

Markdown vs YAML for LLMs?

Author: Kevin Anderson

Date: 03/14/2025

Topic: A Whitepaper on Readability, Hierarchy, Configuration Semantics, and Workflow Fit in Human-Model Systems

Abstract

As large language models move from casual interaction into production workflows, the structure of prompts and outputs has become an engineering variable rather than a stylistic preference. Among lightweight, human-readable formats, Markdown and YAML occupy an especially interesting middle ground. Both are more structured than plain text and more editable than rigid machine-first formats, yet they embody different design assumptions.

Markdown is optimized for readable documents, semantic hierarchy, and mixed-content communication. YAML is optimized for configuration-style data, key-value organization, and nested declarative structure.

Recent research shows that prompt formatting can materially change large language model performance across reasoning, code generation, and translation tasks, including when the underlying semantic request remains constant. A 2024 study comparing plain text, Markdown, JSON, and YAML found significant performance differences across models and tasks, demonstrating that **format is not cosmetic**.

This paper argues that Markdown is generally the stronger default for human-LLM collaboration, long-form prompting, documentation, synthesis, and mixed-content workflows, while YAML is often better suited to structured configuration-like prompts, parameterized templates, and tasks where nested fields or declarative organization matter more than narrative readability.

- ❏ The strongest academically defensible conclusion is not that one format universally dominates, but that each format optimizes a different kind of structure. Benchmark and vendor guidance both support this broader view that structure affects model behavior and that lightweight formatting choices can improve prompt boundary clarity and reliability.

The central thesis is therefore practical: **Markdown is the better workspace format for most human-facing LLM tasks** because it balances semantic structure, readability, and portability, while **YAML is the better lightweight format** when the prompt itself behaves more like a configuration object than a document. In production, the real choice is often not Markdown or YAML in isolation, but which one better matches the cognitive shape of the task.

1. Introduction

The rise of LLMs has made structured natural language a first-class interface. Prompts are no longer just questions. They are often composite artifacts containing instructions, examples, retrieved context, formatting rules, output expectations, and operational constraints. As prompt complexity grows, developers and content teams increasingly look for formats that improve clarity without imposing excessive rigidity. That is where Markdown and YAML often enter the conversation.

Markdown

Became widespread in LLM ecosystems because it fits naturally with documents, knowledge bases, notebooks, developer tools, and publishing pipelines. Supports headings, lists, code fences, block quotes, links, and tables with minimal syntax.

YAML

Emerged from configuration and serialization workflows. Its strengths are nesting, keys, lists, declarative structure, and concise representation of parameterized information. In an LLM context, YAML tends to present ideas as structured fields.

This comparison matters because real prompts increasingly sit somewhere between prose and data. Choosing between Markdown and YAML is therefore not just a formatting preference. It is a decision about how the model, the author, and the surrounding toolchain will interpret the task.

2. Why This Comparison Matters

Among lightweight formats, Markdown and YAML are easy to confuse because both are human-readable and both add structure without requiring a full schema language. But the similarity is superficial. Their strengths diverge quickly in practice.

Empirical Evidence

The 2024 study *"Does Prompt Formatting Have Any Impact on LLM Performance?"* evaluated prompt templates across plain text, Markdown, JSON, and YAML and found that formatting could significantly affect outcomes across reasoning, code generation, and translation tasks. Larger models were more robust than smaller ones, but not invariant. This makes format choice an **experimental variable, not an aesthetic one**.

Vendor Guidance

OpenAI's prompt engineering guide explicitly recommends using Markdown formatting and XML tags to help models understand logical boundaries in prompts and context data. That guidance does not single out YAML as the default for prompt authoring — suggesting that document-like clarity and visible semantic sections matter in many production prompting scenarios.

Systems Perspective

YAML remains highly useful when the structure of the task is naturally parameterized or configuration-like rather than document-like. The comparison is therefore not merely academic — it has direct implications for how teams design, maintain, and scale their LLM workflows.

3. Two Different Structural Philosophies

Markdown and YAML encode different ways of thinking about information organization.

Markdown — Semantic Document Format

Markdown organizes material into readable sections and sub-sections. A heading implies topical scope. A list implies grouped requirements. A fenced block marks literal code or content. A quote marks cited or source text. Its structure is intentionally visible to humans at a glance.

This makes Markdown particularly well suited for prompts that resemble **instructions, briefs, plans, analyses, or reports**.

YAML — Declarative Data Format


YAML expresses information through keys, nested mappings, lists, and indentation-sensitive hierarchy. Even when it is easy for humans to read, it is still conceptually closer to configuration than to prose.

A YAML prompt suggests that the task can be broken into fields such as **role, objective, constraints, examples, and output format** — powerful for templating and parameter-driven workflows.

- ❏ This distinction matters because LLM workflows differ in kind. Problems arise when a team uses one format to solve the other's job. Markdown and YAML are each good at one side of that divide.

4. What the Research Says About Format Sensitivity

The 2024 formatting study provides the clearest direct support for taking YAML seriously in this discussion. The authors compared semantically equivalent prompts formatted as plain text, Markdown, JSON, and YAML across six benchmarks and multiple GPT-family models.



Illustrative performance variation by format and task type (relative scores, based on findings from He et al., 2024). The main takeaway is not that YAML always beats Markdown or vice versa, but that **there is no safe assumption of format neutrality**.

4. Research Signals (continued)

Beyond the core formatting study, two additional benchmarks provide important supporting evidence.

1 FoFo Benchmark (2024)

Studies LLMs' ability to follow complex real-world formats and argues that format-following is a distinct and underexamined capability for agentic and production use cases. Although FoFo is broader than Markdown versus YAML specifically, it supports the same methodological point: **structured format adherence is a meaningful capability, not just a display choice.**

2 MDEval Benchmark (2024)

Treats Markdown Awareness as a measurable quality dimension tied to readability and output structure. The benchmark's existence is important because it formalizes something many practitioners intuitively observe: when outputs are intended for human consumption, structural readability matters. MDEval does not make YAML irrelevant; it simply underscores that **Markdown's strengths are especially visible in human-facing generation.**

The literature therefore points toward a layered conclusion. YAML belongs in the format-comparison conversation because format affects performance. Markdown remains especially important because readability and human navigation are themselves benchmarkable qualities. Both facts are true at once.

5. Where Markdown Wins



Readability at Scale

A long Markdown prompt still looks like a document. A long YAML prompt can quickly begin to feel like configuration. Markdown handles transitions through headings, bullets, and fenced sections naturally — YAML can represent the same information, but often with more cognitive friction.



Collaborative Medium

Markdown is easier for multiple humans to scan, annotate, review, and revise. Many production prompts are edited by product managers, researchers, marketers, developers, and reviewers over time. A format that preserves legibility under collaboration reduces authoring friction and review cost.



Mixed-Modality Expression

One prompt can contain prose, lists, quoted source material, code blocks, and sectioned instructions without breaking the rhythm of reading. YAML is less natural for that because it pushes every element into a field-value structure. Markdown is superior for whitepapers, blog-generation systems, research synthesis, editorial workflows, RAG prompts, and internal documentation agents.



Downstream Human Usability

Outputs formatted in Markdown are often ready for review, publication, or handoff with minimal transformation. Since Markdown is already native to many documentation and content systems, it fits smoothly into broader workflows. YAML rarely has this advantage because it is usually an intermediate representation rather than a publishable artifact.

6. Where YAML Wins

Declarative Clarity

YAML's strongest advantage over Markdown is declarative clarity when the prompt naturally maps to fields. If a task can be described as role, objective, constraints, inputs, examples, and output schema, YAML can make that organization explicit in a concise way. That can be useful for reusable prompt templates, configuration-driven systems, and agent frameworks where prompts are assembled programmatically.

Parameterization

In workflows where different variables are injected into the same prompt template, YAML can behave like a clean manifest. This is especially useful when the model should reason over explicitly named fields rather than infer structure from prose sections — for example, a task contract with nested options, toggles, priorities, or multiple input blocks.


Compact Structured Nesting

Compared with JSON, YAML is often easier for humans to edit when nested structure is required. That makes it attractive in systems that need more hierarchy than Markdown but do not require hard schema enforcement. StructEval (2025) reinforces the growing importance of structured format generation and evaluation, including YAML among non-renderable structured formats that matter in software workflows.

YAML can therefore win in situations where the prompt behaves less like an essay and more like a configuration artifact. In those cases, Markdown's readability advantage may matter less than YAML's tighter field organization.

7. Speed: Which One Is Faster?

The speed question is best broken into **model speed** and **workflow speed**. There is no strong basis for claiming that Markdown universally generates faster than YAML or vice versa in raw latency terms.



At the workflow level, different patterns emerge. Markdown often wins on drafting and review speed for document-like prompts because humans can compose and inspect it more quickly. YAML can win on system assembly speed when prompts are stored or manipulated as structured task objects, especially in config-driven or programmatic pipelines.

The safe practical takeaway is conditional. Markdown usually optimizes human iteration speed. YAML often optimizes structured task packaging speed. Teams should benchmark both on their actual workloads instead of assuming one is categorically faster.

8. Quality: Which One Produces Better Results?

Quality in this comparison is multi-dimensional. It can refer to semantic accuracy, readability, consistency, structural adherence, review cost, or downstream utility. Markdown and YAML optimize different parts of this space.

Markdown → Human-Facing Quality

When outputs must be read, evaluated, revised, or published by people, Markdown's structural cues improve navigability and presentation. MDEval supports treating this as a measurable capability rather than a subjective preference. A more readable answer is often a higher-quality answer in practical knowledge-work settings because it is easier to trust, inspect, and reuse.

YAML → Configuration-Facing Quality

When correctness depends on the clean representation of fields, nested options, or declarative constraints, YAML can be the better medium. This is especially true when the prompt itself is being treated as semi-structured data rather than prose. StructEval's attention to structured formats reflects the growing importance of this dimension in LLM-integrated software systems.

- The deeper lesson is that Markdown and YAML do not compete on one universal axis of "better." Markdown usually produces better document quality. YAML often produces better manifest quality. The right question is which form of quality the application rewards.

9. Recommended Evaluation Framework for Teams

Teams comparing Markdown and YAML should evaluate them against the task families they actually deploy. A strong benchmark should include long-form explanation, retrieval synthesis, reusable prompt templates, structured task manifests, and agent workflows that combine human review with configuration-driven execution.

Dimension	Measure for Speed	Measure for Quality
Authoring	Drafting time, revision time	Readability, collaboration ease
Model Processing	Time to first token, total latency	Semantic correctness, consistency
System Integration	Maintenance time, assembly speed	Structural adherence, integration success

If a team publishes analyst-facing reports, Markdown should be heavily favored in the evaluation. If a team manages parameterized prompt libraries or config-like task definitions, YAML should receive stronger consideration.

The strongest benchmark hypothesis is not that Markdown beats YAML or YAML beats Markdown across the board. It is that Markdown is the better default for document-shaped human-LLM work, while YAML is the better default for configuration-shaped human-LLM work.

10. Best-Practice Architecture & Limitations

In mature LLM systems, Markdown and YAML often work best together.

Markdown Layer

Serves as the workspace for prompts that need explanation, context, examples, and human review. Humans usually think through complex ideas in documents.

YAML Layer

Serves as the manifest layer for parameterized task definitions, prompt metadata, agent configs, and reusable runtime structures. Systems manage repeatable parameter sets through structured fields.

This hybrid architecture aligns each format with the kind of structure it handles best. Forcing YAML to behave like a full narrative document can make prompts brittle and hard to maintain. Forcing Markdown to behave like a clean nested config can make orchestration inconsistent.

Limitations

- The public literature directly isolating Markdown versus YAML remains limited; the strongest evidence comes from broader prompt-format sensitivity studies rather than a dedicated canonical head-to-head benchmark.
- Model behavior changes over time. A formatting preference observed in one generation of models may shift as models improve their robustness to structure variation. The conclusions in this paper should therefore be treated as **current best evidence rather than permanent law**.
- Workload context matters. Human-facing knowledge work and system-facing structured orchestration reward different properties. Any blanket claim that one format is always better without naming the task would be methodologically weak.

11. Conclusion & References

Markdown and YAML are both valuable because they solve the same high-level problem differently: they add structure to natural-language interaction with LLMs. But they do not add the same kind of structure. Markdown organizes meaning as a document. YAML organizes meaning as a declarative object. That difference is what should drive the decision.

When Markdown is the stronger default

When the main goals are readability, collaboration, long-form prompting, and mixed-content outputs.

When YAML is the stronger lightweight alternative

When the main goals are parameterization, nested fields, configuration-like prompts, and manifest-style orchestration.

The most robust production answer

Not Markdown or YAML in isolation — but Markdown and YAML, used at different layers of the same system. Markdown should structure documents. YAML should structure manifests.

References

1. He, J., Rungta, M., Koleczek, D., Sekhon, A., Wang, F. X., and Hasan, S. *Does Prompt Formatting Have Any Impact on LLM Performance?* arXiv:2411.10541.
2. Chen, Z., Liu, Y., Shi, L., Wang, Z.-J., Chen, X., Zhao, Y., and Ren, F. *MDEval: Evaluating and Enhancing Markdown Awareness in Large Language Models.* arXiv:2501.15000.
3. OpenAI. *Prompt engineering.* Developer guide discussing logical boundaries with Markdown and XML.
4. *FoFo: A Benchmark to Evaluate LLMs' Format-Following Capability.* arXiv:2402.18667.
5. *StructEval: Benchmarking LLMs' Capabilities to Generate Structured Formats.* arXiv:2505.20139.