Designing for Cognitive Rhythm: A Three-State Model for Managing Mental Energy in Tools for Thought

Ning Coeva
Independent Researcher
China
coevaxlumen.log@gmail.com

Abstract

Tools for Thought (TfT) promise to augment human cognition, yet current GenAI tools often force knowledge workers into a false choice: productivity or cognitive depth. We propose a three-state model for cognitive rhythm management—*Diffuse*, *Aggregation*, and *Drive*—that reframes this tension. Rather than optimizing for constant engagement, effective TfT should orchestrate transitions between exploratory openness (Diffuse), structural consolidation (Aggregation), and focused execution (Drive). These states are not rigid phases but fluid orientations that may overlap or alternate rapidly. We present implementation-agnostic design patterns, measurable outcome proxies, and an adoption pathway that preserves user agency. Our framework suggests that the key design question is not "how much can AI do?" but "how can AI support the right cognitive mode at the right time?" We invite discussion on operationalizing rhythm-aware design and evaluating cognitive outcomes beyond task performance.

CCS Concepts

• **Human-centered computing** → **HCI theory, concepts and models**; *Interaction design*.

Keywords

tools for thought, cognitive rhythm, human-AI collaboration, generative AI, cognitive load, design patterns

1 Introduction

Generative AI has transformed knowledge work—writing, analysis, research, and creative tasks. Yet this transformation introduces a troubling dilemma: *the more AI does, the less humans may think*—unless tools explicitly protect cognitive agency. We focus on knowledge work contexts where sustained thinking is central to the task, though the framework may extend to other domains.

1.1 The Productivity-Cognition Tension

Current GenAI tools optimize for task completion, but this focus creates friction with cognitive development:

- **Over-reliance** leads to skill atrophy and reduced critical thinking [7]
- **Constant activation** causes cognitive fatigue and attention fragmentation [6]
- **Throughput optimization** crowds out reflection, incubation, and deep learning

The Tools for Thought (TfT) community [1, 3] has begun addressing these concerns, asking: *How do we protect cognition while using AI?* We propose a complementary reframing.

1.2 A Rhythm-Based Perspective

Rather than treating productivity and cognition as competing goals, we suggest shifting from **"how much"** to **"when and how"**:

- Human cognition naturally alternates between modes—exploration, organization, execution
- Effective tools should *support* these rhythms, not override them
- The goal is *sustainable cognitive collaboration*, not maximum output

Our critique is not that productivity is undesirable, but that many tools default to Drive-like interaction, collapsing the cognitive rhythm into a single mode.

1.3 Contributions

This position paper offers:

(1) A **three-state model** for cognitive rhythm (Diffuse / Aggregation / Drive)
(2) **Observable cues** for each state, enabling tool support without surveillance
(3) **Design patterns** aligned with the Workshop's three themes
(4) **Open questions** for operationalization and evaluation

We frame this as a *Structural Energy* perspective: cognition as a metabolic resource that must cycle between stillness, mobilization, and exertion—rather than a constant throughput engine. We intentionally present an implementation-agnostic framework; our goal is to inspire design thinking rather than prescribe specific mechanisms.

2 The Three-State Model

2.1 Core Framing

We conceptualize human-AI collaboration as a *shared cognitive rhythm system*. Users cycle through different cognitive modes; tools can either support or disrupt these cycles. When tools respect cognitive rhythm, users achieve both productivity *and* preserved cognition—not one at the expense of the other. Just as biological systems alternate between rest, activation, and exertion, cognition follows a metabolic pattern: stillness, mobilization, and drive.

2.2 Three States Defined

We propose three primary cognitive states, each serving essential functions:

Diffuse represents open, low-commitment cognition that serves two functions: exploratory intake and cognitive recovery. Users browse, capture fragments, and allow ideas to incubate—but equally, the mind clears residue from prior focused work and restores readiness for the next cycle. Activities include reading broadly, collecting references, and noting fleeting thoughts without pressure to organize.

Aggregation involves organizing, structuring, and resolving ambiguity. Users cluster related ideas, identify patterns, and build coherent frameworks from scattered materials. This is the "sense-making" phase [5] where raw inputs become structured knowledge.

Drive is the highest-energy cognitive state: sustained, deep engagement where multiple cognitive demands converge—drafting while reasoning, deciding while integrating [2]. Distractions become costly; continuity is precious. This state produces tangible outputs but consumes significant cognitive resources and requires active recovery afterward.

Critically, these states are not quality levels—Drive is not "better" than Diffuse. Each serves irreplaceable cognitive functions. Problems arise when tools force users into one mode (typically Drive) regardless of their actual cognitive needs.

Importantly, these states are not rigid phases but *fluid orientations*. Users may occupy blended states (e.g., exploring while partially structuring), switch rapidly between modes, or experience micro-oscillations within a single session. The model describes dominant orientations rather than discrete bins; tools should support graceful blending, not enforce sharp boundaries.

2.3 Recognition Cues

To move from abstract concept to actionable design, we identify observable cues for each state (Table 1). These cues enable tools to support rhythm *without* requiring invasive monitoring.

2.4 Implementation-Agnostic Stance

We deliberately avoid specifying detection algorithms, transition thresholds, or orchestration mechanisms. Our goal is to provide *design-oriented cues* that enable TfT systems to support rhythm-aware cognition. Specific implementations will vary by context, user population, and tool purpose.

3 Design Implications

We now connect the three-state model to the Workshop's three themes: Strategies, Outcomes, and Experience/Adoption.

3.1 TfT Strategies: From Features to Rhythm Patterns

Rather than asking "what features should a TfT have?", rhythm-aware design asks "how should a TfT behave differently across cognitive states?" We propose three design patterns:

Pattern 1: State Declaration. Let users explicitly signal their current mode (e.g., "Explore / Organize / Execute"). This externalizes internal rhythm as a collaborative object, reducing tool-user misalignment. The anti-pattern is tools that treat all interaction as execution mode, interrupting Diffuse exploration with premature structure.

Pattern 2: Friction Shaping. Adjust interaction friction to match cognitive state, while remaining sensitive to task context:

- *Diffuse*: Generally low friction—quick capture, weak structure, tolerance for mess. However, contexts involving critical evaluation (e.g., morally or factually sensitive material) may warrant deliberate friction even during exploration, to support reflective engagement.
- *Aggregation*: Medium friction—guided clustering, alignment prompts, disambiguation aids
- *Drive*: Environmental protection—not friction *against* the user, but friction *around* them. Reduced distraction entry points and clear boundaries protect focus; the execution path itself remains smooth.

Friction levels are thus context-sensitive defaults, not fixed prescriptions. The key is that appropriate friction *at the right time* serves both productivity and cognition.

Pattern 3: Rhythm Transitions. Explicitly support movement between states. While a common progression is Diffuse → Aggregation → Drive, real workflows are iterative: users may loop back from Drive to Diffuse when new questions arise, or oscillate between Aggregation and Diffuse as structure reveals gaps. Key transitions include:

- Diffuse ↔ Aggregation: Transform fragments into structure; return to exploration when gaps emerge
- Aggregation → Drive: Convert structure into action
- Drive → Diffuse: Support active cooldown—clearing cognitive residue and restoring readiness, rather than abruptly stopping

The anti-pattern is tools offering only "start" and "stop," leaving users stuck between half-finished work and cognitive exhaustion. Rhythm-aware tools should make *any* transition lightweight, not just forward ones.

3.1.1 Micro-Vignette. A researcher begins in **Diffuse**: she opens several papers, copies fragments into quick notes, and jumps between browser tabs. *Recognition cues*: frequent switching, fragment-heavy notes, no sustained editing. The tool responds with low-friction capture—a floating note widget, no prompts to organize.

As patterns emerge, she shifts to **Aggregation**: she begins sorting notes, grouping related ideas, and comparing

Table 1: Recognition Cues for Three Cognitive States

State	Experience Cues	Interaction Cues	Risk Cues	Tool Goal
Diffuse	Wandering yet sparked; avoids constraints; clearing from prior load	Frequent switching; fragment-heavy notes; more browsing than editing	Over-drift; low convergence; note hoarding	*Capture & restore*
Aggregation	Wants to consolidate; seeks structure; fears missing pieces	Sorting; grouping; comparing; restructuring	Over-organizing; procrastination	*Structure & align*
Drive	Wants to finish; interruption-sensitive; needs momentum	Sustained editing; fewer switches; hard boundaries	Over-drive; exhaustion; quality collapse	*Protect & recover*

claims across sources. *Recognition cues*: sorting and restructuring behavior, longer dwell time on organizational views. The tool transitions to structure-first mode—offering clustering suggestions, surfacing contradictions, and providing alignment prompts.

When ready to draft, she enters **Drive**: sustained editing in a single document, hard boundaries on notifications. *Recognition cues*: fewer tab switches, continuous text production. The tool protects focus—distractions fade, actionable next steps are highlighted, and interruptions become recoverable rather than derailing.

Midway through drafting, a gap in her argument sends her back to **Diffuse** briefly—the tool recognizes the shift and re-opens the capture interface without losing her draft context. The same tool supports all three rhythms and their non-linear transitions.

3.2 TfT Outcomes: Measurable Proxies

Traditional productivity metrics (words written, tasks completed) miss cognitive dimensions. We propose seven rhythm-aware outcome proxies:

(1) **Rhythm Alignment**: Does the tool feel like it "matches" the user's current state?
(2) **Cognitive Clarity**: Post-session clarity about task and materials (beyond completion)
(3) **Switching Cost**: Number of state transitions and time to recover from interruptions
(4) **Commitment Stability**: Fluctuation in task commitment (reduced back-and-forth)
(5) **Friction Appropriateness**: Was friction "just right"—not always smooth, not always hard?
(6) **Cognitive Debt**: Post-session fatigue or regret ("Was I led into unproductive busy-work?")
(7) **Learning Residue**: Retention and transferability of material after one week (vs. short-term output inflation)

These proxies share a key property: they are *state-relative*. Good outcomes in Diffuse (many fragments captured) differ from good outcomes in Drive (coherent output completed). Traditional metrics conflate these, potentially rewarding tools that force premature Drive at the cost of exploration depth.

Measurement hints: Rhythm Alignment and Friction Appropriateness can be captured via brief post-session Likert prompts; Switching Cost and Commitment Stability via lightweight interaction logging; Cognitive Clarity and Cognitive Debt via end-of-day diary entries; Learning Residue via delayed recall tasks one week post-session.

Punchline: Traditional metrics optimize throughput; rhythm-aware outcomes ask whether a tool optimizes **the right kind of cognition at the right time**.

3.3 Experience & Adoption: A Staged Pathway

Rhythm-aware TfT faces an adoption challenge: users accustomed to productivity-maximizing tools may resist friction or find state-switching unfamiliar. We propose a graduated pathway:

Stage 1: Self-Labeled Rhythm. Users manually select their state at session start (e.g., via a simple three-option toggle). The tool adjusts interface defaults accordingly. This requires no inference, poses no privacy risk, and keeps users in control. To minimize cognitive overhead, labels should use intuitive language ("Explore / Organize / Execute") rather than theoretical terminology, and defaults can be pre-set based on recurring patterns (e.g., always starting morning sessions in Diffuse).

Stage 2: Gentle Suggestions. Based on interaction cues, the tool *suggests* (but does not force) state transitions. Users retain override authority. This builds trust while reducing cognitive burden of self-monitoring.

Stage 3: Shared Rhythm in Collaboration. In team settings, members share current states. "I'm in Drive—please minimize interruptions" becomes an explicit coordination signal. Rhythm becomes collaboration infrastructure.

Key insight: Adoption depends less on inference accuracy and more on whether the tool preserves the user's **sense of agency** over their cognitive rhythm [4].

4 Discussion

4.1 Open Questions

We offer six questions for Workshop discussion:

(1) **Agency vs. Automation**: Should state be user-declared or tool-inferred? At what automation level does support become intrusion?
(2) **Misalignment Harms**: When tools misread state, what causes the most damage—efficiency loss, cognitive burden, or threat to user identity?
(3) **Personalization**: Cognitive rhythm varies substantially across individuals. What is the minimum viable personalization?
(4) **Overfitting to Productivity**: How do we prevent tools from over-reinforcing Drive, causing long-term cognitive debt?
(5) **Ethics & Privacy**: If behavioral cues inform state inference, which cues are acceptable? Where is the line?
(6) **Evaluation Protocol**: What is the minimal viable experiment—short lab study, longitudinal diary, or in-the-wild logging?

4.2 Limitations

This paper presents a conceptual framework requiring empirical validation. Individual differences in cognitive rhythm may be substantial, and integration with existing productivity-focused tools poses non-trivial design challenges. We view this as a starting point for research, not a finished theory.

4.3 Workshop Contribution

We bring a new theoretical lens and design patterns grounded in cognitive rhythm. We seek operationalization strategies, measurement methods, and critical feedback. We hope to learn from participants' empirical experiences with TfT design and evaluation. We look forward to discussing how participants' experiences with TfT tools align with or challenge this three-state framing.

References

[1] Andy Clark and David Chalmers. 1998. The Extended Mind. *Analysis* 58, 1 (1998), 7–19.
[2] Mihaly Csikszentmihalyi. 1990. *Flow: The Psychology of Optimal Experience*. Harper & Row, New York.
[3] Douglas C. Engelbart. 1962. Augmenting Human Intellect: A Conceptual Framework. *Summary Report AFOSR-3223* (1962), 1–89.
[4] Eric Horvitz. 1999. Principles of Mixed-Initiative User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 159–166.
[5] Peter Pirolli and Stuart Card. 2005. The Sensemaking Process and Leverage Points for Analyst Technology. In *Proceedings of International Conference on Intelligence Analysis*, Vol. 5. McLean, VA, 2–4.
[6] John Sweller. 1988. Cognitive Load During Problem Solving: Effects on Learning. *Cognitive Science* 12, 2 (1988), 257–285.
[7] Lev Tankelevitch, Advait Sarkar, Sean Rintel, et al. 2025. Tools for Thought: Research and Design for Understanding, Protecting, and Augmenting Human Cognition with Generative AI. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–6.