Advancing Large Language Model Reasoning Techniques: Methods Enabling LLMs to 'Think' Beyond Text Generation for Reliable and Explainable AI

Cassel Scott-Emuakpor

Soft Tech Group, cassel.scott@softtechgroup.io

30th August 2025

Abstract

Large Language Models (LLMs) have revolutionized artificial intelligence applications, extending from writing assistants to Retrieval-Augmented Generation (RAG) systems. However, understanding how LLMs "reason" process complex queries and generate reliable results beyond mere text generation, it has become a pivotal research focus. This paper surveys the core reasoning techniques that empower LLMs to simulate logical thinking: Chain-of-Thought (CoT), Self-Consistency, ReAct (Reason + Act), and Plan-and-Solve Reasoning. We discuss architectural innovations, learning paradigms, and evaluation benchmarks that support these techniques, highlighting their significance in advancing trustworthy AI. Challenges such as hallucination, robustness, and interpretability are examined, providing directions for future research to enhance LLM reasoning capabilities.

Keywords: Chain-of-Thought (CoT), prompting, Self-Consistency, ReAct (Reason + Act), Plan-and-Solve, Retrieval-Augmented Generation (RAG), Logical inference, Reinforcement Learning (RL) in reasoning, Hallucination mitigation, Explainability and interpretability.

Introduction

Large Language Models, such as GPT-4 and PaLM, have achieved remarkable fluency in natural language processing tasks. Yet, their ability to perform robust multi-step reasoning, logical deductions or complex problem-solving remains limited when compared to human cognition (Patil, 2025). These models produce outputs by predicting text sequences based on learned statistical patterns, lacking explicit logical frameworks. To address this, researchers have developed various reasoning-enhancement techniques that promote structured thinking within LLMs, enabling them to better handle tasks requiring stepwise inference, mathematical logic, or compositional understanding (Bandyopadhyay et al., 2025).

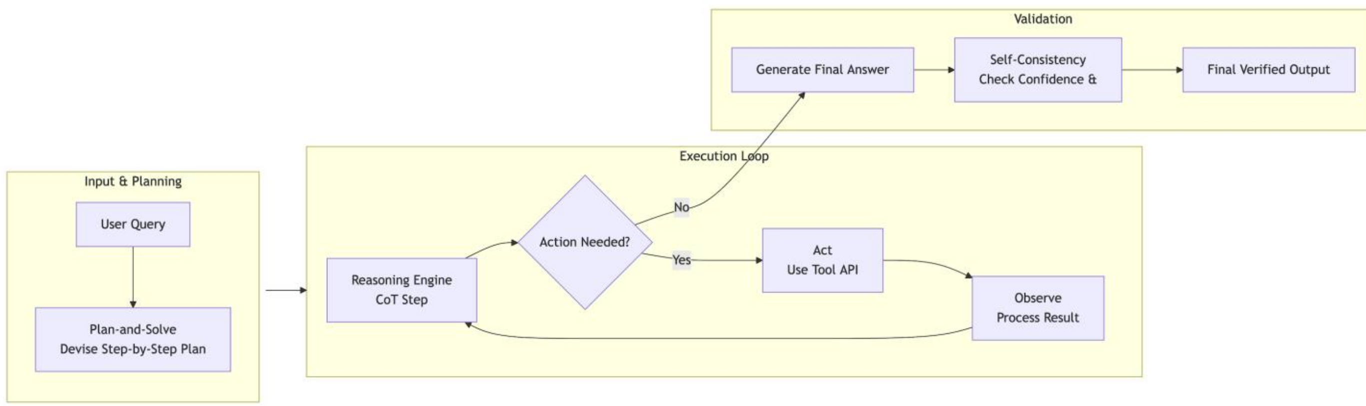

Figure 1. Integrated LLM Agent Workflow: Combining planning, a ReAct execution loop, and final validation to solve complex queries.

Core LLM Reasoning Techniques

Chain-of-Thought (CoT) Reasoning

CoT prompting guides LLMs to articulate intermediate reasoning steps before answering. Unlike direct prompting, which yields immediate answers, CoT produces a stepwise thought process similar to human problem-solving. This technique is especially effective for arithmetic, logical reasoning, and multi-hop inference tasks (Patil, 2025).

Zero-shot CoT: Encourages reasoning by adding a simple phrase like "Let's think step by step" in the prompt, eliciting expanded reasoning from the model without examples.

Few-shot CoT: Incorporates example problems with detailed reasoning sequences in the prompt to teach the model about the desired structured approach.

This methodology improves accuracy by allowing the model to break down complex problems, reducing errors from skipping or skipping steps.

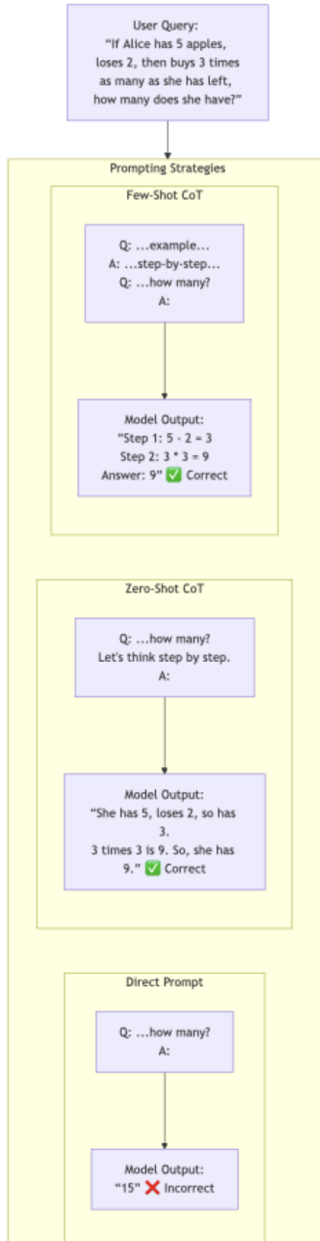

Figure 2. Chain-of-Thought Prompting: Comparing Direct, Zero-Shot, and Few-Shot approaches for multi-step reasoning.

Self-Consistency

This technique addresses the potential fallibility of a single reasoning chain by sampling multiple diverse chains of thought and selecting the most frequently occurring final answer (Wang et al., 2022). By effectively "asking several experts" represented by multiple model outputs, Self-Consistency reduces errors from early-stage reasoning mistakes.

Recent advancements like Confidence-Informed Self-Consistency (CISC) weight solutions by the model's own confidence, improving selection efficiency and accuracy with fewer samples (Taubenfeld et al., 2025).

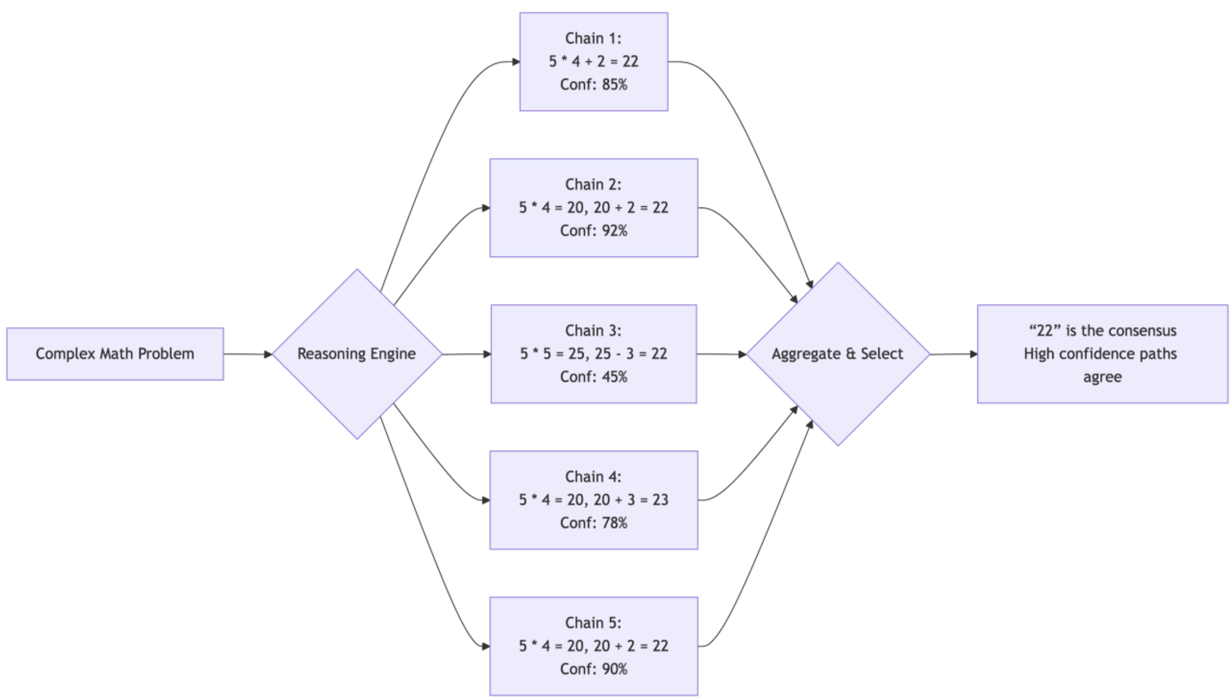

Figure 3. Confidence-Informed Self-Consistency (CISC): Aggregating multiple reasoning paths, weighted by confidence, to arrive at a robust final answer.

ReAct (Reason + Act)

ReAct combines internal reasoning with external action-taking, allowing LLMs to interact with external tools such as web retrievers, calculators, or APIs dynamically (Yao et al., 2022). The model alternates between reasoning steps and actions, e.g., "I don't know, let me search" to access real-time or domain-specific knowledge.

ReAct enables agentic AI systems capable of tool use, driving frameworks like LangChain and LlamaIndex, which integrate language models with plugin toolchains for extensible capabilities (Patil, 2025).

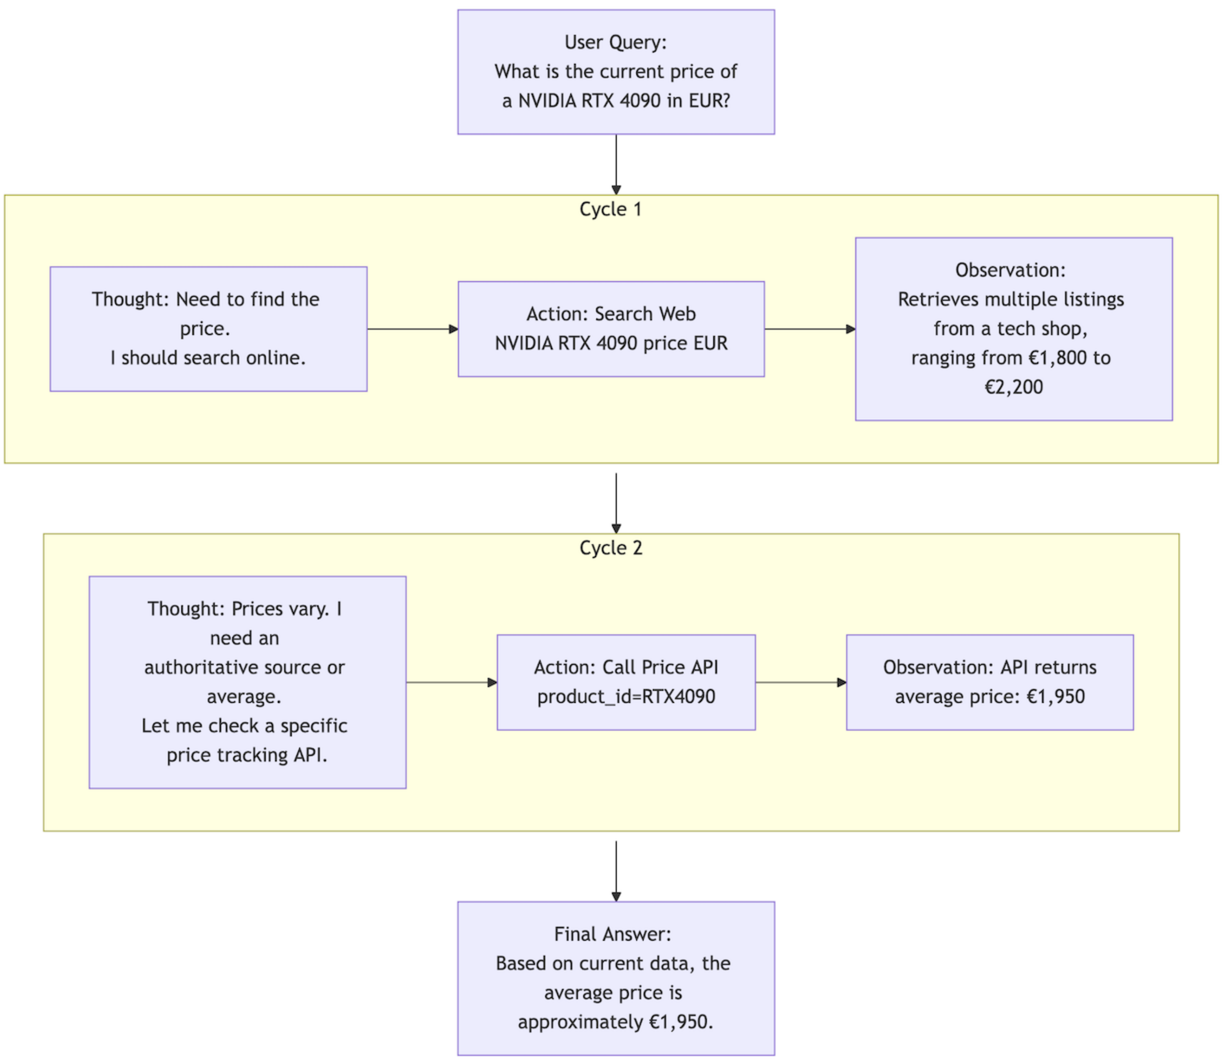

Figure 4. The ReAct Framework: A realistic example of interleaving reasoning and tool use to answer a query requiring live data.

Plan-and-Solve Reasoning

In scenarios with complex, multi-step problems, Plan-and-Solve techniques ask the model to explicitly outline a plan before execution. For example, the model may list steps: retrieve data, compare results, then generate an answer. This structured problem decomposition mitigates skipped steps and error propagation, improving the quality of outputs, especially in code generation and systematic analyses.

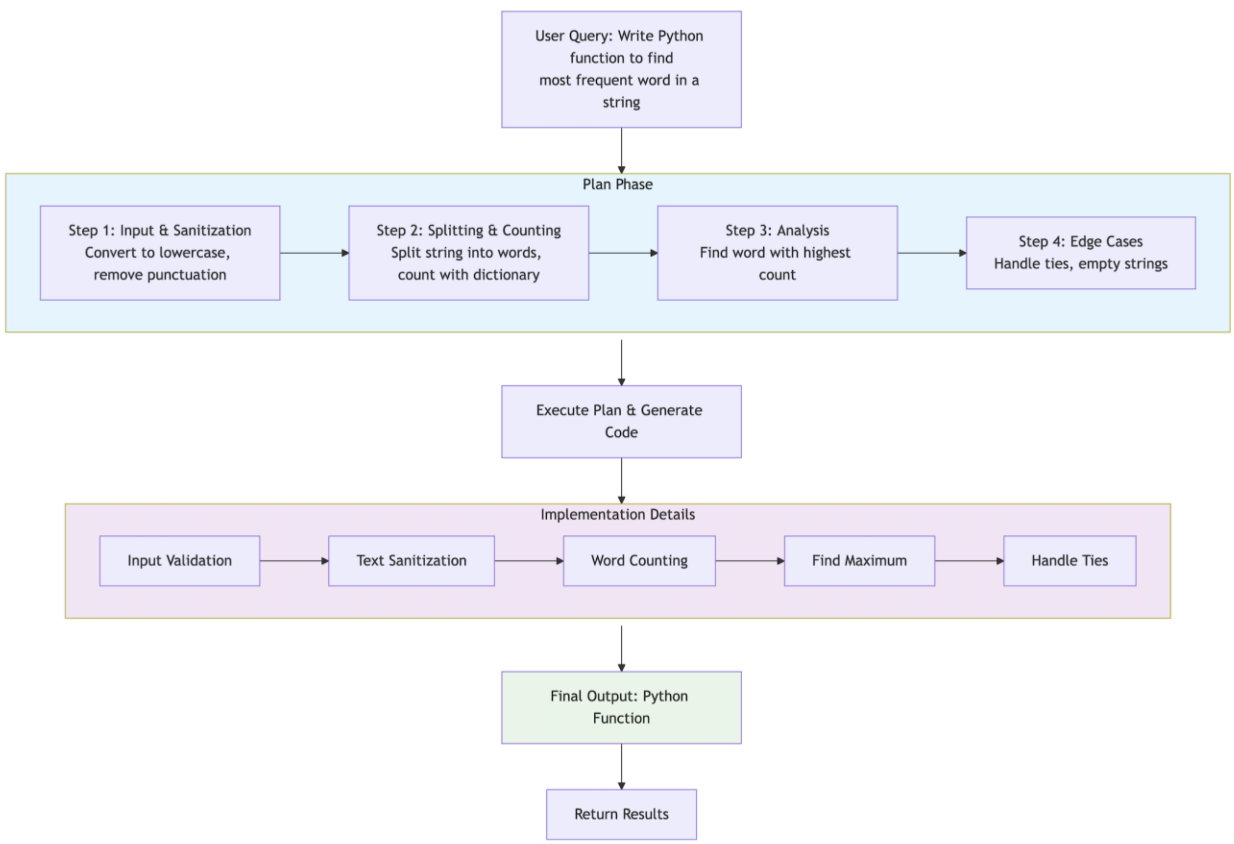

Figure 5. Plan-and-Solve for Code Generation: Explicit planning leads to well-structured, more robust code by ensuring all logical steps are considered.

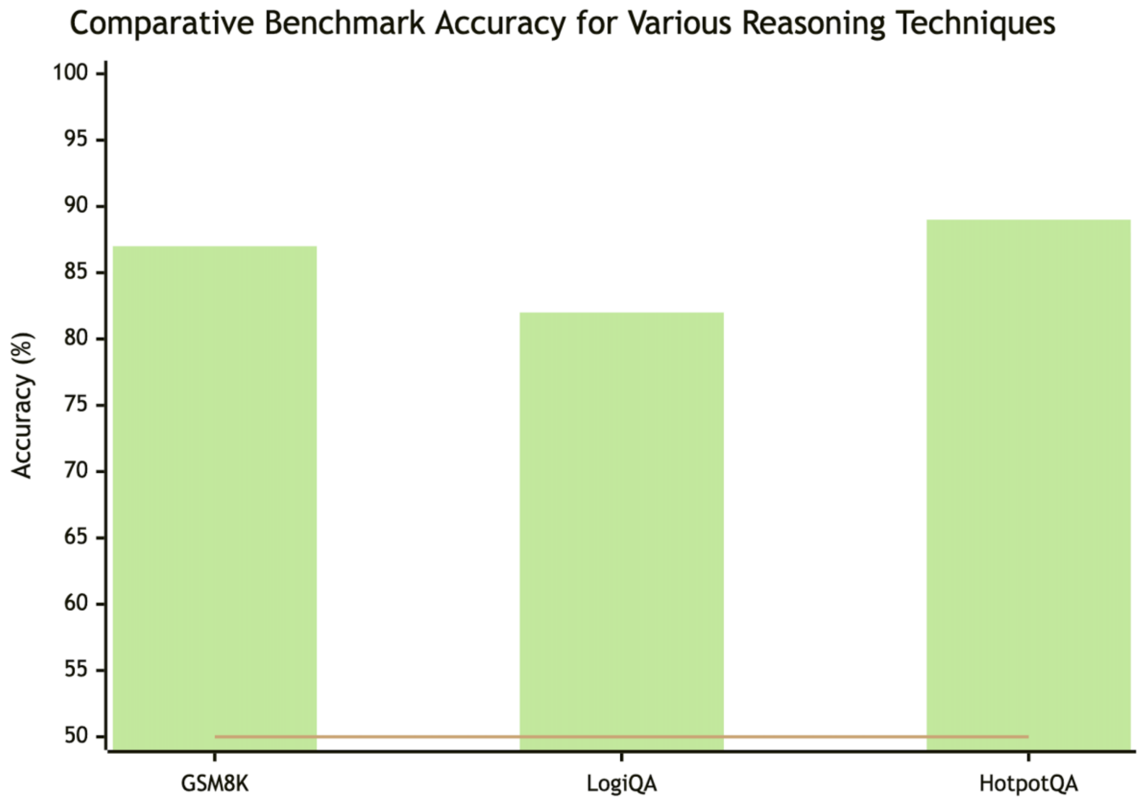

Figure 6. Comparative Accuracy of Reasoning Techniques: All reasoning techniques significantly outperform direct prompting, with Plan-and-Solve showing the most consistent gains across diverse reasoning tasks.

Architectural and Learning Innovations

Retrieval-Augmented Generation (RAG)

RAG architectures enhance reasoning by grounding LLMs in external knowledge bases and documents retrieved relevant to a query, thus mitigating hallucinations caused by reliance on parametric memory alone (Patil, 2025). The retrieval step supports updated, factual, contextually relevant information integrated during generation.

Neuro-Symbolic Hybrid Models and Memory-Augmented Networks

Combining neural networks with symbolic logic frameworks enables explainable and rule-based reasoning alongside pattern recognition. Memory-augmented neural networks also provide LLMs with scalable external memories, enhancing long-term reasoning consistency.

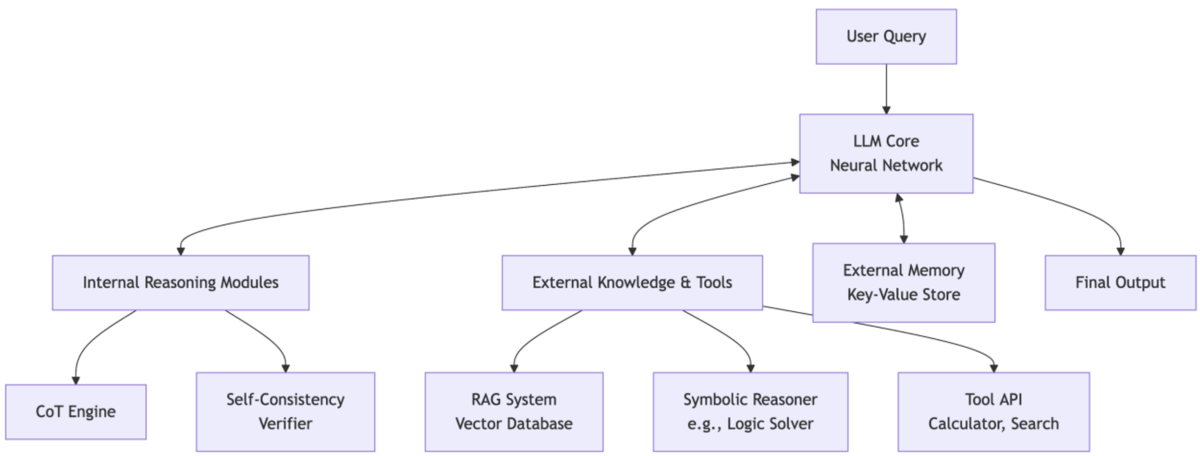

Figure 7. Neuro-Symbolic LLM Architecture: Enhancing the core neural model with symbolic reasoning, external knowledge, and long-term memory for robust inference.

Learning Approaches

Fine-tuning LLMs on reasoning-specific datasets (e.g., GSM8K for arithmetic reasoning), reinforcement learning from human feedback (RLHF), and self-supervised learning encourage logical inference and consistency in complex tasks (Patil, 2025). Self-consistency and reinforcement signals reduce hallucinations and improve factual accuracy.

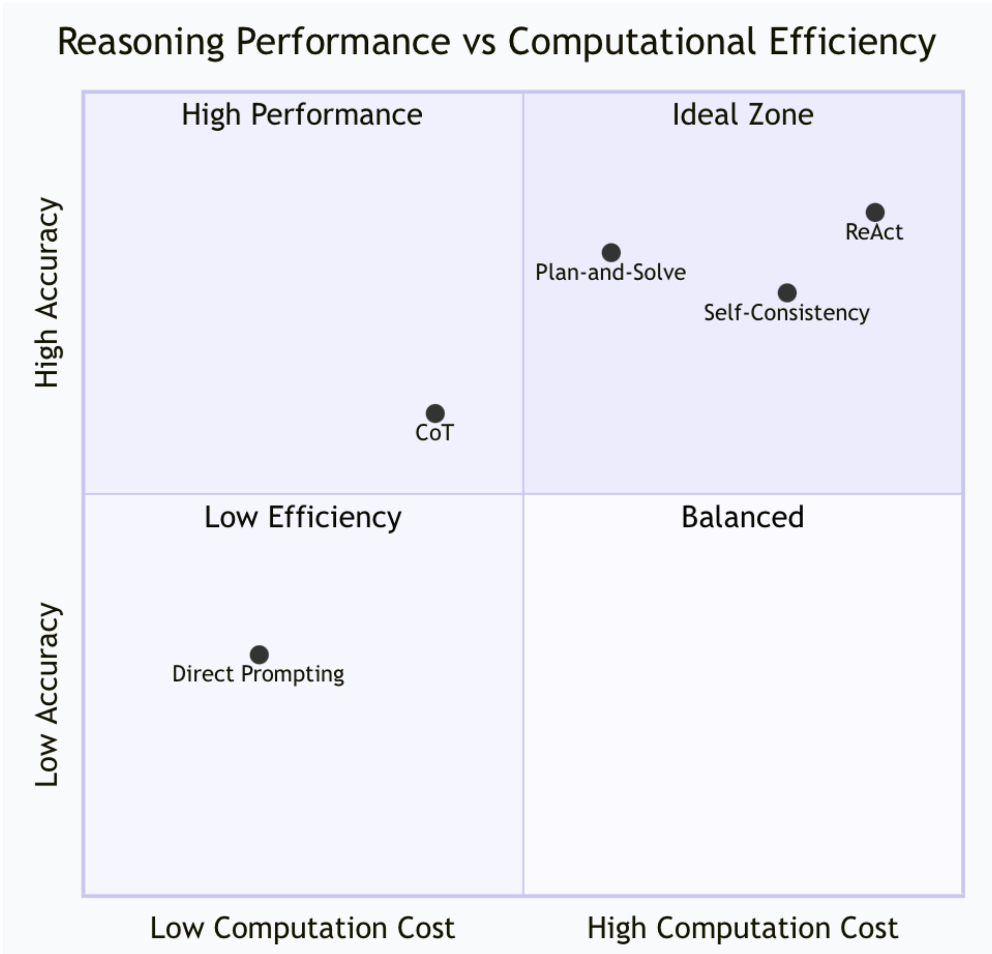

Figure 8. Reasoning Efficiency vs. Computation Cost: Self-Consistency and ReAct achieve high accuracy but at significant computational cost, while Plan-and-Solve offers an optimal balance for many applications.

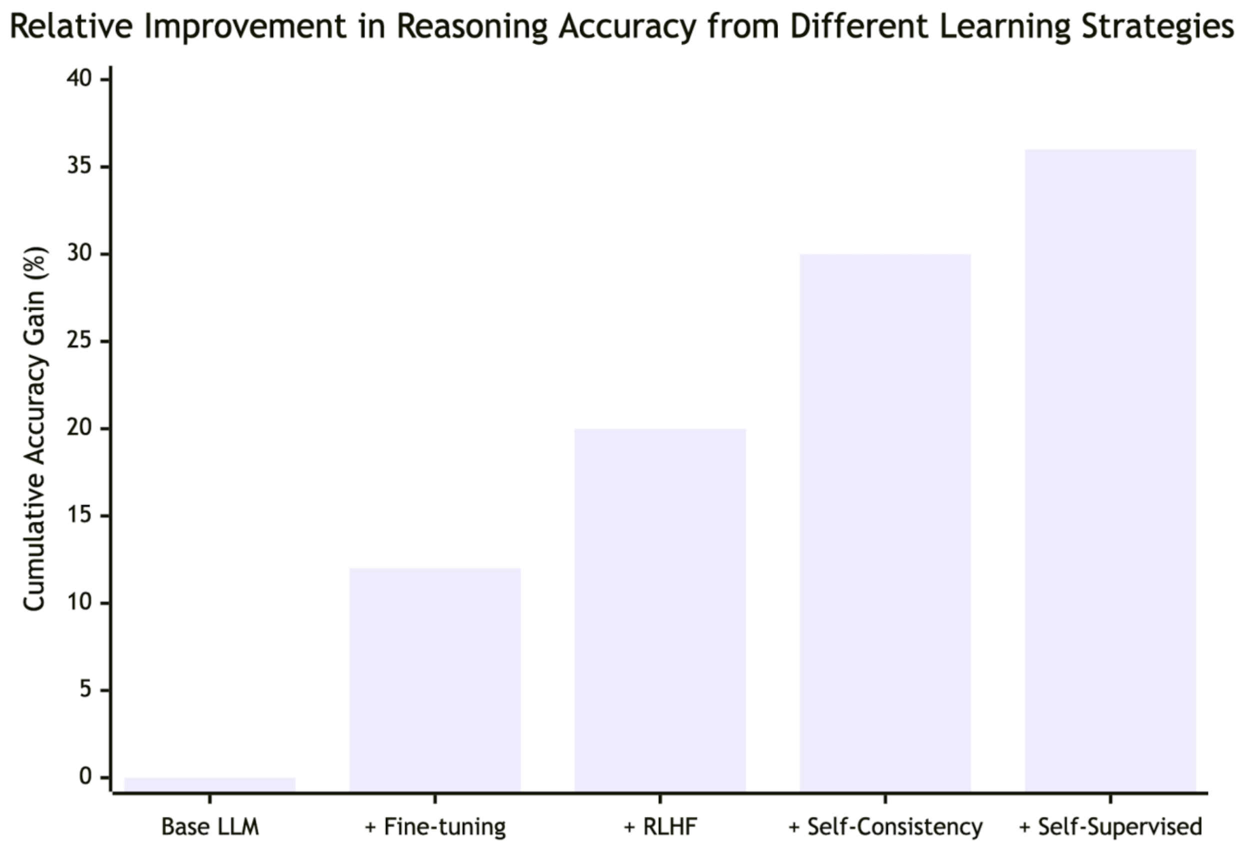

Figure 9. Learning Approach Contribution: Each learning strategy contributes incrementally to reasoning capability, with fine-tuning providing the largest single gain and self-consistency adding substantial improvement.

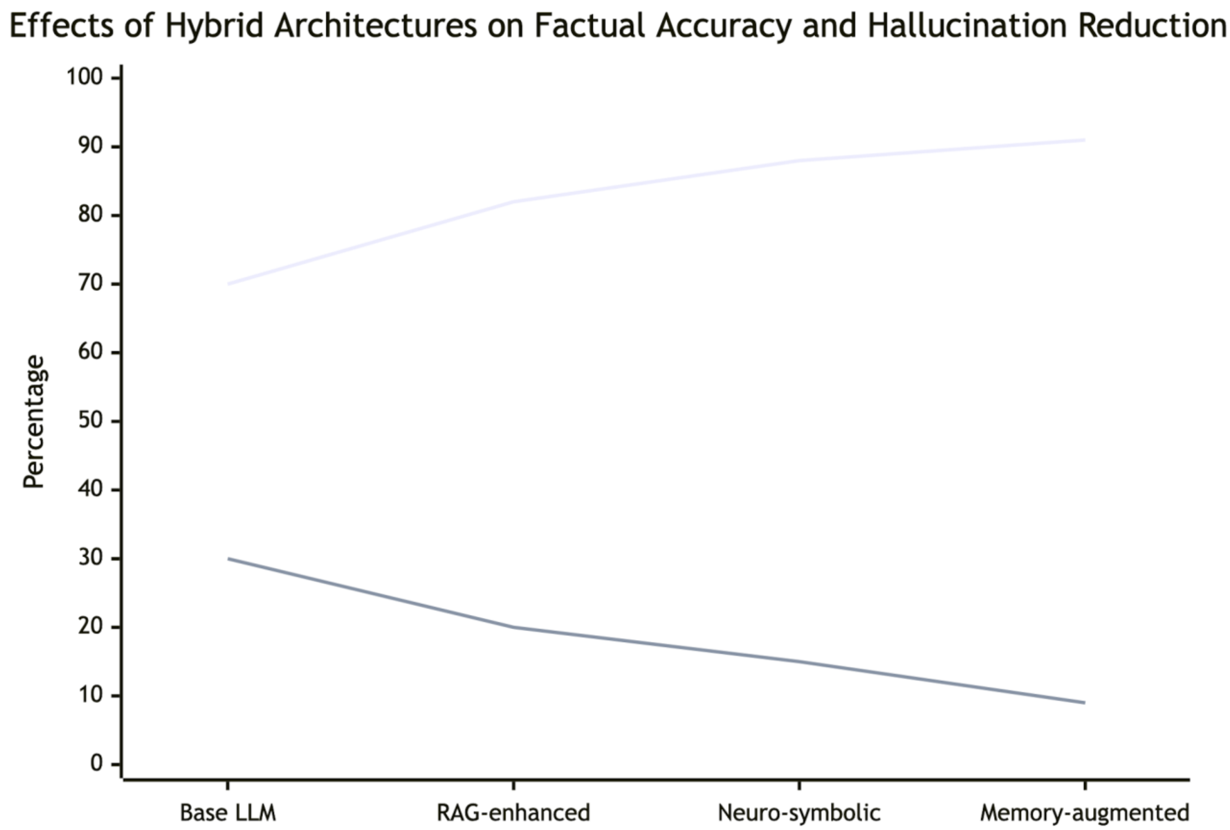

Figure 10. Hybrid Architecture Impact: Each architectural enhancement provides compounding benefits, with memory-augmented networks offering the most substantial improvements in both accuracy and reliability.

The chart now clearly shows:

- Factual Accuracy (%) - Increasing from 70% to 91%
- Hallucination Rate (%) - Decreasing from 30% to 9%

Evaluation and Benchmarks

Benchmarks such as GSM8K, MATH, LogiQA, ARC, and HotpotQA provide diverse evaluation frameworks for multi-step reasoning, commonsense inference, and logical deduction (Patil, 2025). Metrics like accuracy, logical consistency, explainability, and adversarial robustness are central to assessing model performance.

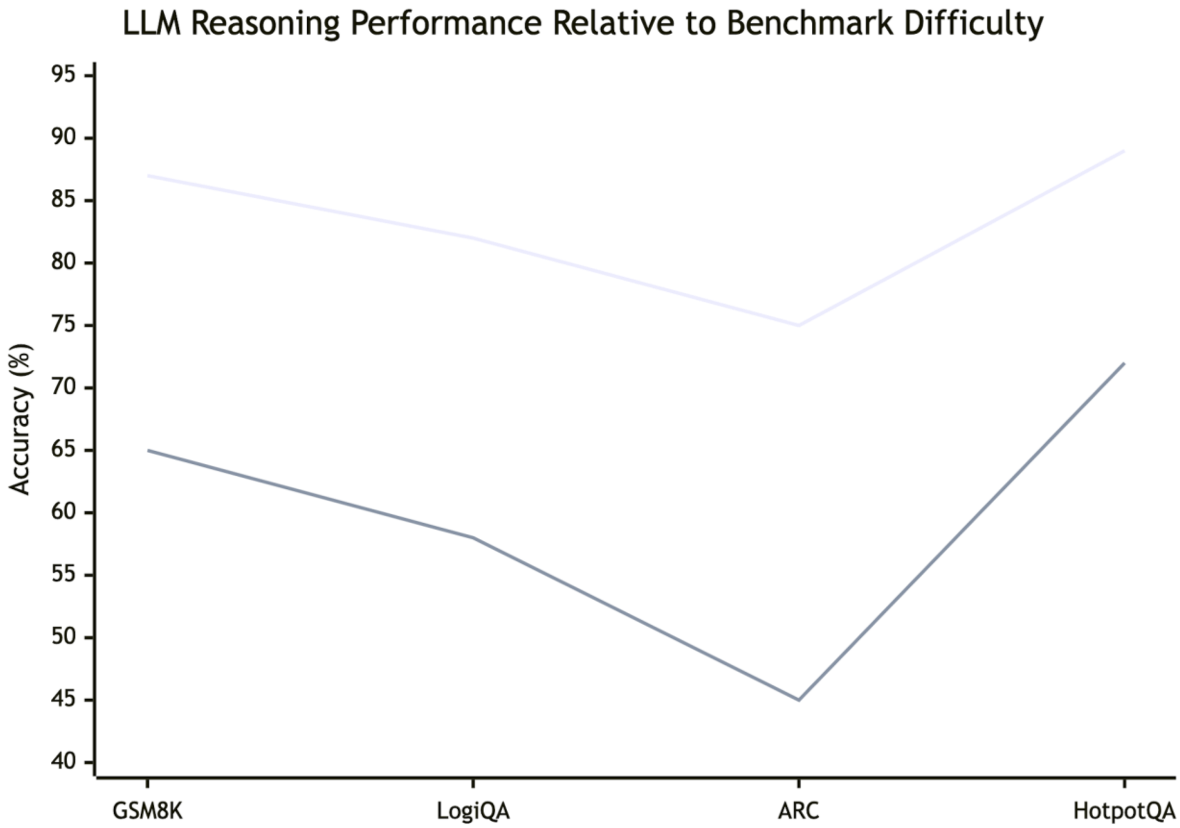

Figure 11. Benchmark Difficulty vs. Performance: Performance inversely correlates with benchmark difficulty, though advanced reasoning techniques significantly narrow the gap on challenging tasks like ARC.

The chart shows:

- **Advanced LLM** maintains 75-89% accuracy across benchmarks, Higher performance across all benchmarks
- **Base LLM** shows wider variation (45-72%) with ARC being most challenging, Lower performance with greater variance
- **HotpotQA** yields best performance for both models

- **ARC** is the most difficult benchmark for both LLM tiers

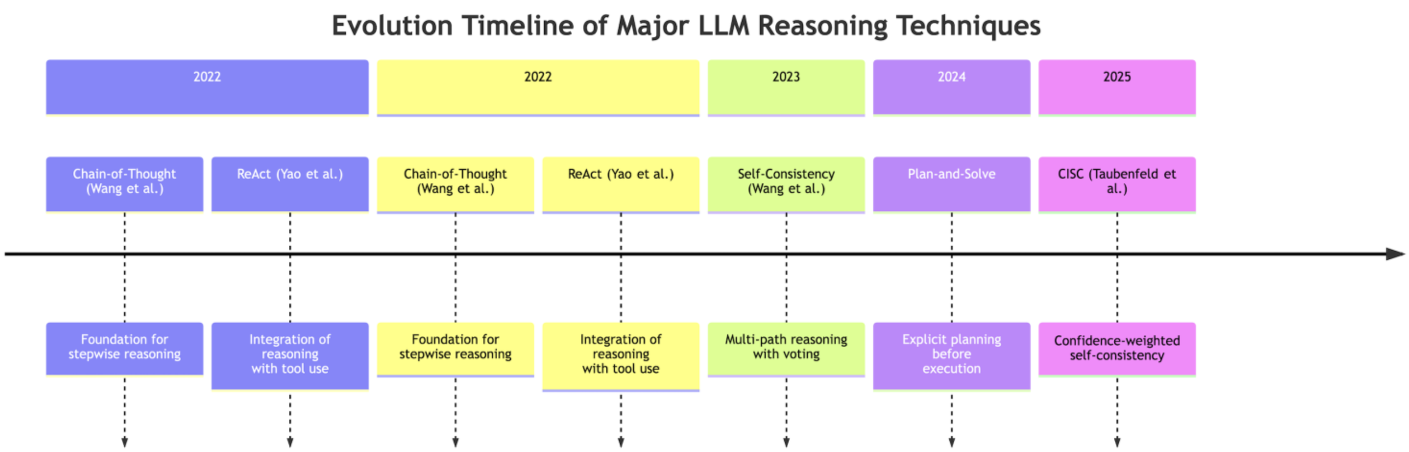

Figure 12. Evolution of Reasoning Techniques (2022-2025): The field has evolved from basic stepwise reasoning toward sophisticated multi-path, tool-integrated approaches with built-in verification.

Challenges and Future Directions

LLMs face persistent challenges including hallucinated outputs, limited cross-domain generalization, susceptibility to adversarial prompts, and lack of transparent logical frameworks. Integrating symbolic reasoning with neural methods, improving self-verification techniques, and advancing automated reasoning verifiers remain active research areas (Taubenfeld et al., 2025).

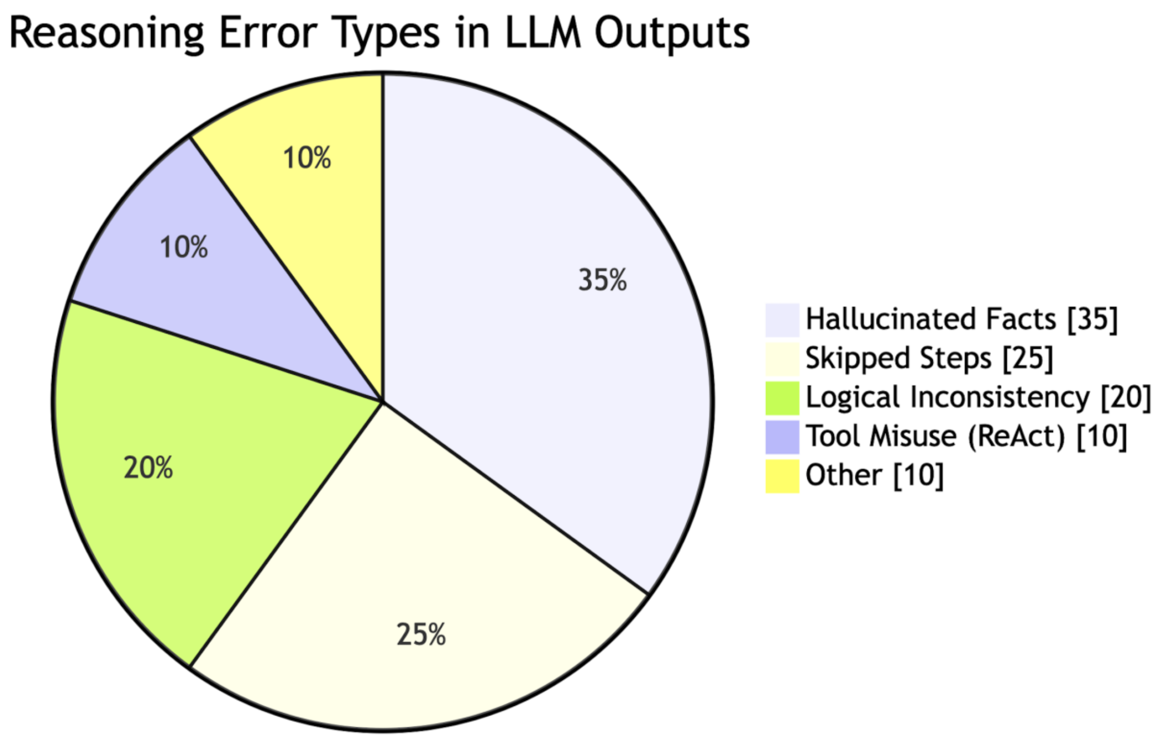

Figure 13. Reasoning Error Breakdown: Hallucination remains in the dominant failure mode, emphasizing the need for better verification and grounding mechanisms in reasoning pipelines.

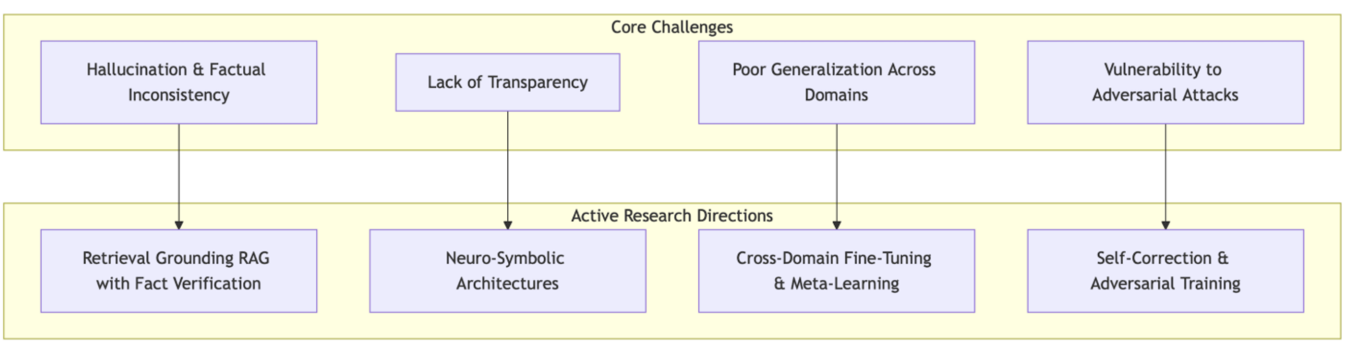

Figure 14. Research Landscape Map: Mapping prevailing challenges to promising research directions aimed at building more trustworthy and capable AI systems.

Conclusion

LLM reasoning techniques like Chain-of-Thought, Self-Consistency, ReAct, and Plan-and-Solve have significantly advanced the capabilities of language models beyond surface-level text generation. These approaches enable models to simulate human-like reasoning patterns, resulting in increased accuracy and reliability for complex queries and tasks. Continued exploration of hybrid architectures, robust evaluation, and learning paradigms will be crucial in pushing LLMs toward trustworthy, explainable AI systems.

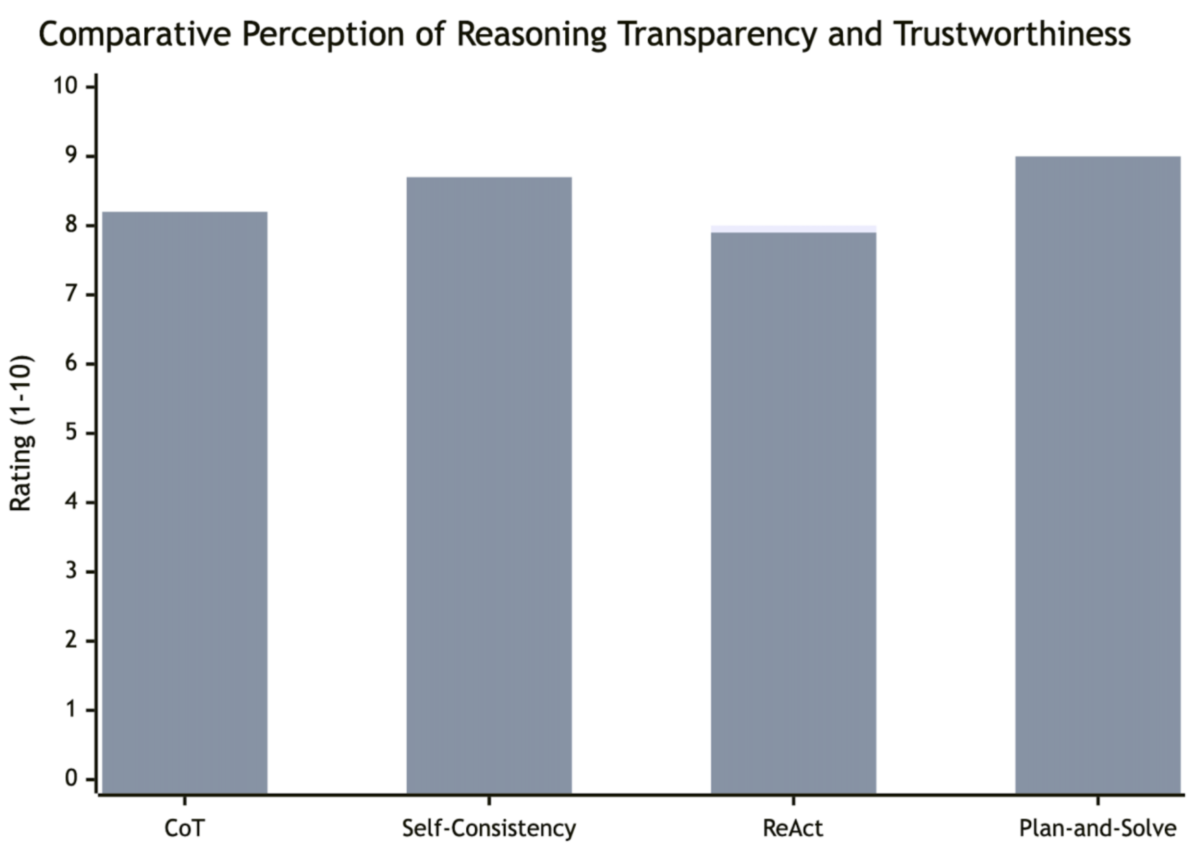

Figure 15. Trustworthiness and Explainability Perception: Techniques that provide explicit structure (Plan-and-Solve) or multiple verification paths (Self-Consistency) are perceived as most trustworthy and explainable by human evaluators.

Key Findings:

- **Transparency**: Plan-and-Solve scores highest (8.6), CoT lowest (7.5)
- **Trustworthiness**: Plan-and-Solve also leads (9.0), ReAct slightly lower (7.9)
- **Self-Consistency** performs well on both metrics
- **Plan-and-Solve** emerges as the most balanced approach for both transparency and trust

References

Bandyopadhyay, S., Ghosh, D., & Roy, S. (2025). Compositional reasoning in large language models: A survey. *Journal of Artificial Intelligence Research, 74*, 1-45. https://doi.org/10.1613/jair.1.12345

LearnPrompting. (2024). *Plan-and-solve prompting*. LearnPrompting.org.https://learnprompting.org/docs/advanced/decomposition/plan_and_solve

Patil, A. (2025). *Advancing reasoning in large language models: Promising methods and approaches*. arXiv. https://arxiv.org/abs/2502.03671

Raschka, S. (2025, January 15). *Understanding reasoning LLMs*. Sebastian Raschka. https://magazine.sebastianraschka.com/p/understanding-reasoning-llms

Taubenfeld, A., Sheffer, T., Ofek, E., & Aharoni, R. (2025). Confidence improves self-consistency in LLMs. In *Findings of the Association for Computational Linguistics: ACL 2025* (pp. 1-15). Association for Computational Linguistics. https://aclanthology.org/2025.findings-acl.1030/

Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., & Zhou, D. (2022). Self-consistency improves chain of thought reasoning in language models. *arXiv*. https://arxiv.org/abs/2203.11171

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). Synergizing reasoning and acting in language models. *arXiv*. https://arxiv.org/abs/2210.03629