result['consistent']:
                inconsistencies.append({
                    'type': 'behavioral_inconsistency',
                    'test': test.name,
                    'result': result
                })

        # Resolve inconsistencies
        if inconsistencies:

```

```

        resolution_plan = self.generate_resolution_plan(inconsistencies)
        self.execute_resolution_plan(resolution_plan)

    return len(inconsistencies) == 0

```

Through these mechanisms, COH/GISMOL ensures that intelligent systems exhibit consistent, reliable behavior across domains, avoiding the "jagged intelligence" problem that plagues many current AI approaches.

## 7. COH as a "Universal World Model"

The concept of a "universal world model" refers to a comprehensive representation that can capture the essential structure and dynamics of any environment or domain [20]. COH provides such a capability through its flexible hierarchical representation and constraint formalism, enabling modeling of diverse worlds from physical environments to abstract conceptual spaces.

**Formal Universality:** The COH 9-tuple provides a formally complete representation that can encode any system with components, attributes, behaviors, adaptive elements, semantics, and constraints. This formal completeness enables COH to serve as a universal representation language.

```

# Example: Universal world model representation
class UniversalWorldModel(COHObject):
    def model_physical_world(self):
        """Model physical world with laws of physics as constraints"""
        world = COHObject(name="PhysicalWorld")

        # Components: physical entities
        world.add_component(self.create_entity("Matter"))
        world.add_component(self.create_entity("Energy"))
        world.add_component(self.create_entity("SpaceTime"))

        # Attributes: physical properties
        world.attributes = {
            'temperature': 2.73, # Kelvin
            'expansion_rate': 73.3, # km/s/Mpc
            'entropy': 'increasing'
        }

        # Methods: physical processes
        world.methods = {
            'evolve': self.physical_evolution,
            'interact': self.entity_interaction,
            'transform': self.energy_matter_conversion
        }

        # Neural components: approximate complex physics
        world.add_neural_component(
            "quantum_approximator",

```

```
        NeuralNetwork(input_dim=256, output_dim=128)
    )

    # Identity constraints: physical laws
    world.add_identity_constraint({
        'name': 'energy_conservation',
        'specification': 'delta_energy == 0',
        'category': 'physical_law'
    })

    world.add_identity_constraint({
        'name': 'causality',
        'specification': 'causes precede effects',
        'category': 'physical_law'
    })

    return world

def model_social_world(self):
    """Model social world with norms and institutions as constraints"""
    world = COHObject(name="SocialWorld")

    # Components: social entities
    world.add_component(self.create_entity("Individuals"))
    world.add_component(self.create_entity("Institutions"))
    world.add_component(self.create_entity("Norms"))

    # Attributes: social properties
    world.attributes = {
        'trust_level': 0.7,
        'cooperation_index': 0.65,
        'inequality_gini': 0.42
    }

    # Methods: social processes
    world.methods = {
        'communicate': self.social_communication,
        'cooperate': self.collective_action,
        'innovate': self.social_innovation
    }

    # Neural components: model complex social dynamics
    world.add_neural_component(
        "social_dynamics_predictor",
```

```

        GraphNeuralNetwork(node_features=16, edge_features=8)
    )

    # Identity constraints: social norms
    world.add_identity_constraint({
        'name': 'reciprocity',
        'specification': 'cooperation requires mutual benefit',
        'category': 'social_norm'
    })

    return world

```

**Multi-Scale Modeling:** COH's hierarchical nature enables modeling at multiple scales simultaneously, from quantum phenomena to cosmological structures, all within a unified framework.

```

# Example: Multi-scale world model
class MultiScaleWorld(COHObject):
    def __init__(self):
        super().__init__(name="MultiScaleUniverse")

        # Create scale hierarchy
        self.add_component(self.create_scale("quantum", scale=1e-35))
        self.add_component(self.create_scale("atomic", scale=1e-10))
        self.add_component(self.create_scale("molecular", scale=1e-9))
        self.add_component(self.create_scale("cellular", scale=1e-6))
        self.add_component(self.create_scale("organismic", scale=1e-2))
        self.add_component(self.create_scale("ecological", scale=1e3))
        self.add_component(self.create_scale("planetary", scale=1e7))
        self.add_component(self.create_scale("stellar", scale=1e12))
        self.add_component(self.create_scale("galactic", scale=1e21))
        self.add_component(self.create_scale("cosmological", scale=1e26))

        # Cross-scale constraints
        self.add_identity_constraint({
            'name': 'scale_consistency',
            'specification': 'aggregate(micro) == macro',
            'description': 'Micro-scale properties must aggregate to macro-scale
properties'
        })

        # Scale-specific neural components
        for scale in self.components:
            scale.add_neural_component(
                f"{scale.name}_dynamics",

```

```

        self.create_scale_appropriate_model(scale)
    )

```

**Abstract World Modeling:** COH can model not only physical worlds but also abstract conceptual spaces, mathematical structures, and informational environments.

```

# Example: Mathematical world model
class MathematicalWorld(COHObject):
    def __init__(self):
        super().__init__(name="MathematicalUniverse")

        # Components: mathematical structures
        self.add_component(self.create_structure("Set"))
        self.add_component(self.create_structure("Group"))
        self.add_component(self.create_structure("Topology"))
        self.add_component(self.create_structure("Category"))

        # Attributes: mathematical properties
        self.attributes = {
            'consistency': True,
            'completeness': 'partial',
            'cardinality': 'infinite'
        }

        # Methods: mathematical operations
        self.methods = {
            'prove': self.mathematical_proof,
            'construct': self.structure_construction,
            'transform': self.mathematical_transformation
        }

        # Neural components: theorem proving and discovery
        self.add_neural_component(
            "theorem_prover",
            TransformerNetwork(input_dim=512, output_dim=256)
        )

        # Identity constraints: mathematical axioms
        self.add_identity_constraint({
            'name': 'law_of_noncontradiction',
            'specification': 'not (P and not P)',
            'category': 'logical_axiom'
        })

        self.add_identity_constraint({

```

```

    'name': 'axiom_of_extensionality',
    'specification': 'sets with same elements are equal',
    'category': 'set_theory'
})

```

**Dynamic World Adaptation:** COH worlds can adapt and evolve over time through neural components and constraint modifications, modeling evolutionary processes and learning.

```

# Example: Evolving world model
class EvolvingWorld(COHObject):
    def __init__(self):
        super().__init__(name="EvolvingUniverse")
        self.generation = 0
        self.fitness_history = []

        # Evolutionary neural component
        self.add_neural_component(
            "world_generator",
            GenerativeAdversarialNetwork(
                generator_dim=1024,
                discriminator_dim=512
            )
        )

        # Fitness constraints (selection criteria)
        self.add_goal_constraint({
            'name': 'evolutionary_fitness',
            'specification': 'maximize complexity * stability',
            'weight': 1.0
        })

    def evolve(self, generations=1000):
        """Evolve the world through multiple generations"""
        for gen in range(generations):
            # Generate candidate modifications
            candidates =
self.neural_components["world_generator"].generate_variations(self)

            # Evaluate fitness
            fitness_scores = []
            for candidate in candidates:
                fitness = self.evaluate_fitness(candidate)
                fitness_scores.append(fitness)

            # Select best candidate

```

```

best_idx = np.argmax(fitness_scores)
best_candidate = candidates[best_idx]

# Incorporate successful variations
self.incorporate_variations(best_candidate)

# Update fitness history
self.fitness_history.append(fitness_scores[best_idx])
self.generation += 1

# Adaptive constraint adjustment
self.adapt_constraints_based_on_fitness(fitness_scores[best_idx])

```

**Inter-World Relationships:** COH can model relationships between different worlds, enabling simulation of multiverse scenarios or parallel reality modeling.

```

# Example: Multiverse modeling
class MultiverseModel(COHObject):
    def __init__(self):
        super().__init__(name="Multiverse")

    # Create multiple worlds
    self.worlds = {
        'physical': self.create_physical_world(),
        'social': self.create_social_world(),
        'mathematical': self.create_mathematical_world(),
        'informational': self.create_informational_world()
    }

    # Inter-world relationships
    self.add_relation(
        self.worlds['physical'],
        self.worlds['informational'],
        "implements",
        "physical_implements_informational"
    )

    self.add_relation(
        self.worlds['social'],
        self.worlds['mathematical'],
        "models_with",
        "social_modeled_with_math"
    )

    # Cross-world constraints

```

```

self.add_identity_constraint({
    'name': 'cross_world_consistency',
    'specification': 'worlds must be mutually consistent where they
intersect',
    'category': 'metaconstraint'
})

```

Through these capabilities, COH serves as a true universal world model, capable of representing any conceivable environment or system while maintaining formal rigor and practical implementability.

## 8. Modelling and Implementing "Agentic Systems"

Agentic systems—autonomous entities that perceive, decide, and act in pursuit of goals—are a fundamental application of AGI. COH provides a natural framework for modeling agents with explicit goal constraints, adaptive behavior, and hierarchical decision-making. This section demonstrates how COH/GISMOL enables the development of sophisticated agentic systems.

**Agent as COHObject:** In COH, agents are simply COHObjects with specific configurations: goal constraints define their objectives, neural components implement their adaptive capabilities, and trigger constraints govern their reactive behaviors.

```

# Example: Basic agent definition
class COHAgent(COHObject):
    def __init__(self, name, capabilities, goals):
        super().__init__(name=name)

        # Agent capabilities as components
        for capability in capabilities:
            self.add_component(COHObject(name=capability))

        # Goal constraints define agent objectives
        for goal in goals:
            self.add_goal_constraint({
                'name': goal['name'],
                'specification': goal['specification'],
                'priority': goal.get('priority', 'medium'),
                'deadline': goal.get('deadline', None)
            })

        # Neural components for learning and adaptation
        self.add_neural_component(
            "policy_network",
            ReinforcementLearning(
                state_dim=64,
                action_dim=16,
                reward_function=self.composite_reward_function
            )
)

```

```
)

self.add_neural_component(
    "world_model",
    PredictiveModel(
        input_dim=128,
        output_dim=64,
        prediction_horizon=10
    )
)

# Identity constraints: agent principles
self.add_identity_constraint({
    'name': 'self_preservation',
    'specification': 'health > 0',
    'category': 'safety',
    'severity': 10
})

# Trigger constraints: reactive behaviors
self.add_trigger_constraint({
    'name': 'threat_response',
    'precondition': 'threat_detected == true',
    'action': self.evade_or_counter,
    'postcondition': 'safe == true'
})

def composite_reward_function(self, state, action, next_state):
    """Reward function combining multiple goal constraints"""
    reward = 0

    # Evaluate progress toward each goal constraint
    for goal in self.goal_constraints:
        progress = self.evaluate_goal_progress(goal, state, next_state)
        reward += goal.priority * progress

    # Penalty for constraint violations
    for constraint in self.identity_constraints:
        if not self.evaluate_constraint(constraint, state):
            reward -= constraint.severity * 10

    return reward
```

**Hierarchical Agent Organizations:** Complex agentic systems often involve hierarchies of agents with different roles and capabilities. COH naturally models such organizations through compositional hierarchies.

```
# Example: Hierarchical multi-agent system
class MultiAgentOrganization(COHObject):
    def __init__(self):
        super().__init__(name="AgentOrganization")

        # Create agent hierarchy
        self.add_component(self.create_agent("Coordinator", "strategic_planning"))
        self.add_component(self.create_agent("Manager", "task_allocation"))

        # Worker agents
        workers = []
        for i in range(10):
            worker = self.create_agent(f"Worker_{i}", "task_execution")
            workers.append(worker)
            self.add_component(worker)

        # Organizational constraints
        self.add_identity_constraint({
            'name': 'coordination_requirement',
            'specification': 'all_agents.aligned == true',
            'category': 'organizational'
        })

        self.add_identity_constraint({
            'name': 'resource_fairness',
            'specification': 'resource_distribution_gini < 0.3',
            'category': 'ethical'
        })

        # Communication and coordination methods
        self.methods['coordinate'] = self.coordinate_agents
        self.methods['resolve_conflict'] = self.conflict_resolution

    def coordinate_agents(self, task):
        """Coordinate multiple agents to accomplish task"""
        # Decompose task hierarchically
        subtasks = self.decompose_task(task)

        # Allocate to appropriate agents
        allocations = self.allocate_subtasks(subtasks)
```

```

# Ensure allocations satisfy constraints
while not self.allocations_valid(allocations):
    # Adjust allocations to satisfy constraints
    allocations = self.adjust_allocations(allocations)

# Execute with monitoring
results = self.execute_allocations(allocations)

return results

def conflict_resolution(self, agent1, agent2, conflict):
    """Resolve conflicts between agents"""
    # Check organizational constraints
    org_constraints = self.get_relevant_constraints(conflict)

    # Find resolution satisfying constraints
    resolution = self.find_constraint_satisfying_resolution(
        agent1, agent2, conflict, org_constraints
    )

    # Implement resolution
    self.implement_resolution(resolution)

    return resolution

```

**Goal-Driven Behavior:** Goal constraints in COH agents drive their behavior through continuous optimization. Agents evaluate potential actions based on their contribution to goal achievement while respecting identity constraints.

```

# Example: Goal-driven decision making
class GoalDrivenAgent(COHAgent):
    def decide_action(self, state):
        """Decide action based on goals and constraints"""
        # Generate candidate actions
        candidates = self.generate_action_candidates(state)

        # Evaluate each candidate
        evaluations = []
        for action in candidates:
            # Predict outcome
            predicted_state = self.predict_outcome(state, action)

            # Evaluate goal progress
            goal_scores = []
            for goal in self.goal_constraints:

```

```
        progress = self.calculate_goal_progress(goal, state,
predicted_state)
        goal_scores.append(progress * goal.priority)

    # Check constraint satisfaction
    constraint_violations = 0
    for constraint in self.identity_constraints:
        if not self.evaluate_constraint(constraint, predicted_state):
            constraint_violations += constraint.severity

    # Composite evaluation
    total_score = sum(goal_scores) - constraint_violations * 10
    evaluations.append((action, total_score, predicted_state))

# Select best action
best_action, best_score, _ = max(evaluations, key=lambda x: x[1])

# Safety check
if not self.safety_check(best_action, state):
    # Fall back to safest action
    safe_actions = [a for a, _, ps in evaluations
                    if self.safety_check(a, state)]
    if safe_actions:
        best_action = safe_actions[0]
    else:
        best_action = self.safe_default_action()

return best_action

def calculate_goal_progress(self, goal, current_state, predicted_state):
    """Calculate progress toward a goal"""
    # Parse goal specification
    if 'maximize' in goal.specification:
        # Extract metric to maximize
        metric = self.extract_metric(goal.specification)
        current_value = self.evaluate_metric(metric, current_state)
        predicted_value = self.evaluate_metric(metric, predicted_state)
        progress = predicted_value - current_value

    elif 'minimize' in goal.specification:
        # Extract metric to minimize
        metric = self.extract_metric(goal.specification)
        current_value = self.evaluate_metric(metric, current_state)
        predicted_value = self.evaluate_metric(metric, predicted_state)
```

```

        progress = current_value - predicted_value

    elif 'achieve' in goal.specification:
        # Binary achievement
        current_achieved = self.evaluate_condition(goal.specification,
current_state)
        predicted_achieved = self.evaluate_condition(goal.specification,
predicted_state)
        progress = 1.0 if (not current_achieved and predicted_achieved) else
0.0

    return progress

```

**Learning and Adaptation:** Neural components in COH agents enable learning from experience while respecting constraints. The ConstraintAwareOptimizer ensures that learning doesn't violate important constraints.

```

# Example: Constraint-aware agent learning
class LearningAgent(COHAgent):
    def learn_from_experience(self, experiences):
        """Learn from experiences while respecting constraints"""
        # Extract training data
        states, actions, rewards, next_states, dones = zip(*experiences)

        # Update world model
        self.neural_components["world_model"].train(
            states, next_states,
            constraints=self.get_learning_constraints()
        )

        # Update policy with constraint-aware optimization
        optimizer = ConstraintAwareOptimizer(
            model=self.neural_components["policy_network"],
            constraints=self.get_training_constraints(),
            learning_rate=0.001,
            penalty_weight=0.1
        )

        # Policy gradient update
        policy_loss = self.calculate_policy_loss(
            states, actions, rewards, next_states, dones
        )

        # Add constraint penalties
        constraint_loss = self.calculate_constraint_violation_penalty(states)

```

```

total_loss = policy_loss + constraint_loss

optimizer.zero_grad()
total_loss.backward()
optimizer.step()

# Update constraint thresholds based on experience
self.adapt_constraint_thresholds(experiences)

def get_training_constraints(self):
    """Get constraints relevant for training"""
    training_constraints = []

    for constraint in self.identity_constraints:
        if constraint.learnable:
            # Convert constraint to differentiable penalty
            penalty = self.constraint_to_penalty(constraint)
            training_constraints.append(penalty)

    return training_constraints

def adapt_constraint_thresholds(self, experiences):
    """Adapt constraint thresholds based on experience"""
    for constraint in self.identity_constraints:
        if constraint.adaptive:
            # Calculate actual violation rate
            violations = sum(1 for s, _, _, _ in experiences
                            if not self.evaluate_constraint(constraint, s))
            violation_rate = violations / len(experiences)

            # Adjust threshold to maintain target violation rate
            if violation_rate > constraint.target_violation_rate:
                # Too many violations, relax constraint
                constraint.threshold *= 1.1
            elif violation_rate < constraint.target_violation_rate * 0.5:
                # Too few violations, tighten constraint
                constraint.threshold *= 0.9

```

**Multi-Agent Coordination:** COH enables modeling of complex multi-agent systems with explicit coordination constraints and communication protocols.

```

# Example: Multi-agent coordination system
class CoordinatedMultiAgentSystem(COHObject):
    def __init__(self):
        super().__init__(name="CoordinatedMAS")

```

```
# Create agent teams
self.teams = {
    'exploration': self.create_agent_team("Explorers", 5),
    'exploitation': self.create_agent_team("Exploiters", 3),
    'monitoring': self.create_agent_team("Monitors", 2)
}

# Team coordination constraints
self.add_identity_constraint({
    'name': 'team_coordination',
    'specification': 'teams.coordination_score > 0.7',
    'category': 'organizational'
})

# Resource sharing constraints
self.add_identity_constraint({
    'name': 'fair_resource_sharing',
    'specification': 'resource_allocation_gini < 0.4',
    'category': 'ethical'
})

# Communication protocol
self.communication_protocol = self.create_communication_protocol()

def coordinate_teams(self, global_goal):
    """Coordinate multiple teams toward global goal"""
    # Decompose global goal into team objectives
    team_objectives = self.decompose_goal(global_goal, self.teams.keys())

    # Assign objectives to teams
    for team_name, objective in team_objectives.items():
        self.teams[team_name].assign_objective(objective)

    # Monitor coordination
    while not self.global_goal_achieved(global_goal):
        # Collect team status
        team_statuses = {}
        for team_name, team in self.teams.items():
            team_statuses[team_name] = team.get_status()

        # Check coordination constraints
        coordination_ok = self.evaluate_coordination_constraints(team_statuses)
```

```
if not coordination_ok:
    # Adjust coordination
    self.adjust_coordination(team_statuses)

# Share information between teams
self.share_information(team_statuses)

# Execute team actions
for team_name, team in self.teams.items():
    team.execute_cycle()

def create_communication_protocol(self):
    """Create constraint-aware communication protocol"""
    protocol = {
        'message_types': ['request', 'offer', 'acknowledge', 'refuse'],
        'constraints': [
            {
                'name': 'response_time',
                'specification': 'response_within < 5.0',
                'category': 'temporal'
            },
            {
                'name': 'truthfulness',
                'specification': 'messages_truthful == true',
                'category': 'ethical'
            }
        ],
        'format': self.message_format,
        'validation': self.validate_message
    }

    return protocol
```

Through these mechanisms, COH/GISMOL enables the development of sophisticated agentic systems that are goal-driven, adaptive, coordinated, and constraint-aware—essential properties for AGI applications.

## 9. Why COH/GISMOL Can Be a Better Framework for AGI

The COH/GISMOL framework offers several advantages over existing approaches to AGI development, addressing key limitations in current AI systems while providing a path toward genuine general intelligence.

**Unified Representation Across Domains:** Unlike domain-specific AI systems or narrowly trained models, COH provides a universal representation that applies equally to all domains. This enables knowledge transfer, consistent reasoning, and true generalization—hallmarks of general intelligence.

```

# Example: Unified representation enables cross-domain reasoning
class CrossDomainReasoner:
    def reason_across_domains(self, problem):
        """Apply consistent reasoning across multiple domains"""
        # Parse problem to identify relevant domains
        domains = self.identify_relevant_domains(problem)

        # Create unified representation
        coh_problem = self.problem_to_coh(problem, domains)

        # Apply domain-appropriate reasoners
        solutions = []
        for domain in domains:
            reasoner = self.get_domain_reasoner(domain)
            solution = reasoner.solve(coh_problem)
            solutions.append(solution)

        # Integrate solutions
        integrated_solution = self.integrate_solutions(solutions,
coh_problem.constraints)

        return integrated_solution

```

**Explicit Constraint Management:** Most AI systems handle constraints implicitly or through post-hoc correction. COH makes constraints first-class citizens, enabling proactive constraint satisfaction, violation prevention, and principled constraint relaxation when necessary.

```

# Example: Proactive constraint management
class ProactiveConstraintManager:
    def plan_with_constraints(self, goal, initial_state):
        """Plan actions while proactively satisfying constraints"""
        # Generate candidate plans
        candidates = self.generate_plan_candidates(goal, initial_state)

        # Filter by constraint satisfaction
        feasible = []
        for plan in candidates:
            # Simulate plan execution
            trajectory = self.simulate_plan(plan, initial_state)

            # Check constraints at each step
            constraint_satisfied = True
            for state in trajectory:
                for constraint in self.constraints:
                    if not self.evaluate_constraint(constraint, state):

```

```

        constraint_satisfied = False
        break

    if not constraint_satisfied:
        break

    if constraint_satisfied:
        feasible.append(plan)

# Optimize among feasible plans
best_plan = self.optimize_plans(feasible, goal)

return best_plan

```

**Hierarchical Abstraction and Composition:** The compositional hierarchy in COH enables modeling at multiple abstraction levels simultaneously, from low-level details to high-level concepts. This mirrors the hierarchical organization of intelligence in biological systems and enables scalable reasoning.

```

# Example: Multi-level reasoning
class HierarchicalReasoner:
    def reason_at_multiple_levels(self, problem, levels=['detailed', 'abstract']):
        """Reason about problem at multiple abstraction levels"""
        solutions = {}

        for level in levels:
            # Create level-appropriate representation
            coh_problem = self.abstract_to_level(problem, level)

            # Apply level-appropriate reasoning
            if level == 'detailed':
                solution = self.detailed_reasoning(coh_problem)
            elif level == 'abstract':
                solution = self.abstract_reasoning(coh_problem)

            solutions[level] = solution

        # Integrate across levels
        integrated = self.integrate_level_solutions(solutions)

        return integrated

```

**Neural-Symbolic Integration:** COH seamlessly integrates neural components (sub-symbolic, pattern-based reasoning) with symbolic constraints and structures. This combines the learning capabilities of neural networks with the precision and explainability of symbolic reasoning.

```

# Example: Neural-symbolic integration

```

```

class NeuralSymbolicIntegrator:
    def solve_with_integration(self, problem):
        """Solve problem using integrated neural-symbolic approach"""
        # Neural component: pattern recognition
        neural_insights = self.neural_component.analyze(problem)

        # Symbolic component: constraint-based reasoning
        symbolic_solution = self.symbolic_reasoner.solve(
            problem,
            constraints=self.constraints
        )

        # Integrate neural and symbolic insights
        integrated = self.integrate_insights(
            neural_insights,
            symbolic_solution,
            integration_constraints=self.integration_constraints
        )

        # Validate against all constraints
        if self.validate_solution(integrated, self.constraints):
            return integrated
        else:
            # Refine using constraint feedback
            return self.refine_with_constraints(integrated)

```

**Explainability and Transparency:** The explicit constraint structure in COH enables detailed explanation of system behavior. When a decision is made or a constraint is violated, the system can provide clear explanations based on its constraint hierarchy.

```

# Example: Constraint-based explanation
class ExplainableSystem:
    def explain_decision(self, decision, context):
        """Generate explanation for decision based on constraints"""
        explanation = {
            'decision': decision,
            'context': context,
            'relevant_constraints': [],
            'constraint_evaluations': [],
            'tradeoffs': []
        }

        # Identify relevant constraints
        for constraint in self.constraints:
            if self.constraint_relevant(constraint, decision, context):

```

```

        explanation['relevant_constraints'].append(constraint.name)

        # Evaluate constraint
        evaluation = self.evaluate_constraint(constraint, context)
        explanation['constraint_evaluations'].append({
            'constraint': constraint.name,
            'satisfied': evaluation['satisfied'],
            'value': evaluation['value'],
            'threshold': constraint.threshold
        })

    # Identify tradeoffs
    for i, cons1 in enumerate(self.constraints):
        for cons2 in self.constraints[i+1:]:
            if self.constraints_conflict(cons1, cons2, decision, context):
                tradeoff = self.analyze_tradeoff(cons1, cons2, decision)
                explanation['tradeoffs'].append(tradeoff)

    return explanation

```

**Safety and Robustness:** The constraint daemon system in COH provides continuous monitoring and enforcement of safety constraints. This is particularly important for AGI systems that must operate reliably in complex, unpredictable environments.

```

# Example: Safety monitoring system
class SafetyMonitoringSystem:
    def __init__(self):
        # Multiple safety daemons with different priorities
        self.daemons = {
            'critical': SafetyDaemon(priority=10, check_interval=0.1),
            'high': SafetyDaemon(priority=7, check_interval=1.0),
            'medium': SafetyDaemon(priority=4, check_interval=5.0),
            'low': SafetyDaemon(priority=1, check_interval=30.0)
        }

        # Safety constraints categorized by criticality
        self.safety_constraints = self.categorize_constraints_by_criticality()

        # Start all daemons
        for level, daemon in self.daemons.items():
            daemon.constraints = self.safety_constraints[level]
            daemon.start()

    def categorize_constraints_by_criticality(self):
        """Categorize constraints by safety criticality"""

```

```

categories = {
    'critical': [], # Life-critical, must never violate
    'high': [],   # Safety-critical, should never violate
    'medium': [], # Important, can briefly violate with monitoring
    'low': []     # Desirable, can violate with justification
}

for constraint in self.constraints:
    if constraint.category == 'safety':
        if constraint.severity >= 9:
            categories['critical'].append(constraint)
        elif constraint.severity >= 7:
            categories['high'].append(constraint)
        elif constraint.severity >= 4:
            categories['medium'].append(constraint)
        else:
            categories['low'].append(constraint)

return categories

```

**Scalability and Compositionality:** The hierarchical nature of COH enables systems to scale from simple agents to complex organizations while maintaining consistent constraint management. New components can be added without disrupting existing functionality.

```

# Example: Scalable system composition
class ScalableCOHSystem:
    def __init__(self):
        self.root = COHObject(name="RootSystem")
        self.component_registry = {}
        self.constraint_propagation_graph = {}

    def add_component(self, component, parent=None):
        """Add component with constraint propagation"""
        if parent is None:
            parent = self.root

        parent.add_component(component)

        # Register component
        self.component_registry[component.name] = component

        # Set up constraint propagation
        self.setup_constraint_propagation(component, parent)

        # Validate integration

```

```

self.validate_integration(component)

def setup_constraint_propagation(self, component, parent):
    """Set up constraint propagation between component and parent"""
    # Inherit relevant constraints from parent
    for constraint in parent.constraints:
        if constraint.inheritable:
            component.add_constraint(constraint.clone())

    # Set up upward propagation for aggregated constraints
    if hasattr(component, 'aggregate_constraints'):
        for agg_constraint in component.aggregate_constraints:
            self.constraint_propagation_graph.setdefault(
                component.name, []
            ).append({
                'type': 'upward',
                'constraint': agg_constraint,
                'target': parent.name
            })

    # Set up downward propagation for distributed constraints
    if hasattr(parent, 'distribute_constraints'):
        for dist_constraint in parent.distribute_constraints:
            self.constraint_propagation_graph.setdefault(
                parent.name, []
            ).append({
                'type': 'downward',
                'constraint': dist_constraint,
                'target': component.name
            })

```

These advantages make COH/GISMOL a superior framework for AGI development, addressing fundamental challenges that have hindered progress in the field while providing a practical path toward building truly general intelligent systems.

## 10. Key Contributions

This paper makes several significant contributions to the field of artificial general intelligence:

1. **Formalization of Constrained Object Hierarchies:** We have presented a complete formalization of COH as a 9-tuple framework that integrates hierarchical composition, neural adaptation, and multi-type constraints into a unified representation for intelligent systems.
2. **Implementation of GISMOL:** We have developed GISMOL as a practical Python implementation of the COH framework, providing developers with tools for building constraint-aware intelligent systems across domains.
3. **Multi-Domain Case Studies:** Through six comprehensive case studies across healthcare, manufacturing, social science, finance, biology, and astronomy, we have demonstrated COH's ability to model complex adaptive systems with explicit constraint management.

4. **Solution to "Jagged Intelligence":** We have shown how COH/GISMOL addresses the "jagged intelligence" problem through unified representation, consistent constraint enforcement, and cross-domain knowledge transfer.
5. **Universal World Modeling Framework:** We have demonstrated that COH can serve as a universal world model, capable of representing any environment or system while maintaining formal rigor.
6. **Agentic System Development Methodology:** We have provided a comprehensive approach to developing agentic systems within the COH framework, with explicit goal constraints, adaptive behavior, and hierarchical coordination.
7. **Constraint-Aware Neural Integration:** We have developed techniques for integrating neural components with constraint systems, enabling learning and adaptation while respecting domain constraints.
8. **Practical Implementation Guidelines:** We have provided detailed implementation examples showing how COH models translate to executable GISMOL systems, bridging the gap between theory and practice.

These contributions advance both the theory and practice of AGI development, providing a comprehensive framework for building intelligent systems that are capable, reliable, and safe across diverse domains.

## 11. Conclusion and Future Research

The Constrained Object Hierarchies framework, implemented through the GISMOL toolkit, represents a significant advance in the pursuit of artificial general intelligence. By providing a unified representation that integrates hierarchical composition, neural adaptation, and explicit constraint management, COH addresses fundamental challenges in AGI development while offering practical tools for implementation.

Our case studies demonstrate COH's versatility across domains, from pandemic response to interstellar exploration. The framework's ability to model complex adaptive systems with verifiable constraint satisfaction makes it particularly suitable for safety-critical applications where reliability is paramount. The integration of neural components enables learning and adaptation while maintaining constraint boundaries, addressing the tension between flexibility and safety that has hindered many AI approaches.

Several directions for future research emerge from this work:

1. **Scalability Optimization:** While COH provides a principled approach to hierarchical modeling, practical deployment in extremely large-scale systems requires optimization of constraint propagation and evaluation algorithms.
2. **Automatic Constraint Learning:** Current implementations require manual specification of constraints. Future work could develop techniques for learning constraints from data or experience while maintaining safety guarantees.
3. **Cross-Domain Generalization:** Further research is needed on mechanisms for automatic knowledge transfer between domains, enabling faster adaptation to new environments.
4. **Human-COH Collaboration:** Developing intuitive interfaces for humans to interact with COH systems, specify constraints, and understand system reasoning will be crucial for real-world deployment.
5. **Formal Verification:** Extending formal verification techniques to COH systems could provide mathematical guarantees of constraint satisfaction under specified conditions.
6. **Neuroscientific Validation:** Further research could investigate the correspondence between COH structures and neural organization in biological intelligent systems, potentially leading to more brain-inspired architectures.
7. **Distributed COH Systems:** Developing techniques for distributed implementation of COH systems could enable larger-scale deployments while maintaining constraint consistency.

The COH/GISMOL framework provides a solid foundation for these research directions while offering immediate practical value for developing intelligent systems across diverse domains. As AI systems become increasingly integrated into critical infrastructure and everyday life, frameworks that prioritize constraint management, explainability, and safety will become increasingly important. COH represents a significant step toward AGI systems that are not only intelligent but also reliable, transparent, and aligned with human values.

## References

1. A. Hendrycks et al., "Measuring Massive Multitask Language Understanding," *International Conference on Learning Representations (ICLR)*, 2021.
2. J. Friston, "The free-energy principle: A unified brain theory?" *Nature Reviews Neuroscience*, vol. 11, no. 2, pp. 127–138, 2010.
3. J. McCarthy, "Programs with common sense," in *Proceedings of the Teddington Conference on the Mechanization of Thought Processes*, London, UK, 1959, pp. 75-91.
4. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436-444, 2015.
5. A. S. d'Avila Garcez, L. C. Lamb, and D. M. Gabbay, *Neural-Symbolic Cognitive Reasoning*. Berlin, Germany: Springer, 2009.
6. J. E. Laird, *The Soar Cognitive Architecture*. Cambridge, MA, USA: MIT Press, 2012.
7. D. Ha and J. Schmidhuber, "World models," *arXiv preprint arXiv:1803.10122*, 2018.
8. G. A. Miller, "The magical number seven, plus or minus two: Some limits on our capacity for processing information," *Psychological Review*, vol. 63, no. 2, pp. 81-97, 1956.
9. J. Hawkins and S. Blakeslee, *On Intelligence*. New York, NY, USA: Times Books, 2004.
10. E. Tsang, *Foundations of Constraint Satisfaction*. London, UK: Academic Press, 1993.
11. K. Marriott and P. J. Stuckey, *Programming with Constraints: An Introduction*. Cambridge, MA, USA: MIT Press, 1998.
12. D. J. Felleman and D. C. Van Essen, "Distributed hierarchical processing in the primate cerebral cortex," *Cerebral Cortex*, vol. 1, no. 1, pp. 1-47, 1991.
13. G. M. Edelman, *Neural Darwinism: The Theory of Neuronal Group Selection*. New York, NY, USA: Basic Books, 1987.
14. J. E. Laird, A. Newell, and P. S. Rosenbloom, "Soar: An architecture for general intelligence," *Artificial Intelligence*, vol. 33, no. 1, pp. 1-64, 1987.
15. J. R. Anderson, *The Adaptive Character of Thought*. Hillsdale, NJ, USA: Lawrence Erlbaum Associates, 1990.
16. T. Brown et al., "Language models are few-shot learners," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 1877-1901.
17. G. Marcus, "Deep learning: A critical appraisal," *arXiv preprint arXiv:1801.00631*, 2018.
18. N. G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*. Cambridge, MA, USA: MIT Press, 2011.
19. W. Battaglia et al., "Relational inductive biases, deep learning, and graph networks," *arXiv preprint arXiv:1806.01261*, 2018.
20. J. B. Hamrick et al., "On the role of planning in model-based deep reinforcement learning," *arXiv preprint arXiv:2011.04021*, 2020.
21. S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. Hoboken, NJ, USA: Pearson, 2020.
22. W. O. Kermack and A. G. McKendrick, "A contribution to the mathematical theory of epidemics," *Proceedings of the Royal Society of London. Series A*, vol. 115, no. 772, pp. 700-721, 1927.
23. J. M. Epstein and R. Axtell, *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge, MA, USA: MIT Press, 1996.

**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.