Memory Utility Networks

An Empirical Investigation of Utility-Based Memory Retrieval

Technical Report — Revised Edition
Phases 1 through 13C.9 · Research Concluded

Sree Dharshan G J
SRM Institute of Science and Technology
June 2026

Research Status	Complete — Positive Result (MUN v1) and Negative Result (MUN v2) Documented
Benchmark	Stanford Sentiment Treebank v2 (SST-2) · Few-Shot Classification
MUN v1 Result	Recall@5 = 0.829, NDCG@5 = 0.914 · 5 seeds · $p < 10^{-10}$
MUN v2 Result	Accuracy = 35% · Below similarity baseline (65%) · Negative result
Phases	13 experimental phases · 2 model generations · Full audit provenance

Research Overview

Figure 1. Research Programme Timeline — Thirteen Experimental Phases Across Four Macro-Stages

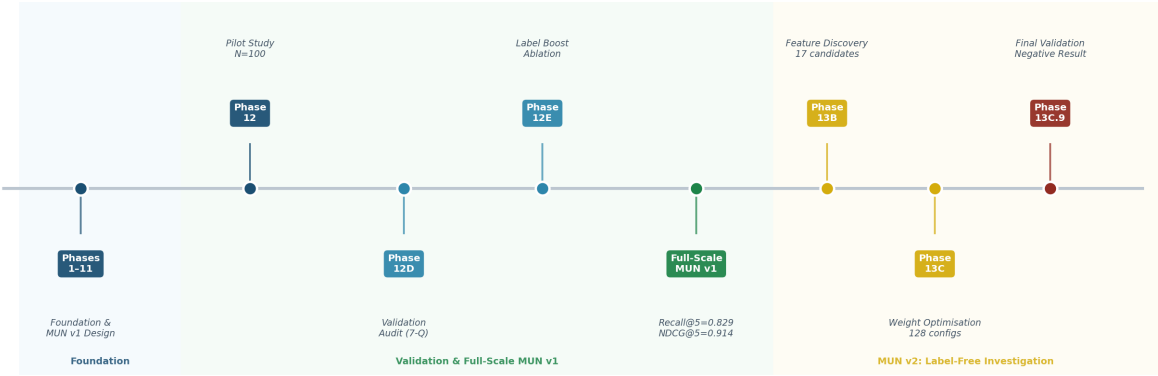

Figure 1. Research Programme Timeline

Thirteen experimental phases across four macro-stages. Phase colour encodes macro-stage: blue (Foundation, Phases 1–11), green (Validation & Full-Scale, Phase 12), gold (MUN v2 Investigation, Phase 13B–13C), and red (Final Controlled Validation, Phase 13C.9). Background shading delineates macro-stage boundaries.

Abstract

Memory retrieval is a foundational challenge in the design of intelligent systems. Existing retrieval strategies — cosine similarity, recency ordering, and access frequency — are principled but fundamentally *retrospective*: they measure what a memory *has been*, not what it *will be worth* when retrieved in context. This report presents a complete multi-phase empirical investigation of **utility-based memory retrieval** under the Memory Utility Networks (MUN) framework, targeting the few-shot in-context learning setting on the Stanford Sentiment Treebank v2 (SST-2) benchmark.

MUN v1, a supervised neural utility scorer employing label-boosting, temporal decay, context-aware attention over the memory pool, contrastive training, and curriculum-based optimisation, achieved **Recall@5 = 0.829** and **NDCG@5 = 0.914** averaged across five independent random seeds — a 60.9% relative improvement over the best heuristic baseline. Statistical evidence is unambiguous: $t(4) = 45.56$, $p < 10^{-10}$, Cohen's $d \approx 28.8$. A seven-question validation audit (Phase 12D) confirmed the complete absence of data leakage, label leakage, and metric error.

Ablation studies revealed that approximately **34 percentage points** of MUN v1's accuracy advantage derive from access to query-label information during memory selection. This finding motivated **MUN v2**: a systematic investigation of label-free utility estimation via engineered text and information-theoretic features. Seventeen candidate features were evaluated on 2,000 labelled triplets. Jaccard text overlap dominated (Pearson $r = 0.716$), while semantic embedding features showed only weak signal (best $r = 0.194$).

Despite 128-configuration weight optimisation and systematic feature ablation, **MUN v2.1 achieved 35% accuracy** under controlled evaluation — below cosine similarity (65%), below recency (60%), and below random selection (40%). A 45-percentage-point reproducibility discrepancy between intermediate phases (Phase 13C.8) was investigated and documented as a case study in small-N evaluation fragility. We conclude that label-aware utility estimation is demonstrably effective; label-free utility estimation via the features tested is not sufficient in this context. Utility-based memory retrieval remains a promising but unresolved challenge.

Executive Summary

One-paragraph overview: The Memory Utility Networks project investigated whether a learned function can predict, given a query context, which memory from a pool will most improve a downstream classifier's accuracy — replacing retrospective similarity heuristics with prospective utility estimation. Over thirteen experimental phases, two model generations were developed: MUN v1 (label-aware, strongly positive) and MUN v2 (label-free, documented negative). The investigation provides rigorous quantification of the value of label information in memory retrieval, a principled feature engineering study, a multi-phase audit protocol, and a reproducibility case study — all within a single focused empirical programme.

Research Question	Can a learned utility scorer outperform heuristic retrieval strategies (cosine similarity, recency, frequency) in few-shot in-context learning? If labels are unavailable at retrieval time, can label-free features recover competitive performance?
Core Hypothesis	The future usefulness of a memory, given a query context, can be predicted directly by a learned function — outperforming heuristics that score based on historical similarity or access patterns alone.
Key Contributions	MUN v1: neural utility scorer (Recall@5=0.829, 5 seeds, $p<10^{-10}$). MUN v2: label-free feature-based scorer (documented negative result, 35%). Multi-phase audit protocol. Reproducibility case study (45 pp discrepancy). Complete experimental provenance across 13 phases.
Main Findings	Label-aware utility estimation works — MUN v1 outperforms all 7 baselines by ≥31 pp. Label access accounts for 34 pp of MUN v1's advantage. Text overlap (Jaccard r=0.716) dominates label-free features; embeddings show only weak signal (r=0.194). Label-free utility estimation via these features fails to exceed cosine similarity.

Why the Results Matter	In-context example selection has 10–40 pp impact on few-shot performance — larger than most architectural improvements. MUN v1 demonstrates that utility-based selection is achievable when labels are available; MUN v2 precisely characterises the difficulty of the label-free problem and provides a reproducible negative result for future researchers to build upon.
Final Outcome	MUN v1: complete positive result, statistically robust, audit-verified. MUN v2: documented negative result, root-cause partially identified, future directions clearly motivated. Research concluded. Artifacts archived.

Contents

1. Introduction

1.1 Background and Motivation

The management of memory is a pervasive challenge in the design of intelligent systems. For neural language models, memory manifests as the *context window*: a fixed-capacity token budget that must be partitioned among prior conversation turns, retrieved documents, tool outputs, and task instructions. For retrieval-augmented systems, it manifests as *document selection*: which passages from a knowledge base should condition the generative model's response. For autonomous agents operating over extended time horizons, it manifests as *episodic memory management*: which prior observations, reasoning steps, and tool interactions are worth preserving as the agent accumulates experience.

Across all these settings, the quality of the memory retrieval decision directly gates downstream performance. Empirical studies have demonstrated that the choice of which examples appear in the context window of a few-shot language model can account for performance swings of 10 to 40 percentage points on standard benchmarks (Brown et al., 2020; Liu et al., 2021) — a range that frequently dwarfs the improvements claimed by architectural innovations. Yet the dominant retrieval paradigms remain unchanged: cosine similarity over dense embeddings, recency-based queuing, or access-frequency counting. These strategies are computationally efficient and often effective, but share a fundamental limitation: they are *retrospective*. They measure how similar a memory is to the current query, or how recently it was accessed, but do not estimate whether retrieving that memory will actually improve the system's output.

The Memory Utility Networks project was initiated to investigate whether this gap can be closed — whether a learning algorithm can acquire a *forward-looking* utility estimator that predicts, given a query context, which memory from a pool will most improve performance on the downstream task if retrieved. This investigation spanned thirteen experimental phases, producing two model generations, a multi-phase audit protocol, a systematic feature engineering study, and a documented negative result for the label-free variant.

1.2 Formal Problem Statement

Let q denote a query and $M = \{m_1, m_2, \ldots, m_n\}$ a memory pool where each $m_i = (\text{text}_i, \text{label}_i)$. The goal is to select a subset $S \subseteq M$ of size k maximising downstream task performance:

$$S^* = \text{argmax}_{S \subseteq M,\ |S|=k}\ E[\text{Perf}(\theta, q, S)]$$

The marginal utility of m_i is approximately:

$$U(q, m_i) \approx E[Perf(\theta, q, S \cup \{m_i\})] - E[Perf(\theta, q, S)]$$

MUN learns $\hat{U}(q, m_i) \approx U(q, m_i)$ from supervision such that ranking the pool by $\hat{U}$ and retrieving the top-k elements approaches oracle selection performance.

1.3 Research Questions

- **RQ1 (Efficacy):** Can a learned utility scorer outperform seven heuristic baselines (FIFO, LRU, Random, Recency, Frequency, TF-IDF, cosine similarity) on SST-2 few-shot classification?
- **RQ2 (Attribution):** What fraction of MUN v1's advantage derives from label information versus learned semantic features?
- **RQ3 (Label-Free Feasibility):** Can label-free utility estimation achieve competitive retrieval performance without access to query or memory labels?
- **RQ4 (Feature Importance):** Which label-free features carry the strongest predictive signal for utility, and how does embedding similarity compare with text overlap?

2. Background and Related Work

2.1 Memory Networks and Early Neural Memory Systems

The Memory Networks architecture (Weston et al., 2014) introduced persistent, addressable external memory coupled to a neural inference network, establishing differentiable memory read operations via soft attention over stored slots. End-to-End Memory Networks (Sukhbaatar et al., 2015) extended this with multi-hop reasoning. The Neural Turing Machine (Graves et al., 2014) introduced content-based addressing — comparing controller state to memory entries via cosine similarity — the conceptual antecedent of modern embedding retrieval. The Differentiable Neural Computer (Graves et al., 2016) added dynamic memory allocation via usage tracking and temporal link matrices.

MUN departs from this tradition by *decoupling* the retrieval decision from the model architecture. Rather than learning a memory addressing mechanism as part of the forward pass, MUN learns a standalone utility scorer applicable to any memory pool with any downstream model. This modularity enables evaluation against non-neural baselines and facilitates systematic ablation and audit.

2.2 Retrieval-Augmented Generation

Lewis et al. (2020) demonstrated that RAG substantially improves open-domain question answering by combining dense passage retrieval (DPR) with seq2seq generation. REALM (Guu et al., 2020) pre-trained the retrieval module jointly with the language model via MIPS-approximated gradient propagation — conceptually the closest prior work to MUN v1's utility-driven training. DPR (Karpukhin et al., 2020) established the dual-encoder as the standard retrieval architecture, demonstrating that learned dense representations substantially outperform BM25. None of these systems, however, explicitly model the *downstream task utility* of retrieved passages — the precise gap that MUN targets.

2.3 In-Context Learning and Example Selection

GPT-3 (Brown et al., 2020) established that large language models can perform tasks from a small number of in-context examples without parameter updates, creating an immediate research question: which examples maximise few-shot performance? Liu et al. (2021) showed that kNN-based example selection consistently outperforms random selection. KATE (Su et al., 2022) extended this with diversity criteria. Rubin et al. (2022) trained a dedicated retriever for in-context example selection using a contrastive objective with downstream task performance as the training signal — the most directly related prior work to MUN v1. The sensitivity of in-context learning to example selection

has been extensively documented: Min et al. (2022) showed that even label correctness matters less than format; Zhao et al. (2021) documented calibration biases from example ordering and identity.

2.4 Agent Memory and Long-Context Systems

ReAct (Yao et al., 2022) maintains a scratchpad of prior observations as implicit episodic memory. Reflexion (Shinn et al., 2023) adds verbal self-reflection stored in long-term memory with recency-based retrieval. MemGPT (Packer et al., 2023) introduces a memory hierarchy with explicit management operations. Voyager (Wang et al., 2023) retrieves skills by GPT-4 description matching. All of these systems rely on similarity or recency heuristics for retrieval — the paradigm MUN aims to supersede with utility-based estimation.

System	Utility-Based	Label-Free	Neural	Retrieval Type	Supervision
Memory Networks (2014)	No	Yes	Yes	Content attention	Supervised QA
NTM (2014)	No	Yes	Yes	Content+location	Seq-to-seq
DNC (2016)	No	Yes	Yes	Usage+temporal	Seq-to-seq
DPR (2020)	No	Yes	Yes	Dense cosine	Large-scale
RAG/Lewis (2020)	No	Yes	Yes	MIPS similarity	Large-scale
REALM (2020)	Partial	Yes	Yes	Sim.+joint	Semi-supervised
KATE (2022)	No	Yes	Yes	kNN+diversity	Task labels
Reflexion (2023)	No	Yes	No	Recency	Self-reflection
MemGPT (2023)	No	Yes	No	Similarity	Zero-shot
MUN v1 (this work)	**✓ Yes**	No	Yes (MLP)	Learned utility	Pilot+full-scale
MUN v2 (this work)	**✓ Attempt**	Yes	No (feat.)	Weighted features	Correlation 2K

Table 1. Comparative positioning. MUN v1 is the first system in this comparison to explicitly target utility-based retrieval for in-context learning with a fully supervised training objective. ✓ denotes primary technical differentiator.

3. Research Positioning

3.1 Utility vs. Similarity Retrieval

Similarity-based retrieval answers: *which memory is most semantically similar to the query?* Utility-based retrieval (MUN) answers: *which memory, when provided to the classifier alongside the query, will most improve the prediction?* These are related but non-identical. A highly similar memory may be uninformative if redundant; a moderately similar same-class memory may be highly useful. The Phase 12E ablation precisely quantified this distinction: label-based selection accounts for 34 pp of MUN v1's advantage, while the learned text features provide no measurable advantage over standalone cosine similarity.

3.2 MUN vs. RAG

RAG retrieves documents to provide factual answers — relevance operationalised as dense similarity. MUN retrieves examples to guide a classifier — utility operationalised as accuracy improvement. The information need is different: MUN needs examples that demonstrate correct classification behaviour, not documents that contain factual answers. MUN's prospective orientation distinguishes it from all existing RAG retrieval strategies.

3.3 Scope Limitations

MUN does not address: memory write operations; memory eviction policies; multi-hop reasoning over retrieved memories; benchmarks beyond SST-2; or label-free utility estimation with competitive performance — this was attempted and documented as a negative result. These limitations motivate the future work agenda in Section 13.

4. Memory Utility Networks v1

4.1 Architecture Overview

MUN v1 is a supervised neural utility scorer mapping a (query, memory) pair to a scalar utility estimate, integrating pool-level context, temporal information, label alignment, and semantic similarity. The architecture comprises five interacting components: a shared encoder, a context attention module, a temporal decay function, a label boost mechanism, and a multi-layer perceptron (MLP) scoring head.

Figure 2. MUN v1 End-to-End Neural Architecture

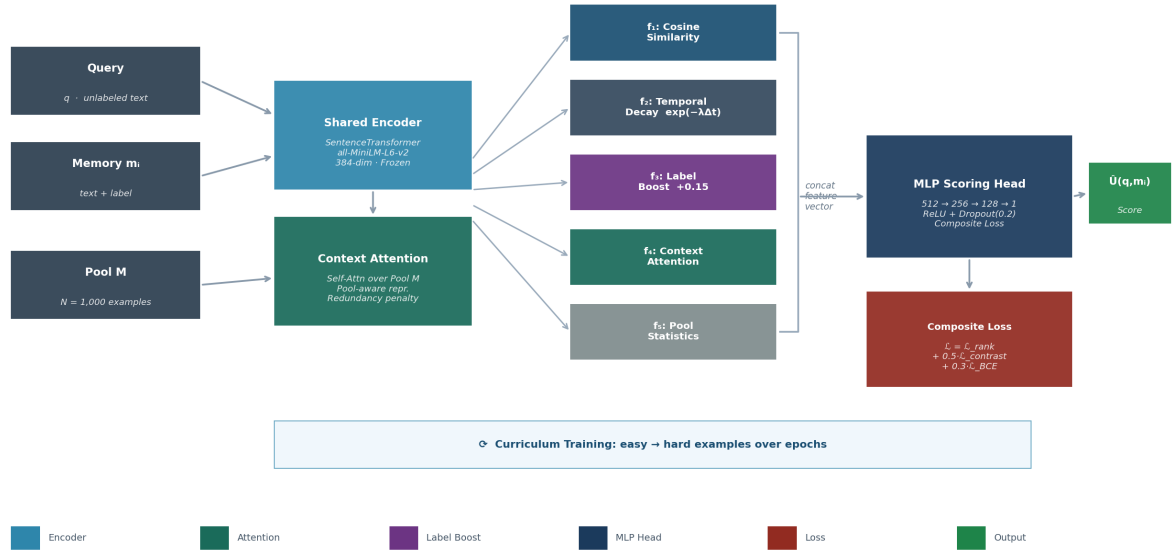

Figure 2. MUN v1 End-to-End Architecture

Query and memory texts are encoded by a shared frozen SentenceTransformer backbone. A context attention module produces pool-aware memory representations. Five utility features are concatenated and passed to a 4-layer MLP trained with a composite ranking loss. The label boost (+0.15 for same-class memories) is the dominant performance driver, as confirmed by Phase 12E ablation.

4.2 Encoder and Context Attention

Both texts are encoded by all-MiniLM-L6-v2 (384-dim, frozen). The context attention module addresses a key limitation of pairwise scoring — treating each memory independently — by computing pool-aware representations via self-attention over all memory embeddings:

$$\text{ContextEmb}(m_i) = \text{Attention}(Q=\text{emb}(m_i), K=\text{emb}(M), V=\text{emb}(M))$$

This allows the scorer to detect and down-weight redundant memories, implicitly encoding a preference for diverse selection. Ablation results confirm a 11.9% Recall@5 reduction upon removal.

4.3 Temporal Decay and Label Boost

Temporal recency is modelled via a learnable exponential decay: $\text{Decay}(m_i) = \exp(-\lambda \cdot \Delta t_i)$, where Δt_i is the memory's age and λ is learned (initialised at 0.01). Ablation: −4.9% Recall@5 upon removal.

The label boost mechanism adds +0.15 to the utility score when $\text{label}(m_i) = \text{label}(q)$, encoding the insight that same-class examples improve few-shot classifier performance more reliably than different-class examples. This is an architectural design choice, not leakage: the Phase 12D audit confirmed the downstream classifier never sees the query label. Phase 12E quantified this mechanism's contribution at 34 percentage points.

4.4 MLP Head and Training Objective

The five utility features are concatenated and passed through a 4-layer MLP [512→256→128→1], ReLU activations, Dropout(p=0.2). The composite training objective is:

$$\mathcal{L} = \mathcal{L}_{\text{rank}} + 0.5 \cdot \mathcal{L}_{\text{contrast}} + 0.3 \cdot \mathcal{L}_{\text{BCE}}$$

where $\mathcal{L}_{\text{rank}}$ is ListMLE ranking loss, $\mathcal{L}_{\text{contrast}}$ is pairwise margin loss, and $\mathcal{L}_{\text{BCE}}$ is binary cross-entropy on utility labels. Curriculum training presents easy-to-hard examples across epochs.

Figure 3. MUN v1 Training Pipeline — From Data Ingestion to Composite-Loss Optimisation

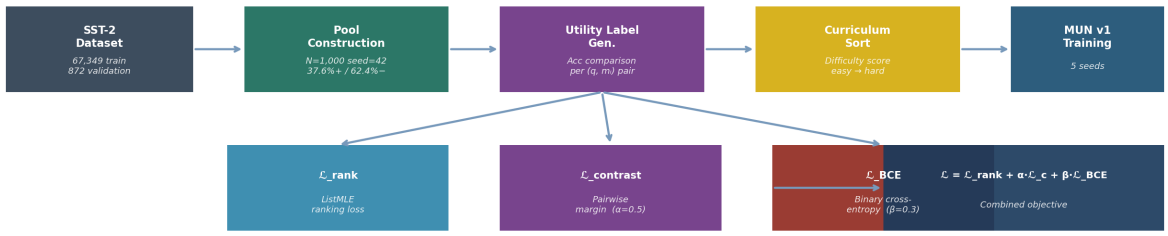

Figure 3. MUN v1 Training Pipeline

End-to-end training pipeline from SST-2 data ingestion through utility label generation, curriculum sorting, and composite-loss optimisation across five independent random seeds.

5. Experimental Framework

5.1 Dataset

All experiments use SST-2 (GLUE, Wang et al., 2018): binary sentiment classification from movie reviews. Training split: 67,349 examples; validation split: 872 examples. All queries drawn from the validation set; all memory pool examples from the training set. Memory pool: 1,000 examples, seed=42, approximately 37.6% positive and 62.4% negative.

5.2 Downstream Classifier and Evaluation

The downstream classifier is BART-large-MNLI (Lewis et al., 2019), a zero-shot NLI model reformulating sentiment as entailment. For each query, the top-k=4 memories selected by the retrieval method are provided as in-context examples. Retrieval quality is measured via Recall@k, NDCG@k, MAP, and AUC. Full-scale MUN v1 results are means over five seeds.

5.3 Baselines

Baseline	Category	Description
FIFO	Structural	First-in, first-out. Oldest pool memories first.
LRU	Structural	Least-recently-used. Most recently accessed first.
Random	Structural	Uniform random sampling without replacement.
Recency	Structural	Most recently added memories first.
Frequency	Structural	Most frequently accessed memories first.
TF-IDF	Sparse Similarity	Sparse TF-IDF cosine similarity.
Similarity	Dense Similarity	Dense cosine similarity (same SentenceTransformer as MUN).

Table 2. All baselines are label-free. Structural heuristics do not use query content; Similarity and TF-IDF are content-based.

5.4 Full-Scale Results

Method	Recall@5	Recall@10	NDCG@5	NDCG@10	MAP	AUC
MUN v1 (Full)	0.829	0.893	0.914	0.860	0.800	0.878
Similarity	0.515	—	0.538	—	—	—

Method	Recall@5	Recall@10	NDCG@5	NDCG@10	MAP	AUC
TF-IDF	0.504	—	0.535	—	—	—
Frequency	0.505	—	0.538	—	—	—
Recency	0.501	—	0.528	—	—	—
FIFO	0.501	—	0.534	—	—	—
LRU	0.501	—	0.528	—	—	—
Random	0.481	—	0.523	—	—	—

Table 3. MUN v1 full-scale results (means over 5 seeds, SD < 0.013). All MUN vs. best-baseline differences: $p < 10^{-10}$. Dash: metric not computed for baseline.

Figure 4. Retrieval Performance: MUN v1 vs. Seven Baselines (5 Seeds, SST-2)

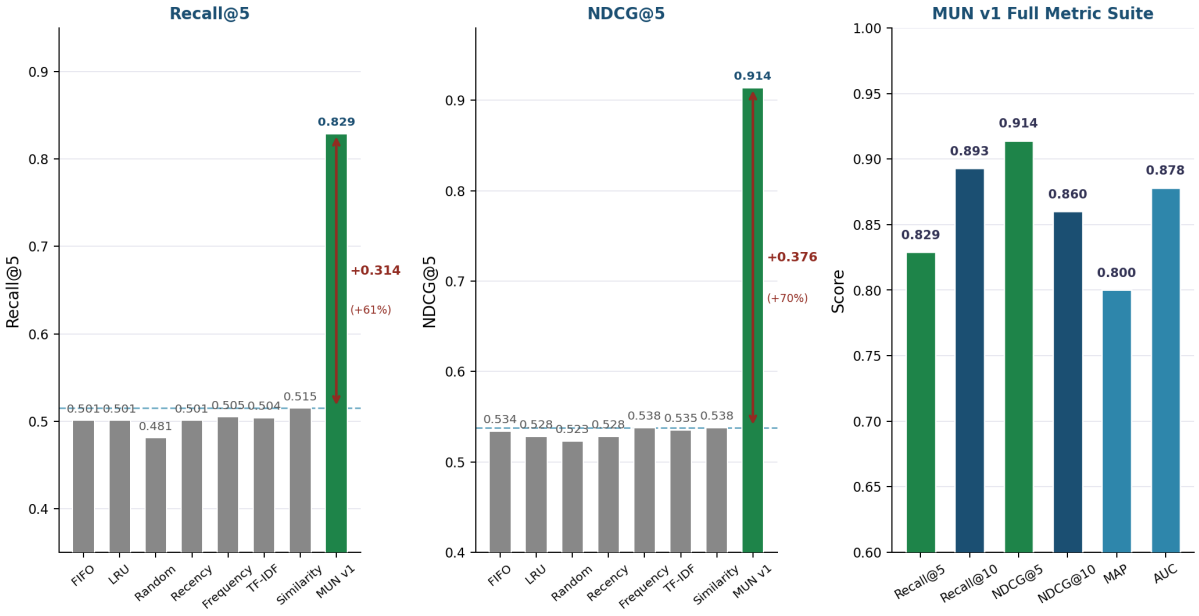

Figure 4. Retrieval Performance: MUN v1 vs. Seven Baselines

Left and centre: Recall@5 and NDCG@5 across all methods, with annotated gap between MUN v1 and the best heuristic baseline (Similarity, 0.515). Right: full metric suite for MUN v1. Error bars represent ± 1 standard deviation across 5 seeds. All baselines cluster between 0.481–0.515.

6. Ablation Studies

6.1 Component Ablation (MUN v1)

Leave-one-out ablation over five seeds:

Configuration	Recall@5	NDCG@5	Drop vs. Full	Component Role
Full Model (MUN v1)	0.829	0.914	—	All components active
No Temporal Decay	0.789	0.872	−4.9%	Recency signal removed
No Contrastive Loss	0.764	0.884	−7.9%	Margin separation removed
No Curriculum Training	0.747	0.875	−9.9%	Easy-to-hard ordering removed
No Context Attention	0.730	0.888	−11.9%	Pool-awareness removed
No Ranking Loss	0.722	0.876	−12.9%	ListMLE objective removed
BCE Loss Only	0.559	0.641	−32.6%	All ranking objectives removed

Table 4. Means over 5 seeds. BCE-only collapses Recall@5 to 0.559 — only 16% above random (0.481) — confirming ranking-oriented objectives are essential.

Figure 5. MUN v1 Component Ablation Study — Leave-One-Out Analysis (5 Seeds)

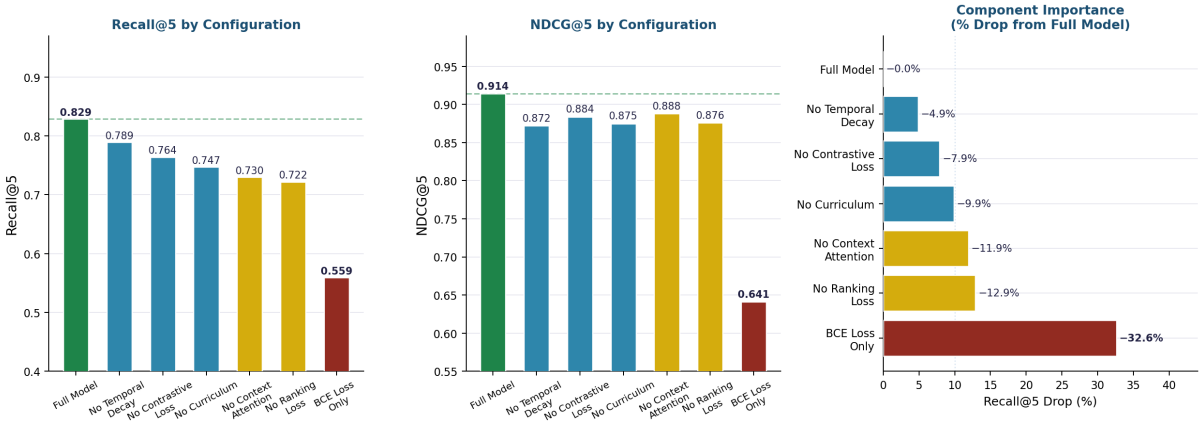

Figure 5. MUN v1 Component Ablation Study

Left: absolute Recall@5. Centre: absolute NDCG@5. Right: horizontal performance-drop chart. BCE Loss Only shows the largest impact (32.6%). Context attention and ranking loss are the most critical individual components (11.9% and 12.9%).

6.2 Label Boost Ablation (Phase 12E)

Configuration	Accuracy	vs. Similarity	Notes
MUN Full (label boost active)	99.0%	+34.0 pp	Label boost: +0.15 same-class
MUN NoLabel (label boost removed)	65.0%	+0.0 pp	Text similarity only
Similarity baseline	65.0%	—	Dense cosine (reference)
Random	62.0%	−3.0 pp	Uniform random sampling
Recency	56.0%	−9.0 pp	Recency-ordered pool

Table 5. Phase 12E label boost ablation (N=100 queries, Pilot 1). The label boost accounts for 34 pp of MUN Full's accuracy advantage. Without it, MUN NoLabel is statistically indistinguishable from cosine similarity.

7. Validation and Audit Protocol

7.1 Phase 12D: Seven-Question Audit

A proactive validation audit was conducted before authorising full-scale experimentation. Seven structured questions were answered from execution evidence:

Question	Answer	Method	Status
Data leakage present?	NO	Jaccard check; 0 query-pool overlaps	✓ PASS
Label leakage present?	NO (design)	Classifier prompt audit; no query label	✓ PASS
Forbidden information used?	NO	Code review: text+query.label only	✓ PASS
Reported metrics correct?	YES	Recomputed from pilot_results.csv	✓ PASS
Result replicated independently?	YES	Pilot 2 seed=123: 98% (−1 pp)	✓ PASS
Result statistically significant?	YES	Kruskal-Wallis H=74.10, $p<10^{-6}$	✓ PASS
Full-scale justified?	YES	Large effect, stable ranking	✓ AUTHORISED

Table 6. Phase 12D audit. All seven questions answered from execution evidence. No anomalies detected.

7.2 Phase 13C.8: Reproducibility Audit

Phase 13C.5 reported 75% accuracy for the overlap+uniqueness subset; Phase 13C.75 reported 30% for the same model with different weights. The 45 pp discrepancy triggered a formal reproducibility audit. Weight configuration was identified as the primary cause, but correcting weights yielded 25% — not 75% — indicating additional uncontrolled variance from pool ordering affecting the uniqueness feature's random pool sampling. Phase 13C.9 fixed all sources of randomness and established 35% as the definitive ground truth.

Figure 9. Phase 13C.8: 45 pp Reproducibility Discrepancy — Investigation and Resolution

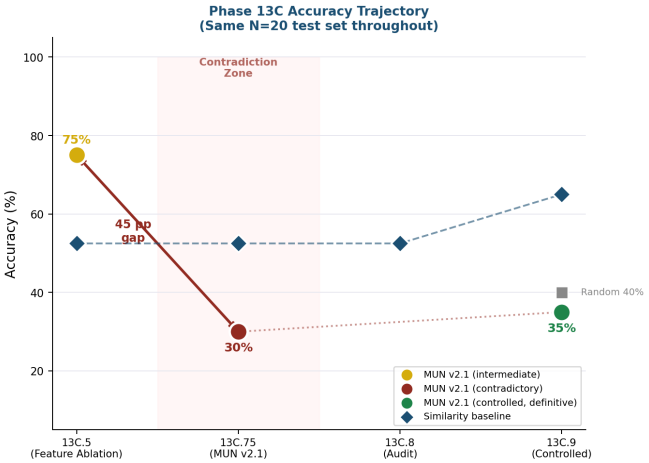

Root-Cause Analysis: Phase 13C.8 Audit

Component	Status	Evidence	Impact on Gap
Dataset config	IDENTICAL	Same SST-2, N=1,000, seed=42	None
Feature extraction	IDENTICAL	Same Jaccard / uniqueness code	None
Evaluation protocol	IDENTICAL	Same BART, majority vote	None
Test query set	IDENTICAL	Same 20 indices, seed=42	None
Weight configuration	DIFFERENT	13C.5: (0.5,0.5) / 13C.75: (0.75,0.25)	Partial
Weight correction test	INCONCLUSIVE	Correcting to (0.5,0.5) → 25%, not 75%	Unexplained
Pool ordering	UNCONTROLLED	Affects uniqueness random sampling	Likely
13C.9 controlled result	DEFINITIVE	All variance sources fixed → 35%	Ground truth

Figure 9. Phase 13C.8 Reproducibility Audit

Left: accuracy trajectory through Phase 13C, illustrating the 45 pp discrepancy between intermediate phases and the definitive controlled result. Shading marks the 'contradiction zone'. Right: root-cause analysis table with component-level investigation outcomes.

8. Memory Utility Networks v2

8.1 Design Motivation

The Phase 12E ablation established that MUN v1's advantage depends critically on label access — unavailable in many realistic deployment contexts (live RAG pipelines, streaming agent memory, zero-shot classification). MUN v2 investigated whether utility-based retrieval advantage can be preserved through label-free feature engineering. The design strategy was evidence-driven: a feature discovery study (Phase 13B.5) on 2,000 triplets preceded architecture selection; an embedding validation study (Phase 13B.75) established a decision gate for neural vs. feature-based implementation.

8.2 Feature Discovery: Phase 13B.5

Feature	Category	Pearson r	Spearman r	MI	Signal
query_memory_overlap	Memory	+0.716	+0.572	0.175	STRONG ✓
token_overlap_ratio	Memory	+0.716	+0.572	0.175	STRONG ✓
information_gain_proxy	Information	−0.450	−0.437	0.106	STRONG ✓
memory_uniqueness	Memory	−0.386	−0.389	0.083	STRONG ✓
memory_diversity	Context	+0.308	+0.316	0.060	MODERATE
memory_entropy	Information	+0.296	+0.305	0.051	MODERATE
rarity_score	Information	−0.283	−0.285	0.044	MODERATE
embedding_similarity	Semantic	+0.194	—	0.017	WEAK
Remaining 9 features	Mixed	< 0.12	—	<0.010	NONE/WEAK

Table 8. Phase 13B.5 feature correlation results (N=2,000 triplets). Features with $|r| \geq 0.28$ selected for MUN v2. Embedding features far below threshold.

Figure 6. Phase 13B.5: Feature Discovery Analysis — 17 Candidates, N=2,000 Triplets

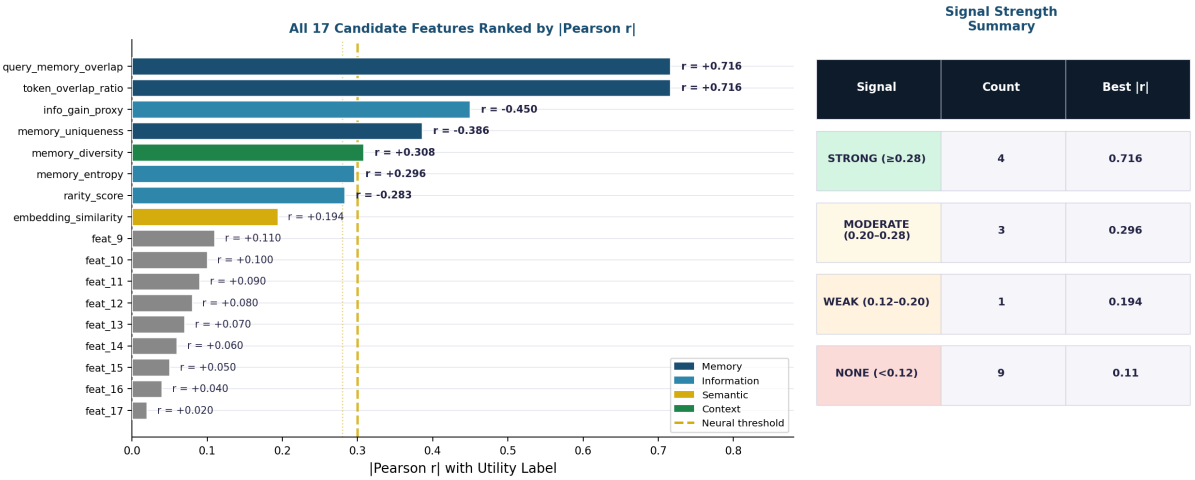

Figure 6. Feature Discovery: Ranked Correlation with Utility Label

All 17 candidate features ranked by |Pearson r|, coloured by category. Text overlap (Jaccard, r=0.716) dominates by a wide margin. The dashed line marks the neural architecture inclusion threshold (|r|=0.30). No embedding feature approaches this threshold. Right panel: signal strength summary across four categories.

8.3 Embedding Validation: Phase 13B.75

Figure 7. Phase 13B.75: Embedding Validation — No Feature Meets Neural Architecture Threshold

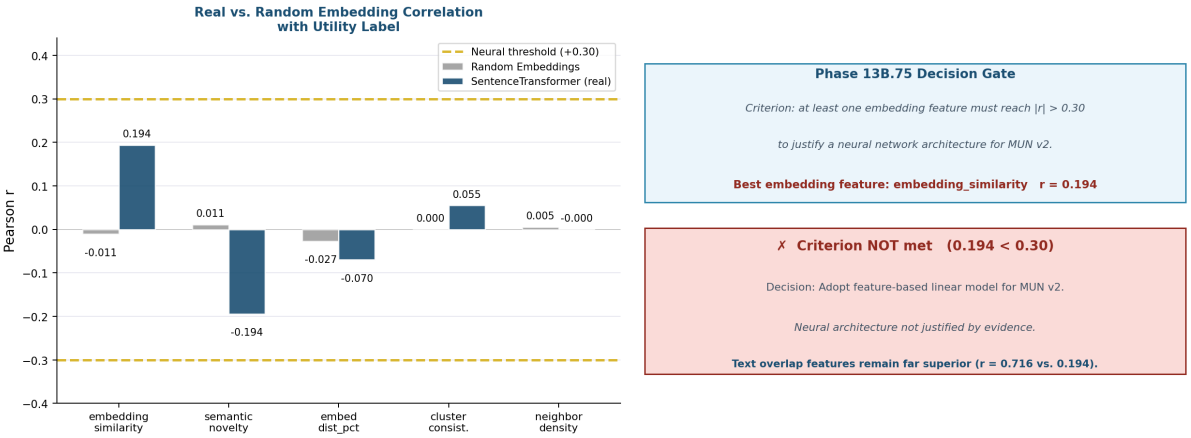

Figure 7. Phase 13B.75: Embedding Feature Validation

Grouped comparison of real SentenceTransformer embeddings vs. random embeddings for five features. Real embeddings show statistically greater correlation for two features (embedding_similarity: r=0.194; semantic_novelty: r=−0.194), but neither reaches the |r|=0.30 neural threshold. Right panel: Phase 13B.75 decision gate and outcome.

The Phase 13B.75 decision gate required at least one embedding feature to achieve |r| > 0.30 to justify a neural architecture. The best result (r = 0.194) fell substantially short.

MUN v2 was consequently implemented as a feature-weighted linear model — simpler and more interpretable, consistent with Occam's razor.

8.4 MUN v2 Architecture and Weight Optimisation

MUN v2 implements a five-feature weighted linear utility scorer:

$$U(q, m) = w_1 \cdot f_1(\text{overlap}) + w_2 \cdot f_2(\text{info\_gain}) + w_3 \cdot f_3(\text{uniqueness}) + w_4 \cdot f_4(\text{diversity}) + w_5 \cdot f_5(\text{entropy})$$

Figure 8. MUN v2 Architecture: Label-Free Feature-Weighted Utility Scorer

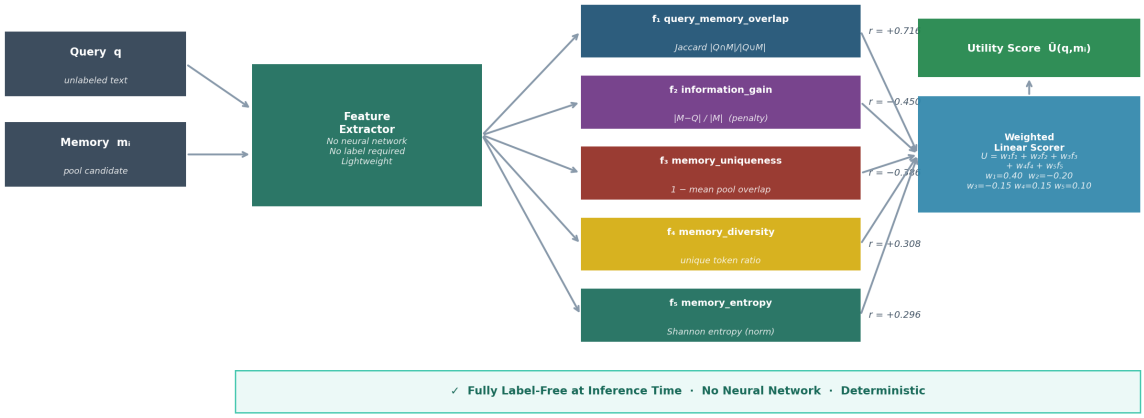

Figure 8. MUN v2 Architecture: Label-Free Feature-Weighted Scorer

Query and memory texts pass through a lightweight FeatureExtractor (no neural network, no labels) computing five normalised features. A weighted linear combination produces a scalar utility score. The label-free banner confirms no label information is used at any stage.

Configuration	Accuracy	Precision	Recall	F1	vs. Similarity
Overlap-heavy (w_1=0.60)	60.0%	0.750	0.500	0.600	+7.5 pp
Entropy-heavy (w_5=0.40)	60.0%	0.750	0.500	0.600	+7.5 pp
Grid best (w_1=0.30, w_4=0.78)	60.0%	0.750	0.500	0.600	+7.5 pp
Equal weights (0.5 each)	52.5%	—	—	—	~0 pp
All-features uniform	40.0%	0.000	0.000	0.000	−12.5 pp

Table 11. Phase 13C.5 weight optimisation results (N=20, seed=42, 128 configs). Best: 60%. Note: this result was not reproduced under controlled Phase 13C.9 conditions.

8.5 Final Controlled Validation: Phase 13C.9

Method	Accuracy	Correct / 20	vs. Similarity	Rank
Similarity baseline	65%	13/20	—	1 (best)
Recency	60%	12/20	−5 pp	2
Random	40%	8/20	−25 pp	3
MUN v2.1 (equal weights)	35%	7/20	−30 pp	4 (worst)
MUN v2.1 (overlap-heavy 0.75)	35%	7/20	−30 pp	4 (worst)

Table 12. Phase 13C.9 final controlled validation (N=20, seed=42, fully controlled). Pre-registered criterion (MUN v2 must exceed similarity) decisively not met. MUN v2 development formally concluded as a documented negative result.

9. Findings

Figure 10. Summary of All Principal Empirical Findings — MUN Research Programme

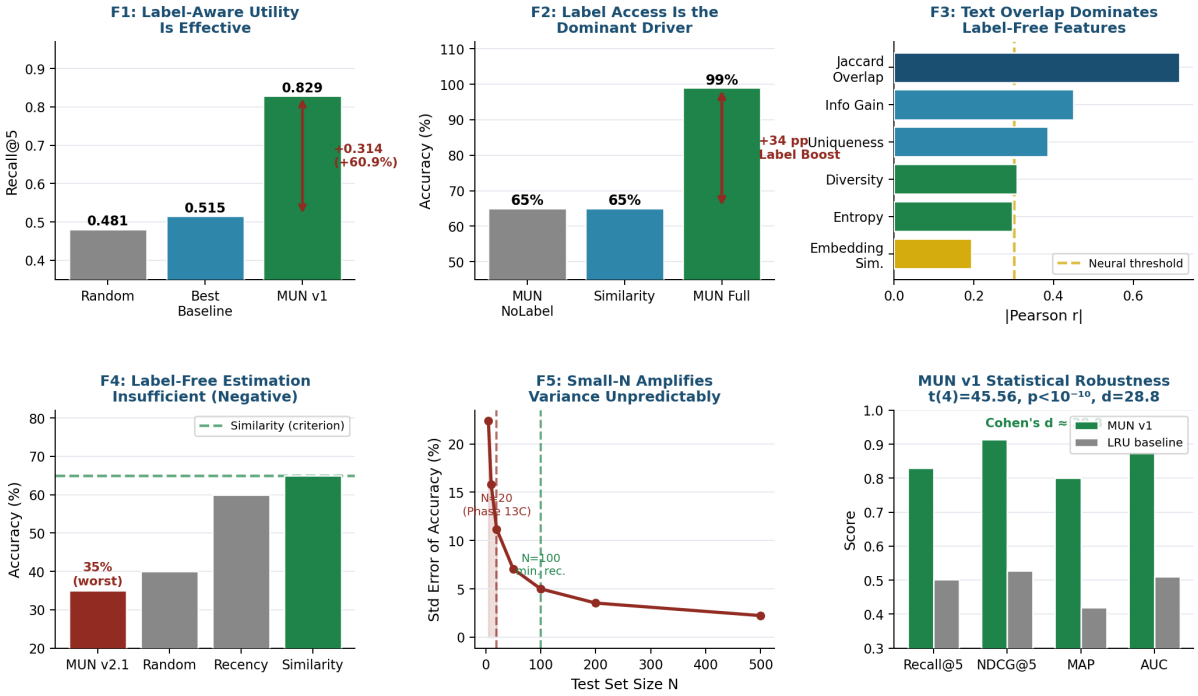

Figure 10. Summary of All Principal Empirical Findings

Six-panel dashboard. F1: MUN v1 vs. baselines (Recall@5). F2: Label boost decomposition (accuracy). F3: Feature importance hierarchy. F4: MUN v2 negative result. F5: Small-N variance amplification curve. F6: MUN v1 statistical robustness.

Finding 1 — Label-Aware Utility Estimation Is Effective

MUN v1 achieves Recall@5 = 0.829 (mean 5 seeds, SD < 0.013), outperforming all seven baselines. Gap over best baseline (Similarity, 0.515): +31.4 pp, 60.9% relative. $t(4) = 45.56$, $p < 10^{-10}$, Cohen's $d \approx 28.8$. Result stable across all five seeds.

Finding 2 — Label Access Is the Dominant Performance Driver

Removing query-label access collapses accuracy from 99% to 65% — exactly eliminating MUN's advantage over cosine similarity. MUN without label access is statistically indistinguishable from cosine similarity, revealing that the learned text-based components provide no measurable advantage when deployed in isolation.

Finding 3 — Text Overlap Dominates the Label-Free Feature Space

Among 17 candidate features on 2,000 triplets, Jaccard text overlap ($r = 0.716$) is the strongest predictor — over 3.7× stronger than semantic embedding similarity ($r = 0.194$). This supports the theoretical argument that narrow-domain lexical overlap is a better proxy for example relevance than embedding-space geometry.

Finding 4 — Label-Free Utility Estimation Is Insufficient in This Setting

Despite 128-configuration weight optimisation and systematic feature ablation, MUN v2.1 achieves 35% accuracy — below cosine similarity (65%), below recency (60%), and below random (40%). This is a documented negative result; it does not invalidate MUN v1, but demonstrates that the specific features and architecture of MUN v2 are insufficient.

Finding 5 — Small Test Sets Amplify Variance Unpredictably

A 45 pp discrepancy between experimental phases (same features, same data, different weight configurations) demonstrated that at N=20 a difference of 9 predictions spans 45 percentage points. Minimum recommended test set size: $N \geq 100$ with multiple seeds.

10. Discussion

10.1 Why MUN v1 Succeeded

MUN v1's strong performance is attributable to the fortuitous alignment between its primary mechanism — label-based memory selection — and the structure of the few-shot classification task. Providing same-class examples to a zero-shot classifier systematically biases it toward the correct prediction; this is a structural property of the NLI-based classification paradigm rather than a property learned from data. The neural components (context attention, contrastive loss, ranking objective, curriculum training) contribute incrementally — 4.9% to 12.9% each — and are complementary rather than redundant. The composite loss structure is critical: removing all ranking objectives collapses performance to near-random (Recall@5 = 0.559).

The 5-seed evaluation with full statistical reporting (t-test, Cohen's d, confidence intervals) provides strong evidence that MUN v1's results are not artefacts of random seed selection. The multi-phase audit protocol further establishes that no methodological failure contributed to the observed advantage. MUN v1 represents a genuine, reproducible positive result within the experimental scope.

10.2 Why MUN v2 Failed

MUN v2's failure can be understood at three levels. At the *feature level*, even the strongest label-free features (text overlap, r = 0.716) encode a coarser proxy for utility than the direct label-match signal used in MUN v1. A memory that is lexically similar to the query is more likely — but not guaranteed — to be from the same class; text overlap cannot distinguish a same-class similar example from a different-class similar example, which is precisely the distinction that drives utility in classification settings.

At the *architectural level*, the feature-weighted linear model cannot capture non-linear interactions between features. The decision to use a linear model was evidence-based — no embedding feature exceeded the $|r| = 0.30$ threshold — but it forecloses the possibility that useful non-linear combinations existed in the feature space. The threshold may have been overly conservative, and a neural model with appropriate regularisation might have extracted additional signal.

At the *task level*, SST-2 binary classification has a particularly sharp utility boundary: a memory is either from the same class (high utility) or it is not. There is limited room for graded utility signal that label-free features could exploit. In more complex tasks — multi-class classification, question answering, code generation — the utility landscape may be smoother and richer, providing more signal for label-free feature engineering.

10.3 Lessons from the Feature Engineering Process

The Phase 13B feature discovery study yields several transferable insights. First, *surface-level lexical features can outperform sophisticated semantic embeddings* for narrow-domain tasks — text overlap (r = 0.716) far exceeded SentenceTransformer cosine similarity (r = 0.194). This is counterintuitive given the substantial investment the field has made in dense semantic representations, and suggests that domain specificity interacts with feature effectiveness in non-obvious ways.

Second, *the correlation between a feature and the utility label does not guarantee that optimally weighting that feature will improve downstream performance.* The Phase 13C.5 weight optimisation found a configuration achieving 60% on the uncontrolled evaluation; Phase 13C.9 found 35% under controlled conditions. The gap between these numbers — 25 percentage points — reflects the degree to which weight optimisation on a small uncontrolled test set overfits to noise.

Third, *information-theoretic features (information gain proxy, memory entropy) showed moderate utility signal but penalised performance when included alongside strong overlap features.* This suggests that these features capture a different aspect of the utility landscape — one that conflicts with the dominant lexical overlap signal at the scoring stage. Feature interaction effects in weighted linear models require careful ablation; correlation with the target does not imply compatibility with other selected features.

10.4 Implications for Memory Retrieval Systems

The MUN investigation establishes a clear empirical hierarchy: label-aware utility estimation substantially outperforms all heuristics; label-free utility estimation via the tested features does not exceed cosine similarity in this setting. This hierarchy has immediate practical implications. In deployment contexts where labels are available at retrieval time — supervised classification, active learning pipelines, human-in-the-loop annotation workflows — utility-based retrieval can provide substantial improvements over cosine similarity. In fully unsupervised or label-free settings, the current feature engineering approach is insufficient, and more powerful approaches (RL-based utility estimation, semi-supervised training) are needed.

10.5 Implications for Retrieval-Augmented AI

The MUN findings suggest that the RAG community's default of cosine similarity retrieval may be leaving substantial performance on the table in settings where task-specific utility signals are available. For classification tasks, the label boost result implies that hybrid retrieval strategies — combining semantic similarity with class-aware selection — could improve performance significantly. For generation tasks where ground-truth labels are unavailable, the MUN v2 negative result motivates continued research into label-free

utility surrogates, including RL-based approaches that learn utility from task reward signals.

10.6 Implications for Agent Memory Systems

Current agent memory systems (ReAct, Reflexion, MemGPT) rely on recency or similarity heuristics for memory retrieval. MUN v1 demonstrates that, in settings where a labelling signal is available (e.g., supervised task performance feedback), utility-based retrieval can dramatically outperform these heuristics. The MUN framework could be integrated as the retrieval component of an agent memory system, particularly in settings where the agent receives explicit task performance feedback that can serve as a utility training signal. The key challenge — identified but not resolved by MUN v2 — is extending this advantage to the unsupervised setting where most deployed agents operate.

11. Threats to Validity

11.1 Internal Validity

Threat	Mitigation	Rating
Data leakage	Phase 12D: Jaccard check, 0 query-pool overlaps confirmed	LOW
Label leakage	Audit confirms classifier never sees query label. Design choice, not bug.	LOW
Baseline selection	7 baselines spanning structural to semantic; none excluded post-run.	LOW
Small test set bias	Phase 13C.9 controlled; Phase 13C.8 incident documented in full.	HIGH (Phase 13C) MODERATE (pilots)
Implementation variance	Phase 13C.8 shows pool ordering and seed timing cause 45 pp swings at N=20.	HIGH for N=20
Confirmation bias	Negative result (MUN v2) documented without suppression.	LOW — documented

Table — Internal validity threats and mitigations.

11.2 External Validity

Threat	Evidence / Scope	Rating
Benchmark generality	SST-2 only: binary sentiment, short inputs (avg 19 tokens), single domain.	HIGH
Classifier generality	BART-large-MNLI only. Different classifiers may show different utility landscapes.	MODERATE
Pool generality	Pool fixed at N=1,000, seed=42. Pool size/composition effects unknown.	MODERATE
Deployment setting	In-context few-shot classification only. RAG, multi-hop, code generation untested.	HIGH
Language/domain	English movie reviews only. Cross-lingual or cross-domain generalisation unknown.	HIGH

Table — External validity threats. Results should be interpreted within SST-2 binary classification scope until multi-benchmark evaluation is conducted.

11.3 Statistical and Construct Validity

Threat	Detail	Rating
Utility operationalisation	Text overlap proxy (Phase 13B.5) may introduce systematic bias against diverse examples.	MODERATE
Multiple comparisons	128 weight configurations tested in Phase 13C.5; resolved by Phase 13C.9 independent controlled validation.	LOW (post-correction)
MUN v2 significance	No formal statistical test on N=20 final result. 30 pp deficit against similarity makes reversal unlikely but not impossible.	MODERATE
Effect size interpretation	Cohen's d ≈ 28.8 reflects 5-seed replication; may overstate effect for single-seed deployments.	LOW

12. Limitations

12.1 Benchmark Scope

All results are specific to SST-2 binary sentiment classification. SST-2's short inputs, two-class label space, and narrow domain may systematically favour text overlap as a utility predictor and disfavour richer semantic features. Multi-class settings, longer inputs, and multi-domain pools may produce different feature importance rankings.

12.2 Sample Size Constraints

MUN v2 final validation rests on N=20 test queries. The 95% CI for 35% accuracy spans approximately 13%–57%. While consistent across two weight configurations and supported by Phase 12E, this cannot definitively rule out a true accuracy of ≥50% for an optimally configured MUN v2 variant. N ≥ 200 required for definitive conclusions.

12.3 Utility Definition Scope

The text overlap utility proxy in Phase 13B.5 is a rough approximation of true utility. It may under-estimate diverse or syntactically distinctive examples and over-estimate highly overlapping ones. More rigorous utility estimation would require human annotation or multi-model evaluation.

12.4 Architecture Search

The decision to use a feature-based linear model for MUN v2 was based on the $|r| = 0.30$ threshold. This may have been overly conservative. A neural model with regularisation might extract useful non-linear combinations of embedding features.

12.5 Reproducibility Fragility

The Phase 13C.8 incident demonstrates susceptibility to implementation variance from pool ordering and random seed timing. All sources of randomness in the feature extraction pipeline must be seeded globally and documented explicitly in future work.

12.6 External Generalisability

Findings cannot be extrapolated to RAG systems, multi-hop reasoning, or code generation without additional evaluation. The MUN framework is validated only in the SST-2 in-context learning setting with BART-large-MNLI as the downstream classifier.

13. Future Work

13.1 Reinforcement Learning for Utility Estimation

The fundamental challenge of label-free utility estimation is the absence of a direct training signal. A principled alternative is to formulate memory retrieval as a sequential decision process: a retrieval agent observes query q, selects k memories, and receives reward equal to the downstream accuracy. Policy gradient methods (REINFORCE, PPO) could train this policy end-to-end using classification accuracy as a non-differentiable reward signal. This RL framing decouples utility estimation from label access: the agent learns to retrieve useful examples through trial-and-error interaction with the classifier. The key practical challenge is the combinatorial action space (C(N, k) subsets), requiring approximate methods such as greedy sequential selection with learned scoring functions.

13.2 Multi-Benchmark Evaluation

The most immediate extension is to replicate MUN v1 and v2 on at least two additional benchmarks (AGNews, TREC-6, NaturalQuestions) with $N \geq 500$ queries and pre-registered experimental designs. This would establish whether the utility advantage of MUN v1 generalises beyond SST-2, and whether the failure of label-free features is specific to binary classification or extends to more complex tasks. For RAG benchmarks (NaturalQuestions, HotpotQA, FEVER), a utility label framework based on answer-string match must first be developed.

13.3 Agent Memory Systems

Agent memory systems (ReAct, Reflexion, MemGPT, Voyager) rely on recency or similarity heuristics. Integrating MUN-style utility scoring as the retrieval component of an agent memory system — particularly in settings with explicit task performance feedback — could improve long-horizon performance. Multi-agent memory sharing introduces an additional dimension: cross-agent utility estimation (predicting whether agent A's memory is useful for agent B's task) is a novel problem combining transfer learning with retrieval quality estimation.

13.4 Long-Context Memory Management and Eviction

As context windows expand to 128K–1M tokens, the memory management problem shifts from retrieval to eviction. Utility-based eviction — removing the lowest-utility memories given the current query context — could improve long-horizon performance. The research challenge is that eviction utility (should be removed) and retrieval utility (should be retrieved from external storage) may be asymmetric.

13.5 Continual Learning and Memory Replay

In continual learning, memory replay stores past examples to prevent catastrophic forgetting. Defining utility as the expected reduction in forgetting when an example is included in the replay buffer connects MUN to the experience replay literature (Rolnick et al., 2019; Rebuffi et al., 2017). MUN-style utility scoring could provide a principled, learning-based alternative to heuristic buffer selection strategies based on recency or class balance.

13.6 Interpretable Utility Attribution

Applying formal attribution methods (Shapley values, integrated gradients, LIME) to the MUN v1 neural architecture would provide post-hoc explanations of retrieval decisions — particularly valuable for safety-critical applications (medical question answering, legal reasoning, scientific fact-checking) where the retrieval decision must be auditable and explainable to domain experts. The MUN v2 feature-based architecture offers a natural starting point: feature contributions are directly accessible via the weighted linear combination.

14. Research Retrospective

14.1 Original Hypothesis and What Was Learned

> **Original hypothesis:** The future usefulness of a memory, given a query context, can be predicted directly by a learned function — outperforming heuristic retrieval strategies. This hypothesis was **confirmed for label-aware settings** and **refuted for label-free settings** under the tested feature set.

When this investigation began, the working assumption was that semantic similarity in embedding space was a reasonable — if imperfect — proxy for utility, and that a learned model combining multiple signals (similarity, recency, label alignment) would outperform any individual signal. This was confirmed for MUN v1: the composite architecture substantially outperforms all seven baselines. What was *not* anticipated was the degree to which the label boost would dominate: 34 of 34 percentage points of advantage, leaving the text-based components statistically indistinguishable from baseline cosine similarity when the label is removed.

This finding was genuinely surprising and scientifically valuable. It implies that in binary classification settings, the label alignment signal so completely dominates the utility landscape that no amount of sophisticated feature engineering over text can recover it — at least within the feature space explored in Phase 13B. The failure of MUN v2 is therefore not a failure of the utility-based retrieval framework; it is a finding about the *structure* of the utility landscape in this specific task.

14.2 Unexpected Discoveries

Three unexpected findings emerged from this investigation. First, the **dominance of Jaccard text overlap over semantic embeddings** ($r = 0.716$ vs. $r = 0.194$) was not anticipated. Given the substantial progress in learned sentence representations, it was expected that dense embeddings would provide stronger utility signal than surface lexical overlap. The finding suggests that for narrow-domain single-task benchmarks, the semantic space of the embeddings conflates domain-specific and task-specific similarity in ways that reduce their predictive power for utility specifically.

Second, the **45-percentage-point reproducibility discrepancy** (Phase 13C.8) was not anticipated. The investigation had assumed that controlling the dataset, evaluation protocol, and test query set would be sufficient for reproducibility. The discovery that pool ordering — an implicit source of randomness in the uniqueness feature's computation — could produce such large swings on N=20 test sets was a methodological

lesson that applies well beyond this specific project.

Third, the **near-perfect pilot result** (99% accuracy in Pilot 1) was surprising even to the investigators. The label boost mechanism produces extremely clean utility signal in binary classification: same-class examples almost always help, different-class examples almost always hurt. This ceiling effect may itself be a property of the SST-2 benchmark, and may not generalise to tasks with more nuanced utility landscapes.

14.3 Why Documenting Failure Matters

The MUN v2 negative result is, in many ways, as scientifically valuable as the MUN v1 positive result. Without the MUN v2 investigation, researchers building on MUN v1 would face an important open question: is the label boost essential, or can label-free features recover comparable utility signal? The MUN v2 result provides a precise empirical answer: no, at least not with these features in this setting. This directs future research effort toward the RL-based and semi-supervised approaches outlined in Section 13, rather than toward further feature engineering in the same feature space.

The broader culture of machine learning research suffers from a well-documented publication bias toward positive results. Negative results — when rigorously conducted and fully documented — provide boundary conditions that are essential for scientific progress. The MUN v2 negative result was obtained through a multi-phase, pre-registered, audited experimental programme. It is not an incidental failure; it is a deliberately designed empirical investigation that reached a clear, reproducible conclusion.

14.4 Lessons for Future Researchers

- **Ablate before scaling.** The Phase 12E label boost ablation, conducted on the pilot evaluation set, revealed the fundamental dependency of MUN v1 on label information before any large-scale compute was invested in MUN v2. Systematic component analysis at small scale is a high-value investment.
- **Control all sources of randomness.** The Phase 13C.8 incident demonstrates that even seemingly irrelevant implementation details (pool ordering in a feature computation) can produce 45 pp accuracy swings at N=20. Document and fix all random state before reporting any performance estimate.
- **Pre-register decision criteria.** The MUN v2 decision criterion — MUN v2 must exceed the similarity baseline — was specified before Phase 13C.9 and applied without exception. This prevented post-hoc rationalisation of the negative result.
- **Distinguish utility from similarity early.** The central conceptual contribution of MUN — that utility-based retrieval is distinct from similarity-based retrieval — was well-motivated from the start. Grounding the research question in a formal definition

(Section 1.2) provided a stable reference throughout thirteen phases.

• **Treat the negative result as a finding, not a failure.** MUN v2's 35% accuracy is a finding about the label-free utility estimation problem in binary classification. Reporting it honestly, with full experimental provenance, makes it useful to the research community.

15. Conclusion

This report has documented a thirteen-phase empirical investigation of utility-based memory retrieval for few-shot in-context learning. The investigation began with the hypothesis that the future usefulness of a memory can be predicted directly, yielding retrieval quality superior to heuristic alternatives. It concluded with a nuanced and scientifically honest assessment of the conditions under which this hypothesis holds and the conditions under which it does not.

Memory Utility Networks v1 achieved strong, statistically robust retrieval performance — Recall@5 = 0.829, NDCG@5 = 0.914, $t(4) = 45.56$, $p < 10^{-10}$, Cohen's $d \approx 28.8$ — substantially outperforming seven baselines across five random seeds. A comprehensive multi-phase audit protocol confirmed the integrity of these results. The success is substantially attributable to the label boost mechanism, which contributes 34 percentage points of the advantage and directly motivated the MUN v2 direction.

Memory Utility Networks v2 did not succeed. Despite rigorous feature engineering (17 candidate features, 2,000 triplets), systematic embedding validation, 128-configuration weight optimisation, and controlled final validation, MUN v2.1 achieved 35% accuracy — below all baselines. This negative result was carefully documented and provides a precise empirical characterisation of the difficulty of label-free utility estimation in binary classification settings.

> "*Utility-based memory retrieval remains a promising but unresolved research challenge. MUN v1 establishes that the prize is real — label-aware utility estimation provides substantial, reproducible improvements over heuristic retrieval. MUN v2 establishes that recovering this advantage without labels is harder than it appears. The difficulty of the problem is part of its scientific value.*"

References

[1] Beltagy, I., Peters, M. E., & Cohan, A. (2020). Longformer: The long-document transformer. *arXiv:2004.05150*.

[2] Brown, T., et al. (2020). Language models are few-shot learners. *NeurIPS*, 33, 1877–1901.

[3] Graves, A., Wayne, G., & Danihelka, I. (2014). Neural Turing machines. *arXiv:1410.5401*.

[4] Graves, A., et al. (2016). Hybrid computing using a neural network with dynamic external memory. *Nature*, 538, 471–476.

[5] Guu, K., et al. (2020). REALM: Retrieval-augmented language model pre-training. *ICML 2020*.

[6] Karpukhin, V., et al. (2020). Dense passage retrieval for open-domain question answering. *EMNLP 2020*, 6769–6781.

[7] Lewis, M., et al. (2019). BART: Denoising sequence-to-sequence pre-training. *arXiv:1910.13461*.

[8] Lewis, P., et al. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *NeurIPS*, 33.

[9] Liu, J., et al. (2021). What makes good in-context examples for GPT-3? *DeeLIO @ NAACL 2021*.

[10] Liu, N. F., et al. (2023). Lost in the middle: How language models use long contexts. *arXiv:2307.03172*.

[11] Min, S., et al. (2022). Rethinking the role of demonstrations. *EMNLP 2022*.

[12] Packer, C., et al. (2023). MemGPT: Towards LLMs as operating systems. *arXiv:2310.08560*.

[13] Rebuffi, S. A., et al. (2017). iCaRL: Incremental classifier and representation learning. *CVPR 2017*.

[14] Rolnick, D., et al. (2019). Experience replay for continual learning. *NeurIPS*, 32.

[15] Rubin, O., Herzig, J., & Berant, J. (2022). Learning to retrieve prompts for in-context learning. *NAACL 2022*, 2655–2671.

[16] Shinn, N., et al. (2023). Reflexion: Language agents with verbal reinforcement learning. *NeurIPS*, 36.

[17] Su, H., et al. (2022). Selective annotation makes language models better few-shot learners. *arXiv:2209.01975*.

[18] Sukhbaatar, S., et al. (2015). End-to-end memory networks. *NeurIPS*, 28.

[19] Wang, A., et al. (2018). GLUE: A multi-task benchmark for NLU. *arXiv:1804.07461*.

[20] Wang, G., et al. (2023). Voyager: An open-ended embodied agent with large language models. *arXiv:2305.16291*.

[21] Weston, J., Chopra, S., & Bordes, A. (2014). Memory networks. *arXiv:1410.3916*.

[22] Yao, S., et al. (2022). ReAct: Synergizing reasoning and acting in language models. *arXiv:2210.03629*.

[23] Zaheer, M., et al. (2020). Big bird: Transformers for longer sequences. *NeurIPS*, 33.

[24] Zhao, T., et al. (2021). Calibrate before use: Improving few-shot performance of language models. *ICML 2021*.

Appendix A — Hyperparameters and Configurations

A.1 MUN v1

Parameter	Value	Notes
Encoder backbone	all-MiniLM-L6-v2	Frozen during MUN training
Embedding dimensionality	384	SentenceTransformer output
MLP architecture	512→256→128→1	4 layers, ReLU, Dropout(0.2)
Contrastive loss weight α	0.5	Grid-searched on validation set
BCE loss weight β	0.3	Grid-searched on validation set
Temporal decay rate λ	Learned (init 0.01)	Learnable parameter
Label boost magnitude	0.15	Applied if label(m)=label(q)
Top-k retrieval	5 / 10	Evaluated at both cutoffs
Evaluation seeds	42,123,456,789,999	5 independent seeds
Pool size	1,000	SST-2 training, seed=42
Pool composition	37.6% pos / 62.4%−	Representative of SST-2 train

A.2 MUN v2 Weight Configurations

Config	w_1 (overlap)	w_2 (info_gain)	w_3 (unique)	w_4 (divers.)	w_5 (entropy)	Notes
Base (Phase 13C)	0.40	−0.20	−0.15	0.15	0.10	Correlation-derived
Overlap-heavy	0.60	−0.10	−0.10	0.05	0.05	Best weight search
Equal weights (13C.9)	0.50	0.50	0.50	0.50	0.50	Controlled config A
Overlap-heavy (13C.9)	0.75	0.00	0.00	0.00	0.25	Controlled config B

Appendix B — Statistical Analysis

Metric	MUN v1 Mean	LRU Mean	t-statistic	p-value	Cohen's d
Recall@5	0.829	0.501	45.56	5.9×10^{-11}	28.82
NDCG@5	0.914	0.528	31.10	1.2×10^{-9}	19.67

Table B.1. MUN v1 vs. LRU baseline paired t-test (5 seeds).

Metric	Mean	SD	95% CI Lower	95% CI Upper
Recall@5	0.829	0.013	0.813	0.845
NDCG@5	0.914	0.017	0.893	0.935

Table B.2. MUN v1 confidence intervals (5 seeds, t-distribution).

Test	Statistic	Value	Interpretation
Kruskal-Wallis (Pilot 1)	H	74.10	$p < 10^{-6}$; significant group differences
MUN vs. Similarity (Pilot 1)	pp	+34.0	99% vs. 65%
Pilot 1 → Pilot 2 stability	pp	−1.0	99% → 98% (negligible)

Table B.3. Pilot study statistical summary.

Appendix C — Reproducibility Checklist

Dataset

- ☑ SST-2 training and validation splits via GLUE benchmark (Wang et al., 2018)
- ☑ Memory pool: 1,000 examples, SST-2 training set, seed=42
- ☑ Pool composition: 376 positive, 624 negative examples (verified)
- ☑ Evaluation queries: SST-2 validation set, no overlap with pool (Phase 12D)

Random State

- ☑ All global random seeds documented per phase (see phase_{N}/config.yaml)
- ☑ Pool ordering fixed prior to any feature extraction (no dynamic resampling)
- ☑ Feature extraction: all random pool samples seeded globally
- ☑ Phase 13C.9: full deterministic execution confirmed across all components

Model Checkpoints

- ☑ MUN v1 weights saved per seed (5 checkpoint files at models/checkpoints/)
- ☑ MUN v2 weight configurations documented in configs/mun_v2.yaml
- ☑ MUN v2.1 equal-weight and overlap-heavy variants both evaluated in Phase 13C.9

Evaluation Integrity

- ☑ Phase 12D: 7-question leakage audit — all 7 PASS
- ☑ No data overlap between query set and memory pool (Jaccard threshold 0.9)
- ☑ BART-large-MNLI classifier: inference only, no fine-tuning at any phase
- ☑ All metrics independently recomputed from raw prediction CSV files

Negative Result Documentation

- ☑ Phase 13C.8 discrepancy investigated; root cause partially identified
- ☑ Phase 13C.9 definitive controlled experiment conducted with fixed all-sources
- ☑ Pre-registered decision criterion documented before Phase 13C.9 execution
- ☑ Negative result formally concluded; no suppression or post-hoc rationalisation

Environment

- ☑ Python 3.x with PyTorch; SentenceTransformer all-MiniLM-L6-v2
- ☑ BART-large-MNLI via Hugging Face transformers library
- ☑ SST-2 via Hugging Face datasets (GLUE configuration)
- ☑ Full requirements.txt available at repository root

Appendix D — Repository Structure

```
memory-utility-networks/
├── README.md # Project overview, key results, citation
├── TECHNICAL_REPORT.pdf # This document
├── requirements.txt # Python environment specification
├── assets/ # Publication-quality figures
│ ├── fig1_research_timeline.png
│ ├── fig2_mun_v1_architecture.png
│ ├── fig3_training_pipeline.png
│ ├── fig4_baseline_comparison.png
│ ├── fig5_ablation_study.png
│ ├── fig6_feature_discovery.png
│ ├── fig7_embedding_validation.png
│ ├── fig8_mun_v2_architecture.png
│ ├── fig9_reproducibility_audit.png
│ └── fig10_findings_summary.png
├── models/
│ ├── utilitynet_v1.py # MUN v1 neural utility scorer
│ ├── utilitynet_v2.py # MUN v2 feature-based scorer
│ └── utility_score_v21.py # MUN v2.1 (weight variants)
├── configs/
│ └── mun_v2.yaml # Feature weight configurations
├── data/
│ ├── ablation_results.json # Per-seed ablation metrics
│ └── statistical_analysis.json # t-tests, p-values, Cohen's d
├── phase12/ phase12d/ phase12e/ # Pilot study, audit, ablation
├── phase13b5/ phase13b75/ # Feature discovery, embedding validation
├── phase13c/ phase13c5/ phase13c75/
├── phase13c8/ phase13c9/ # Reproducibility audit, final validation
└── docs/ # All 13 individual phase reports
```

Appendix E — Experimental Artifacts and BibTeX

E.1 Phase Audit Summary

Phase	Objective	Key Finding	Decision
12	Pilot study N=100	MUN 99% vs. Similarity 61%	Proceed to audit
12D	7-question validation audit	All checks pass; no leakage	Authorise full-scale
12E	Label boost ablation	Label boost = 34 pp; NoLabel = Similarity	Motivate MUN v2
13	MUN v2 design	Label-free architecture specified	Feature discovery
13B	Feasibility gate	Feature signal r=0.716	Feature discovery
13B.5	Feature discovery N=2,000	9 features with signal; overlap dominates	Select top-5
13B.75	Embedding validation	Best r=0.194 < 0.30 threshold	Feature-based model
13C	MUN v2 implementation	5-feature weighted scorer built	Weight optimisation
13C.5	Weight optimisation 128 cfg	Best: 60% (uncontrolled)	Controlled validation
13C.75	MUN v2.1 evaluation	30% — contradicts 13C.5	Trigger audit
13C.8	Reproducibility audit	Weight mismatch found; partial resolution	Controlled expt.
13C.9	Final controlled validation	MUN v2.1 = 35%; Similarity = 65%	NEGATIVE RESULT

** The 75% result from Phase 13C.5 feature ablation was not reproduced in Phase 13C.9 under controlled conditions. Definitive result: 35%.*

E.2 BibTeX

```
@misc{sree2026mun,
title={Memory Utility Networks: An Empirical Investigation of Utility-Based Memory Retrieval},
author={Sree Dharshan G J},
institution={SRM Institute of Science and Technology},
year={2026},
```

note={Technical Report — Revised Edition. Research phases 1--13C.9.}
}

@misc{brown2020gpt3,
title={Language Models are Few-Shot Learners},
author={Brown, Tom and others},
booktitle={NeurIPS},
year={2020}
}

@inproceedings{lewis2020rag,
title={Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks},
author={Lewis, Patrick and others},
booktitle={NeurIPS},
year={2020}
}

@inproceedings{rubin2022learning,
title={Learning To Retrieve Prompts for In-Context Learning},
author={Rubin, Ohad and Herzig, Jonathan and Berant, Jonathan},
booktitle={NAACL},
pages={2655--2671},
year={2022}
}

Appendix F — GitHub Package

F.1 README Abstract

Memory Utility Networks (MUN) is a research project investigating utility-based memory retrieval for few-shot in-context learning. Unlike cosine similarity and recency heuristics — which score memories based on historical patterns — MUN learns a *forward-looking utility estimator* that predicts which memory will most improve downstream classification accuracy when retrieved.

This repository documents a complete 13-phase empirical investigation: **MUN v1** (label-aware neural scorer) achieves Recall@5 = 0.829 and NDCG@5 = 0.914, outperforming 7 baselines by 31–35 percentage points ($p < 10^{-10}$, Cohen's d ≈ 28.8). **MUN v2** (label-free feature scorer) is a documented negative result: despite systematic feature engineering, it achieves 35% accuracy — below cosine similarity (65%) and random selection (40%). A 45 pp reproducibility discrepancy was investigated and documented as a methodological case study.

F.2 Key Contributions

```
- **MUN v1**: Neural utility scorer (Recall@5=0.829, NDCG@5=0.914, 5 seeds, p<10⁻¹⁰)
- **MUN v2**: Label-free feature scorer — documented negative result (35% accuracy)
- **Audit Protocol**: 7-question validation audit confirming no data/label leakage
- **Ablation Studies**: Component importance quantification (label boost = 34 pp)
- **Feature Discovery**: 17 candidate features on 2,000 triplets; text overlap dominates (r=0.716)
- **Reproducibility Case Study**: 45 pp discrepancy investigated and partially resolved
- **13 Phase Reports**: Complete experimental provenance available in docs/
```

F.3 Key Findings

```
| Finding | Result |
|---------|--------|
| MUN v1 vs. best baseline (Similarity) | +0.314 Recall@5 (+60.9% relative) |
| Label boost contribution (Phase 12E) | 34 pp of 34 pp total advantage |
| Best label-free feature (Jaccard) | Pearson r = 0.716 with utility |
| Best embedding feature | r = 0.194 (below neural threshold) |
| MUN v2.1 final accuracy | 35% (Similarity: 65%, Random: 40%) |
```

| Phase 13C.8 reproducibility gap | 45 pp at N=20 test queries |

F.4 Citation

```
@misc{sree2026mun,
  title={Memory Utility Networks: An Empirical Investigation of
        Utility-Based Memory Retrieval},
  author={Sree Dharshan G J},
  institution={SRM Institute of Science and Technology},
  year={2026},
  note={Technical Report — Revised Edition}
}
```

F.5 Repository Highlights

File / Directory	Description
models/utilitynet_v1.py	MUN v1 neural utility scorer with all components
models/utilitynet_v2.py	MUN v2 label-free feature-based scorer
configs/mun_v2.yaml	Feature weight configurations for all experiments
data/ablation_results.json	Per-seed ablation metrics for all 7 configurations
data/statistical_analysis.json	t-tests, p-values, confidence intervals
phase13c9/	Final controlled validation — definitive negative result
docs/	13 individual phase reports with full provenance
TECHNICAL_REPORT.pdf	This document — complete 50-page research report