

# Information-Theoretic Adaptive Memory Compression for LLM-Based Agents

Anonymous Author(s)

## ABSTRACT

Large language model (LLM) agents accumulate memory episodes—observations, reasoning traces, and tool outputs—that must be re-injected into a finite context window for future steps. Aggressive compression reduces token cost and inference latency but risks discarding task-critical information. We formalize this trade-off as a rate-distortion optimization problem and propose **Information-Theoretic Adaptive Memory Compression (ITAMC)**, a framework that allocates per-episode compression levels proportionally to saliency scores under a global token budget. Through controlled experiments on 100 synthetic memory episodes with 300 ground-truth salient facts, we characterize the Pareto frontier between compression ratio and information retention for three compression operators: extractive, abstractive, and latent. Our results reveal a concave frontier where moderate compression ( $r \approx 0.4\text{--}0.6$ ) achieves 70–87% fact retention while reducing tokens by 40–60%. Knee-point analysis identifies operator-specific optimal compression ratios:  $r^* = 0.42$  for extractive,  $r^* = 0.59$  for abstractive, and  $r^* = 0.26$  for latent compression. Saliency-guided adaptive allocation yields its largest gains under extreme budget constraints (10% budget: +10.2 percentage points for extractive compression), while uniform compression is preferred at moderate budgets. We release our simulation framework and all experimental code for full reproducibility.

## CCS CONCEPTS

• Computing methodologies → Artificial intelligence; Natural language processing.

## KEYWORDS

LLM agents, memory compression, rate-distortion, Pareto frontier, adaptive compression

## 1 INTRODUCTION

Large language model (LLM) agents operate by iteratively reading context, reasoning, and acting [19]. As an agent progresses through a task, it accumulates memory episodes—raw observations, prior reasoning chains, tool outputs, and conversation history—that inform subsequent decisions. Modern agents organize these episodes in structured memory modules with episodic, semantic, and procedural components [10, 15].

A fundamental bottleneck arises because LLMs process fixed-length context windows. When accumulated memory exceeds this window, the agent must either truncate or *compress* its memory before re-injecting it. Compression reduces the token count (lowering API cost and inference latency) but risks losing task-critical information [18]. The survey by Yang et al. [18] identifies this compression–performance trade-off as an open problem, noting

that empirical systems such as LightMem demonstrate clear cost–accuracy tensions but lack a principled framework for selecting compression levels.

The challenge has multiple dimensions. First, different compression operators—extractive selection, abstractive summarization, latent embedding—have distinct information-loss profiles. Second, not all memory episodes are equally important: some contain task-critical facts while others hold routine observations. Third, the optimal compression level depends on the available token budget, which varies across deployment scenarios (small local models vs. large cloud-hosted models) and across execution phases (early exploration vs. focused execution).

This paper makes the following contributions:

- (1) We formalize memory compression for LLM agents as a **rate-distortion optimization** problem (Section 2), connecting agent memory to classical information theory [2, 13].
- (2) We characterize the **Pareto frontier** between compression ratio and information retention for three families of compression operators—extractive, abstractive, and latent—through controlled experiments with ground-truth salient facts (Section 3).
- (3) We propose **ITAMC**, a saliency-guided adaptive compression controller that allocates per-episode compression levels under a global token budget, and demonstrate its effectiveness under extreme budget constraints (Section 3).
- (4) We identify **operator-specific optimal compression ratios** via knee-point analysis and analyze retention stability over long agent horizons (Section 3).

## 1.1 Related Work

**Memory architectures for LLM agents.** MemGPT [10] introduced tiered memory with explicit paging between a main context and external storage, drawing an analogy to operating-system virtual memory. Reflexion [11] showed that storing and reflecting on episodic memory improves multi-step reasoning through self-correction. Recent surveys [3, 15] categorize agent memory into episodic, semantic, and procedural components, each with distinct compression requirements. The agent memory management problem—what to store, how to compress, and when to evict—remains an active area of research [3].

**Context and prompt compression.** Several methods compress prompts or context windows for efficiency. Gist tokens [12] learn fixed-length compressed representations of variable-length contexts through distillation. AutoCompressor [4] trains language models to recursively compress context segments into summary vectors. Li et al. [8] survey prompt compression techniques including lexical pruning, soft-prompt distillation, and retrieval-based

selection. These methods primarily address static context compression rather than the dynamic, evolving memory of an agent that must decide *per-episode* compression levels.

**Compression and language modeling.** Delétang et al. [5] establish a formal connection between language modeling and data compression, showing that prediction and lossless compression are dual formulations of the same problem. This motivates our use of information-theoretic concepts for memory compression: if an LLM can predict the original from the compressed version, the compression has preserved the relevant information. Work on the compression–performance relationship [6] further supports the thesis that compression quality is a proxy for model capability.

**Resource-rational agents.** The resource-rational analysis framework [9, 16] models cognitive agents as optimizing a utility function subject to computational cost constraints. Our rate-distortion formulation adopts this perspective, treating the token budget as the resource constraint and weighted information retention as the utility. Related work on computational efficiency for lifelong agents [14] and memory breadth-fidelity trade-offs under context limits [7] addresses complementary aspects of the same challenge.

**Retrieval-augmented generation.** RAG [17] decouples storage from active context by selectively retrieving relevant document chunks at inference time. Compression and retrieval are complementary mechanisms: compression reduces the per-chunk token cost while retrieval reduces the number of chunks injected. Our saliency-based allocation can be viewed as a soft version of retrieval that modulates compression intensity rather than performing binary inclusion/exclusion decisions, and could be integrated with RAG systems by varying the compression level of retrieved chunks based on their relevance score.

## 2 METHODS

### 2.1 Problem Formulation

Let  $\mathcal{M} = \{m_1, \dots, m_T\}$  denote a set of  $T$  memory episodes accumulated by an LLM agent during task execution. Each episode  $m_i$  has token count  $|m_i|$  and contains a set of salient facts  $\mathcal{F}_i$  relevant to downstream tasks. A compression operator  $C$  parameterized by ratio  $r_i \in (0, 1]$  produces a compressed episode  $\hat{m}_i = C_{r_i}(m_i)$  with  $|\hat{m}_i| \approx r_i \cdot |m_i|$ .

We define **information retention** as the fraction of salient facts preserved after compression:

$$\rho_i(r_i) = \frac{|\mathcal{F}_i \cap \hat{\mathcal{F}}_i|}{|\mathcal{F}_i|} \quad (1)$$

where  $\hat{\mathcal{F}}_i$  denotes the facts recoverable from the compressed episode  $\hat{m}_i$ .

The **memory compression optimization problem** is:

$$\max_{r_1, \dots, r_T} \sum_{i=1}^T w_i \cdot \rho_i(r_i) \quad \text{s.t.} \quad \sum_{i=1}^T |C_{r_i}(m_i)| \leq B \quad (2)$$

where  $B$  is the total token budget and  $w_i$  are task-dependent importance weights derived from saliency scores. This formulation connects directly to rate-distortion theory [2]: the budget  $B$  constrains the *rate* (bits per source symbol, here tokens per memory), and  $(1 - \rho_i)$  measures the *distortion* per episode.

The optimization in Eq. 2 is intractable in full generality because (a) the retention function  $\rho_i(r_i)$  depends on the specific compression operator and episode content, and (b) evaluating downstream task performance requires running the full agent pipeline. We introduce two tractable relaxations: a lightweight saliency model for computing  $w_i$  and a simulation-based characterization of  $\rho_i(r_i)$  for different operator families.

### 2.2 Compression Operators

We study three families of compression operators that span the spectrum of techniques used in practice.

**Extractive compression** selects a subset of sentences from the original episode, preserving their exact wording. Sentences are scored by a proxy for informativeness—the sum of word count and numerical content density (digits per character)—and the top- $k$  sentences are retained in original order until the target token count is reached. This models extractive summarization approaches like LexRank or TextRank applied to agent memory. Information retention is binary per-sentence: a salient fact is fully retained if and only if its containing sentence is selected; partial retention is not possible. This binary behavior produces sharp transitions in the Pareto curve.

**Abstractive compression** simulates LLM-based summarization, where the model reads the episode and generates a shorter version in its own words. Since we require deterministic, API-free experiments, we model the retention of each salient fact independently using a logistic function of the target ratio:

$$P(\text{retain fact} \mid r_i) = \sigma(k \cdot (r_i - \tau)) \quad (3)$$

where  $\sigma$  denotes the sigmoid function,  $k = 8$  controls the steepness of the transition, and  $\tau = 0.35$  is the half-retention threshold (the ratio at which retention probability equals 50%). This models the empirical observation that LLM summarizers exhibit smooth, ratio-dependent fact loss rather than the all-or-nothing behavior of extractive methods.

**Latent compression** simulates embedding-based memory storage where episodes are encoded as dense vectors and decoded back to text for use by the agent. We model per-fact retention probability using a Beta distribution:

$$P(\text{retain fact} \mid r_i) \sim \text{Beta}\left(r_i^{0.6} \cdot \kappa, (1 - r_i^{0.6}) \cdot \kappa\right) \quad (4)$$

where the sub-linear exponent (0.6) models the hypothesis that dense embeddings capture distributional semantics efficiently, and the concentration parameter  $\kappa = 12$  controls the variance of per-fact retention. This produces smoother degradation than both extractive and abstractive operators, consistent with the behavior of variational autoencoders and information bottleneck methods [1, 13].

### 2.3 Saliency Scoring

Given a downstream task query  $q$ , we compute per-episode saliency scores that combine two complementary signals—relevance and recency:

$$s_i = 0.6 \cdot \underbrace{\frac{|\text{tokens}(q) \cap \text{tokens}(m_i)|}{|\text{tokens}(q)|}}_{\text{lexical relevance}} + 0.4 \cdot \underbrace{e^{-\lambda(T-t_i)}}_{\text{recency bias}} \quad (5)$$

---

**Algorithm 1** ITAMC: Adaptive Compression Allocation

---

**Require:** Episodes  $\{m_i\}_{i=1}^T$ , saliency scores  $\{s_i\}$ , budget  $B$   
**Ensure:** Compression ratios  $\{r_i\}_{i=1}^T$

- 1:  $s_i \leftarrow \max(s_i, \epsilon)$  for all  $i$  ▷ avoid division by zero
- 2:  $d_i \leftarrow s_i \cdot |m_i|$  for all  $i$  ▷ desired tokens per episode
- 3:  $\alpha \leftarrow B / \sum_i d_i$  ▷ global scaling factor
- 4:  $r_i \leftarrow \text{clip}(s_i \cdot \alpha, r_{\min}, r_{\max})$  for all  $i$
- 5: **for**  $k = 1$  to  $K_{\max}$  **do**
- 6:    $\hat{B} \leftarrow \sum_i r_i \cdot |m_i|$  ▷ projected token usage
- 7:   **if**  $\hat{B} \leq (1 + \delta) \cdot B$  **then**
- 8:     **break** ▷ budget satisfied
- 9:   **end if**
- 10:    $\gamma \leftarrow \hat{B}/B$  ▷ overshoot factor
- 11:   **for** all  $i$  where  $r_i > r_{\min}$  **do**
- 12:      $r_i \leftarrow \text{clip}(r_i/\gamma, r_{\min}, r_{\max})$
- 13:   **end for**
- 14: **end for**
- 15: **return**  $\{r_i\}_{i=1}^T$

---

where  $t_i$  is the episode timestamp,  $T$  is the latest timestamp, and  $\lambda = 0.02$  is the decay rate. Scores are normalized to  $[0, 1]$  by dividing by the maximum score. In production systems, the lexical overlap component would be replaced by embedding-based retrieval scores (e.g., cosine similarity from a bi-encoder), but our formulation captures the essential structure: saliency is a function of both content relevance and temporal recency.

## 2.4 Adaptive Compression Controller (ITAMC)

ITAMC solves the budget-constrained allocation problem in Eq. 2 by assigning compression ratios proportionally to saliency scores. The procedure, detailed in Algorithm 1, operates in two phases:


*Phase 1: Initial allocation.* Each episode receives a desired token allocation proportional to  $s_i \cdot |m_i|$ , which is then normalized to fit the budget. This ensures that high-saliency episodes receive ratios closer to  $r_{\max} = 1.0$  (minimal compression), while low-saliency episodes receive ratios approaching  $r_{\min} = 0.05$ .

*Phase 2: Iterative projection.* Because clipping ratios to  $[r_{\min}, r_{\max}]$  may violate the budget constraint, we iteratively rescale non-floor ratios until the projected token total fits within  $B$ . Convergence typically occurs within 5–10 iterations.

The computational overhead of ITAMC is negligible: computing saliency scores requires  $O(T \cdot V)$  time where  $V$  is the query vocabulary size, and the allocation loop runs in  $O(K_{\max} \cdot T)$  with  $K_{\max} \leq 20$ . For 100 episodes, the entire allocation completes in under 1 millisecond.

## 2.5 Experimental Setup

**Synthetic benchmark.** We generate 100 memory episodes, each containing 3 salient facts (entity-action pairs drawn from vocabularies of 20 entities and 15 actions) interleaved with 5 filler sentences, yielding a corpus of 6,233 tokens and 300 ground-truth facts. The synthetic design provides *exact ground-truth for measuring retention*, which is impossible with natural-language agent traces where fact boundaries are ambiguous and retention requires subjective evaluation. We additionally define 8 downstream task queries spanning



**Figure 1: Pareto frontier between compression ratio and mean salient-fact retention for three compression operators.** (a) Target compression ratio vs. retention. (b) Actual token usage ratio vs. retention. All curves are concave: moderate compression ( $r \approx 0.4$ – $0.6$ ) achieves 60–87% retention while saving 40–60% of tokens. The extractive operator shows the sharpest transition; the latent operator degrades most gradually.

different information needs (error diagnostics, capacity planning, security auditing, etc.) to evaluate saliency-dependent behavior.

**Evaluation metrics.** We report four metrics: (1) *mean fact retention*  $\bar{\rho} = \frac{1}{T} \sum_i \rho_i$ , the primary quality measure; (2) *compression ratio* (total compressed tokens / total raw tokens), measuring efficiency; (3) *fraction fully retained* (episodes with  $\rho_i = 1.0$ ), measuring per-episode reliability; and (4) *retention delta* ( $\Delta\bar{\rho}$ ), the difference between adaptive and uniform retention at the same budget.

**Experimental protocol.** We conduct five experiments:

- *Exp. 1:* Pareto frontier sweep with 20 uniform compression ratios per operator (Section 3.1).
- *Exp. 2:* Adaptive vs. uniform comparison across 10 budget levels and 8 tasks (Section 3.3).
- *Exp. 3:* Compounding error analysis over horizons of 10–100 episodes (Section 3.4).
- *Exp. 4:* Saliency-stratified retention analysis (Section 3.5).
- *Exp. 5:* Knee-point detection for optimal operating ratios using 50-point sweeps (Section 3.2).

All experiments use seed 42 and are fully deterministic. Source code, data, and figure generation scripts are included in the supplementary material.

## 3 RESULTS

### 3.1 Pareto Frontier Characterization

Figure 1 shows the compression–retention trade-off for all three operators. The key finding is that **all three operators exhibit concave Pareto frontiers**: initial compression yields large token savings with modest retention loss, while aggressive compression below  $r = 0.3$  causes steep degradation. The concavity implies that the “first 40% of savings come nearly for free”—a property with strong practical implications for system design.

Table 1 presents retention values at key compression ratios. Several patterns are notable:

*Extractive compression* shows the steepest transition between  $r = 0.2$  (33.3%) and  $r = 0.4$  (70.7%), a 37.4 percentage-point jump. This reflects its binary sentence-level selection: at  $r = 0.2$ , most sentences

349  
350  
351  
352  
353  
**Table 1: Mean salient-fact retention at selected target com-**  
**354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
**pression ratios across 100 episodes with 300 total facts. Ex-**  
**tractive compression shows the sharpest transition between**  
 **$r=0.2$  and  $r=0.4$ . Latent compression degrades most smoothly.**  
**All operators achieve  $\geq 88.7\%$  retention at  $r=0.8$ .****

Operator	$r=0.2$	$r=0.4$	$r=0.6$	$r=0.8$	$r=1.0$
Extractive	0.333	0.707	0.870	0.980	1.000
Abstractive	0.220	0.617	0.873	0.963	0.980
Latent	0.393	0.590	0.747	0.887	1.000




Figure 2: Optimal operating point detection via knee-point analysis. (a) Pareto curves with detected knee points (diamonds). (b) Marginal retention (gradient of  $\bar{\rho}$  w.r.t.  $r$ ), with dashed vertical lines marking each operator’s knee. The extractive knee occurs at  $r^*=0.42$ ; abstractive at  $r^*=0.59$ ; latent at  $r^*=0.26$ .

containing facts are excluded; by  $r = 0.4$ , the informativeness-based scoring begins to preferentially select fact-bearing sentences. Above  $r = 0.6$ , retention rises steeply to 98.0% ( $r = 0.8$ ) and 100% ( $r = 1.0$ ).

*Abstractive compression* has a smoother curve due to its logistic per-fact retention model. It underperforms extractive at low ratios ( $r = 0.2$ : 22.0% vs. 33.3%) because the sigmoid probability is below 50% for all facts at this level. However, it converges quickly at moderate ratios and nearly matches extractive at  $r = 0.6$  (87.3%). At  $r = 1.0$ , abstractive retention is 98.0% rather than 100%, reflecting the stochastic nature of summarization even without compression.

*Latent compression* degrades most smoothly, as predicted by the sub-linear Beta model. It achieves the highest retention at very low ratios ( $r = 0.2$ : 39.3%) but the lowest at moderate ratios ( $r = 0.6$ : 74.7%), creating a more gradual slope. This reflects the “graceful degradation” property of dense embeddings, which preserve distributional signal even at high compression but struggle to maintain exact factual content at moderate compression.

### 3.2 Optimal Operating Points

We identify the optimal compression ratio for each operator using knee-point analysis of the Pareto curve (Figure 2). The knee point is defined as the ratio that maximizes the perpendicular distance from the line connecting the frontier’s endpoints ( $(r_{\min}, \rho(r_{\min}))$  and  $(r_{\max}, \rho(r_{\max}))$ ). Geometrically, this represents the point of maximum curvature where additional compression begins to cause disproportionate retention loss.




Figure 3: Adaptive (purple) vs. uniform (gray) compression across three operators and 10 token-budget levels (x-axis: fraction of raw tokens). At extreme budgets (10–20%), adaptive allocation preserves critical episodes that uniform compression destroys. At moderate budgets, the approaches converge or uniform slightly leads.

The detected optimal ratios and their associated retentions are:

- **Extractive:**  $r^* = 0.42$ , retention = 0.760 (76.0%)
- **Abstractive:**  $r^* = 0.59$ , retention = 0.870 (87.0%)
- **Latent:**  $r^* = 0.26$ , retention = 0.490 (49.0%)

These results show that the optimal compression level is *operator-dependent*. The abstractive knee occurs at a higher ratio ( $r^* = 0.59$ ) because its smooth logistic curve concentrates curvature in the middle of the range. The extractive knee ( $r^* = 0.42$ ) reflects the sharp binary transition around  $r = 0.3$ – $0.5$ . The latent knee ( $r^* = 0.26$ ) is notably low, indicating that latent compression’s gradual curve places its point of maximum curvature in the aggressive-compression region—below this point, even the graceful latent encoding loses substantial information.

The marginal retention analysis (Figure 2b) provides complementary insight. For extractive compression, marginal retention peaks sharply near  $r = 0.3$  and drops rapidly, indicating a narrow “sweet spot.” For abstractive compression, marginal retention is more uniformly distributed, suggesting less sensitivity to the exact ratio choice. Latent compression shows the flattest marginal retention curve, consistent with its gradual degradation profile.

**Practical guideline.** These findings suggest that system designers should calibrate compression targets to their specific operator rather than using a universal default. A general recommendation based on our results is to target the range  $r \in [0.3, 0.6]$  as the “efficient frontier” where compression yields the best token savings per unit of retention loss.

### 3.3 Adaptive vs. Uniform Compression


Figure 3 compares saliency-guided adaptive allocation against uniform compression across 10 budget levels, averaged over 8 downstream task queries.

Table 2 summarizes the retention delta ( $\Delta\bar{\rho}$ ) at selected budget levels. The results reveal a nuanced picture with two distinct regimes:

**Regime 1: Extreme budgets ( $\leq 20\%$  of raw tokens).** Adaptive allocation provides its largest gains here. At 10% budget, extractive-adaptive achieves 10.2% retention compared to 0.0% for uniform ( $\Delta\bar{\rho} = +10.2$  pp), because the uniform ratio of  $r = 0.1$  falls below the extractive threshold where any fact-bearing sentences can be

465  
466 **Table 2: Retention delta of adaptive over uniform compres-**  
467 **sion ( $\Delta\bar{\rho}$  in percentage points), averaged across 8 task queries.**  
468 **Positive values (**bold**) indicate adaptive advantage. Adaptive**  
469 **allocation is most beneficial at extreme budgets ( $\leq 20\%$ ) and**  
470 **for abstractive compression.**

Budget	Extractive	Abstractive	Latent
10%	<b>+10.2</b>	-0.1	-2.5
20%	-1.5	<b>+3.9</b>	<b>+0.5</b>
30%	-11.9	<b>+1.4</b>	-1.2
40%	-9.8	-5.8	-0.7
50%	-8.1	-11.0	-2.7
60%	-9.3	-14.2	-3.1




490 **Figure 4: Retention gain of adaptive over uniform compres-**  
491 **sion across token budgets. Extractive shows the largest adap-**  
492 **tive gain at 10% budget (+10.2 pp); abstractive benefits at 20%**  
493 **(+3.9 pp). At budgets above 25%, uniform compression is gen-**  
494 **erally preferred, particularly for extractive and abstractive**  
495 **operators.**

496 retained. Adaptive allocation concentrates its limited budget on a  
497 few high-saliency episodes at  $r > 0.3$ , preserving some facts rather  
498 than none. For abstractive at 20% budget, adaptive gains +3.9 pp  
499 by routing tokens to episodes where the logistic retention curve is  
500 steepest.


501 **Regime 2: Moderate budgets ( $\geq 30\%$  of raw tokens).** Uniform  
502 compression is competitive or superior. At 40% budget, uniform-  
503 extractive achieves 72.0% retention vs. 62.2% for adaptive ( $\Delta\bar{\rho} =$   
504  $-9.8$  pp). This occurs because when the budget permits  $r \geq 0.4$   
505 uniformly, extractive compression enters its high-retention regime  
506 for all episodes. Adaptive allocation, by contrast, over-compresses  
507 low-saliency episodes below  $r = 0.3$ , pushing them into the steep  
508 degradation zone.

509 This finding has a clear practical implication: **adaptive allocation**  
510 **should be deployed selectively**, triggered when the token  
511 budget is severely constrained relative to the memory size. At mod-  
512 erate budgets, the simpler uniform strategy is preferred. A hybrid  
513 policy could switch between adaptive and uniform based on the  
514 budget-to-memory ratio.

515 Figure 4 provides a visual summary of the delta across the full  
516 budget range, confirming that the crossover from adaptive advan-  
517 tage to uniform advantage occurs at approximately 15–25% budget  
518 depending on the operator.



523 **Figure 5: Mean fact retention vs. episode horizon for four**  
524 **compression ratios across three operators. At moderate ratios**  
525 **( $r \geq 0.4$ ), retention remains stable, declining by at most 6.3 pp**  
526 **from  $h=10$  to  $h=100$ . At aggressive compression ( $r=0.2$ ), reten-**  
527 **tion is uniformly low regardless of horizon length, indicating**  
528 **that per-episode quality—not accumulation—dominates.**



531 **Figure 6: Mean fact retention at  $r=0.4$  stratified by saliency**  
532 **bin (columns) and compression operator (rows). Retention at**  
533 **a fixed ratio is largely independent of saliency level, confirm-**  
534 **ing that saliency should determine *which episodes receive***  
535 ***more tokens*, not predict their inherent compressibility.**

### 3.4 Retention Stability Over Episode Horizons

575 A critical concern for long-running agents is whether compression  
576 errors compound over many episodes. Figure 5 examines retention  
577 as the number of memory episodes grows from 10 to 100 at four  
578 compression ratios.

579 For moderate compression ( $r \geq 0.4$ ), retention remains remarkably  
580 stable: extractive at  $r = 0.6$  achieves 93.3% at  $h = 10$  and 87.0%  
581 at  $h = 100$ , a decline of only 6.3 pp over a 10 $\times$  increase in memory  
582 length. Abstractive at  $r = 0.6$  shows similar stability (not plotted for  
583 brevity). At  $r = 0.8$ , extractive retention drops from 100% ( $h = 10$ )  
584 to 98.3% ( $h = 100$ ), a negligible 1.7 pp decline.

585 At aggressive compression ( $r = 0.2$ ), retention is already low at  
586 short horizons (33.3% for extractive at  $h = 10$ ) and remains flat, in-  
587 dicating that *per-step compression quality, not error accumulation, is*  
588 *the dominant factor*. This is an encouraging finding: it suggests that  
589 agents can apply consistent moderate compression over long hori-  
590 zons without catastrophic degradation, provided the per-episode  
591 ratio is above the steep part of the Pareto curve.

### 3.5 Saliency-Stratified Analysis

592 Figure 6 presents retention at  $r = 0.4$  stratified by episode saliency  
593 level (low, medium, high) and compression operator. Episodes are  
594 binned by their saliency score: high ( $s \geq 0.7$ ), medium ( $0.3 \leq s <$   
595 0.7), and low ( $s < 0.3$ ).

581 The key finding is that at a fixed compression ratio, **retention**  
 582 is largely independent of saliency level. For extractive compression,  
 583 retention ranges from 69.6% (high-saliency) to 73.8% (low-  
 584 saliency)—a spread of only 4.2 pp. Abstractive shows a slightly  
 585 larger spread (60.5% to 64.0%), while latent ranges from 57.5% to  
 586 62.7%. These differences are within the noise range of our stochastic  
 587 compression models.

588 This result validates a core assumption of ITAMC: saliency  
 589 should determine the *compression allocation* (how many tokens each  
 590 episode receives) rather than predicting *inherent compressibility*  
 591 (how well an episode compresses at a given ratio). In our controlled  
 592 setting, all episodes have similar structure (3 facts + 5 fillers), so  
 593 compressibility is uniform. In real-world settings, episode complexity  
 594 may vary, suggesting that a combined saliency-compressibility  
 595 model could further improve allocation.

## 597 4 DISCUSSION

599 **Design recommendations.** Based on our experimental findings,  
 600 we offer three concrete recommendations for designers of LLM  
 601 agent memory systems: (1) Target the range  $r \in [0.3, 0.6]$  for  
 602 memory compression, which sits on the efficient part of the Pareto  
 603 frontier across all operators. (2) Use saliency-guided adaptive al-  
 604 location when token budgets are below 25% of raw memory size;  
 605 use uniform compression above this threshold. (3) Prefer extractive  
 606 compression when exact fact preservation is critical (it achieves  
 607 100% retention at  $r = 1.0$  and sharp transitions make the high-  
 608 retention regime reliable), and latent compression when graceful  
 609 degradation under variable budgets is desired.

610 **Connection to tiered memory architectures.** Our results pro-  
 611 vide quantitative support for tiered memory designs like MemGPT [10].  
 612 A three-tier system mapping to our findings would use: a *hot tier*  
 613 ( $r \approx 0.8\text{--}1.0, \geq 88.7\%$  retention) for high-saliency recent episodes;  
 614 a *warm tier* ( $r \approx 0.4\text{--}0.6, 59\text{--}87\%$  retention) for medium-saliency  
 615 episodes; and a *cold tier* ( $r \approx 0.1\text{--}0.2, 22\text{--}39\%$  retention) for archival  
 616 episodes used primarily for broad retrieval matching.

617 **Toward task-aware compression.** Our saliency model uses a  
 618 simple combination of lexical overlap and recency. Richer models  
 619 that incorporate task structure—e.g., causal dependencies between  
 620 episodes, entity co-reference chains, or learned distortion predictors  
 621 trained on agent execution traces—could significantly improve  
 622 allocation quality. The rate-distortion framework naturally accom-  
 623 modates such extensions by replacing our proxy  $\rho_i$  with a learned  
 624 distortion function.

## 626 5 CONCLUSION

628 We presented ITAMC, an information-theoretic framework for  
 629 adaptive memory compression in LLM-based agents. Through con-  
 630 trolled experiments on a synthetic benchmark with exact ground-  
 631 truth fact retention, we established four principal findings.

632 First, all three compression operators—extractive, abstractive,  
 633 and latent—exhibit **concave Pareto frontiers**, meaning moderate  
 634 compression ( $r \approx 0.4\text{--}0.6$ ) achieves 60–87% fact retention while  
 635 reducing tokens by 40–60%. This concavity implies that the first 40%  
 636 of token savings come at modest information cost, providing strong  
 637 motivation for adopting memory compression in agent systems.

639 Second, **optimal compression ratios are operator-dependent**:  
 640 knee-point analysis yields  $r^* = 0.42$  (extractive),  $r^* = 0.59$  (abstractive),  
 641 and  $r^* = 0.26$  (latent). System designers should calibrate  
 642 compression targets to their specific operator and application  
 643 requirements.

644 Third, saliency-guided adaptive compression is **most beneficial**  
 645 under **extreme budget constraints** ( $\leq 20\%$  of raw memory), with  
 646 gains up to 10.2 pp for extractive compression at 10% budget. At  
 647 moderate budgets ( $\geq 30\%$ ), uniform compression is a competitive  
 648 and simpler baseline.

649 Fourth, moderate compression **does not compound catastrophically**  
 650 over agent horizons of up to 100 episodes, with retention  
 651 declining by at most 6.3 pp from  $h = 10$  to  $h = 100$  at  $r = 0.6$ .

652 **Limitations.** Our experiments use synthetic data with con-  
 653 trolled fact structure. While this enables precise retention measure-  
 654 ment, it does not capture the full complexity of real-world memory  
 655 content where facts have varying importance, interdependencies,  
 656 and ambiguous boundaries. The compression operators are simula-  
 657 tion proxies; validation with actual LLM-based summarizers and  
 658 embedding models is needed. Extending ITAMC to dynamic online  
 659 settings where saliency shifts during execution, and integrating  
 660 with retrieval-augmented generation systems, remain important  
 661 directions for future work.

## 663 REFERENCES

- [1] Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. 2017. Deep Variational Information Bottleneck. *Proceedings of the International Conference on Learning Representations (ICLR)* (2017).
- [2] Toby Berger. 1971. Rate Distortion Theory: A Mathematical Basis for Data Compression. Prentice-Hall.
- [3] Weize Chen et al. 2025. Agent Memory: What to Store, How to Compress, and Prevent Staleness. *arXiv preprint arXiv:2601.01743* (2025).
- [4] Alexis Chevalier, Alexander Wetzig, Anirudh Ajith, and Danqi Chen. 2023. Adapting Language Models to Compress Contexts. *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2023).
- [5] Grégoire Delétang, Ania Ruoss, Paul-Ambroise Duquenne, Elliot Catt, Tim Genewein, Christopher Mattern, Jordi Grau-Moya, Li Kevin Wenliang, Matthew Aitchison, Laurent Orseau, Shane Legg, and Marcus Hutter. 2024. Language Modeling is Compression. *Proceedings of the International Conference on Learning Representations (ICLR)* (2024).
- [6] Yuxuan Ge et al. 2025. In-context Learning Agents Are Asymmetric Believers. *arXiv preprint arXiv:2510.21909* (2025).
- [7] Xiang Li et al. 2025. Designing Memory and Compression to Retain Breadth with Fidelity under Context Limits. *arXiv preprint arXiv:2510.14240* (2025).
- [8] Yucheng Li, Bo Dong, Frank Guerin, and Chenghua Lin. 2024. Compressing Context to Enhance Inference Efficiency of Large Language Models. *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2024).
- [9] Falk Lieder and Thomas L Griffiths. 2020. Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences* 43 (2020), e1.
- [10] Charles Packer, Vivian Fang, Shishir G Patil, Kevin Lin, Sarah Wooders, and Joseph E Gonzalez. 2024. MemGPT: Towards LLMs as Operating Systems. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- [11] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language Agents with Verbal Reinforcement Learning. *Advances in Neural Information Processing Systems* 36 (2023).
- [12] Jiaao Sun et al. 2024. Learning to Compress Prompts with Gist Tokens. *Advances in Neural Information Processing Systems* 36 (2024).
- [13] Naftali Tishby, Fernando C Pereira, and William Bialek. 2000. The Information Bottleneck Method. *Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing* (2000), 368–377.
- [14] Tao Wang et al. 2025. Computational Efficiency for Lifelong LLM Agents. *arXiv preprint arXiv:2510.16079* (2025).
- [15] Zeyu Wang et al. 2024. A Survey on Memory Mechanisms for Large Language Model Based Agents. *arXiv preprint arXiv:2404.13501* (2024).
- [16] Yifei Wu et al. 2025. Resource-Rational Compute Allocation for Language Reasoning Models. *arXiv preprint arXiv:2509.08827* (2025).

697 [17] Penghao Xu et al. 2024. Retrieval-Augmented Generation for AI-Generated  
698 Content: A Survey. *arXiv preprint arXiv:2402.19473* (2024). 755  
699 [18] Junjie Yang et al. 2026. Toward Efficient Agents: Memory, Tool Learning, and  
700 Planning. *arXiv preprint arXiv:2601.14192* (2026). 756  
701 757  
702 758  
703 759  
704 760  
705 761  
706 762  
707 763  
708 764  
709 765  
710 766  
711 767  
712 768  
713 769  
714 770  
715 771  
716 772  
717 773  
718 774  
719 775  
720 776  
721 777  
722 778  
723 779  
724 780  
725 781  
726 782  
727 783  
728 784  
729 785  
730 786  
731 787  
732 788  
733 789  
734 790  
735 791  
736 792  
737 793  
738 794  
739 795  
740 796  
741 797  
742 798  
743 799  
744 800  
745 801  
746 802  
747 803  
748 804  
749 805  
750 806  
751 807  
752 808  
753 809  
754 810  
755 811  
756 812