

Communication

Not peer-reviewed version

---

# Reverse Prompting: A Memory-Efficient Paradigm for LLM Agents Through Structural Regeneration

---

[Ünver Çiftçi](#)\*

Posted Date: 20 August 2025

doi: 10.20944/preprints202508.1424.v1

Keywords: large language models; prompt compression; memory



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Communication

# Reverse Prompting: A Memory-Efficient Paradigm for LLM Agents Through Structural Regeneration

Ünver Çiftçi

Math & AI Institute; unverciftci@gmail.com

## Abstract

Large language model (LLM) agents face a fundamental challenge: how to maintain memory and context across extended interactions when constrained by limited context windows. Current approaches rely on external storage systems, vector databases, or retrieval mechanisms that are often complex and opaque. We introduce *reverse prompting*, a simple alternative where agents store compact, human-readable recipes instead of raw outputs. These recipes capture the essential structure and intent of generated content—such as code, documents, or plans—and can be used to regenerate functionally equivalent outputs when needed. For example, a 50-line machine learning script can be compressed into a 15-line recipe specifying the model architecture, dataset, and training parameters. When fed back to an LLM, this recipe produces new code that accomplishes the same task. We demonstrate the concept with practical examples across different domains and discuss its potential applications, limitations, and advantages over existing memory mechanisms. Reverse prompting offers a transparent, model-agnostic approach to agent memory that maintains human interpretability while enabling efficient storage and cross-session continuity.

**Keywords:** large language models; prompt compression; memory

## 1. Introduction

Large language models are increasingly being deployed as autonomous agents that need to maintain context and memory across extended interactions [1,2]. Whether helping with software development, writing documents, or planning complex tasks, these agents often need to recall and build upon their previous work. However, current LLMs face a fundamental constraint: they can only process a limited amount of text at once, and they cannot inherently remember what happened in previous conversations or sessions.

This memory limitation creates practical problems. A coding agent that writes a function in one session cannot easily recall that code in the next session. A research assistant that drafts an outline cannot seamlessly continue writing the full document later. Current solutions typically involve storing conversation histories in external databases, using retrieval-augmented generation (RAG) systems [3], or employing vector embeddings to retrieve relevant past content [4], but these approaches are often complex, opaque, and unreliable.

We propose a different approach called **reverse prompting**. Instead of storing the actual outputs that an LLM generates (like complete code files or full documents), we extract and store compact *recipes*—structured, human-readable instructions that can be used to regenerate the original content when needed. Think of it as creating a cooking recipe from a finished meal: the recipe is much shorter than describing every ingredient and step in detail, but it contains enough information to recreate the dish.

For example, instead of storing a 50-line machine learning training script, an agent might store a 15-line recipe that specifies the model architecture, dataset, hyperparameters, and training procedure. When the agent needs that code again, it can feed the recipe to an LLM (itself or another model) and get back functionally equivalent code.

This approach offers several potential advantages: it's memory-efficient, human-readable, works across different LLM models, and maintains the essential functionality while allowing for natural variations in implementation details. The purpose of this paper is to introduce the concept, demonstrate its viability with concrete examples, and discuss where it might be most useful for LLM agents.

## 2. What is Reverse Prompting?

Reverse prompting is the process of taking an LLM-generated output and distilling it into a structured, compact recipe that can be used to regenerate similar content. The key insight is that many LLM outputs—especially structured ones like code, formatted documents, or step-by-step plans—contain patterns and essential elements that can be captured more efficiently than storing the full text.

A reverse prompt (or recipe) has several important characteristics:

- **Compact:** The recipe is significantly shorter than the original output
- **Structured:** Information is organized in a clear, systematic format
- **Human-readable:** Anyone can understand what the recipe specifies
- **Regenerative:** When used as a prompt, it produces functionally equivalent output
- **Model-agnostic:** The same recipe can work with different LLMs

The process works in two directions. In the forward direction, an LLM generates content from a prompt. In the reverse direction, we analyze that generated content and extract a recipe that captures its essential structure and requirements. This recipe then becomes a new prompt that can regenerate similar content.

Unlike traditional compression or summarization, reverse prompting aims for *functional equivalence* rather than preserving every detail. A recipe for code should produce code that works the same way, even if variable names or comments differ. A recipe for a document should generate text that conveys the same information and follows the same structure, even if the exact wording varies.

This approach is particularly well-suited for outputs that have clear structural patterns: software code, technical documents, procedural instructions, configuration files, and similar content where the underlying logic and organization matter more than specific stylistic choices.

## 3. A Detailed Example: Python Code

To illustrate reverse prompting concretely, let's walk through a complete example with Python code for training a neural network.

### 3.1. Original Generated Code

Suppose an LLM generates the following 38-line Python script for training an MNIST classifier:

```
import torch
import torch.nn as nn
import torch.optim as optim
from torchvision import datasets, transforms

class Net(nn.Module):
    def __init__(self):
        super(Net, self).__init__()
        self.fc1 = nn.Linear(784, 128)
        self.relu = nn.ReLU()
        self.fc2 = nn.Linear(128, 10)

    def forward(self, x):
        x = x.view(-1, 784)
        x = self.relu(self.fc1(x))
```

```
        return self.fc2(x)

train_loader = torch.utils.data.DataLoader(
    datasets.MNIST('.', train=True, download=True,
                  transform=transforms.ToTensor()),
    batch_size=64, shuffle=True
)

model = Net()
criterion = nn.CrossEntropyLoss()
optimizer = optim.SGD(model.parameters(), lr=0.01)

for epoch in range(5):
    for data, target in train_loader:
        optimizer.zero_grad()
        output = model(data)
        loss = criterion(output, target)
        loss.backward()
        optimizer.step()

print("Training finished.")
```

### 3.2. Extracted Recipe

Instead of storing these 38 lines, we extract the following recipe:

Task: Train neural network for MNIST digit classification  
Framework: PyTorch

#### Model Architecture:

- Input: 784 features (28x28 flattened)
- Hidden layer: 128 units with ReLU activation
- Output: 10 classes (digits 0-9)
- Architecture type: Simple feedforward network

#### Dataset:

- Name: MNIST
- Split: Training set
- Preprocessing: Convert to tensor
- Batch size: 64
- Shuffle: Yes

#### Training Configuration:

- Epochs: 5
- Loss function: Cross-entropy
- Optimizer: SGD with learning rate 0.01
- Training loop: Standard PyTorch pattern

Output: Print "Training finished." when complete

Requirements: Generate complete, runnable PyTorch code

### 3.3. Regeneration

When this recipe is provided to an LLM with a prompt like "Generate Python code that implements this specification," the model produces new code that:

- Implements the same neural network architecture
- Uses the same dataset and preprocessing
- Applies the same training configuration
- Produces functionally identical results

The regenerated code might have different variable names, slightly different organization, or additional comments, but it accomplishes exactly the same task. This demonstrates the key principle: preserving function while allowing stylistic variation.

### 3.4. Compression and Benefits

In this example:

- **Original:** 38 lines of code
- **Recipe:** 22 lines of structured specification
- **Compression ratio:** About 42% reduction in size
- **Functional preservation:** 100% - the regenerated code works identically

The recipe is also more readable for humans who want to understand what the code does, and it can be easily modified (e.g., changing the learning rate or number of epochs) without rewriting the entire program.

## 4. Applications and Use Cases

Reverse prompting can be applied across various domains where LLM agents generate structured content. Here are some promising application areas:

### 4.1. Software Development

Programming tasks are natural candidates for reverse prompting because code has clear structure and well-defined functionality.

**Function libraries:** Instead of storing complete function implementations, agents can maintain recipes that specify the algorithm, input/output types, and key logic. This allows for code regeneration in different programming languages or with updated syntax.

**Configuration files:** Database schemas, API configurations, and deployment scripts can be captured as recipes that specify the essential settings and relationships, making them easier to modify and adapt to new environments.

**Test suites:** Testing patterns and scenarios can be stored as recipes that specify the test cases, expected behaviors, and edge conditions, allowing regeneration of tests for different codebases.

### 4.2. Document Generation

Technical and structured documents often follow predictable patterns that recipes can capture effectively.

**Technical reports:** Research papers, project summaries, and analysis documents can be reduced to recipes specifying the structure, key findings, methodology, and conclusions. This enables consistent formatting while allowing content adaptation.

**Documentation:** API documentation, user guides, and tutorials can be stored as recipes that capture the information hierarchy, examples, and key concepts, making them easier to update as systems evolve.

**Communication templates:** Email templates, meeting agendas, and project proposals can be abstracted into recipes that maintain professional structure while allowing customization for specific contexts.

#### 4.3. Planning and Workflows

Multi-step processes and planning tasks benefit from recipe-based representation.

**Project plans:** Complex projects can be broken down into recipes that specify milestones, dependencies, resource requirements, and success criteria, enabling plan adaptation for similar future projects.

**Troubleshooting procedures:** Diagnostic workflows and problem-solving steps can be captured as recipes that maintain the logical flow while allowing adaptation to specific technical contexts.

**Educational content:** Lesson plans, course outlines, and learning exercises can be stored as recipes that preserve pedagogical structure while enabling content updates and customization.

#### 4.4. Cross-Domain Benefits

Reverse prompting offers several advantages that apply across these application areas:

**Version control:** Recipes can be easily tracked, compared, and merged, making it simpler to manage evolving content and collaborate across teams.

**Customization:** Recipes can be modified for specific needs without starting from scratch, enabling rapid adaptation of proven patterns to new contexts.

**Knowledge transfer:** Recipes capture institutional knowledge in a readable format that can be shared, taught, and improved over time.

**Quality consistency:** By standardizing the essential elements while allowing implementation flexibility, recipes help maintain quality standards across different outputs.

### 5. Limitations and Challenges

While reverse prompting shows promise, it's important to acknowledge its limitations and the challenges that need to be addressed.

#### 5.1. What Doesn't Work Well

**Highly creative content:** Poetry, creative writing, marketing copy, and other content where style, tone, and specific word choices matter cannot be effectively captured in recipes. The regenerated content may lose the original's creative essence or emotional impact.

**Context-dependent outputs:** Content that relies heavily on external context, recent events, or specific cultural references is difficult to compress into self-contained recipes. The regenerated version may miss important nuances or become outdated.

**Highly irregular or random content:** Outputs with little underlying structure or pattern resist recipe extraction. Data dumps, random lists, or content with no clear organizational logic cannot be meaningfully compressed.

**Precision-critical content:** Legal documents, medical prescriptions, financial calculations, or other content where exact wording and precise details matter cannot tolerate the functional equivalence approach of reverse prompting.

#### 5.2. Technical Challenges

**Recipe extraction:** Currently, creating good recipes requires manual effort and domain expertise. Automating this process while maintaining quality is a significant challenge that needs research attention.

**Quality control:** Determining when a regenerated output is "good enough" compared to the original is non-trivial. Different use cases may require different fidelity standards.

**Model variability:** Different LLMs may interpret the same recipe differently, potentially leading to inconsistent results. Ensuring reliable cross-model performance requires careful recipe design.

**Recipe maintenance:** As underlying technologies, libraries, or standards evolve, recipes may become outdated and need updating to remain effective.

### 5.3. Practical Considerations

**Learning curve:** Users need to understand how to write effective recipes and recognize what content is suitable for this approach. This requires training and experience.

**Domain specificity:** Effective recipe formats may vary significantly across domains, requiring specialized knowledge and potentially different tools for different applications.

**Storage trade-offs:** While recipes are more compact than full content, they still require storage and management. For very simple or short outputs, the overhead may not be worthwhile.

### 5.4. Open Questions

Several important questions remain unanswered and require further research:

- How can we automate recipe extraction while maintaining quality?
- What are the optimal recipe formats for different domains?
- How do we measure and ensure adequate fidelity for different use cases?
- Can recipes be made more robust to model changes and updates?
- How does this approach scale to very large or complex outputs?

These limitations don't invalidate the concept, but they do suggest that reverse prompting is best viewed as a complementary technique rather than a universal solution to LLM memory challenges.

## 6. Future Directions and Research Opportunities

Reverse prompting opens several promising avenues for future research and development. Here are key areas that warrant investigation:

### 6.1. Automated Recipe Extraction

The most immediate need is developing systems that can automatically extract recipes from generated content without manual intervention.

**Pattern recognition:** Machine learning approaches could identify common structural patterns in different domains and learn to extract the essential elements automatically. This might involve training specialized models on examples of content-recipe pairs.

**Interactive tools:** Semi-automated systems could guide users through recipe creation by suggesting key elements to capture and providing templates for different content types.

**Domain-specific extractors:** Specialized tools for different domains (code, documents, plans) could leverage domain knowledge to create more effective recipes.

### 6.2. Recipe Standards and Formats

Establishing standard formats for recipes across different domains would improve consistency and tool interoperability.

**Schema development:** Creating standardized schemas for common content types would make recipes more portable across different systems and models.

**Recipe languages:** Developing domain-specific languages or markup formats for recipes could provide more structure while maintaining human readability.

**Validation frameworks:** Tools for testing and validating recipe quality could help ensure that recipes produce acceptable regenerations.

### 6.3. Integration with Agent Architectures

Exploring how reverse prompting fits into existing and future LLM agent systems presents important research questions.

**Hybrid memory systems:** Combining reverse prompting with other memory mechanisms (RAG, vector databases, conversation histories) could provide comprehensive memory solutions that leverage the strengths of each approach.

**Dynamic recipe management:** Systems that automatically decide when to create recipes, how to organize them, and when to update or retire them could make the approach more practical for real-world deployment.

**Multi-agent scenarios:** Investigating how agents can share and adapt recipes created by other agents could enable collaborative knowledge building.

#### 6.4. Evaluation and Metrics

Developing better ways to measure the effectiveness of reverse prompting is crucial for progress.

**Fidelity metrics:** Creating reliable measures of how well regenerated content preserves the essential qualities of the original across different domains.

**Efficiency benchmarks:** Establishing standards for comparing the memory efficiency and retrieval speed of reverse prompting against other memory mechanisms.

**User studies:** Understanding how humans interact with recipe-based systems and what factors contribute to successful adoption.

#### 6.5. Advanced Applications

Several more sophisticated applications could emerge as the field develops:

**Recipe evolution:** Systems that can modify and improve recipes over time based on usage patterns and feedback.

**Cross-domain transfer:** Investigating whether recipes from one domain can be adapted for use in related domains.

**Compositional recipes:** Developing ways to combine and modify existing recipes to create new ones, enabling more flexible content generation.

**Personalization:** Adapting recipes to individual user preferences and working styles while maintaining their core functionality.

These research directions could transform reverse prompting from an interesting concept into a practical, widely-used approach for LLM agent memory management.

## 7. Conclusion

We have introduced reverse prompting as a simple yet potentially powerful approach to LLM agent memory management. By storing compact, human-readable recipes instead of full outputs, this method offers an alternative to complex retrieval systems and external databases that currently dominate the field.

The core insight is straightforward: many LLM outputs contain structural patterns that can be captured more efficiently than storing the complete text. Our detailed example with Python code demonstrates that functional equivalence can be preserved while achieving significant compression, and our survey of applications shows potential value across diverse domains.

However, reverse prompting is not a universal solution. It works best for structured, pattern-rich content where functional equivalence matters more than exact reproduction. Creative content, precision-critical outputs, and highly irregular text remain challenging for this approach. Understanding these boundaries is crucial for effective application.

The most significant barrier to adoption is the current need for manual recipe extraction. Until this process can be automated reliably, reverse prompting will require human expertise and effort that may limit its practical deployment. This represents the most important near-term research challenge.

Despite these limitations, we believe reverse prompting offers unique advantages that make it worth pursuing. Its transparency and human readability distinguish it from opaque vector embedding approaches. Its model-agnostic nature provides flexibility as LLM technology evolves. Its emphasis on structural understanding rather than verbatim storage aligns well with how humans naturally organize and recall information.

The path forward involves developing automated extraction tools, establishing standard recipe formats, and integrating this approach with existing agent architectures. As these technical challenges

are addressed, reverse prompting could become a valuable component of the toolkit for building more capable and reliable LLM agents.

We encourage the research community to explore this direction, develop better tools and methods, and identify new applications where structured regeneration offers advantages over traditional memory approaches. The concept is simple enough to be broadly accessible yet rich enough to support substantial research and development efforts.

Ultimately, reverse prompting represents one possible solution to the fundamental challenge of giving LLM agents better memory. Whether it proves broadly useful will depend on continued research, practical experimentation, and honest assessment of its capabilities and limitations in real-world applications.

**Acknowledgments:** We thank Eren Ünlü for reading the first draft and providing valuable feedback that helped improve this work.

## References

1. Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K.; Yao, S. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems* **2023**, *36*, 8634–8652.
2. Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.; Cao, Y.; Narasimhan, K. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems* **2023**, *36*, 11809–11822.
3. Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.t.; Rocktäschel, T.; et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems* **2020**, *33*, 9459–9474.
4. Borgeaud, S.; Mensch, A.; Hoffmann, J.; Cai, T.; Rutherford, E.; Millican, K.; Van Den Driessche, G.B.; Lespiau, J.B.; Damoc, B.; Clark, A.; et al. Improving language models by retrieving from trillions of tokens. In Proceedings of the International conference on machine learning. PMLR, 2022, pp. 2206–2240.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.