

# ENTROPIC MEMORY

## A thermodynamics–inspired consolidation mechanism for lifelong agent learning

Jing Du · Hang Zhao

Khoury College of Computer Sciences, Northeastern University

Entropy–aware consolidation improves robustness to distractors under a fixed memory budget.

### Problem

Long interaction streams fill agent memory with noisy observations. Retrieval quality drops even when the memory budget is fixed.

- Relevant evidence becomes harder to surface
- Utility–only heuristics can preserve distractors

### Free–energy objective

$$F = E + \lambda S$$

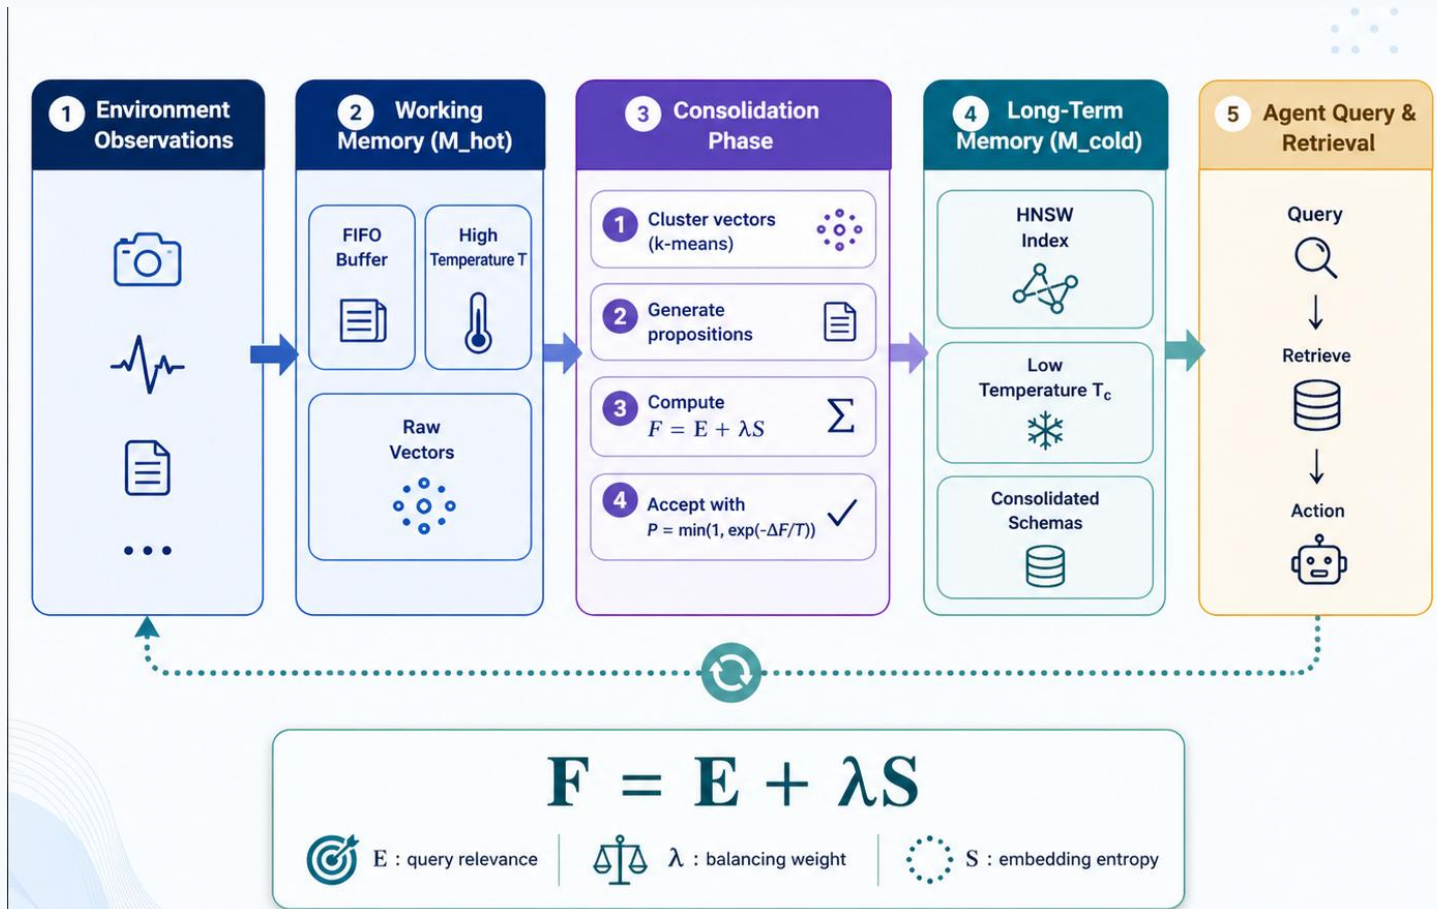
$$E(m) = -\text{Utility}(m) \quad S(m) = H(e)$$
$$P = \min(1, \exp(-\Delta F / T))$$

### Core contribution

A two–tier memory system: a hot working buffer periodically crystallizes into a cold long–term store via a thermodynamic objective.

- Objective–driven consolidation
- Temperature–controlled stochastic replacement

# Method: hot buffer → free-energy consolidation → cold store



## Consolidation rule

At every sleep phase,  $M_{hot}$  is clustered and summarized into candidate propositions. Each candidate competes against the weakest memory in  $M_{cold}$ .

**Free-energy objective**

$$F = E + \lambda S$$

$E(m) = -\text{Utility}(m)$      $S(m) = H(e)$   
 $P = \min(1, \exp(-\Delta F / T))$

Temperature  $T$  controls plasticity: high  $T$  explores more replacements; low  $T$  favors stable, low-entropy memories.

- k-means clustering
- LLM proposition generation
- Metropolis acceptance

# Setup and results

## Infinite Room setup

- 5,000-step horizon
- Concept drift rate  $\delta = 0.005$
- Retrieval threshold  $\tau = 0.7$
- Retrieve top-3 memories
- Capacity  $C = 20$ ; concepts  $N = 50$
- Metric: survival rate (SR)

$$\text{Info Density} = \text{SR} \times (N / C)$$

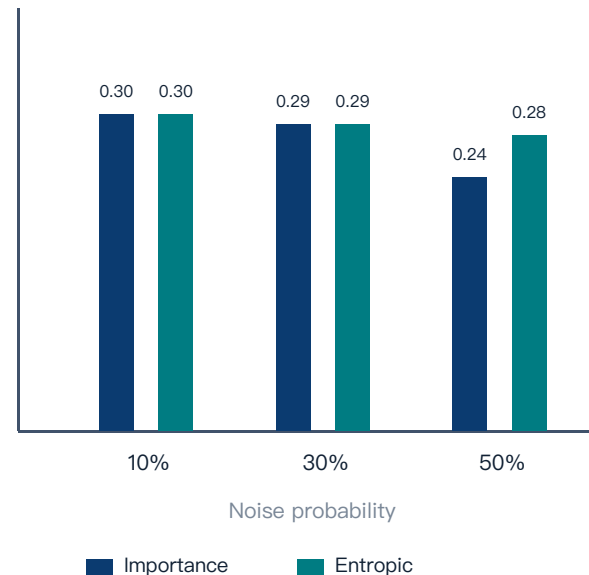
## 30% noise: matches Importance

Method	Survival	Hit@3	Info
Random	0.02 ± 0.00	0.02	0.05
FIFO	0.26 ± 0.04	0.31	0.65
LRU	0.27 ± 0.04	0.33	0.68
Importance	0.29 ± 0.04	0.35	0.73
<b>Entropic</b>	<b>0.29 ± 0.04</b>	<b>0.36</b>	<b>0.73</b>

Table 1. Mean ± std over 5 seeds. Capacity is fixed at 20 for all methods.

**50% noise: 0.24 → 0.28**  
+15% relative improvement

## Noise resilience



The entropy term helps suppress high-frequency noisy embeddings that utility-only memory rules may preserve.

# Takeaways and next steps

## What Entropic Memory changes

- 1 Objective > heuristic**  
Consolidation is posed as free-energy minimization rather than FIFO/LRU pruning.
- 2 Stochasticity helps**  
Metropolis acceptance can escape local utility maxima.
- 3 Entropy filters noise**  
Ambiguous embeddings are penalized even when frequently retrieved.

**Result: +15% relative survival at 50% noise**

## Limitations and future work

- Synthetic Infinite Room only
- Sleep phase adds about 2s per consolidation cycle
- Fixed  $T = 1.0$  and fixed consolidation hyperparameters

---

## Next

- End-to-end evaluation with full LLM agents
- Adaptive temperature and consolidation schedules
- Shared memory for multi-agent systems