

Codette: Multi-Perspective Reasoning as a Convergent Dynamical System with Meta-Cognitive Strategy Evolution

Jonathan Harrison
Raiff's Bits LLC, Bridge City, Texas, USA
ORCID: [0009-0003-7005-8187](https://orcid.org/0009-0003-7005-8187)
jonathan@raiffsbits.com

April 2026

Abstract

We present CODETTE, a modular cognitive architecture that models multi-perspective reasoning as a constrained dynamical system. The system integrates six heterogeneous reasoning agents coordinated through a meta-cognitive layer with reflective memory (cocoons). Our theoretical foundation, RC+ ξ (Recursive Convergence + Epistemic Tension), formalizes cognitive state evolution and proves convergence under Lyapunov stability analysis. Empirical evaluation across 17 problem domains demonstrates a **93.1%** composite quality improvement over single-agent baselines ($p < 0.0001$), with reasoning depth increasing from 0.402 to 0.855.

1 Introduction

Modern Large Language Models (LLMs) suffer from hallucination and a lack of persistent reasoning structure [1]. While Chain-of-Thought (CoT) [2] improves performance, it remains a single-perspective linear process. CODETTE addresses this by formalizing multi-perspective reasoning as a convergent dynamical system where heterogeneous agents produce coherent outputs through state-space stabilization.

2 System Architecture

The architecture separates the cognitive load into the Perspective Plane, the Memory Substrate (Cocoon traces), and the Meta-Cognitive Strategy Engine.

Codette Architecture (Functional Flow)




Figure 1: The CODETTE Architecture: Interaction between specialized reasoning agents and the meta-cognitive layer.

3 Experimental Evaluation

We conducted 68 evaluations across 17 problems in 6 categories. CODETTE outperformed all baselines in reasoning depth and diversity.

Table 1: Benchmark Results by Condition (0-1 Scale)

Condition	Composite	Depth	Diversity	Coherence	Ethics	Novelty	Grounding	Turing
SINGLE	0.338	0.402	0.237	0.380	0.062	0.327	0.456	0.412
MULTI	0.632	0.755	0.969	0.503	0.336	0.786	0.604	0.180
MEMORY	0.636	0.770	0.956	0.500	0.340	0.736	0.599	0.291
CODETTE	0.652	0.855	0.994	0.490	0.366	0.852	0.575	0.245




Figure 2: Dimension-level performance of CODETTE across evaluation metrics.

3.1 Performance and Latency

Validation testing shows a 100% pass rate across safety and retention categories. However, we note a significant latency tradeoff for the increased reasoning depth.




Figure 3: Composite quality improvement.




Figure 4: Latency vs performance tradeoff.

4 Conclusion

CODETTE proves that multi-perspective stability leads to significant depth improvements. The trade-off between depth and conversational naturalness remains a key area for future research.

References

- [1] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623, 2021.
- [2] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, 2022.