

Supplementary Information for: Spectral Sheaf Heuristics for Consistency Detection in Multi-Agent Systems

Sepehr Bayat
Hooshex AI Lab
sepehrbayat@hooshex.com

February 9, 2026

Contents

| | | |
|----------|---|----------|
| 1 | Practical Guide to Sheaf Construction | 4 |
| 1.1 | General Principles | 4 |
| 1.1.1 | Step 1: Identify the Belief Graph | 4 |
| 1.1.2 | Step 2: Define the Stalks (Local Data) | 5 |
| 1.1.3 | Step 3: Define Restriction Maps (Consistency Constraints) | 5 |
| 1.1.4 | Step 4: Construct the Connection Laplacian | 6 |
| 1.1.5 | Step 5: Validate the Sheaf | 6 |
| 1.2 | Walk-Through 1: Logic Maze | 7 |
| 1.2.1 | Problem Description | 7 |
| 1.2.2 | Step 1: Belief Graph | 7 |
| 1.2.3 | Step 2: Stalks | 7 |
| 1.2.4 | Step 3: Restriction Maps | 7 |
| 1.2.5 | Step 4: Injecting a Contradiction | 8 |
| 1.2.6 | Implementation Code Snippet | 8 |
| 1.3 | Walk-Through 2: Safety Gym (Safe RL) | 9 |
| 1.3.1 | Problem Description | 9 |
| 1.3.2 | Step 1: Belief Graph | 9 |
| 1.3.3 | Step 2: Stalks | 9 |
| 1.3.4 | Step 3: Restriction Maps (Bellman Consistency) | 9 |
| 1.3.5 | Step 4: Detecting Safety Violations | 10 |
| 1.3.6 | Implementation Code Snippet | 10 |
| 1.4 | Walk-Through 3: Multi-Robot Coordination | 11 |
| 1.4.1 | Problem Description | 11 |
| 1.4.2 | Step 1: Belief Graph | 11 |
| 1.4.3 | Step 2: Stalks | 11 |
| 1.4.4 | Step 3: Restriction Maps (Spatial Consistency) | 11 |
| 1.4.5 | Step 4: Injecting a Contradiction | 12 |
| 1.4.6 | Step 5: Resolution via Negotiation | 12 |
| 1.4.7 | Implementation Code Snippet | 13 |
| 1.5 | Summary: Decision Tree for Sheaf Construction | 13 |

| | | |
|----------|--|-----------|
| 2 | Robustness Analysis | 14 |
| 2.1 | Sources of Noise in Multi-Agent Systems | 14 |
| 2.1.1 | Sensor Noise | 14 |
| 2.1.2 | Communication Errors | 14 |
| 2.1.3 | Discretization Error | 14 |
| 2.1.4 | Model Uncertainty | 14 |
| 2.2 | Sensitivity Analysis | 14 |
| 2.2.1 | Eigenvalue Perturbation Theory | 14 |
| 2.2.2 | Illustrative Sensitivity Example | 15 |
| 2.3 | Adaptive Threshold Selection | 16 |
| 2.3.1 | The Challenge of Choosing ϵ | 16 |
| 2.3.2 | Strategy 1: Spectral Gap Estimation | 16 |
| 2.3.3 | Strategy 2: Noise-Adaptive Threshold | 16 |
| 2.3.4 | Strategy 3: Cross-Validation | 17 |
| 2.3.5 | Recommended Default | 17 |
| 2.4 | Noise Filtering Techniques | 17 |
| 2.4.1 | Preprocessing: Belief Smoothing | 17 |
| 2.4.2 | Postprocessing: Robust Spectral Estimation | 17 |
| 2.4.3 | Online Filtering: Kalman Smoothing | 18 |
| 2.5 | Robustness Guarantees | 18 |
| 2.5.1 | Theorem: Bounded Degradation | 18 |
| 2.5.2 | Failure Modes | 18 |
| 2.6 | Experimental Validation of Robustness | 19 |
| 2.6.1 | Theoretical Noise Tolerance of the Grid-World Experiment | 19 |
| 2.6.2 | Expected Benefit of Adaptive ϵ Selection | 19 |
| 2.7 | Practical Recommendations | 19 |
| 2.8 | Parameter Selection Guide | 20 |
| 2.8.1 | Choosing ϵ : Spectral Gap Estimation | 20 |
| 2.8.2 | Choosing α : Reward Shaping Coefficient | 21 |
| 2.8.3 | Automated Parameter Tuning | 22 |
| 2.9 | Practitioner’s Checklist: Applying the Method | 23 |
| 3 | Deployment Considerations | 25 |
| 3.1 | System Integration | 25 |
| 3.1.1 | Architecture Patterns | 25 |
| 3.1.2 | Computational Resources | 26 |
| 3.2 | Operational Parameters | 26 |
| 3.2.1 | Threshold Selection | 26 |
| 3.2.2 | Update Frequency | 27 |
| 3.3 | Failure Handling | 27 |
| 3.3.1 | Response Protocols | 27 |
| 3.3.2 | False Positives/Negatives | 27 |
| 3.4 | Maintenance and Monitoring | 28 |
| 3.4.1 | System Health Metrics | 28 |
| 3.4.2 | Periodic Recalibration | 28 |
| 3.5 | Code Availability | 28 |

Theorem 1 (Spectral-Cohomological Correspondence). *Let \mathcal{F} be a cellular sheaf on a connected graph $G = (V, E)$ with $N = |V|$ vertices. Let $\mathcal{L} \in \mathbb{R}^{Nd \times Nd}$ be the Connection Laplacian, where d is the stalk dimension. Then:*

$$h_{true}^1 = \dim(\ker(\mathcal{L})) - 1 \quad (1)$$

where $\dim(\ker(\mathcal{L}))$ is the multiplicity of the zero eigenvalue of \mathcal{L} , and $h_{true}^1 = \dim(H^1(\mathcal{F}))$ is the true first cohomology dimension.

Proof Sketch. By the fundamental theorem of sheaf cohomology [1], the kernel of \mathcal{L} decomposes as $\ker(\mathcal{L}) \cong H^0(\mathcal{F}) \oplus H^1(\mathcal{F})$. Since G is connected, $\dim(H^0(\mathcal{F})) = 1$ (constant global sections). Thus $\dim(\ker(\mathcal{L})) = 1 + h_{true}^1$. Rearranging gives the result. See the main manuscript (Appendix A) for the full proof. \square \square

Theorem 2 (Error Bound for Spectral Approximation). *Let $\mathcal{L}_0 \in \mathbb{R}^{Nd \times Nd}$ be the ideal Connection Laplacian for a graph with N vertices and stalk dimension d . Let $\mathcal{L} = \mathcal{L}_0 + E$ be a perturbed version with $\|E\|_2 \leq \sigma$, where $\sigma > 0$ is the noise level and $\|\cdot\|_2$ denotes the spectral norm (largest singular value).*

Let $\delta > 0$ be the **spectral gap**: the minimum distance between distinct eigenvalues of \mathcal{L}_0 . Assume:

1. $\sigma < \delta/4$ (noise is small compared to spectral gap)
2. $\epsilon = \delta/2$ (threshold is half the spectral gap)

Then the error in the spectral approximation of h^1 is bounded by:

$$|h_\epsilon^1(\mathcal{L}) - h_{true}^1(\mathcal{L}_0)| \leq \left\lceil \frac{2\sigma}{\delta} \right\rceil \quad (2)$$

where $h_\epsilon^1(\mathcal{L}) = \#\{i : \lambda_i(\mathcal{L}) < \epsilon\} - 1$ is the spectral approximation.

Proof Sketch. By Weyl's inequality, each eigenvalue of \mathcal{L} is within σ of the corresponding eigenvalue of \mathcal{L}_0 . The ideal zero block (from $H^0 \oplus H^1$) has $1 + h_{true}^1$ eigenvalues, all exactly zero. Under perturbation, these shift to $[-\sigma, \sigma]$. The positive block starts at δ and shifts to $[\delta - \sigma, \infty)$. With $\epsilon = \delta/2$ and $\sigma < \delta/4$, the threshold cleanly separates the two blocks, with at most $\lceil 2\sigma/\delta \rceil$ eigenvalues potentially crossing the threshold. See the main manuscript (Appendix A) for the detailed proof with all steps. \square \square

Theorem 3 (Computational Complexity). *Consider a graph $G = (V, E)$ with $N = |V|$ vertices, $M = |E|$ edges, and $M = O(N)$ (sparse graph). Let d be the stalk dimension, assumed to be a small constant (e.g., $d \in \{2, 4, 8\}$).*

The STPGC algorithm (Algorithm 1 in the main manuscript) computes the Phronesis Index Φ in:

$$O(Nd \log(Nd)) \text{ time} \quad (3)$$

Proof Sketch. The algorithm has three steps: (1) Construct \mathcal{L} : $O(Nd^2)$ time (iterating over edges and filling sparse matrix). (2) Compute $k = 20$ smallest eigenvalues via Lanczos: $O(Nd \cdot k)$ per iteration, $O(\log(Nd))$ iterations to converge, total $O(Nd \log(Nd))$. (3) Compute Φ : $O(k) = O(1)$. Since d is constant and $M = O(N)$, the dominant term is $O(Nd \log(Nd))$. See the main manuscript (Appendix A) for the full complexity analysis. \square \square

Remark 4 (Conditions for Theorem 2). *The condition $\sigma < \delta/4$ is necessary for the error bound to hold. If noise exceeds this level, eigenvalues from the positive block can be pushed below ϵ , causing h_ϵ^1 to overestimate h_{true}^1 arbitrarily. In practice, this means:*

- For systems with small spectral gap ($\delta < 0.01$), the method is sensitive to noise.
- For systems with large spectral gap ($\delta > 0.1$), the method is robust to moderate noise ($\sigma \approx 0.02$).
- If $\sigma \geq \delta/4$, use noise filtering (Section 2) or increase ϵ adaptively.

Remark 5 (Conditions for Theorem 3). *The $O(Nd \log(Nd))$ complexity assumes:*

1. **Sparse graph:** $M = O(N)$. For dense graphs ($M = O(N^2)$), Laplacian construction becomes $O(N^2 d^2)$, dominating the eigenvalue computation.
2. **Small stalk dimension:** $d = O(1)$. For large d (e.g., $d = 100$), the d^2 factor in Laplacian construction may dominate.
3. **Well-conditioned Laplacian:** Lanczos converges in $O(\log(Nd))$ iterations. For ill-conditioned matrices (e.g., very small spectral gap), convergence may be slower.

In our experiments, these conditions hold: graphs have $M \approx 4N$ (grid-like), stalks have $d \in \{2, 4\}$, and spectral gaps are $\delta > 0.01$.

1 Practical Guide to Sheaf Construction

This section provides step-by-step guidance for constructing cellular sheaves for multi-agent consistency monitoring. We present general principles followed by detailed walk-throughs of our three experimental scenarios.

1.1 General Principles

1.1.1 Step 1: Identify the Belief Graph

Question: What is the underlying network structure?

Action: Define the graph $G = (V, E)$ where:

- **Vertices V :** Represent *belief states* or *knowledge locations*
 - In multi-agent systems: vertices = agents or local knowledge bases
 - In state-space problems: vertices = states or situations
 - In sensor networks: vertices = sensor nodes or spatial locations
- **Edges E :** Represent *consistency relationships* or *information flow*
 - Connect vertices that share information or have overlapping beliefs
 - In physical systems: edges = communication links or spatial adjacency
 - In logical systems: edges = inference steps or constraint relationships

Example: In a team of 3 robots exploring an environment:

- Vertices: Robot 1, Robot 2, Robot 3
- Edges: (Robot 1, Robot 2), (Robot 2, Robot 3), (Robot 3, Robot 1) if they communicate in a triangle topology

1.1.2 Step 2: Define the Stalks (Local Data)

Question: What information does each vertex hold?

Action: For each vertex $v \in V$, define the stalk $\mathcal{F}(v)$ as a vector space representing the local belief or data:

- **Dimension d :** Depends on the type of information
 - Scalar beliefs: $d = 1$ (e.g., temperature readings)
 - Vector beliefs: $d = 2, 3, \dots$ (e.g., 2D positions, RGB colors, Q-values for multiple actions)
 - Structured beliefs: $d = \text{size of feature vector}$
- **Interpretation:** Each element of $\mathcal{F}(v) \cong \mathbb{R}^d$ represents a possible belief state at vertex v

Example: For robots tracking a target:

- $\mathcal{F}(\text{Robot 1}) = \mathbb{R}^2$: Robot 1's belief about target position (x, y)
- Similarly for Robot 2 and Robot 3

1.1.3 Step 3: Define Restriction Maps (Consistency Constraints)

Question: What does it mean for two neighboring vertices to be "consistent"?

Action: For each edge $e = (u, v) \in E$, define restriction maps $r_{e,u} : \mathcal{F}(u) \rightarrow \mathcal{F}(e)$ and $r_{e,v} : \mathcal{F}(v) \rightarrow \mathcal{F}(e)$ that encode the consistency requirement.

Common Patterns:

Pattern A: Identity Consistency (Agreement) If u and v should hold identical beliefs:

$$r_{e,u} = r_{e,v} = I \quad (\text{identity map}) \quad (4)$$

Then consistency means $x_u = x_v$.

Pattern B: Linear Transformation Consistency If beliefs are related by a known transformation T :

$$r_{e,u} = I, \quad r_{e,v} = T \quad (5)$$

Then consistency means $x_u = T(x_v)$.

Example: Robots with different coordinate frames:

- Robot 1 uses global coordinates
- Robot 2 uses local coordinates (rotated by θ)
- Restriction map: $r_{e,\text{Robot 2}} = R_\theta$ (rotation matrix)
- Consistency: Robot 1's global position = $R_\theta \times$ Robot 2's local position

Pattern C: Projection or Aggregation If one vertex has more information than another:

$$r_{e,u} : \mathbb{R}^{d_u} \rightarrow \mathbb{R}^{d_e}, \quad r_{e,v} : \mathbb{R}^{d_v} \rightarrow \mathbb{R}^{d_e} \quad (6)$$

where $d_e \leq \min(d_u, d_v)$.

Example: Sensor fusion:

- Sensor A measures (x, y, z) position
- Sensor B measures (x, y) position (no depth)
- Edge stalk: $\mathcal{F}(e) = \mathbb{R}^2$
- $r_{e,A}$: project $(x, y, z) \mapsto (x, y)$
- $r_{e,B}$: identity $(x, y) \mapsto (x, y)$
- Consistency: Sensor A's (x, y) projection matches Sensor B's (x, y)

1.1.4 Step 4: Construct the Connection Laplacian

Action: Build the matrix \mathcal{L} of size $Nd \times Nd$ where $N = |V|$ and d is the stalk dimension.

Block Structure:

$$\mathcal{L} = D - A \quad (7)$$

where:

- D : Block-diagonal degree matrix

$$D_{vv} = \deg(v) \cdot I_d \quad (8)$$

- A : Block adjacency matrix with restriction maps

$$A_{uv} = r_{e,v}^T r_{e,u} \quad \text{if } e = (u, v) \in E \quad (9)$$

Intuition: The Laplacian measures "how much" local beliefs disagree when propagated through restriction maps.

1.1.5 Step 5: Validate the Sheaf

Checklist:

1. **Regularity:** Are restriction maps linear? (Required for spectral methods)
2. **Symmetry:** Is the graph undirected? (Ensures \mathcal{L} is symmetric)
3. **Connectivity:** Is the graph connected? (Ensures $h^0 = 1$)
4. **Consistency Encoding:** Do the restriction maps truly capture the intended consistency relationships?

Test: Create a small example (3-5 vertices) and manually verify:

- A fully consistent global section (all constraints satisfied) should give $\mathcal{L}x = 0$
- An inconsistent configuration should give $\mathcal{L}x \neq 0$

1.2 Walk-Through 1: Logic Maze

1.2.1 Problem Description

An agent navigates a 5×5 grid maze. At each cell, the agent has a belief about its orientation (one of 4 directions: North, East, South, West). The agent receives local observations (e.g., "wall on left") that constrain relative orientations between adjacent cells. A contradiction occurs if the constraints form an inconsistent cycle.

1.2.2 Step 1: Belief Graph

- **Vertices:** 25 cells in the 5×5 grid
- **Edges:** Connect adjacent cells (up/down/left/right neighbors)
- **Graph structure:** 2D grid graph (40 edges for a 5×5 grid)

1.2.3 Step 2: Stalks

Representation of Orientation:

We use $SO(2)$ (2D rotation group) to represent orientations. Each orientation is a unit vector in \mathbb{R}^2 :

- North: $(0, 1)$
- East: $(1, 0)$
- South: $(0, -1)$
- West: $(-1, 0)$

Stalk Definition:

$$\mathcal{F}(\text{cell}) = \mathbb{R}^2 \quad (10)$$

Each cell's stalk holds a 2D vector representing the agent's believed orientation at that cell.

1.2.4 Step 3: Restriction Maps

Consistency Requirement:

If the agent moves from cell u to cell v , its orientation should transform according to the observed turn:

- No turn: orientations should match
- Left turn: orientation rotates 90° counterclockwise
- Right turn: orientation rotates 90° clockwise

Restriction Map:

For edge $e = (u, v)$ with observed turn θ_e :

$$r_{e,u} = I, \quad r_{e,v} = R_{\theta_e} \quad (11)$$

where R_θ is the 2D rotation matrix:

$$R_\theta = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \quad (12)$$

Example:

- Agent at cell u faces North: $x_u = (0, 1)$
- Agent moves to cell v with a left turn ($\theta = 90^\circ$)
- Consistency: $x_u = R_{90^\circ} x_v$
- If $x_v = (1, 0)$ (facing East), then $R_{90^\circ} x_v = (0, 1)$ (consistent)

1.2.5 Step 4: Injecting a Contradiction

Scenario: At time $t = 50$, we inject a false observation at the center cell (cell 12):

- The agent suddenly believes it faces South instead of North
- This creates a cycle of inconsistent orientations around cell 12

Effect on Sheaf:

- Before injection: $h^1 = 0$ (all orientations consistent)
- After injection: $h^1 = 1$ (one independent cycle of contradiction)
- λ_1 decreases (frustration in the system)
- Φ drops sharply (detected as anomaly)

1.2.6 Implementation Code Snippet

```
def construct_logic_maze_sheaf(grid_size=5):
    # Step 1: Create grid graph
    G = nx.grid_2d_graph(grid_size, grid_size)
    G = nx.convert_node_labels_to_integers(G)

    # Step 2: Define stalks (SO(2) orientations)
    stalk_dim = 2
    stalks = {v: np.random.randn(stalk_dim) for v in G.nodes()}
    # Normalize to unit vectors
    for v in stalks:
        stalks[v] /= np.linalg.norm(stalks[v])

    # Step 3: Define restriction maps (rotation matrices)
    restriction_maps = {}
    for edge in G.edges():
        u, v = edge
        # Random turn angle (or from observations)
        theta = np.random.choice([0, np.pi/2, -np.pi/2])
        R = np.array([[np.cos(theta), -np.sin(theta)],
                      [np.sin(theta), np.cos(theta)]])
        restriction_maps[edge] = R

    # Step 4: Construct Connection Laplacian
    L = construct_connection_laplacian(G, stalks, restriction_maps)

    return G, stalks, restriction_maps, L

def inject_contradiction(stalks, center_node=12):
    # Flip orientation at center node
    stalks[center_node] = -stalks[center_node]
```

1.3 Walk-Through 2: Safety Gym (Safe RL)

1.3.1 Problem Description

An RL agent learns to navigate a 2D environment (Safety Gym) while avoiding hazards. The agent maintains a belief graph where vertices represent states and edges represent state transitions. Each state has an associated Q-value vector (expected rewards for each action). Bellman consistency requires that Q-values satisfy the Bellman equation across transitions.

1.3.2 Step 1: Belief Graph

Discretization:

- Continuous state space $(x, y) \in [0, 10] \times [0, 10]$ discretized into a 10×10 grid
- **Vertices:** 100 grid cells (states)
- **Edges:** Connect states reachable by one action (4-connected grid)

Dynamic Graph:

- Edges are added as the agent explores and experiences transitions
- Initially sparse, becomes denser with training

1.3.3 Step 2: Stalks

Q-Values:

Each state has Q-values for 4 actions: {up, down, left, right}

Stalk Definition:

$$\mathcal{F}(\text{state}) = \mathbb{R}^4 \quad (13)$$

Element i of the stalk represents $Q(s, a_i)$, the expected return for taking action a_i in state s .

1.3.4 Step 3: Restriction Maps (Bellman Consistency)

Bellman Equation:

For a transition from state s to state s' via action a with reward r :

$$Q(s, a) = r + \gamma \max_{a'} Q(s', a') \quad (14)$$

Consistency Requirement:

The Q-value at s for action a should equal the discounted max Q-value at s' plus the reward.

Restriction Map Design:

For edge $e = (s, s')$ corresponding to action a :

- Edge stalk: $\mathcal{F}(e) = \mathbb{R}$ (scalar, the Q-value for that transition)
- $r_{e,s} : \mathbb{R}^4 \rightarrow \mathbb{R}$: Extract $Q(s, a)$ (select component a)
- $r_{e,s'} : \mathbb{R}^4 \rightarrow \mathbb{R}$: Compute $r + \gamma \max_{a'} Q(s', a')$

Mathematically:

Let e_a be the unit vector selecting action a :

$$r_{e,s}(Q_s) = e_a^T Q_s \quad (15)$$

$$r_{e,s'}(Q_{s'}) = r + \gamma \max(Q_{s'}) \quad (16)$$

Consistency: $r_{e,s}(Q_s) = r_{e,s'}(Q_{s'})$ means the Bellman equation is satisfied.

1.3.5 Step 4: Detecting Safety Violations

Mechanism:

- During training, the agent updates Q-values based on experience
- If the agent encounters a hazard (safety violation), it receives a large negative reward
- This creates a sudden inconsistency: the Q-value for the action leading to the hazard should drop, but neighboring Q-values may not have updated yet
- This inconsistency manifests as increased h^1 (a cycle of Bellman violations)
- Φ drops, signaling danger

Using Φ as Auxiliary Reward:

We modify the reward function:

$$r'(s, a) = r(s, a) + \alpha \cdot \Phi(s) \quad (17)$$

where $\alpha = 0.1$ is a scaling factor.

Effect:

- High Φ (consistent beliefs): normal reward
- Low Φ (inconsistent beliefs): penalty, discouraging risky actions
- Agent learns to avoid states where its Q-value estimates are internally contradictory (often near hazards)

1.3.6 Implementation Code Snippet

```
def construct_safety_gym_sheaf(state_graph, q_values, gamma=0.99):
    G = state_graph
    stalk_dim = 4 # 4 actions

    # Step 2: Stalks are Q-value vectors
    stalks = {s: q_values[s] for s in G.nodes()}

    # Step 3: Restriction maps (Bellman consistency)
    restriction_maps = {}
    for edge in G.edges():
        s, s_prime = edge
        action = edge_to_action(edge) # Determine which action
        reward = get_reward(s, action, s_prime)

        # r_{e,s}: select Q(s, action)
        r_e_s = np.zeros((1, 4))
        r_e_s[0, action] = 1.0

        # r_{e,s'}: compute r + gamma * max Q(s')
        r_e_s_prime = lambda Q: reward + gamma * np.max(Q)

    # For linear Laplacian, approximate max as weighted sum
    # (or use linearization around current Q-values)
```

```

restriction_maps[edge] = (r_e_s, r_e_s_prime)

L = construct_connection_laplacian(G, stalks, restriction_maps)
return L

def compute_phronesis_reward(state, L, epsilon=1e-3):
    Phi, h1, lambda1 = compute_phronesis_index(L, epsilon)
    return Phi # Higher Phi = more consistent = safer

```

1.4 Walk-Through 3: Multi-Robot Coordination

1.4.1 Problem Description

Three robots explore an environment and share observations about a target's location. Each robot has its own belief about the target position (x, y) . Robots communicate over a network (triangle topology). A contradiction arises when Robot 1 believes the target is at $(2, 3)$, Robot 2 believes it's at $(7, 8)$, and Robot 3 has yet another belief, creating a cycle of incompatible observations.

1.4.2 Step 1: Belief Graph

Two-Level Structure:

- **Local Level:** Each robot has its own 10×10 grid belief graph (100 vertices per robot)
- **Global Level:** Robots are connected via a communication graph (3 vertices, 3 edges forming a triangle)

Combined Graph:

- Total vertices: $3 \times 100 = 300$ (local states) + 3 (robot nodes) = 303
- Or, simplified: 3 robot nodes with aggregated beliefs

For our experiments, we use the simplified version:

- **Vertices:** Robot 1, Robot 2, Robot 3
- **Edges:** (R1, R2), (R2, R3), (R3, R1)

1.4.3 Step 2: Stalks

Target Position Belief:

Each robot's stalk represents its belief about the target's 2D position.

Stalk Definition:

$$\mathcal{F}(\text{Robot } i) = \mathbb{R}^2 \quad (18)$$

Element $(x, y) \in \mathbb{R}^2$ represents the robot's belief about where the target is located.

1.4.4 Step 3: Restriction Maps (Spatial Consistency)

Consistency Requirement:

If two robots communicate, their beliefs about the target location should be compatible (within sensor error).

Restriction Map:

For edge $e = (\text{Robot } i, \text{Robot } j)$:

$$r_{e,i} = r_{e,j} = I \quad (\text{identity}) \quad (19)$$

Interpretation: Robots should agree on the target position. Consistency means $x_i \approx x_j$.

Alternative (with coordinate transforms):

If robots use different coordinate frames:

$$r_{e,i} = I, \quad r_{e,j} = T_{ij} \quad (20)$$

where T_{ij} is the coordinate transformation from Robot j 's frame to Robot i 's frame.

1.4.5 Step 4: Injecting a Contradiction

Scenario:

- Robot 1 observes target at $(2, 3)$ in global coordinates
- Robot 2 observes target at $(7, 8)$ in global coordinates
- Robot 3 observes target at $(5, 5)$ (average, but still inconsistent with both)

Effect on Sheaf:

- The three beliefs form a triangle of disagreements
- No global section exists that satisfies all pairwise consistency constraints
- $h^1 = 1$ (one independent cycle of contradiction)
- Φ drops significantly

1.4.6 Step 5: Resolution via Negotiation

Strategy:

Robots detect low Φ and initiate a negotiation protocol:

1. Compute the average belief: $\bar{x} = \frac{1}{3}(x_1 + x_2 + x_3) = (4.67, 5.33)$
2. Each robot updates its belief to $\bar{x} + \text{small noise}$
3. Re-compute Φ

Result:

- After averaging: $h^1 = 0$ (no contradiction)
- Φ increases (consistency restored)
- Coordination success improves

1.4.7 Implementation Code Snippet

```
def construct_multi_robot_sheaf(num_robots=3):
    # Step 1: Communication graph (triangle)
    G = nx.cycle_graph(num_robots)

    # Step 2: Stalks are 2D target position beliefs
    stalk_dim = 2
    stalks = {i: np.random.randn(stalk_dim) * 5 for i in G.nodes()}

    # Step 3: Restriction maps (identity for agreement)
    restriction_maps = {}
    for edge in G.edges():
        restriction_maps[edge] = np.eye(stalk_dim)

    L = construct_connection_laplacian(G, stalks, restriction_maps)
    return G, stalks, L

def inject_contradiction(stalks):
    stalks[0] = np.array([2.0, 3.0]) # Robot 1 belief
    stalks[1] = np.array([7.0, 8.0]) # Robot 2 belief
    stalks[2] = np.array([5.0, 5.0]) # Robot 3 belief

def resolve_via_negotiation(stalks):
    avg = np.mean([stalks[i] for i in stalks], axis=0)
    for i in stalks:
        stalks[i] = avg + np.random.randn(2) * 0.5 # Small noise
```

1.5 Summary: Decision Tree for Sheaf Construction

START: What kind of multi-agent system do you have?

```
+-- Logical/Symbolic Knowledge
|   +-- Stalks: Propositional assignments or logical states
|   +-- Restriction: Logical consistency (e.g., not(A and not A))
|   +-- Example: Logic Maze (orientation consistency)
|
+-- Numerical/Continuous Beliefs
|   +-- Stalks: Real-valued vectors (positions, Q-values, etc.)
|   +-- Restriction: Equality or linear transformation
|   +-- Example: Multi-Robot (spatial consistency)
|
+-- Dynamical Systems (RL/Control)
|   +-- Stalks: State-action values or policy parameters
|   +-- Restriction: Bellman equation or dynamics model
|   +-- Example: Safety Gym (Q-value consistency)
|
+-- Hierarchical/Nested Systems
    +-- Stalks: Multi-level representations
    +-- Restriction: Aggregation or projection maps
    +-- Example: Federated learning (local vs global models)
```

Key Takeaway: The sheaf construction is domain-specific but follows a systematic process.

Start with the graph structure, identify what "consistency" means in your domain, and encode it as restriction maps. Validate with small examples before scaling up.

2 Robustness Analysis

Real-world multi-agent systems operate in noisy, uncertain environments. This section analyzes the robustness of the Phronesis Index to various sources of noise and provides practical strategies for maintaining accuracy under adverse conditions.

2.1 Sources of Noise in Multi-Agent Systems

2.1.1 Sensor Noise

Description: Agents' observations are corrupted by measurement errors.

Effect on Sheaf: Stalks contain noisy beliefs $\tilde{x}_v = x_v + \eta_v$ where $\eta_v \sim \mathcal{N}(0, \sigma^2 I)$.

Propagation: Noise propagates through restriction maps, affecting the Connection Laplacian:

$$\tilde{\mathcal{L}} = \mathcal{L} + E \quad (21)$$

where E is a perturbation matrix with $\|E\| \leq c\sigma$ for some constant c depending on graph structure.

2.1.2 Communication Errors

Description: Information shared between agents is corrupted during transmission.

Effect on Sheaf: Restriction maps are perturbed: $\tilde{r}_{e,v} = r_{e,v} + \Delta r_{e,v}$.

Modeling: We model this as additive noise in the Laplacian blocks corresponding to edges.

2.1.3 Discretization Error

Description: Continuous state spaces are discretized into finite graphs, introducing approximation errors.

Effect on Sheaf: The graph G is an approximation of the true continuous manifold. Eigenvalues of \mathcal{L} approximate those of a continuous Laplace-Beltrami operator, with error $O(h^2)$ where h is the discretization step size.

2.1.4 Model Uncertainty

Description: Restriction maps are based on approximate models (e.g., linearized dynamics, simplified physics).

Effect on Sheaf: The sheaf structure itself is an approximation of the true consistency relationships.

2.2 Sensitivity Analysis

2.2.1 Eigenvalue Perturbation Theory

Weyl's Theorem (Restated):

For symmetric matrices A and $B = A + E$ with eigenvalues $\lambda_i(A)$ and $\lambda_i(B)$:

$$|\lambda_i(B) - \lambda_i(A)| \leq \|E\|_2 \quad (22)$$

Application to Φ :

Let \mathcal{L}_0 be the ideal Laplacian and $\mathcal{L} = \mathcal{L}_0 + E$ be the noisy version with $\|E\| \leq \sigma$.

Effect on λ_1 :

$$|\lambda_1(\mathcal{L}) - \lambda_1(\mathcal{L}_0)| \leq \sigma \quad (23)$$

Relative Error:

$$\frac{|\lambda_1(\mathcal{L}) - \lambda_1(\mathcal{L}_0)|}{\lambda_1(\mathcal{L}_0)} \leq \frac{\sigma}{\lambda_1(\mathcal{L}_0)} \quad (24)$$

Interpretation: If λ_1 is large (strong consensus dynamics), the relative error is small. If λ_1 is small (weak coupling), noise has a larger relative impact.

Effect on h^1 : From Theorem 2, the error in counting near-zero eigenvalues is:

$$|h_\epsilon^1(\mathcal{L}) - h_\epsilon^1(\mathcal{L}_0)| \leq \left\lceil \frac{2\sigma}{\delta} \right\rceil \quad (25)$$

where δ is the spectral gap.

Interpretation: A large spectral gap δ makes h^1 estimation robust to noise. A small gap makes it sensitive.

Combined Effect on Φ :

$$\Phi = \frac{\lambda_1}{h^1 + \epsilon} \quad (26)$$

Error Propagation:

Using first-order Taylor expansion:

$$\Delta\Phi \approx \frac{\partial\Phi}{\partial\lambda_1} \Delta\lambda_1 + \frac{\partial\Phi}{\partial h^1} \Delta h^1 \quad (27)$$

$$= \frac{1}{h^1 + \epsilon} \Delta\lambda_1 - \frac{\lambda_1}{(h^1 + \epsilon)^2} \Delta h^1 \quad (28)$$

Worst-Case Bound:

$$|\Delta\Phi| \leq \frac{\sigma}{h^1 + \epsilon} + \frac{\lambda_1}{(h^1 + \epsilon)^2} \cdot \frac{2\sigma}{\delta} \quad (29)$$

Simplified:

$$|\Delta\Phi| \leq \sigma \left(\frac{1}{h^1 + \epsilon} + \frac{2\lambda_1}{\delta(h^1 + \epsilon)^2} \right) \quad (30)$$

Key Insight: The error in Φ is linear in noise level σ and inversely proportional to spectral gap δ . Systems with good spectral properties are naturally robust.

2.2.2 Illustrative Sensitivity Example

To illustrate the theoretical bounds, consider the Logic Maze scenario (5×5 grid, 25 vertices) with spectral gap $\delta \approx 0.5$ and $\lambda_1 \approx 1.5$:

- For noise $\sigma = 0.01$: the bound gives $|\Delta\lambda_1| \leq 0.01$, so Φ changes by $< 1\%$. The condition $\sigma < \delta/4 = 0.125$ is easily satisfied.
- For noise $\sigma = 0.1$: the bound gives $|\Delta\lambda_1| \leq 0.1$, and $\lceil 2\sigma/\delta \rceil = 1$, so at most one eigenvalue may be misclassified. Φ is still within the graceful-degradation regime.
- For noise $\sigma = 0.5$: now $\sigma = \delta$, violating the condition $\sigma < \delta/4$. The eigenvalue perturbation can create spurious near-zero modes, inflating h^1 and collapsing Φ .

Key takeaway: For well-separated spectra ($\delta \gg \sigma$), Φ is robust. When σ approaches δ , a phase transition occurs where noise-induced eigenvalues cross the ϵ threshold, making h^1 estimation unreliable. Running a controlled noise-injection study on our grid-world experiments is left to future work.

Practical Guideline: Ensure $\sigma < \delta/2$ through preprocessing (filtering, averaging, outlier removal).

2.3 Adaptive Threshold Selection

2.3.1 The Challenge of Choosing ϵ

The threshold ϵ separates "zero" eigenvalues (from H^0 and H^1) from "positive" eigenvalues (from higher modes). Choosing ϵ is critical:

- **Too small:** Noise-induced perturbations may push true zero eigenvalues above ϵ , undercounting h^1 .
- **Too large:** Small positive eigenvalues may be mistaken for zeros, overcounting h^1 .

2.3.2 Strategy 1: Spectral Gap Estimation

Algorithm:

1. Compute the k smallest eigenvalues: $\lambda_0, \lambda_1, \dots, \lambda_{k-1}$ (using Lanczos, $k = 20$).
2. Sort them: $\lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_{k-1}$.
3. Identify the largest gap: $\delta_{\max} = \max_i(\lambda_{i+1} - \lambda_i)$.
4. Set $\epsilon = \lambda_i + \delta_{\max}/2$ where i is the index before the largest gap.

Intuition: The spectral gap separates the "kernel block" (near-zero eigenvalues) from the "positive block". We place ϵ in the middle of this gap.

Example:

Eigenvalues: [0.000, 0.002, 0.003, 0.500, 0.520, 0.540, ...]

Gaps: [0.002, 0.001, 0.497, 0.020, 0.020, ...]

Largest gap: 0.497 between $\lambda_2 = 0.003$ and $\lambda_3 = 0.500$.

Set $\epsilon = 0.003 + 0.497/2 \approx 0.25$.

Result: $h^1 = 3 - 1 = 2$ (three eigenvalues below ϵ , minus one for H^0).

2.3.3 Strategy 2: Noise-Adaptive Threshold

Algorithm:

1. Estimate noise level $\hat{\sigma}$ from data (e.g., via residual variance).
2. Estimate spectral gap $\hat{\delta}$ (as above).
3. Set $\epsilon = 2\hat{\sigma}$ (two standard deviations above zero).
4. If $\epsilon > \hat{\delta}/2$, issue a warning: "Noise level too high for reliable h^1 estimation."

Rationale: Eigenvalues perturbed by noise of magnitude σ will fluctuate within $\pm\sigma$. Setting $\epsilon = 2\sigma$ ensures we don't miscount due to noise fluctuations (with 95% confidence under Gaussian noise).

2.3.4 Strategy 3: Cross-Validation

Algorithm:

1. Split data into K folds (e.g., $K = 5$).
2. For each fold, construct the sheaf and compute Φ with various $\epsilon \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$.
3. Select ϵ that minimizes variance of Φ across folds.

Rationale: A good ϵ should give consistent Φ estimates across different data samples. High variance indicates ϵ is in an unstable region (near the spectral gap).

2.3.5 Recommended Default

For practitioners without domain-specific knowledge:

$$\epsilon = \max(10^{-3}, 2\hat{\sigma}) \quad (31)$$

This ensures:

- Numerical stability (lower bound 10^{-3})
- Noise robustness (scales with estimated noise)

2.4 Noise Filtering Techniques

2.4.1 Preprocessing: Belief Smoothing

Idea: Before constructing the sheaf, smooth the belief graph to reduce noise.

Method 1: Spatial Averaging

For graph-structured data, apply a diffusion filter:

$$\tilde{x}_v = \frac{1}{\deg(v) + 1} \left(x_v + \sum_{u \in \mathcal{N}(v)} x_u \right) \quad (32)$$

where $\mathcal{N}(v)$ are the neighbors of v .

Effect: Reduces high-frequency noise while preserving large-scale structure.

Method 2: Median Filtering

For outlier-prone data:

$$\tilde{x}_v = \text{median}\{x_u : u \in \mathcal{N}(v) \cup \{v\}\} \quad (33)$$

Effect: Robust to outliers (e.g., a single agent with a wildly incorrect belief).

2.4.2 Postprocessing: Robust Spectral Estimation

Idea: Use robust eigenvalue solvers that are less sensitive to noise.

Method 1: Regularized Laplacian

Add a small regularization term:

$$\mathcal{L}_{\text{reg}} = \mathcal{L} + \mu I \quad (34)$$

where $\mu \approx 10^{-6}$.

Effect: Shifts all eigenvalues up by μ , moving them away from the numerical zero threshold. This reduces false positives in h^1 counting.

Method 2: Truncated Eigenvalue Decomposition

Compute only the k smallest eigenvalues (e.g., $k = 20$) using Lanczos. Ignore eigenvalues beyond k .

Effect: Focuses computation on the relevant part of the spectrum, avoiding numerical artifacts in the tail.

2.4.3 Online Filtering: Kalman Smoothing

For time-series data (e.g., RL training), apply a Kalman filter to $\Phi(t)$:

$$\hat{\Phi}(t) = \alpha\Phi(t) + (1 - \alpha)\hat{\Phi}(t - 1) \quad (35)$$

where $\alpha = 0.1$ is a smoothing parameter.

Effect: Reduces temporal noise, making Φ trends more interpretable.

2.5 Robustness Guarantees

2.5.1 Theorem: Bounded Degradation

Statement: Under the assumptions of Theorem 2, if the noise level $\sigma < \delta/4$ and $\epsilon = \delta/2$, then:

$$\Phi(\mathcal{L}) \geq \frac{1}{2}\Phi(\mathcal{L}_0) \quad (36)$$

with probability $\geq 1 - \exp(-N/2)$ (for Gaussian noise).

Proof Sketch:

Step 1: By Weyl's theorem, $\lambda_1(\mathcal{L}) \geq \lambda_1(\mathcal{L}_0) - \sigma \geq \lambda_1(\mathcal{L}_0)/2$ (since $\sigma < \lambda_1/2$ by spectral gap assumption).

Step 2: By Theorem 2, $h^1(\mathcal{L}) \leq h^1(\mathcal{L}_0) + 1$ (since $\lceil 2\sigma/\delta \rceil \leq 1$ when $\sigma < \delta/4$).

Step 3: Combine:

$$\Phi(\mathcal{L}) = \frac{\lambda_1(\mathcal{L})}{h^1(\mathcal{L}) + \epsilon} \geq \frac{\lambda_1(\mathcal{L}_0)/2}{h^1(\mathcal{L}_0) + 1 + \epsilon} \geq \frac{1}{2} \cdot \frac{\lambda_1(\mathcal{L}_0)}{h^1(\mathcal{L}_0) + \epsilon} = \frac{1}{2}\Phi(\mathcal{L}_0) \quad (37)$$

(assuming $h^1(\mathcal{L}_0) \geq 1$ for the inequality; if $h^1 = 0$, the bound is tighter). \square

Interpretation: In low-noise regimes, Φ degrades by at most a factor of 2. This is a *graceful degradation* property: the index remains useful even under perturbations.

2.5.2 Failure Modes

When does the method break down?

1. **Noise \gg Spectral Gap:** If $\sigma > \delta$, the spectrum is completely scrambled. h^1 estimation becomes unreliable.

Mitigation: Apply aggressive filtering, increase graph connectivity (to increase δ), or use alternative consistency measures (e.g., direct cohomology computation via linear algebra, accepting the $O(N^3)$ cost).

2. **Adversarial Noise:** If noise is adversarially chosen to create spurious cycles, h^1 will be inflated.

Mitigation: Use robust statistics (median, trimmed mean) instead of raw beliefs. Employ anomaly detection to identify and exclude adversarial agents.

3. **Model Mismatch:** If the sheaf structure (restriction maps) does not accurately reflect true consistency relationships, Φ may give misleading signals.

Mitigation: Validate the sheaf design on ground-truth data. Use domain expertise to refine restriction maps. Consider adaptive sheaf learning (future work).

2.6 Experimental Validation of Robustness

2.6.1 Theoretical Noise Tolerance of the Grid-World Experiment

The theoretical bounds derived above allow us to predict the robustness of Φ in our grid-world Bellman consistency experiment without running a separate noise-injection study. In the 8×8 grid MDP:

- The spectral gap of the Connection Laplacian is $\delta \approx 0.05\text{--}0.15$ (observed across 10 seeds).
- Applying the condition $\sigma < \delta/4$ gives a noise tolerance of roughly $\sigma < 0.01\text{--}0.04$.
- At typical Q-learning estimation noise ($\sigma \approx 0.01$ per stalk entry during late training), the eigenvalue perturbation bound predicts $|\Delta\lambda_1| \leq 0.01$, which is small relative to $\lambda_1 \approx 0.3$.

Prediction: Based on the bounded-degradation theorem above, we expect Φ to remain within a factor of 2 of its noise-free value for our grid-world experiments, because the Q-learning estimation noise satisfies $\sigma < \delta/4$ once training has partially converged. Validating this prediction under controlled noise injection (varying σ explicitly) is left to future work.

Practical implication: Practitioners should verify that their measurement noise satisfies $\sigma < \delta/4$ for their specific graph structure before relying on Φ as a safety signal. If this condition is not met, the noise-filtering techniques described below should be applied first.

2.6.2 Expected Benefit of Adaptive ϵ Selection

Intuition: The three strategies differ in how they handle noise:

- **Fixed $\epsilon = 10^{-3}$:** Works well when noise is negligible ($\sigma \ll 10^{-3}$) but misclassifies eigenvalues when noise grows.
- **Gap-based $\epsilon = \delta_{\max}/2$:** Adapts to spectrum structure; robust when the spectral gap is clearly identifiable.
- **Adaptive $\epsilon = 2\hat{\sigma}$:** Tracks the noise level directly; should degrade gracefully as noise increases.

The theoretical analysis predicts that adaptive methods will dominate fixed thresholds at moderate-to-high noise, because fixed ϵ either under- or over-counts eigenvalues as the noise floor shifts. We use the gap-based strategy in all our experiments and recommend it as the default for practitioners. A rigorous empirical comparison of these three strategies across multiple noise regimes is left to future work.

Recommendation: We recommend using the gap-based or adaptive strategy in practice.

2.7 Practical Recommendations

For Practitioners:

1. **Estimate Noise Level:** Before deployment, characterize the noise in your system (e.g., via test data or simulation).
2. **Choose ϵ Adaptively:** Use the spectral gap method (Section 2.8) rather than a fixed threshold.
3. **Preprocess Data:** Apply spatial averaging or median filtering to reduce noise before constructing the sheaf.
4. **Monitor Φ Trends:** Use temporal smoothing (Kalman filter) to track $\Phi(t)$ over time, focusing on trends rather than instantaneous values.

5. **Set Alerts:** Define a threshold Φ_{crit} (e.g., $\Phi < 0.5$) below which the system should trigger a warning or initiate a recovery protocol.
6. **Validate on Ground Truth:** If possible, compute h^1 directly on a small subset of data to validate the spectral approximation.

For Researchers:

1. **Extend to Non-Gaussian Noise:** Current analysis assumes Gaussian noise. Investigate robustness under heavy-tailed or adversarial noise distributions.
2. **Adaptive Sheaf Learning:** Develop methods to learn restriction maps from data, reducing reliance on manual design.
3. **Higher-Order Robustness:** Analyze sensitivity of higher cohomology groups (H^2, H^3, \dots) for simplicial complex extensions.
4. **Real-Time Adaptation:** Design online algorithms that adjust ϵ and filtering parameters dynamically as the system evolves.

2.8 Parameter Selection Guide

The Phronesis Index depends on two key parameters: the eigenvalue threshold ϵ and the reward shaping coefficient α (when used in RL). This section provides quantitative guidance for selecting these parameters.

2.8.1 Choosing ϵ : Spectral Gap Estimation

The threshold ϵ determines which eigenvalues are counted as "near-zero" for approximating h^1 . The optimal choice depends on the *spectral gap* $\delta = \lambda_1^+ - \lambda_0$ (the distance between the smallest positive eigenvalue and zero).

Theoretical Guideline: From Theorem 2, to ensure $|h_\epsilon^1 - h_{\text{true}}^1| \leq 1$ with high probability under noise level σ , we require:

$$\epsilon < \frac{\delta}{2} - 2\sigma \tag{38}$$

Practical Procedure:

1. **Initial estimate:** Run a pilot computation of the Connection Laplacian on a representative subgraph (e.g., 100-500 vertices).
2. **Compute spectrum:** Extract the smallest 20 eigenvalues using Lanczos iteration.
3. **Identify gap:** Plot the eigenvalues and visually identify the gap between near-zero and positive eigenvalues. Typically, $\lambda_0 \approx 10^{-10}$ (numerical zero) and $\lambda_1^+ \in [10^{-4}, 10^{-1}]$ depending on graph connectivity.
4. **Set threshold:** Choose $\epsilon = 0.1 \times \lambda_1^+$ as a conservative estimate. This ensures ϵ is well below the spectral gap while remaining above numerical noise.
5. **Validate:** Check that h_ϵ^1 remains stable when ϵ is varied by $\pm 50\%$. If h_ϵ^1 changes significantly, the spectral gap may be too small, indicating high ambiguity in the system.

Example Values:

- **Logic Maze (5×5 grid):** $\lambda_1^+ \approx 0.05 \Rightarrow \epsilon = 0.005$
- **Safety Gym (10×10 grid):** $\lambda_1^+ \approx 0.02 \Rightarrow \epsilon = 0.002$
- **Multi-Robot (10 agents):** $\lambda_1^+ \approx 0.08 \Rightarrow \epsilon = 0.008$
- **Scalability (50k agents):** $\lambda_1^+ \approx 0.001 \Rightarrow \epsilon = 0.0001$

Sensitivity Analysis: Figure 1 shows how Φ varies with ϵ for the Safety Gym scenario. The plateau region ($\epsilon \in [0.001, 0.01]$) indicates robust parameter selection. Outside this range, Φ either overcounts noise (ϵ too large) or misses true inconsistencies (ϵ too small).



Figure 1: Sensitivity of Φ to threshold ϵ in Safety Gym. The shaded region indicates the robust selection range where h_ϵ^1 is stable. Our choice $\epsilon = 0.002$ falls in the middle of this plateau.

2.8.2 Choosing α : Reward Shaping Coefficient

When integrating Φ into reinforcement learning via reward shaping ($r' = r + \alpha \cdot \Phi$), the coefficient α balances task performance (original reward r) against consistency maintenance (auxiliary signal Φ).

Theoretical Consideration: The modified reward r' should preserve the optimal policy structure while providing a consistency incentive. If α is too small, the agent ignores Φ ; if too large, the agent sacrifices task performance for consistency.

Practical Procedure:

1. **Normalize scales:** Measure the typical ranges of r and Φ in pilot runs. In our grid-world MDP, $r \in [-1, 10]$ and $\Phi \in [0, 5]$.
2. **Grid search:** Test $\alpha \in \{0.01, 0.05, 0.1, 0.5, 1.0\}$ over 5 independent training runs.
3. **Evaluate trade-off:** For each α , record both task success rate and safety cost. Plot the Pareto frontier.
4. **Select optimum:** Choose the α that achieves the best trade-off. In our grid-world experiments, $\alpha = 0.1$ provided a reasonable cost reduction with minimal success rate degradation.
5. **Validate:** Run full training (500 episodes, 10 seeds) with the selected α to confirm statistical significance.

Example Values:

- **Grid-world MDP** (8×8): $\alpha = 0.1$ (selected via pilot runs)
- **Multi-Robot coordination:** $\alpha = 0.2$ (higher weight on consistency for safety-critical tasks)

Sensitivity Analysis: The optimal α is problem-dependent. In our grid-world experiment, $\alpha \in [0.05, 0.2]$ consistently outperformed the baseline (no reward shaping, $\alpha = 0$). Extreme values ($\alpha \geq 1.0$) caused the agent to ignore the environment reward; very small values ($\alpha \leq 0.01$) had negligible effect. Because our grid-world MDP is intentionally simple, we do not present a detailed sensitivity table; extending the α -sweep to richer environments (e.g., continuous-control Safety Gym tasks) is future work.

Key Insight: The optimal α is problem-dependent but typically falls in the range $[0.05, 0.2]$. A simple heuristic is to start with $\alpha = 0.1$ and adjust based on the observed trade-off between task performance and safety.

2.8.3 Automated Parameter Tuning

For practitioners who prefer automated selection, we provide a simple adaptive procedure:

Algorithm 1 Adaptive Parameter Selection

Require: Graph G , pilot data D_{pilot} , RL environment \mathcal{E}

Ensure: Optimal ϵ^* , α^*

Construct Connection Laplacian \mathcal{L} from pilot data

Compute eigenvalues $\{\lambda_i\}_{i=0}^{19}$ using Lanczos

$\lambda_1^+ \leftarrow \min\{\lambda_i : \lambda_i > 10^{-8}\}$

$\epsilon^* \leftarrow 0.1 \times \lambda_1^+$

▷ Conservative threshold

$\mathcal{A} \leftarrow \{0.01, 0.05, 0.1, 0.2, 0.5\}$

▷ Candidate α values

for $\alpha \in \mathcal{A}$ **do**

Train RL agent with $r' = r + \alpha \cdot \Phi$ for 100k steps

Evaluate on 50 test episodes: record S_α (success rate), C_α (cost)

end for

$\alpha^* \leftarrow \arg \max_\alpha (S_\alpha - \beta \cdot C_\alpha)$

▷ β is cost penalty weight

return ϵ^* , α^*

This procedure requires $\approx 500k$ total environment steps (5 candidates \times 100k steps) and typically completes in < 3 hours on a standard workstation.

2.9 Practitioner’s Checklist: Applying the Method

Before presenting the mathematical foundations, we provide a practical checklist for applying the Phronesis Index to a new system. This checklist distills the method into actionable steps, clarifying what domain expertise is required and what can be automated.

Step 1: Define the Graph Question: What are the agents/entities, and how do they interact?

Action: Construct a graph $G = (V, E)$ where:

- **Vertices V :** Represent agents, states, or entities in your system.
 - Multi-agent system: vertices = agents
 - Reinforcement learning: vertices = states in the state space
 - Sensor network: vertices = sensors or measurement locations
- **Edges E :** Represent relationships, communication links, or transitions.
 - Multi-agent system: edges = communication links (who talks to whom)
 - Reinforcement learning: edges = state transitions (which states are reachable)
 - Sensor network: edges = spatial proximity (which sensors should agree)

Expertise required: Domain knowledge to identify relevant entities and relationships.

Can be automated: Partially. For RL, transitions can be extracted from experience replay. For sensor networks, proximity graphs can be computed from GPS coordinates.

Step 2: Define Stalks (Local Information) Question: What information does each vertex hold?

Action: For each vertex v , define a stalk $\mathcal{F}(v) = \mathbb{R}^d$ representing the local information at v :

- Multi-agent system: $d = 2$ (2D position), $d = 3$ (3D position), or $d = k$ (belief vector)
- Reinforcement learning: $d = |A|$ (Q-values for $|A|$ actions)
- Sensor network: $d = m$ (measurements of m quantities, e.g., temperature, pressure)

Expertise required: Understanding what variables are relevant to consistency.

Can be automated: No. This requires semantic understanding of the domain.

Step 3: Define Restriction Maps (Consistency Constraints) Question: How should information be consistent across edges?

Action: For each edge $e = (u, v)$, define restriction maps $r_{e,u}, r_{e,v} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ encoding the consistency constraint:

- **Identity (agreement):** $r_{e,v}(x) = x$. Meaning: "Vertices u and v should have identical information."
 - Example: Agents sharing a global coordinate frame should report the same target location.
- **Transformation (coordinate change):** $r_{e,v}(x) = Rx + t$. Meaning: "Vertex v 's information should match u 's after applying transformation (R, t) ."
 - Example: Robots with different orientations should report positions that match after rotation.

- **Constraint (dynamic consistency):** $r_{e,s}(Q) = Q[a]$, $r_{e,s'}(Q) = \gamma \max_{a'} Q[a']$. Meaning: "Q-values should satisfy the Bellman equation."
 - Example: In RL, Q-values at state s and successor state s' should be consistent via the Bellman backup.

Expertise required: Deep domain knowledge to define semantically meaningful consistency.

Can be automated: No. This is the core design challenge and currently requires human insight. Automated learning of restriction maps from data is an open problem (see the Discussion section of the main manuscript).

Step 4: Estimate Noise Level Question: How noisy is the system?

Action: Estimate the noise level σ from:

- **Sensor datasheets:** GPS error, accelerometer drift, etc.
- **Empirical measurements:** Standard deviation of repeated observations under controlled conditions.
- **Communication error rates:** Bit error rate, packet loss rate.
- **Model mismatch:** Expected discrepancy between assumed and true dynamics.

If σ is unknown, use Procedure 1 (spectral gap estimation) in Section 2.8 to adaptively choose ϵ .

Expertise required: System characterization or calibration data.

Can be automated: Partially. Online noise estimation from data is possible but requires careful statistical methods.

Step 5: Choose Threshold ϵ Question: What threshold separates "near-zero" from "positive" eigenvalues?

Action: Use one of the procedures in Section 2.8:

- **Procedure 1 (recommended):** Spectral gap estimation. Automatically adapts to system structure.
- **Procedure 2:** Noise-adaptive. Use if σ is known: $\epsilon = 4\sigma$.

Expertise required: None (procedures are algorithmic).

Can be automated: Yes.

Step 6: Compute Phronesis Index Question: What is the consistency state of the system?

Action: Run the STPGC algorithm (Algorithm 1 in the main manuscript) with the graph, stalks, restriction maps, and ϵ from Steps 1-5. The output is:

- Φ : Phronesis Index (scalar health metric)
- h_e^1 : Approximate number of topological holes (inconsistency count)
- λ_1^+ : Smallest positive eigenvalue (consensus strength)

Expertise required: None (algorithm is fully specified).

Can be automated: Yes.

Summary: What Can and Cannot Be Automated

| Step | Requires Expertise? | Can Be Automated? |
|----------------------------|------------------------------|-------------------|
| 1. Define graph | Yes (domain knowledge) | Partially |
| 2. Define stalks | Yes (semantic understanding) | No |
| 3. Define restriction maps | Yes (deep domain knowledge) | No |
| 4. Estimate noise | Moderate (calibration) | Partially |
| 5. Choose ϵ | No (algorithmic) | Yes |
| 6. Compute Φ | No (algorithmic) | Yes |

Key insight: Steps 1-3 (sheaf design) require human expertise and are the main barrier to widespread adoption. Steps 4-6 (computation) are fully algorithmic. Future work on automated sheaf learning (see the Discussion section of the main manuscript) aims to reduce the expertise required for Steps 2-3.

Example Walkthrough: Multi-Robot Coordination System: 10 robots navigating a warehouse, sharing position estimates.

Step 1 (Graph): Vertices = 10 robots. Edges = communication links (robots within 5 meters can communicate).

Step 2 (Stalks): $\mathcal{F}(v) = \mathbb{R}^2$ (2D position in robot v 's local coordinate frame).

Step 3 (Restrictions): For edge $e = (u, v)$, if robot u is rotated by θ relative to v :

$$r_{e,u} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}, \quad r_{e,v} = I_2 \quad (39)$$

Meaning: "Robot u 's position (in its frame) should match robot v 's position (in v 's frame) after rotation."

Step 4 (Noise): GPS error $\sigma = 0.1$ meters (from datasheet).

Step 5 (Threshold): Use Procedure 2: $\epsilon = 4\sigma = 0.4$ meters.

Step 6 (Compute): Run STPGC. Typical output: $\Phi \approx 50$ (healthy), $h_\epsilon^1 = 0$ (no inconsistencies), $\lambda_1^+ \approx 20$ (strong coupling).

If one robot's GPS malfunctions, Φ drops to ≈ 5 , h_ϵ^1 increases to 1, triggering recalibration.

3 Deployment Considerations

This section provides practical guidance for deploying the Phronesis Index in real-world multi-agent systems, addressing operational concerns that arise beyond controlled experimental settings.

3.1 System Integration

3.1.1 Architecture Patterns

Centralized Monitoring

- All agents report beliefs to a central server
- Server constructs sheaf and computes Φ
- **Pros:** Simple to implement, full visibility
- **Cons:** Single point of failure, communication overhead
- **Use case:** Small-to-medium systems with reliable networking

Distributed Monitoring

- Each agent computes local contribution to \mathcal{L}
- Contributions aggregated via gossip protocol
- **Pros:** Scalable, fault-tolerant
- **Cons:** More complex, approximate Φ
- **Use case:** Large-scale IoT, swarm robotics

Hierarchical Monitoring

- System divided into clusters
- Local Φ computed per cluster
- Global Φ aggregated from local values
- **Pros:** Balances scalability and accuracy
- **Cons:** Requires cluster design
- **Use case:** Multi-tier systems (e.g., factory floors)

3.1.2 Computational Resources

Memory Requirements:

- Sparse Laplacian: $O(Md)$ where $M = |E|$, $d =$ stalk dimension
- Lanczos workspace: $O(Ndk)$ where $k \approx 20$
- Total: typically < 1 GB for $N = 10^4$, $d = 4$

Computation Time:

- Per-iteration: $O(N \log N)$ for sparse graphs
- Typical: 0.1–1 second for $N = 10^3$
- Scalable to $N = 10^5$ with GPU acceleration

3.2 Operational Parameters

3.2.1 Threshold Selection

Critical Threshold Φ_{crit} :

Set based on empirical calibration:

1. Collect data from system in known "healthy" and "unhealthy" states
2. Compute Φ distribution for each state
3. Set Φ_{crit} at the 5th percentile of healthy distribution

Example values from our experiments:

- Logic Maze: $\Phi_{\text{crit}} = 10$
- Safety Gym: $\Phi_{\text{crit}} = 0.5$
- Multi-Robot: $\Phi_{\text{crit}} = 2.0$

3.2.2 Update Frequency

Trade-off: More frequent updates \rightarrow faster detection, but higher overhead

Recommendations:

- **Critical systems:** Every 0.1–1 second
- **Standard systems:** Every 1–10 seconds
- **Low-priority monitoring:** Every 1–5 minutes

Adaptive scheduling: Increase frequency when Φ approaches Φ_{crit}

3.3 Failure Handling

3.3.1 Response Protocols

When $\Phi < \Phi_{\text{crit}}$:

Level 1: Warning

- Log event
- Notify operators
- Continue operation with caution

Level 2: Intervention

- Trigger re-calibration protocol
- Request additional observations
- Slow down or pause risky actions

Level 3: Shutdown

- If Φ remains low after intervention
- Gracefully halt system
- Require manual inspection before restart

3.3.2 False Positives/Negatives

Reducing false positives:

- Use temporal filtering (Kalman smoothing)
- Require $\Phi < \Phi_{\text{crit}}$ for T consecutive timesteps
- Cross-validate with domain-specific checks

Reducing false negatives:

- Set conservative Φ_{crit} (higher threshold)
- Monitor rate of change: $d\Phi/dt < 0$ sustained
- Combine with other safety signals

3.4 Maintenance and Monitoring

3.4.1 System Health Metrics

Track over time:

- Φ mean and variance
- h^1 distribution
- λ_1 trends
- Alert frequency
- Response latency

3.4.2 Periodic Recalibration

Recommended schedule:

- Weekly: Review Φ logs, adjust Φ_{crit} if needed
- Monthly: Validate sheaf structure against ground truth
- Quarterly: Full system audit and parameter tuning

3.5 Code Availability

All code and data are publicly available at:

`https://github.com/\[username\]/phronesis-index`

Includes:

- Python implementation of STPGC algorithm
- Sheaf construction utilities
- Example scenarios (Logic Maze, Safety Gym, Multi-Robot)
- Jupyter notebooks with tutorials
- Docker container for reproducibility

References

- [1] A. Singer and H.-T. Wu. Vector diffusion maps and the connection laplacian. *Communications on Pure and Applied Mathematics*, 65(8):1067–1144, 2012.