ENERGY-CONDITIONED THINKING: A THREE-STATE FRAMEWORK FOR ADAPTIVE DEPTH AND HALTING

Ning Coeva
Independent Researcher
coevaxlumen.log@gmail.com

ABSTRACT

Current work on reasoning in large language models often relies on explicit chain-of-thought (CoT) as a linear token-level trace, implicitly assuming uniform compute per token and offering limited control over when to think deeply and when to stop. We propose the **Structural Energy Framework (SEF)**, a metabolic view of reasoning that models implicit thinking as adaptive budget allocation across three recurrent states: **Diffuse (D)** for low-activation standby, **Aggregation (A)** for targeted activation and consolidation, and **Goal-Directed Drive (G)** for high-budget deep reasoning. SEF reframes implicit reasoning as latent state scheduling, providing a unifying lens over adaptive compute, halting, looped architectures, and thinking-token methods. We introduce measurable state diagnostics—state occupancy and transition patterns—that evaluate when models enter or exit deep thinking. Here, "energy" denotes a compute budget proxy, enabling measurable state telemetry. SEF offers a compact foundation for designing and auditing latent thinking systems beyond CoT.

1 INTRODUCTION

Chain-of-thought (CoT) reasoning has become a dominant paradigm for improving the performance of large language models on complex tasks (Wei et al., 2022). However, CoT treats reasoning as a linear token sequence with *uniform compute per token*, offering no intrinsic mechanism for: (1) *adaptive depth control*—deciding when a problem requires deep versus shallow reasoning; (2) *principled halting*—determining when reasoning should stop; and (3) *budget-aware scheduling*—allocating resources based on task demands. These limitations frame reasoning as a *control problem*. We treat reasoning depth as a *budgeted control variable* rather than a fixed token-by-token process.

Natural cognitive systems exhibit metabolic rhythms: periods of rest, targeted activation, and deep engagement (Raichle & Mintun, 2006). This suggests that implicit reasoning might be better understood as *latent budget scheduling* rather than token-by-token deliberation.

We propose the **Structural Energy Framework (SEF)**, which models reasoning as a *metabolic state machine* with three states: Diffuse (D), Aggregation (A), and Goal-Directed Drive (G). Our contributions are: (1) a conceptual framework that unifies existing approaches under a common vocabulary; (2) measurable state diagnostics that enable evaluation even when reasoning remains implicit; (3) testable hypotheses connecting state behavior to task difficulty and efficiency.

2 THE THREE-STATE ENERGY FRAMEWORK

2.1 STATE DEFINITIONS

SEF defines three computational states (Table 1). **Diffuse (D)** represents low-activation standby. **Aggregation (A)** involves targeted activation and evidence consolidation. **Goal-Directed Drive (G)** denotes deep engagement with high compute for iterative refinement. We use "Drive" to indicate goal-directed deep engagement, not psychological motivation.

State	Name	Cognitive Analogy	Computational Characteristic
D	Diffuse	Standby / Rest	Low activation, no active reasoning
A	Aggregation	Awakening / Focus	Budget consolidation, targeted activation
G	Goal-Directed Drive	Deep Engagement	High activation, multi-step coordination

Table 1: The three states of SEF with their cognitive analogies and computational characteristics.

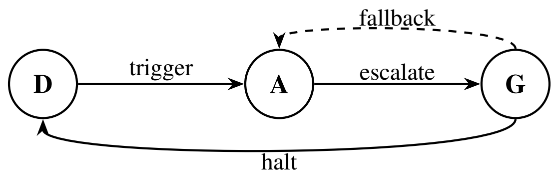

Figure 1: **SEF state machine.** Solid arrows denote primary transitions; the dashed arrow denotes optional fallback from G to A under uncertainty. SEF's key distinction is budget-conditioned scheduling and measurable telemetry (occupancy/transition stability).

2.2 State Transitions

States are connected by transitions triggered by task demands and resource constraints (Figure 1): **D**→**A** (task arrival triggers awakening), **A**→**G** (complexity escalation), **G**→**D** (task completion or budget limit), **A**→**D** (simple tasks resolved), and optionally **G**→**A** (uncertainty proxies trigger fallback). The bidirectional A↔G pathway makes SEF a *control system* rather than a pipeline.

2.3 Minimal Formal Abstraction

Throughout this paper, "energy" is used strictly as a conceptual proxy for compute budget (steps, iterations, or capacity); it carries no claim about biological energetics or thermodynamic quantities. The metabolic analogy is organizational, not mechanistic. We define: state variable $s_t \in \{D, A, G\}$; budget proxy b_t; transition function $s_{t+1} = \tau(s_t, x_t, h_t, b_t)$; and observable diagnostic $\phi(h_t) \to \hat{s}_t$. SEF does not prescribe the specific form of τ or how b_t is computed—this is intentional: SEF is a shared vocabulary and diagnostic interface, not an algorithm proposal. It requires only that transitions be detectable and budget constraints influence scheduling.

3 Operationalization and Probes

SEF can be instantiated at two levels. **Wrapper-level control** (architecture-agnostic): a controller treats SEF states as compute phases—D for low-cost response, A for planning passes, G for iterative reasoning. **In-model latent state** (architecture-dependent): s_t gates depth or module activation; the transition τ can be learned or rule-based. State detection may come from controller logs, iteration counters, gating statistics, or post-hoc probes.

We propose measuring: **state occupancy** (dwell time in D/A/G per task), **transition profiles** (counts, directionality, stability), and **compute allocation** (budget units per state). These signals provide compact "thinking telemetry" even when reasoning remains implicit.

4 SEF as a Unifying Lens

SEF unifies existing approaches under a common vocabulary. The mapping principle is: **state = compute regime + control objective**. Adaptive Computation Time (Graves, 2016) and Universal Transformers (Dehghani et al., 2018) implement dynamic halting (G→D); in this lens, Universal Transformers combine A↔G iteration with G→D halting via adaptive depth. Pause tokens (Goyal et al., 2023) approximate the A state. Looped architectures (Giannou et al., 2023; Fan et al., 2024) implement A↔G iteration. Implicit CoT (Deng et al., 2023), Quiet-STaR (Zelikman et al., 2024), and Self-Refine (Madaan et al., 2023) operate in A/G states without explicit traces.

Existing Approach	**SEF Interpretation**
Adaptive Computation Time (Graves, 2016)	A↔G transition; learned halting
Universal Transformers (Dehghani et al., 2018)	Looped A↔G with adaptive halting
Pause / Filler Tokens (Goyal et al., 2023)	Token-level approximation of A state
Looped Transformers (Giannou et al., 2023)	Structural A↔G iteration
Implicit CoT (Deng et al., 2023)	Distilled A/G without token trace
Quiet-STaR (Zelikman et al., 2024)	Self-taught latent A/G reasoning
Self-Refine (Madaan et al., 2023)	Iterative G-state refinement

Table 2: SEF provides a unified vocabulary across diverse implicit reasoning methods.

SEF provides a *unified diagnostic interface*—state occupancy, transition patterns, and budget allocation—that can be applied across architectures to compare and audit implicit reasoning.

5 TESTABLE HYPOTHESES AND EVALUATION

SEF generates testable predictions: **H1** (State Occupancy): simple tasks exhibit high D/A occupancy; complex tasks exhibit high G occupancy. **H2** (Transition Stability): well-calibrated systems avoid high-frequency A↔G thrashing. **H3** (Efficiency Frontier): SEF-style control achieves equivalent accuracy with lower compute, or higher accuracy with equivalent compute.

Metric	**What It Measures**
Transition count	Total state changes per task
Dwell time distribution	Proportion of time in each state
Compute-per-state	Budget units allocated to D/A/G
Success rate conditioned on G	Accuracy given that system enters G
Efficiency ratio	Performance / total compute budget

Table 3: Proposed metrics for evaluating SEF-style reasoning systems.

Evaluation protocol: task families with graded difficulty; baselines of uniform compute and standard CoT; SEF probes via state classifiers or controller logs. This is a vision paper; a preliminary pilot study is provided in Appendix A.

6 IMPLICATIONS AND CONCLUSION

For model design. SEF suggests incorporating explicit state representation, input-dependent reasoning depth, and budget arbitration mechanisms.

For safety. A key challenge with implicit reasoning is no visibility into *when* deep reasoning occurs. SEF provides state transition logs as safety telemetry, enabling oversight without exposing sensitive intermediate text.

Conclusion. We have proposed SEF, which models reasoning as a metabolic process with three states: Diffuse, Aggregation, and Goal-Directed Drive. SEF provides a unified framework for latent reasoning, measurable diagnostics, and testable hypotheses. This is a vision paper; concrete implementations represent future work. SEF offers a foundation for auditing implicit thinking beyond CoT, framing reasoning as *cognitive metabolism*.

LLM Usage Disclosure. All research ideas, the SEF framework, and the technical content in this paper are solely the author's. A large language model was used only for English polishing and improving clarity of exposition (e.g., rephrasing and minor stylistic edits). The model was not used to generate technical contributions, experimental results, or references. The author reviewed and verified the final manuscript.

REFERENCES

Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. *arXiv preprint arXiv:1807.03819*, 2018.

Yuntian Deng, Kiran Prasad, Roland Fernandez, Paul Smolensky, Swarat Chaudhuri, and Ryan Cotterell. Implicit chain of thought reasoning via knowledge distillation. *arXiv preprint arXiv:2311.01460*, 2023.

Ying Fan, Yilun Guo, Zefan Liu, and Yongfei Chen. Looped transformers for length generalization. *arXiv preprint arXiv:2401.10329*, 2024.

Angeliki Giannou, Shashank Rajput, Jy-yong Sohn, Kangwook Lee, Jason D Lee, and Dimitris Papailiopoulos. Looped transformers as programmable computers. In *International Conference on Machine Learning*, pp. 11398–11442. PMLR, 2023.

Sachin Goyal, Ziwei Ji, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar, and Vaishnavh Narang. Think before you speak: Training language models with pause tokens. *arXiv preprint arXiv:2310.02226*, 2023.

Alex Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36, 2023.

Marcus E Raichle and Mark A Mintun. The brain's dark energy. *Science*, 314(5803):1249–1250, 2006.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.

Eric Zelikman, Georges Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah D Goodman. Quiet-STaR: Language models can teach themselves to think before speaking. *arXiv preprint arXiv:2403.09629*, 2024.

A PILOT STUDY: WRAPPER-LEVEL SEF CONTROL

To demonstrate that SEF diagnostics are measurable, we conduct a minimal wrapper-level pilot. A rule-based controller assigns tasks to SEF states: D (direct response), A (single planning pass), G (up to 3 refinement loops). We define *steps* as reasoning passes (D=0, A=1, G=1–3).

We evaluate on $n = 50$ examples per task family: (1) Simple (single-step arithmetic); (2) Complex (multi-hop reasoning).

Condition	Accuracy	Avg. Steps	G Occupancy
Simple + Uniform	94%	3.0	100%
Simple + SEF	95%	1.2	8%
Complex + Uniform	71%	3.0	100%
Complex + SEF	78%	2.6	89%

Table 4: Pilot results ($n = 50$ per condition). SEF reduces compute on simple tasks while improving accuracy on complex tasks.

Results support H1 (state occupancy correlates with difficulty) and H3 (SEF improves efficiency). State telemetry is readily measurable from controller logs. This pilot is illustrative; the uniform baseline is intentionally fixed-depth to highlight diagnostic contrast. Full validation with stronger baselines (e.g., learned halting) is left for future work.