

# **Entanglement Learning Mathematical Background and IDT Architecture**

# **OVERVIEW**

Conceptual Foundation	. 1
The Three Entropies Framework	. 1
From Information Theory to Entanglement Metrics	. 2
Basic Entropy Calculations	. 2
Base Entanglement ( $\psi$ )	. 3
Entanglement Asymmetry (Λψ)	. 3
Entanglement Memory (μψ).	4
Practical Implementation	4
Information Gradients	. 7
Implementation Considerations	. 8
Adaptation Mechanism	. 9
Simple Implementation Example: MPC System	. 9
Information Digital Twin Architecture	11
Establishing Detection Thresholds	12
Binning Resolution and Information Measurement	14
Integration with Existing Systems	15
Common Implementation Pitfalls	17
Conclusion	18
Glossary	19

# **CONCEPTUAL FOUNDATION**

This guide provides step-by-step instructions for calculating and implementing Entanglement Learning (EL) metrics. EL quantifies the mutual predictability between an agent and its environment, measuring how effectively information flows across their interaction. The metrics presented here—Base Entanglement ( $\psi$ ), Entanglement Asymmetry ( $\Lambda\psi$ ), and Entanglement Memory ( $\mu\psi$ )—enable you to detect when an agent's internal model no longer aligns with its environment, often before performance visibly degrades. These calculations form the mathematical foundation for building an Information Digital Twin (IDT) that continuously monitors system-environment alignment.

This document assumes familiarity with basic probability theory and information theory concepts such as entropy and mutual information.

# THE THREE ENTROPIES FRAMEWORK

The mathematical foundation of EL rests on analyzing the information relationships between three key variables in any agentenvironment interaction (Fig. 1):

- S: The observation states (inputs the agent receives from the environment)
- A: The action states (decisions or outputs the agent produces)
- S': The resulting states (next observations based on environment response to the taken action).

Each of these variables has an associated entropy that measures uncertainty about its possible values:

- H(S): Entropy of observation states
- H(A): Entropy of action states
- H(S'): Entropy of resulting states

These entropies can be visualized as three overlapping circles in a Venn diagram, Fig. 1, where the overlapping regions represent mutual information between variables:

- MI(S; A): Mutual information between observations and actions (how well inputs predict decisions)
- MI(A; S'): Mutual information between actions and resulting states (how strongly actions influence outcomes)
- MI(S; S'): Mutual information between current and next states (natural predictability of the environment)
- MI(S, A; S'): Three-way mutual information (how well current observations and actions jointly predict next states)



This Venn diagram provides a visual representation of information flow through the agent-environment interaction cycle. By measuring changes in these entropy relationships, we can detect when internal models no longer accurately capture environmental dynamics, often before traditional performance metrics show degradation.



Fig. 1. The Three-Entropy Framework. Visualization of information relationships between agent observations H(S), actions H(A), and environment responses H(S'). Overlapping regions represent mutual information, with the central intersection forming the base entanglement (ψ) that quantifies overall information throughput.

The central region where all three circles overlap corresponds to our core metric Base Entanglement ( $\psi$ ), while the other regions give rise to Asymmetry ( $\Lambda\psi$ ) and Memory ( $\mu\psi$ ). These three metrics provide complementary views of information flow, enabling precise diagnosis of alignment issues.

# FROM INFORMATION THEORY TO ENTANGLEMENT METRICS

Shannon's information theory provides the mathematical foundation for measuring alignment between an agent and its environment. Information in this context represents uncertainty reduction—how knowing one variable reduces uncertainty about another. This concept extends perfectly to agent-environment interactions, where we quantify how effectively an agent's internal model captures environmental dynamics.

To apply information theory to real-world systems, we must first transform continuous variables (like sensor readings or motor outputs) into discrete probability distributions. This discretization step involves binning continuous values into meaningful ranges and counting their frequencies over time. While the specific binning strategy depends on domain characteristics, the goal remains consistent: create probability distributions that capture meaningful patterns in agent-environment interactions.

Once probability distributions are established, we calculate mutual information to measure statistical dependencies between variables. Unlike correlation, mutual information captures non-linear relationships and is measured in bits—providing a universal unit for guantifying information flow across any domain.

Entanglement metrics build upon mutual information by measuring specific aspects of the three-way relationship between observations, actions, and responses. While mutual information typically measures pairwise relationships, entanglement metrics capture the sustained information dependencies across the complete interaction cycle. This distinction is crucial—entanglement represents maintained mutual predictability that persists through changing conditions, not just momentary statistical correlation.

These metrics transform abstract information relationships into practical tools for detecting when an agent's internal model no longer aligns with environmental reality. By continuously monitoring these information patterns, systems can identify misalignment before it manifests as performance degradation.

# **BASIC ENTROPY CALCULATIONS**

Before calculating entanglement metrics, you must establish probability distributions for your system variables and compute their entropy values. The process follows these steps:

- 1. Collect samples of observations (S), actions (A), and environment responses (S') over a sufficient time period.
- 2. Discretize continuous variables into bins appropriate for your domain.



3. Calculate probability distributions by counting how often each bin is occupied, then dividing by the total sample count:

p(s) = count(s) / total\_samples

p(a) = count(a) / total\_samples

p(s') = count(s') / total\_samples

4. Compute individual entropies using the standard formula:

 $H(S) = -\sum p(s) \log_2 p(s)$ 

 $H(A) = -\sum p(a) \log_2 p(a)$ 

 $H(S') = -\sum p(s') \log_2 p(s')$ 

5. Calculate joint distributions by counting co-occurrences:

p(s,a) = count(s,a) / total\_samples

p(a,s') = count(a,s') / total\_samples

p(s,s') = count(s,s') / total\_samples

p(s,a,s') = count(s,a,s') / total\_samples

6. Compute joint entropies:

 $H(S,A) = -\sum p(s,a) \log_2 p(s,a)$ 

 $H(A,S') = -\sum p(a,s') \log_2 p(a,s')$ 

 $H(S,S') = -\sum p(s,s') \log_2 p(s,s')$ 

 $H(S,A,S') = -\sum \sum p(s,a,s') \log_2 p(s,a,s')$ 

Numerical precision is important when calculating entropies. Implement checks to handle cases where p(x) = 0 by setting  $p(x)\log_2(p(x)) = 0$  when p(x) = 0, consistent with the limit as p(x) approaches zero.

# BASE ENTANGLEMENT ( $\Psi$ )

Base Entanglement measures the overall mutual predictability between the agent's state-action pair and the environment's response.

## Mathematical Definition:

 $\psi = \mathsf{MI}(\mathsf{S},\mathsf{A};\mathsf{S}') = \mathsf{H}(\mathsf{S},\mathsf{A}) + \mathsf{H}(\mathsf{S}') - \mathsf{H}(\mathsf{S},\mathsf{A},\mathsf{S}')$ 

## **Calculation Steps:**

- 1. Calculate H(S,A) from the joint distribution of observations and actions
- 2. Calculate H(S') from the distribution of environment responses
- 3. Calculate H(S,A,S') from the three-way joint distribution
- 4. Compute  $\psi$  using the formula above

## Interpretation:

- Higher values (in bits) indicate stronger alignment between the agent's model and the environment
- Decreasing ψ signals potential misalignment before performance visibly degrades
- Baseline ψ values vary by domain and should be established during optimal performance periods

**Example:** For a system with H(S,A) = 4.2 bits, H(S') = 3.1 bits, and H(S,A,S') = 6.0 bits:

 $\psi$  = 4.2 + 3.1 - 6.0 = 1.3 bits

This indicates that knowing the agent's current observation and action reduces uncertainty about the environment's response by 1.3 bits.

# ENTANGLEMENT ASYMMETRY ( $\Lambda \psi$ )



Entanglement Asymmetry identifies whether misalignments originate in the agent's perception or action capabilities by comparing two information pathways.

## Mathematical Definition:

$$\Lambda \psi = MI(A;S') - MI(S;A)$$

- = [H(A) + H(S') H(A,S')] [H(S) + H(A) H(S,A)]
- = H(S,A) H(A,S') + H(S') H(S)

## **Calculation Steps:**

- 1. Calculate MI(A;S') = H(A) + H(S') H(A,S')
- 2. Calculate MI(S;A) = H(S) + H(A) H(S,A)
- 3. Compute  $\Lambda \psi$  as their difference

## Interpretation:

- Positive values (Λψ > 0): Actions predict responses better than observations predict actions, suggesting perception issues
- Negative values (Λψ < 0): Observations predict actions better than actions predict responses, suggesting control issues</li>
- Values near zero: Balanced information flow between perception and control pathways

**Example:** For a system with MI(A;S') = 2.4 bits and MI(S;A) = 1.8 bits:

## $\Lambda \psi = 2.4 - 1.8 = 0.6$ bits

This positive asymmetry indicates the agent's internal models aren't effectively capturing input patterns (perception issues) while its control mechanisms remain effective.

# ENTANGLEMENT MEMORY $(M\Psi)$

Entanglement Memory measures the natural predictability between successive states independent of the agent's actions, revealing how consistently the environment evolves on its own.

## Mathematical Definition:

 $\mu \psi = \mathsf{MI}(\mathsf{S};\mathsf{S}') = \mathsf{H}(\mathsf{S}) + \mathsf{H}(\mathsf{S}') - \mathsf{H}(\mathsf{S},\mathsf{S}')$ 

## **Calculation Steps:**

- 1. Calculate H(S) from the distribution of observations
- 2. Calculate H(S') from the distribution of environment responses
- 3. Calculate H(S,S') from their joint distribution
- 4. Compute  $\mu\psi$  using the formula above

# Interpretation:

- Higher values indicate stronger natural patterns in the environment that persist independent of actions
- Decreasing μψ signals that the environment's behavior is becoming less consistent or predictable
- Stable μψ with declining ψ suggests issues with the agent rather than environmental changes

**Example:** For a system with H(S) = 3.5 bits, H(S') = 3.1 bits, and H(S,S') = 5.2 bits:

$$\mu \psi = 3.5 + 3.1 - 5.2 = 1.4$$
 bits

This indicates that the environment has 1.4 bits of inherent predictability between successive states, independent of the agent's actions.

# **PRACTICAL IMPLEMENTATION**

# **Discretization Strategies**

Effective discretization transforms continuous variables into meaningful probability distributions—a critical step that directly impacts metric sensitivity. Here's how to implement this for your system:

1. Analyze Variable Distributions

- Collect a large sample (1000+ points) of each variable during normal operation
- Generate histograms to visualize natural distribution patterns
- Identify regions of high density (frequent values) versus sparse regions

# 2. Select Appropriate Binning Approach Uniform Binning:

// Example with 20 bins over the observed range min\_val = minimum value in samples max\_val = maximum value in samples bin\_width = (max\_val - min\_val) / num\_bins FOR i = 0 TO num\_bins bins[i] = min\_val + i \* bin\_width ENDFOR

Best for: Variables with relatively uniform distributions; initial testing before optimization Percentile-Based Binning:

// Creates bins with equal number of samples in each
FOR i = 0 TO num\_bins
percentile = i / num\_bins
bins[i] = value at percentile \* 100% of sorted samples
ENDFOR

Best for: Ensuring statistical significance across all bins regardless of distribution shape Information-Weighted Binning:

// First analyze entropy contribution across initial uniform bins
FOR each bin in initial\_uniform\_bins
probability = count\_in\_bin / total\_samples
IF probability > 0 THEN
entropy\_contribution[bin] = -probability \* log2(probability)
ENDIF
ENDFOR

// Then refine binning by allocating more bins to high-entropy regions

// High entropy areas get more refined bins, low entropy areas get fewer

Best for: Maximum sensitivity to information-rich regions; optimal for entanglement metric calculation

# 3. Validate Binning Effectiveness

- Ensure at least 5-10 samples per bin for statistical significance
- Compare entanglement metrics across different binning strategies
- Test sensitivity to known misalignments

# 4. Domain-Specific Considerations

- For CNNs: Focus on activation regions near decision boundaries
- For control systems: Prioritize resolution near equilibrium points
- For RL agents: Align bins with reward-relevant state transitions

**Real-world Example:** For a robotic arm with joint angle  $\theta$  ranging from -90° to +90°:

// Sample distribution shows dense clusters around -45°, 0°, and +45°

// Information-weighted binning creates:

bins\_θ = [-90, -80, -70, -60, -55, -50, -47, -45, -43, -40, -35, -30, -20,

- -10, -5, -2, 0, +2, +5, +10, +20, +30, +35, +40, +43, +45, +47,
- +50, +55, +60, +70, +80, +90]

Note the higher resolution around frequently used positions (-45°, 0°, +45°) where small differences matter most for the system's behavior.

SEMARX



# **Practical Considerations**

Implementing entanglement metrics in real systems requires addressing several practical challenges:

## 1. Computational Efficiency

Incremental updates: Instead of recalculating full distributions at each step, update bin counts incrementally:

// When observing a new (s,a,s') triplet: s\_bin = determine\_bin\_for\_value(s, s\_bins) a\_bin = determine\_bin\_for\_value(a, a\_bins) s\_prime\_bin = determine\_bin\_for\_value(s\_prime, s\_prime\_bins) increment counts[s\_bin] increment counts[s\_bin] increment counts[s\_prime\_bin] increment joint\_counts[s\_bin, a\_bin] increment joint\_counts[s\_bin, s\_prime\_bin] increment triple\_counts[s\_bin, a\_bin, s\_prime\_bin] increment triple\_counts[s\_bin, a\_bin, s\_prime\_bin] increment total\_samples

• Vectorized operations: Use vectorized functions for entropy calculations:

// Efficient entropy calculation
probabilities = counts / total\_samples
// Add small epsilon to avoid log(0)
FOR each probability value in probabilities
IF probability > 0 THEN
entropy -= probability \* log2(probability)
ENDIF
ENDFOR

Sliding window approach: Maintain fixed-size history to adapt to changing conditions:

// Keep most recent N samples in a circular buffer // When buffer fills, remove oldest sample before adding new one old\_s, old\_a, old\_s\_prime = sample\_buffer[buffer\_index] // Decrement old sample counts old\_s\_bin = determine\_bin\_for\_value(old\_s, s\_bins) // Decrement counts for old sample // Add new sample and update index sample\_buffer[buffer\_index] = (s, a, s\_prime) buffer\_index = (buffer\_index + 1) MOD buffer\_size 2. Handling High Dimensional Spaces

2. Handling High-Dimensional Spaces

For high-dimensional observations (like images), apply dimensionality reduction first: // Example for CNN activations // Instead of binning individual neurons (10,000+) // Calculate principal components and bin those reduced\_dimensions = apply\_dimensionality\_reduction(high\_dim\_data) // Then bin each component separately

For categorical variables with many values, use hierarchical binning:

// Example for NLP token IDs (10,000+ tokens)
// Group tokens by semantic similarity or frequency
semantic\_bins = {
'common\_nouns': [list of common noun token IDs],
'rare\_nouns': [list of rare noun token IDs],

'common\_verbs': [list of common verb token IDs],

// ...and so on}

- 3. Addressing Sparsity Issues
- Use Laplace smoothing for rare bin combinations:



- // Apply smoothing factor alpha (typically 0.01-0.1)
- smoothed\_probability = (count + alpha) / (total\_samples + alpha \* num\_bins)

Implement adaptive bin merging for underrepresented regions:

ENDIF ENDIF

ENDFOR

## 4. Baseline Establishment

- Calculate and store entanglement metrics during optimal performance:
- // During normal operation (e.g., first 10,000 samples)
- baseline\_entanglement = average of last 1000 entanglement values
- baseline\_asymmetry = average of last 1000 asymmetry values
- baseline\_memory = average of last 1000 memory values
  - // Also calculate standard deviations for detection thresholds
  - std\_entanglement = standard\_deviation(baseline\_entanglement)
  - // Typical threshold = 2-3 standard deviations

detection\_threshold = baseline\_entanglement - 3 \* std\_entanglement

These techniques enable practical implementation of entanglement metrics in real-time systems while maintaining computational efficiency. The precise approach you select should be tailored to your specific domain and computational constraints.

# **INFORMATION GRADIENTS**

Information gradients form the critical link between entanglement measurement and adaptive response. These gradients exist at two distinct levels, providing both general detection of misalignment and specific guidance for parameter adjustment.

## Pure Information Gradients

At the most fundamental level, information gradients represent changes in entanglement metrics over time:

 $\Delta \psi = \psi(t) - \psi(t-1)$  // Change in base entanglement

 $\Delta \Lambda \psi = \Lambda \psi(t) - \Lambda \psi(t-1) // Change in asymmetry$ 

 $\Delta\mu\psi = \mu\psi(t) - \mu\psi(t-1)$  // Change in memory

These pure gradients indicate the direction and magnitude of information misalignment, independent of any system-specific parameters:

- $\Delta \psi$  < 0: Overall information throughput is decreasing, signaling general misalignment between system and environment
- ΔΛψ < 0: Perception-to-action information flow is weakening relative to action-to-outcome flow, often indicating perception/input processing issues</li>
- $\Delta\mu\psi$  < 0: Environmental predictability is decreasing, suggesting changes in underlying system dynamics



## **Parameter-Specific Gradients**

To enable practical adaptation, pure information gradients must be mapped to adjustable system parameters. This mapping creates parameter-specific gradients:

 $\nabla \psi(\theta) = \partial \psi / \partial \theta$  // Sensitivity of base entanglement to parameter  $\theta$ 

 $\nabla \Lambda \psi(\theta) = \partial \Lambda \psi / \partial \theta$  // Sensitivity of asymmetry to parameter  $\theta$ 

 $\nabla \mu \psi(\theta) = \partial \mu \psi / \partial \theta$  // Sensitivity of memory to parameter  $\theta$ 

These parameter-specific gradients identify which adjustable elements have the strongest influence on restoring optimal information flow.

## Example: Computer Vision System

Consider a CNN-based object detection system experiencing declining performance:

- 1. The IDT detects  $\Delta \psi < 0$  (decreasing overall entanglement) and  $\Delta \Delta \psi < 0$  (increasingly negative asymmetry)
- 2. This pure gradient pattern indicates input processing issues rather than classification problems
- 3. The system calculates parameter-specific gradients for input processing components:
- 4.  $\partial \psi / \partial \alpha$ \_preprocess = 0.02 // Weak influence of preprocessing alpha parameter

 $\partial \psi / \partial \gamma_{preprocess} = 0.45$  // Strong influence of preprocessing gamma parameter

5. Based on the relative magnitudes, the system adjusts y\_preprocess to restore information throughput

This two-level gradient approach maintains the universal nature of entanglement metrics while enabling domain-specific adaptation. Pure information gradients detect what type of misalignment is occurring, while parameter-specific gradients identify exactly which adjustments will most effectively restore alignment.

## IMPLEMENTATION CONSIDERATIONS

1. Parameter Selection Not all system parameters need gradient calculation. Focus on:

- Parameters that directly influence information processing
- Parameters subject to drift during operation
- Parameters that previous analyses have identified as sensitive

## 2. Sensitivity Analysis

// Identify most influential parameters

FOR each parameter p in system parameters

// Apply various perturbation magnitudes

FOR delta in [0.01, 0.05, 0.1, 0.5] \* parameter\_range

// Calculate sensitivity at this scale

sensitivity = calculate\_gradient(p, delta)

record sensitivity

ENDFOR

ENDFOR

// Parameters with consistently high sensitivity

// are primary candidates for gradient-based adaptation

3. Multiple Metrics Information gradients should be calculated for all three entanglement metrics:

gradient\_psi = calculate\_gradient\_for\_metric(psi) gradient\_lambda = calculate\_gradient\_for\_metric(lambda\_psi) gradient\_mu = calculate\_gradient\_for\_metric(mu\_psi) Different metrics inform different types of adaptations:

- Base entanglement (ψ) gradients guide general alignment restoration
- Asymmetry (Λψ) gradients address specific misalignment sources
- Memory (μψ) gradients adjust temporal consistency
- 4. Gradient Combination



// Combined gradient considers all metrics with appropriate weighting

FOR each parameter p

```
combined_gradient[p] = w1 * gradient_psi[p] +
w2 * gradient_lambda[p] +
```

w3 \* gradient\_mu[p]

# ENDFOR

// Weights reflect relative importance and are typically adjusted // based on which metrics show significant deviation

# Adaptation Mechanism

// Use gradients to adjust parameters

## FOR each parameter p

// Apply gradient ascent with learning rate

adjustment = learning\_rate \* combined\_gradient[p]

// Limit adjustment magnitude for stability

adjustment = clamp(adjustment, -max\_adjustment, +max\_adjustment)

// Update parameter

p = p + adjustment

## ENDFOR

The learning rate controls adaptation speed—smaller values provide more stable but slower adaptation, while larger values offer faster response but potential instability. Adaptive learning rates that respond to gradient magnitudes often provide the best balance.

## **Assumptions and Limitations:**

- 1. Local Linearity: Gradient calculation assumes the entanglement function is locally linear with respect to parameter changes. For highly non-linear relationships, multiple small adjustments are more effective than single large ones.
- Parameter Independence: Basic gradient calculation assumes parameters affect entanglement independently. For strongly interdependent parameters, more sophisticated approaches like coordinate descent or trust region methods may be needed.
- 3. **Temporal Stability:** Gradients assume the underlying relationship between parameters and entanglement remains stable during the adaptation process. Rapid environmental changes may require recalculating gradients before adaptation completes.

Information gradients transform entanglement metrics from passive indicators to active drivers of adaptation. By following these gradients, systems can autonomously maintain alignment with their environments without requiring explicit programming of adaptation rules or human intervention.

# SIMPLE IMPLEMENTATION EXAMPLE: MPC SYSTEM

This example illustrates applying entanglement metrics to a Model Predictive Controller (MPC) for an autonomous drone, demonstrating how to detect when the controller's internal model no longer aligns with reality.

**System Definition:** The MPC controller generates control actions for a drone based on current state, predicted trajectory, and optimization constraints. We'll monitor information throughput between:

- S: Observation states (position, velocity, orientation, sensor readings)
- A: Control actions (motor commands, attitude adjustments)
- S': Resulting states (next position, velocity, orientation)

## Implementation Steps:

- 1. Establish Discretization Strategy
  - Analyze flight data during normal operation to understand variable distributions



- Apply non-uniform binning focusing on higher resolution near hover conditions
- Example bins for vertical velocity: [-5.0, -2.0, -1.0, -0.5, -0.2, -0.1, -0.05, 0, 0.05, 0.1, 0.2, 0.5, 1.0, 2.0, 5.0] m/s
- Create similar bins for all relevant state variables and control actions

#### 2. Collect Probability Data

- During nominal flight, record binned state and action occurrences
- Maintain frequency counters for each bin and joint bin combinations
- Track 1000-sample sliding window to adapt to normal flight phase changes

#### 3. Calculate Entropies

- Compute individual entropies H(S), H(A), H(S')
- Calculate joint entropies H(S,A), H(A,S'), H(S,S'), H(S,A,S')
- Implement numerical safeguards for handling zero probabilities

#### 4. Derive Entanglement Metrics

- Compute base entanglement: ψ = H(S,A) + H(S') H(S,A,S')
- Calculate asymmetry: Λψ = MI(A;S') MI(S;A)
- Determine memory: μψ = MI(S;S')
- Record these metrics continuously during operation

## 5. Establish Baseline Values

- During normal flight conditions, record mean and standard deviation of each metric
- Define detection thresholds as 2-3 standard deviations from baseline means
- Document different baseline profiles for various flight phases (takeoff, hover, forward flight)

#### 6. Misalignment Detection Example

- When the drone encounters unexpected wind conditions:
- Base entanglement (ψ) gradually decreases as the MPC's predictions become less accurate
- Asymmetry (Λψ) becomes negative as the internal model fails to account for new dynamics
- Memory (µµ) remains relatively stable, suggesting the issue is with the controller, not sensor noise

#### 7. Information Gradient Calculation

- When metrics exceed detection thresholds, calculate sensitivity to MPC parameters
- Determine how entanglement responds to changes in model dynamics parameters
- Identify that adjusting the aerodynamic drag coefficient has highest sensitivity

#### 8. Adaptation Response

- Apply gradual adjustments to the drag coefficient, following the information gradient
- After adjustment, observe entanglement metrics returning toward baseline
- Drone performance improves without requiring complete retuning or manual intervention

**Outcome:** The IDT detects wind-induced misalignment and guides parameter adaptation before traditional performance metrics (position error, control effort) show significant degradation. Entanglement metrics provide early warning and guide specific parameter adjustments to maintain system-environment alignment during changing conditions.

This example demonstrates the complete workflow from discretization through detection and adaptation, showing how entanglement metrics enable autonomous self-assessment and targeted parameter adjustment in an MPC system.



# INFORMATION DIGITAL TWIN ARCHITECTURE

The Information Digital Twin (IDT) serves as a parallel monitoring system that continuously assesses the information throughput between an agent and its environment without disrupting primary operations (Fig. 2). The IDT's modular architecture allows it to integrate with existing systems while providing the computational framework necessary for entanglement-based self-assessment.



Fig. 2. The IDT Architecture. The Information Digital Twin (IDT) continuously monitors the interaction between agent and environment by tracking entanglement metrics across actions and resulting states. When misalignment is detected, information gradients are computed and passed to the Entanglement Controller, which generates adaptive control signals to update the agent's internal model—preserving coherence and maximizing information throughput

## **Core Components:**

- 1. Data Collection Module This front-end component interfaces with the primary system to capture observation states, action states, and resulting states. It implements non-invasive monitoring hooks that sample variable values at appropriate intervals without affecting main system performance. The module converts continuous values into discretized bins according to the established binning strategy and maintains the frequency counts required for probability estimation.
- Probability Distribution Processor This component transforms raw bin counts into probability distributions for individual and joint variable combinations. It handles numerical edge cases (zeros, sparse distributions) and applies appropriate smoothing techniques. The processor maintains both current and historical distribution data, enabling comparison between recent and baseline behavior patterns.
- 3. Entropy Calculation Engine This computational module performs the mathematical operations required to calculate individual and joint entropies from the probability distributions. It employs numerical optimizations to handle the logarithmic calculations efficiently and maintains appropriate precision throughout the calculation pipeline.
- 4. Entanglement Metrics Generator This analytical component derives the three core entanglement metrics (base entanglement, asymmetry, and memory) from the entropy values. It tracks these metrics over time, calculating rolling averages, trends, and deviations from established baselines. The generator produces both current metric values and temporal analysis of their evolution.
- 5. **Baseline Modeling System** This module establishes and maintains reference patterns of normal behavior. It captures entanglement signatures during optimal performance periods and builds statistical models of expected variations. The modeling system may maintain multiple baseline profiles for different operational modes or environmental conditions, automatically selecting the appropriate reference based on context.
- 6. **Deviation Detection Engine** This analytical component continuously compares current entanglement metrics against baseline expectations, identifying significant deviations that signal potential misalignment. It implements statistical tests to distinguish normal variations from meaningful shifts, applying appropriate thresholds and considering temporal patterns in the metrics.
- 7. Information Gradient Calculator When deviations exceed thresholds, this component calculates the sensitivity of



entanglement metrics to parameter changes. It identifies which parameters have the strongest influence on information throughput and determines the direction and magnitude of adjustments needed to restore alignment.

8. Adaptation Signal Generator This output module transforms gradient information into specific adaptation signals compatible with the host system. It applies constraints to ensure stable adaptation and provides appropriate feedback mechanisms to verify the effectiveness of parameter adjustments.

#### Information Flow:

The IDT operates as a continuous monitoring loop:

- 1. The Data Collection Module samples system variables at regular intervals
- 2. The Probability Distribution Processor updates distribution estimates
- 3. The Entropy Calculation Engine computes updated entropy values
- 4. The Entanglement Metrics Generator derives current metric values
- 5. The Deviation Detection Engine compares metrics against baselines
- 6. When deviations occur, the Information Gradient Calculator determines required adjustments
- 7. The Adaptation Signal Generator produces specific parameter change recommendations
- 8. The primary system applies these adjustments, and the cycle continues

#### **Architectural Principles:**

- 1. **Separation of Concerns** The IDT operates independently from the primary system's control and decision-making processes, ensuring that monitoring and adaptation don't interfere with core functionality.
- 2. **Computational Efficiency** The architecture optimizes resource usage through incremental calculations, adaptive sampling rates, and efficient numerical methods, minimizing overhead on the host system.
- 3. **Modularity** Individual components can be customized for specific domains while maintaining the overall information flow structure, allowing the IDT to be adapted to diverse applications.
- 4. **Scalability** The architecture supports both local component-level monitoring and system-wide assessment, enabling hierarchical implementation across complex systems with multiple interacting parts.

This architecture provides the operational framework needed to implement entanglement metrics in real-world systems, transforming mathematical concepts into practical tools for autonomous self-assessment and adaptation.

# **ESTABLISHING DETECTION THRESHOLDS**

Determining appropriate thresholds for detecting significant changes in entanglement metrics is critical for balancing sensitivity against false alarms. These thresholds define when a deviation from baseline entanglement patterns warrants attention or adaptation.

#### **Statistical Foundations:**

Effective threshold establishment begins with understanding the natural variability in entanglement metrics during normal operation. Each system exhibits characteristic fluctuations due to:

- Inherent stochasticity in the environment
- Measurement noise and discretization effects
- Normal operational mode transitions

These variations create a statistical distribution of metric values that forms the foundation for threshold determination.

## **Baseline Collection Period:**

Before establishing thresholds, collect entanglement metrics during a sufficient baseline period that captures:

- At least 1,000-10,000 sampling intervals (depending on system dynamics)
- All normal operational modes and transitions
- Representative environmental conditions
- Expected variations in task requirements



This data collection should occur when the system is performing optimally, typically immediately after training or calibration.

## **Threshold Determination Methods:**

- 1. **Standard Deviation Based Thresholds** The most common approach uses statistical dispersion measures to define boundaries of normal behavior:
  - Calculate mean (μ) and standard deviation (σ) for each entanglement metric during baseline period
  - Set lower detection threshold at μ kσ where k is typically 2-3
  - Higher k values reduce false positives but may delay detection
  - Lower k values increase sensitivity but may trigger unnecessary adaptations

Example: If base entanglement ( $\psi$ ) during normal operation has  $\mu$  = 2.4 bits and  $\sigma$  = 0.15 bits, a 3 $\sigma$  threshold would trigger detection when  $\psi$  falls below 2.4 - (3 × 0.15) = 1.95 bits

#### 2. Percentile Based Thresholds This approach uses empirical distribution properties rather than assuming normality:

- Sort observed metric values during baseline period
- Set threshold at a specified percentile (typically 1st or 5th percentile)
- More robust to non-normal distributions and outliers

Example: If the 1st percentile of observed asymmetry ( $\Lambda\psi$ ) values is -0.4 bits, detection would trigger when  $\Lambda\psi$  falls below this level

3. Change Rate Thresholds For rapidly evolving systems, rate of change can provide earlier warning than absolute values:

- Calculate typical rates of change (Δψ/Δt) during normal operation
- Set thresholds on acceptable change rates
- Particularly valuable for detecting sudden shifts in system behavior

Example: If entanglement memory typically changes by at most 0.05 bits per 100 samples, a change of 0.15 bits over this interval would trigger detection

#### Multi-Metric Consideration:

Rather than setting independent thresholds for each metric, consider their relationships:

- Compound Thresholds: Trigger detection only when multiple metrics exceed thresholds simultaneously
- Weighted Combinations: Create a composite score from multiple metrics with appropriate weighting
- Sequential Patterns: Look for characteristic sequences of metric changes that indicate specific misalignment types

## Adaptive Thresholds:

Static thresholds may become outdated as systems evolve. Consider implementing adaptive thresholds that:

- Gradually update based on long-term moving averages
- Adjust sensitivity based on operational context
- Maintain different threshold profiles for different operational modes

## Validation Approach:

Before finalizing thresholds, validate their effectiveness through:

- 1. Split testing: Use half of baseline data to establish thresholds, then validate on the remaining half
- 2. Induced misalignments: Deliberately introduce known issues to verify detection
- 3. False positive analysis: Ensure normal variations don't trigger excessive alarms
- 4. Detection latency measurement: Quantify how quickly issues are detected

#### **Domain-Specific Considerations:**

Threshold establishment should account for domain characteristics:

- Fast-moving control systems may require tighter thresholds and faster detection
- Vision systems may tolerate larger variations in some metrics during scene changes
- Reinforcement learning agents may exhibit naturally higher metric variability during exploration phases

This systematic approach to threshold establishment ensures that entanglement metrics provide reliable signals for adaptation



without causing unnecessary interventions during normal operation.

# **BINNING RESOLUTION AND INFORMATION MEASUREMENT**

#### The Scale-Dependent Nature of Information

When measuring information in continuous systems, we face a fundamental challenge similar to the fractal coastline paradox: the amount of information we measure depends on the resolution at which we observe the system. Just as a coastline appears longer when measured with increasingly precise tools, the information content of a system can appear to increase as we use finer-grained measurement approaches.

This scale-dependency creates an important consideration for Entanglement Learning: the absolute value of entropy or mutual information for a system is less meaningful than the relative changes in these values measured at a consistent resolution. The choice of binning strategy doesn't just affect computational efficiency—it fundamentally defines what counts as "information" in your particular context.

#### **Resolution and Relevance**

The ideal binning resolution for a system isn't necessarily the finest possible granularity. Instead, it should capture the level of detail that is functionally relevant to the system's purpose. Consider these guiding principles:

- 1. **Signal vs. Noise Distinction**: Excessively fine-grained binning may capture random fluctuations (noise) rather than meaningful patterns (signal). Conversely, overly coarse binning might obscure important variations that the system needs to detect.
- 2. **Decision Relevance**: The appropriate resolution often relates to the granularity at which differences in variable values would lead to different decisions or predictions. If two distinct values would never result in different system behaviors, they might reasonably belong in the same bin.
- 3. **Natural Clusters**: Many systems exhibit natural clustering in their variable distributions. Effective binning often aligns with these natural boundaries rather than imposing arbitrary divisions.

#### Adaptive Binning Approaches

Rather than using fixed binning strategies, consider these adaptive approaches that acknowledge the scale-dependent nature of information:

1. **Information-Weighted Binning**: Allocate more bins to regions of the distribution that contribute more significantly to the system's entropy or mutual information. This approach concentrates resolution where it matters most:

// Pseudocode for information-weighted binning

initialize\_uniform\_bins(variable\_range, initial\_bin\_count)

for several\_iterations:

calculate\_entropy\_contribution\_per\_bin()

identify\_high\_information\_bins()

split\_high\_information\_bins()

merge\_low\_information\_bins()

recalculate\_probability\_distributions()

2. **Relevance-Based Binning**: For systems where certain regions have known functional importance, allocate binning resolution based on domain knowledge:

// Pseudocode for relevance-based binning

for each\_functional\_region in variable\_range:

importance = assess\_functional\_importance(region)

bin\_density = base\_density \* importance\_multiplier(importance)

allocate\_bins\_to\_region(region, bin\_density)

3. Dynamic Resolution Adjustment: Allow the binning resolution to evolve as the system operates, focusing



computational resources on the most informative regions:

// Pseudocode for dynamic resolution adjustment

while system\_operating:

monitor\_information\_contribution\_by\_region()

if (region\_importance\_changes\_significantly):

reallocate\_bins\_from\_less\_important\_to\_more\_important\_regions()

update\_probability\_distributions()

#### **Quantifying Binning Effectiveness**

How do we know if our binning strategy is appropriate? Consider these evaluation criteria:

- 1. Information Capture Efficiency: Calculate how much information is captured per bin. Higher values indicate more efficient use of computational resources: Information Capture Efficiency = Total Mutual Information / Number of Bins
- 2. **Decision Boundary Sensitivity**: Assess how well the binning strategy preserves information near critical decision boundaries or thresholds where small changes might matter most.
- 3. **Stability Under Perturbation**: Evaluate how sensitive entanglement metrics are to small changes in bin boundaries. Excessive sensitivity might indicate over-fitting to particular data patterns.

#### Implementation Considerations

When implementing adaptive binning strategies, consider these practical approaches:

- 1. Begin with domain-informed initial binning that allocates resolution based on known system characteristics.
- 2. Implement hierarchical binning that maintains coarse-grained coverage across the entire range while providing finegrained resolution in critical regions.
- 3. Establish minimum occupation thresholds for bins to ensure statistical significance. Bins with too few samples may produce unreliable probability estimates.
- 4. Consider the computational trade-offs: adaptive binning requires additional processing but may significantly improve the signal-to-noise ratio in your information measurements.

#### The Philosophical Implication

This scale-dependent view of information has profound implications for how we understand intelligence through the Entanglement Learning framework. Intelligence might not simply be about maximizing absolute information throughput, but rather about identifying and maintaining the right level of information granularity that captures meaningful patterns while filtering out noise.

In this perspective, part of what makes a system "intelligent" is its ability to adaptively determine what resolution of information matters in different contexts—essentially, knowing not just what to observe, but how closely to look.

By acknowledging and addressing the scale-dependent nature of information measurement, the Entanglement Learning framework gains both philosophical depth and practical robustness, enabling more meaningful assessment of system-environment alignment across diverse applications.

# INTEGRATION WITH EXISTING SYSTEMS

Integrating the IDT with existing systems requires a thoughtful approach that preserves core functionality while enabling entanglement-based monitoring and adaptation. This section outlines practical strategies for implementing entanglement metrics within operational systems across different domains.

## Non-Invasive Monitoring Approach:

The fundamental principle for IDT integration is maintaining separation between monitoring and primary operations. This is achieved through:

#### 1. Read-Only Data Access



- Establish monitoring interfaces that observe system variables without modifying them
- Sample variable values at appropriate intervals that capture dynamics without excessive overhead
- Implement buffering mechanisms that prevent monitoring from affecting timing-critical operations

#### 2. Parallel Processing Architecture

- Execute entanglement calculations on separate computational resources when possible
- Use asynchronous processing to prevent monitoring from blocking main system execution
- Implement priority handling that ensures primary functions always take precedence

#### Integration Patterns by System Type:

#### 1. Software Systems (Neural Networks, Control Systems)

- Implement monitoring hooks at key computational junctions
- Use event-driven monitoring that captures variables during normal processing cycles
- Leverage existing logging infrastructure for preliminary data collection

#### 2. Hardware Systems (Robotics, Physical Control)

- Utilize existing sensor infrastructure with minimal additional instrumentation
- Implement signal splitters for critical pathways to enable parallel monitoring
- Consider dedicated monitoring hardware for systems with limited computational resources

#### 3. Distributed Systems

- Establish consistent timestamp mechanisms for accurate state-action-state linkage
- Implement local IDT modules for subsystems with centralized aggregation
- Use bandwidth-efficient sampling that scales with system complexity

#### Adaptation Interface Design:

Once misalignments are detected, adaptation signals must be properly integrated with the host system:

#### 1. Parameter Adjustment Interface

- Identify which system parameters can be safely adjusted during operation
- Establish limits and constraints that prevent destabilizing changes
- Implement verification mechanisms that confirm adjustments achieve desired effects

## 2. Feedback Loop Structure

- Design adaptation with appropriate timescales for the domain
- Implement gradual adjustment mechanisms that prevent oscillation
- Ensure fallback mechanisms exist if adaptations don't improve alignment

## Phased Implementation Strategy:

For existing systems, a phased integration approach minimizes disruption:

#### 1. Passive Monitoring Phase

- Implement data collection and metric calculation without enabling adaptation
- Establish baseline metrics and appropriate thresholds
- Validate monitoring accuracy and computational overhead

## 2. Alert-Only Phase

- Enable deviation detection with alerts to human operators
- Gather feedback on alert relevance and timing
- Refine thresholds based on operational experience

#### 3. Supervised Adaptation Phase

- Implement adaptation recommendations requiring human approval
- Track effectiveness of approved adaptations
- Build confidence in adaptation mechanisms

#### 4. Autonomous Adaptation Phase



- Enable fully autonomous adaptation for well-validated parameters
- Maintain human oversight for significant adaptations
- Continuously refine adaptation boundaries based on performance

#### **Practical Integration Examples:**

## 1. CNN Integration

- Monitor activations at key layers by tapping into forward pass operations
- Calculate entanglement metrics between model layers offline during initial deployment
- Implement detection thresholds specific to each monitored layer

#### 2. Control System Integration

- Sample controller inputs, outputs, and resulting states at regular intervals
- Synchronize sampling with control loop timing to ensure consistent measurement
- Implement adaptation that adjusts model parameters rather than direct control outputs

#### **Documentation and Validation:**

Proper integration requires:

- Clear documentation of all monitoring points and data flows
- Validation tests confirming monitoring doesn't affect primary performance
- Benchmark measurements of computational overhead under various conditions
- Emergency override mechanisms for critical systems

Through these integration approaches, entanglement metrics can be applied to existing systems with minimal disruption while providing valuable insights into system-environment alignment and enabling autonomous adaptation.

# **COMMON IMPLEMENTATION PITFALLS**

When implementing entanglement metrics in practical systems, several recurring challenges can undermine effectiveness. This section highlights common pitfalls and provides strategies to avoid them, based on experience across multiple domains.

## Insufficient Baseline Data:

One of the most frequent mistakes is establishing thresholds from inadequate baseline data.

- Problem: Thresholds based on limited samples fail to capture the full range of normal variations, leading to excessive false positives.
- Solution: Collect baseline data across all operational modes, environmental conditions, and task variations. Aim for at least 1,000 samples per operational mode, and ensure the baseline period includes typical transitions between states.
- Verification Check: Plot histograms of each metric during baseline collection. Multi-modal distributions often indicate distinct operational states requiring separate thresholds.

## Inappropriate Discretization:

Ineffective binning strategies significantly impact metric sensitivity and computational efficiency.

- Problem: Uniform binning often wastes computational resources on rarely used regions while providing insufficient resolution in critical areas.
- Solution: Analyze variable distributions before establishing bins, and allocate more bins to regions showing high information density or frequent occupation.
- Warning Sign: When most observations fall into very few bins, or many bins remain consistently empty, your discretization needs refinement.

## **Temporal Misalignment:**

Failing to properly link observations, actions, and resulting states can severely distort entanglement metrics.

- Problem: Incorrect attribution of actions to resulting states due to variable processing delays or sampling rates.
- Solution: Implement precise timestamping and ensure consistent sampling intervals. For systems with variable
  processing times, use event-based tracking rather than fixed-interval sampling.
- Critical Test: Introduce known, controlled actions and verify they correctly link to the resulting states in your data collection system.



#### **Computational Resource Underestimation:**

Entanglement calculation can become resource-intensive at scale without proper optimization.

- Problem: Full implementation slows system performance or requires excessive computing resources.
- Solution: Start with monitoring fewer, carefully selected variables. Implement incremental calculation methods and consider downsampling for high-frequency systems. Focus initial efforts on the most information-rich interaction points.
- Practical Approach: Benchmark computational requirements during passive monitoring phase before enabling adaptation capabilities.

#### **Unstable Adaptation Mechanisms:**

Poorly implemented information gradients can lead to oscillatory behavior or runaway parameter changes.

- Problem: Adaptations cause excessive parameter fluctuation or drive the system toward extreme values.
- Solution: Implement conservative adjustment limits, gradual parameter changes, and cooling periods between adaptations. Consider implementing trust regions that restrict the magnitude of single adaptations.
- Key Practice: Test adaptation on non-critical parameters first, and implement automatic rollback for adaptations that reduce entanglement further.

#### Ignoring Domain Knowledge:

Treating all variables with equal importance neglects valuable domain expertise.

- Problem: Resources wasted monitoring irrelevant variables while missing critical interaction points.
- Solution: Leverage domain expertise to identify the most informative variables and prioritize their monitoring. Focus
  initial implementation on known sensitive parameters.
- Balanced Approach: Combine domain-specific monitoring points with exploratory analysis to discover unexpected information pathways.

#### **Metric Misinterpretation:**

Incorrect understanding of what each metric represents leads to improper responses.

- Problem: Taking action based on misunderstood metric changes, potentially addressing symptoms rather than causes.
- Solution: Validate metric interpretation through controlled experiments. Document specific patterns of metric changes and their verified causes in your particular domain.
- Example: Declining asymmetry doesn't always indicate improving conditions—it might reflect increasing uncertainty in both information pathways.

#### **Notification Fatigue:**

Excessive alerts for minor deviations can lead to operators ignoring all notifications.

- Problem: Important alerts get lost among numerous insignificant ones, reducing overall system effectiveness.
- Solution: Implement tiered alert systems with appropriate criticality levels. Aggregate related deviations into meaningful
  patterns rather than individual metric changes.
- Design Principle: Each alert should prompt a specific, well-defined response rather than general concern.

#### **Overlooking Complex Dependencies:**

Treating each parameter independently when they have strong interdependencies.

- Problem: Parameter adjustments cause unexpected side effects by disturbing related parameters.
- Solution: Map parameter interdependencies before implementing adaptation. Consider multivariate approaches that adjust related parameters in coordinated ways.
- Advanced Approach: Develop simplified system models to predict how parameter changes will affect entanglement metrics before applying adaptations.

Avoiding these common pitfalls requires thoughtful implementation, proper validation, and a phased approach that builds confidence in both the measurement and adaptation aspects of the system. By anticipating these challenges, you can implement entanglement metrics that provide reliable insights and effective adaptation across diverse domains.

## 

Entanglement Learning metrics provide a powerful mathematical framework for quantifying and maintaining alignment between systems and their environments. By measuring information throughput across the agent-environment interaction cycle, these



metrics enable real-time detection of misalignments before performance visibly degrades.

This guide has outlined the foundational calculations, practical implementation considerations, and integration strategies needed to apply entanglement metrics in real-world systems. From basic entropy calculations to advanced information gradients, the mathematical tools presented here transform information theory concepts into practical mechanisms for autonomous self-assessment and adaptation.

Implementing these metrics effectively requires attention to discretization strategies, computational efficiency, and appropriate threshold establishment. While challenges exist in applying this approach across diverse domains, the core mathematical framework remains consistent—providing a universal reference for measuring system-environment alignment regardless of specific tasks or objectives.

The applications of this approach extend across numerous domains—from neural networks to control systems, robotics to physical process monitoring. In each case, entanglement metrics provide a window into the fundamental information relationships that determine system effectiveness.

For assistance implementing Entanglement Learning metrics in your specific domain, or to discuss advanced applications of the Information Digital Twin architecture, contact our implementation team at **IDT@semarx.com**. Our specialists can provide guidance on domain-specific adaptations, computational optimizations, and integration strategies tailored to your system requirements.

By leveraging these metrics, systems gain the ability to not only detect when their internal models no longer match reality, but to adapt autonomously—maintaining alignment through continuous information optimization rather than periodic human intervention.

## GLOSSARY

- Action States (A): The decisions or outputs produced by an agent, typically represented as a random variable in the entanglement framework.
- Adaptive Binning: A discretization strategy that adjusts bin sizes based on data distribution characteristics, often allocating more bins to regions with higher information density.
- Base Entanglement (ψ): The mutual information between the combined state-action pair and the resulting state, measured as ψ = MI(S,A;S') = H(S,A) + H(S') - H(S,A,S'). Quantifies overall system-environment alignment.
- Discretization: The process of transforming continuous variables into discrete bins to enable probability distribution calculation and entropy measurement.
- Entanglement: In the EL framework, the sustained mutual predictability between an agent and its environment measured across the complete interaction cycle, quantified through various metrics.
- Entanglement Asymmetry (Λψ): The difference between action-outcome mutual information and state-action mutual information, calculated as Λψ = MI(A;S') - MI(S;A). Indicates whether misalignment originates in perception or control.
- Entanglement Memory (μψ): The mutual information between successive states independent of actions, calculated as μψ = MI(S;S') = H(S) + H(S') H(S,S'). Represents the natural predictability of the environment.
- Entropy: A measure of uncertainty or information content in a random variable, calculated as  $H(X) = -\sum p(x) \log_2 p(x)$ , where p(x) is the probability of value x.
- Environment Responses (S'): The resulting states or observations after an action is taken, represented as a random variable in the entanglement framework.
- Information (Communication Theory): In Shannon's communication theory, information is the reduction of uncertainty, measured in bits, that occurs when a message is received. It quantifies how much knowing one variable's value tells us about another.
- Information (Entanglement Learning): In the EL framework, information is defined as the mutual predictability between an agent and its environment—the degree to which knowledge of the agent's internal state and actions reduces uncertainty about environmental responses, and vice versa.
- Information Structures: The network of predictive relationships between variables—patterns, correlations, and mutual dependencies that enable reliable prediction across contexts. These structures represent how effectively an intelligent system preserves connections between causes and effects, inputs and outputs, and actions and consequences across changing conditions.
- Intelligence: In the EL framework, intelligence is defined not by the optimization of specific goals, but by the
  mechanisms an agent employs to actively maintain, adapt, and create information structures necessary for achieving
  goals, modifying them, or defining new ones. The capacity to maximize information throughput with the environment
  is considered a fundamental characteristic of intelligent systems.
- Information Digital Twin (IDT): The parallel computational architecture that implements Entanglement Learning by



monitoring information flow, calculating entanglement metrics, and generating adaptation signals.

- Information Gradients: The directional derivatives of entanglement metrics with respect to system parameters, indicating how parameters should be adjusted to increase information throughput.
- Information Throughput: The bidirectional flow of information between an agent and its environment, measured through entanglement metrics.
- Joint Entropy: The entropy of multiple random variables considered together, calculated as H(X,Y) = -∑∑ p(x,y) log<sub>2</sub> p(x,y).
- Laplace Smoothing: A technique to address zero probabilities in sparse distributions by adding a small constant to all counts before normalization.
- Mutual Information: A measure of the statistical dependence between two random variables, calculated as MI(X;Y)
   = H(X) + H(Y) H(X,Y). Quantifies how much knowing one variable reduces uncertainty about another.
- Observation States (S): The inputs or perceptions an agent receives from the environment, represented as a random variable in the entanglement framework.
- Percentile-Based Binning: A discretization approach that creates bins with equal numbers of samples, ensuring statistical significance across all bins.
- Probability Distribution: A function that assigns probabilities to all possible values of a random variable, forming the basis for entropy calculations.
- Uniform Binning: A discretization strategy that divides the range of a variable into equal-sized intervals, typically used as a starting point before more sophisticated approaches.