

Affective Tools for Thought: Towards Shared Attention and Affective Reorienting in AI-Supported Thinking

Yifu Liu
UCL Interaction Centre
London, United Kingdom
yifu.liu.22@ucl.ac.uk

Raffaele Andrea Buono
UCL Interaction Centre
London, United Kingdom
raffaele.buono.18@ucl.ac.uk

Nadia Bianchi-Berthouze
UCL Interaction Centre
London, United Kingdom
n.berthouze@ucl.ac.uk

Abstract

Current Tools for Thought (TfTs) treat affect as either friction that slows cognitive progress or a signal to optimise it. Drawing on enactive cognitive science, we argue that affect is constitutive of cognition: it reshapes the trajectory of thinking, not just the speed. We identify two core barriers for Affective TfTs—the lack of *Shared Attention* (caring, directed attention to the user’s mode of engagement) and the lack of *Affective Reorienting* (the capacity to use emotional moments to open new trajectories rather than reinforcing predetermined ones)—and propose three design strategies that address both: Chain of Emotion × Chain of Thought, Affective Mirror, and Prompted Reorienting. The strategies are grounded in empirical findings from a study of a touch-aware conversational agent for embodied craft learning, and are oriented as provocations for future design.

Keywords

Affective Computing, Tools for Thought, Enactivism, Generative AI, Learning Science

ACM Reference Format:

Yifu Liu, Raffaele Andrea Buono, and Nadia Bianchi-Berthouze. 2026. Affective Tools for Thought: Towards Shared Attention and Affective Reorienting in AI-Supported Thinking. In . ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

Current Tools for Thought (TfT)—AI systems designed to augment human cognition rather than merely automate tasks [18]—operate on an implicit assumption: that thinking proceeds along a roughly linear trajectory toward a goal. Within this framework, AI functions as an instrumental support, a tool to aid the user along *the* trajectory more effectively. The destination, and oftentimes the path itself, is a given. Within such framework (*instrumental progress model*), affect is either understood as friction and noise that slows down progress, or an informative signal to leverage to improve adherence.

This paper challenges this model, and the role of affect within it, from an enactivist standpoint [28]. Affect does not merely speed up, or slow down, thinking along a fixed path. We argue that affect,

as an enactive scaffold to cognition, has the potential to reshape paths and user journeys, beyond instrumentalist impositions which frame tasks as ready-made objectives. For instance, a frustrated learner is not just experiencing a cognitive burden to be smoothed over; similarly, an eager learner is not merely moving along a path quicker, with tools being designed to support such increase in speed. Rather, we argue that affect foregrounds affordances and potentialities [29] which radically reshape the activity space itself, and allow tasks to move along new gradients. We argue that by reframing AI as Affective Tools for Thought could support *divergent trajectories* for thinking. These emerge when affect is taken as a scaffold to cognition, reframing progress away from the execution of a single, predetermined outcome and loop.

We make two contributions. First, we articulate a theoretical grounding for why affect is constitutive of, not incidental to, thinking—and what this means for TfT design (§2). Second, we identify two core barriers for Affective TfTs (*Shared Attention* and *Affective Reorienting*), and propose three design strategies that address both, grounded in empirical findings from the study of a touch-aware conversational agent [16] (§3). The strategies are provocations for future design, with their limitations discussed in (§4).

2 Theoretical Grounding: Affect and Thinking

Research on Tools for Thought has drawn on ideas of reflective thinking to ground the design of AI systems that augment cognition [17, 18]. While informative, the affective implications behind these conceptualisations have been often set aside. For instance, Singh et al. [17] cite Dewey’s understanding of perplexity as a prerequisite for critical thinking. Perplexity is however treated as a purely cognitive state, rather than as an embodied and extended affective experience that may trigger and sustain reflective enquiry. The TfT synthesis provided by Tankelevitch et al. [18] acknowledges this gap, noting that, while most current work augments deliberate, high-level, ‘System 2’ reasoning, leveraging emotions and other ‘System 1’ processes remains an underdeveloped frontier. We argue however that such dichotomised framing of cognition and affect, System 2 and System 1, presents the two as ultimately separate channels intersecting *a posteriori*.

This assumption has shaped how existing TfT systems are designed. Zhang and Reicherts’ *ExtendAI* [21] asks users to externalise their reasoning before receiving AI feedback—a form of ‘process-oriented support’ that reads the user’s thinking rather than delivering end-to-end solutions. Cheung’s [3] ‘ignorant co-learner’ deliberately withholds knowledge to create cognitive friction, fostering moments of uncertainty, dissonance, or pause that compel users to think critically. Gmeiner et al.’s [7] metacognitive support agents proactively ask reflective questions during design tasks

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Tools for Thought Workshop, CHI26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

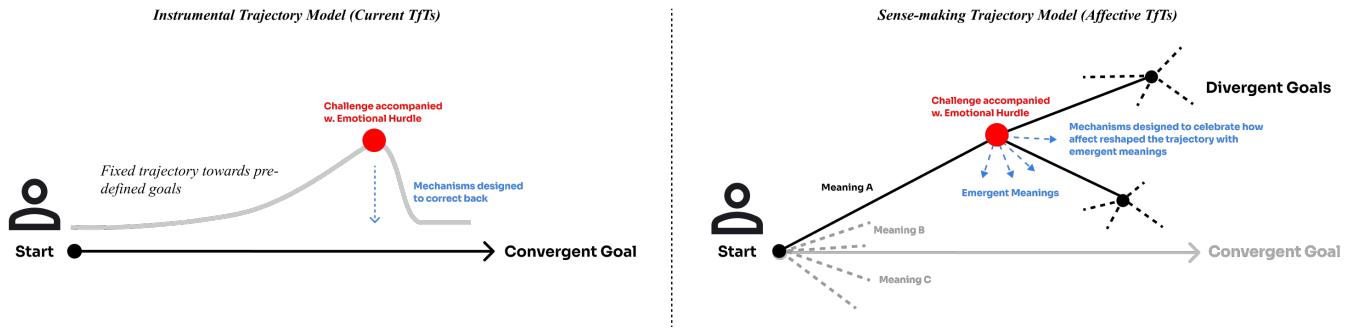


Figure 1: (Left) The Instrumental Trajectory Model common in current TfTs, where affect is treated as a hurdle to be corrected to reach a predefined goal. (Right) The proposed Sense-making Trajectory Model, where affective hurdles act as scaffolds that reshape thinking toward emergent, divergent goals.

to leverage affect. And while these all understand cognition as something AI can directly engage with, as a Tool for Thought, the studies presented all understand reasoning and cognition as the main objects of support: affect is understood as merely an external parameter, either as noise to be filtered out, or a signal to adjust to aid the user in task completion.

Three positions on affect and cognition. We see the approaches to TfT as responding to two main ways of positioning affect. Against these, we propose a third position.

Position 1—Affect as friction. The Cartesian view which frames affect and bodies as pitted against rational thinking and cognition. Here, TfT should optimise information processing, and minimise emotional disruption.

Position 2—Affect as instrument. This view, epitomised by most affect automatic recognition systems, treats emotions as detectable and tractable signals that can be leveraged to optimise performance by building semi-adaptable systems [4, 11]. A system here detects frustration and reduces task difficulty; it designs *around* affect, rather than *through* it.

We see both positions as *instrumentalising* affect, as subservient to cognition. We thus argue that both positions frame a system of individual, self-contained cognizers, with AI functioning as an optimising tool for adherence to predefined paths of reasoning. Cognitive reasoning follows linear, individual trajectories which are either hampered, or supported, by affect as an externality.

Position 3—Affect as constitutive scaffold. We argue that affective modes of engagement are constitutive of thinking, thus opening up ways to rethink tasks and goals. We argue that affect is not a variable within thinking, or an externality to plans: rather, it is the very embodied and situated resource through which worlds are brought about and provisionally constituted. In this sense, we draw on the enactivist idea of sense-making, as an agent’s embodied and affectively-charged activity of bringing forth a world of significance through their ongoing coupling with an environment and other agents [20, 22, 23]. Affect scaffolds embodiment and cognition [24]: it produces attentional orientations, interpretive possibilities, significant potentialities, multiple relationships to a world and the tasks emergent from it. Affect thus does not simply *influence* how a task is performed, but actively *shapes* what can count as a task, what can be perceived as relevant, and what kinds of actions

become intelligible or viable in a given situation. It is through affective attunement that certain features of the environment come to matter at all, and that particular problem framings, goals, and criteria of success take form.

From instrumental progress to sense-making. We then propose a rethinking of TfT away from the ideal of instrumental progress, and instead fostering affect-based multiplicity. We do this by arguing that affect is not merely a factor ‘in’ cognition: instead, it is a central engine to cognising, as the capacity to construct worlds through sense-making. Here, agents (humans and AI itself) and the world are not pre-given entities externally interacting, but they are rather mutually constituted through engagements [19] which intrinsically link them, based on forms of shared attention and care—i.e., intersubjectivity [25, 26]. From this premise, we conclude that tasks themselves, their trajectories and benchmarks for success, are emergent through affect and affect-driven sense-making. When a learner’s frustration transforms the meaning of a task from “an interesting challenge” into “evidence of my inadequacy,” this opens up a space to precisely support the user where the task could lead them, rethinking the entire landscape of what is cognitively possible. In our empirical study [16], we observed this directly: chatbot-led reflections did not simply restore participants to their original task trajectory. Rather, affect re-pivoted thoughts. During guided reflection, one participant arrived at the experience of repair as “almost meditative” (P12); another, prompted to reconsider tactile assumptions, discovering unexpected relationships with the material (P15); a third, when helped by the AI to notice they were rushing, began to reconceive the activity as something to attend to, rather than merely complete (P6). In each case, affect-aware interaction opened a new trajectory, rather than correcting a deviation. This reframes the role of a TfT from an instrument that optimises progress and productivity along predefined trajectories, to a participant in sense-making that can open, rather than close down, alternative ways of understanding the task, and oneself within it.

From these premises, we start drawing out some preliminary design stakes and implications: (1) TfTs that assume affect to be a constitutive scaffold should treat emotional moments—frustration, confusion, surprise, doubt—not as obstacles to overcome, but as potential pivot points where the meaning of the task might productively shift, and with it, perhaps, the task itself and the world

within which the task comes to be configured; (2) TtTs as relational systems that support shared attention and concern, treating interaction not as the exchange of inputs and outputs but as a process through which both the user and the system can reshape how the task is understood. In this view, meaningful support arises from sustaining intersubjective engagement—moments of noticing, pausing, and reorienting together—through which divergent trajectories of thinking can become possible.

3 Framework: Affective Tools for Thought

The framework has two components. First, we identify two core barriers that current TtTs do not address: *Shared Attention* and *Affective Reorienting* (§3.1). Second, we propose three design strategies, Chain of Emotion × Chain of Thought, Affective Mirror, and Prompted Reorienting, each of which addresses both barriers differently (§3.2, Figure 5). We discuss their limitations in §4.

Our framework builds on the theoretical foundations drawn in (§2) by specifically designing for “care”, and against “Loop-in-the-Human”. The first idea comes from Heidegger’s concept of *Sorge* [27]: here, attention is not merely a form of detection, monitoring and response, as current affect recognition systems provide. Rather, care is a form of intersubjective orientation: carer and cared for interlock each other, and transform each other towards the constitution of a common trajectory. For instance, a human tutor who notices a student struggling does not merely register the struggle and adjust parameters; they attend with concern, drawing on their own experience to pose questions that are personally authentic and oriented toward who this person is becoming. This caring attention is constitutive of a kind of support that enables collective (re)making, and is precisely what current AI systems lack.

We borrow the second idea from the work of Gambarotto and colleagues to argue that current Human-in-the-Loop strategies reduce human purposiveness to quantifiable costs and benefits. The result is that the “Human-in-the-Loop increasingly risks becoming a Loop-in-the-Human” [6]: a configuration where human thought is shaped to resemble the algorithmic systems it was meant to oversee. In TtTs, this means systems designed to support thinking may constrain it, channelling users along predetermined trajectories and treating deviations as errors.

3.1 Affective TtT Barriers

The identified barriers builds on the scenarios and empirical findings from our study of human-AI interaction in embodied craft learning [16], a Wizard-of-Oz system using a touch-aware chatbot to support participants repairing garments.

Barriers 1: The Lack of Shared Attention. By *Shared Attention* we mean not merely that the system monitors the user, but that it attends to the user’s *affective mode of engagement*—the way they are relating to the task emotionally in the moment. This involves three components: (i) *noticing how* the person is engaging, not just *what* they are doing; (ii) *interpreting* that engagement within the developing trajectory—reading the present moment in light of what has come before; and (iii) *making that attention felt*—communicating, through the quality of its responses, that the person’s mode of engagement matters.

Instead of capturing or sensing the entire context, Participants described a human tutor as someone who not only “sense the critical points” and offer proactive support, but “makes the action feel seen”(P4)[16]. This names the quality of *Sorge*: a directed, caring attention that differs from tracking or monitoring. Current TtTs provide sophisticated cognitive support, but not this quality of **shared** attention: for instance, current TtTs *track behaviour*, such as keystrokes, reasoning patterns, task completion [21] [3] [7]—without attending to how the person *relates* to what they are doing. Without Shared Attention, the system operates on a model of cognition that is disconnected from the person’s lived experience: it sees the task but not the person doing it.

Barriers 2: The Lack of Affective Reorienting. By *Affective Reorienting* we mean the capacity to use moments of emotional intensity—frustration, resistance, self-doubt, but also unexpected delight or curiosity—as pivot points where the person’s relationship to the task might shift. The key distinction is between loop-reinforcing responses, which treat affect as incidental and adjust parameters within the existing trajectory (e.g., reducing difficulty when frustration is detected), and reorienting responses, which attend to what the emotional moment reveals about how the person relates to the task. Current TtTs lack this second capacity.

In our study [16], participants experienced self-reinforcing cascades where frustration produced self-doubt, self-doubt produced rushing, and rushing produced more errors—a spiral that led some to disengage entirely, reaching what they described as “saturation points” or declaring themselves “not gifted at this.” This was not merely an emotional event, but a shift in what the task meant for the person: from “can I learn this?” to “am I the kind of person who can do this?”. Our chatbot responses remained loop-reinforcing: when P7 expressed resistance, we provided easier tasks—adjusting difficulty but preserving the same evaluative frame. A **reorienting response** would have treated resistance not as an obstacle but as an opportunity, inviting a reconsideration of what the difficulty signified, and what it had opened up.

More broadly, current proactive support mechanisms [10] demonstrate the value of anticipatory assistance, but operate within the existing trajectory, smoothing the path before difficulty arrives. They make the loop more efficient; they do not question whether the loop is the right one. What is missing is the capacity to recognise when emotional intensity signals that the trajectory itself needs to change—and to stay with the person at that moment rather than resolving it. Without this capacity, the system risks becoming what Gambarotto et al. call a Loop-in-the-Human [6].

3.2 Affective TtT Strategies

Affective Mirror. The Affective Mirror directs the system’s reading of the user’s engagement back to the user, creating occasions for affective self-awareness. Rather than the system acting on what it senses, it makes the user’s own mode of engagement visible to them, not as a diagnostic label, but as material for reflection.

We arrive at this strategy through empirical observation. In our study, moments where the system reflected aspects of engagement back to participants, even imperfectly, generated unexpected self-awareness. P7 found incorrect chatbot feedback useful because it prompted self-examination. When our touch-aware system offered

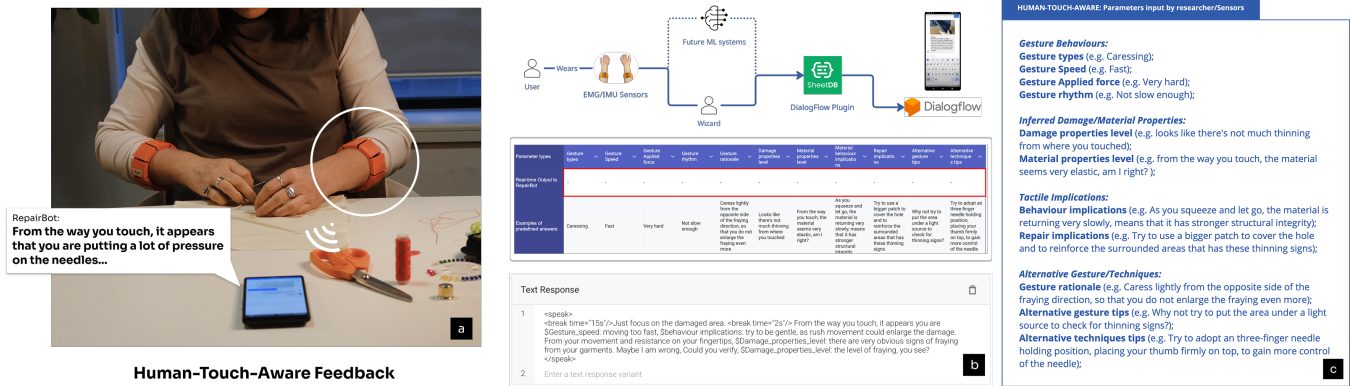


Figure 2: The RepairBot Feedback System; an overview of touch-aware interaction where the system analyzes gesture behaviors to provide inferred material properties and alternative techniques as affective mirrors.

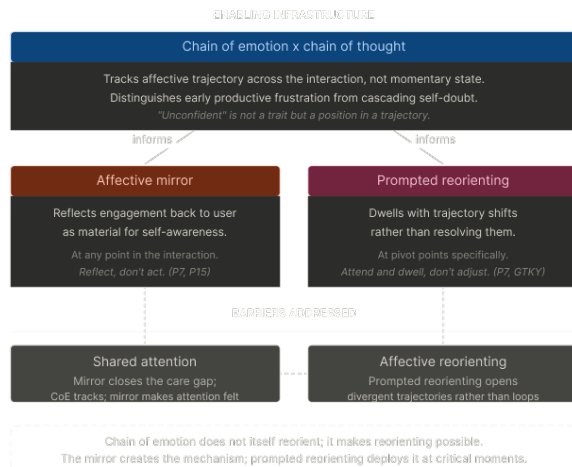


Figure 3: The Affective TfT Framework; mapping the two core identified barriers to three proposed design strategies.




Figure 4: Shared Attention in Context; comparing the inter-subjective orientation of a Human Tutor (A) and the Chatbot (B).

an inference about P15's material interaction, the encounter led P15 to re-examine her own tactile perception and discover an unexpected relationship with the material. This happened not because the system was right, but because the reflection opened a new way

of attending. Through facilitated noticing, P12 arrived at the experience of repair as "a very calm, almost meditative experience," and P6 came to see they were "rushing ahead instead of enjoying the process." None of these outcomes were pre-designed; they emerged through the encounter between system and person.

These observations align with existing work: research on body movement and touch has demonstrated that embodied behaviour expresses not merely task actions but affective engagement; the body's involvement shapes experience itself [2]. Yet users are often unaware of what their engagement reveals. Höök's *soma design* [9] argues that technology can cultivate this bodily self-awareness, and Sanches et al. find that *ambiguous* bio-feedback prompts more reflection than precise diagnostic labels [15]. For future TfTs, the Affective Mirror implies a specific design orientation: when the system senses a shift in engagement, its first response should be to *reflect* rather than *act*, surfacing what it has noticed in open-ended, non-diagnostic language ("I notice you've slowed down and are handling the material differently") rather than classifying an emotion or adjusting a parameter. This addresses Shared Attention by closing the care gap: if the system cannot attend with genuine concern, it can help the *user attend to themselves*. It also creates the conditions for Affective Reorienting. As P12's and P6's shifts suggest, the act of seeing one's own engagement can itself become the moment where the meaning of the task changes.

Chain of Emotion × Chain of Thought. Chain of Emotion proposes that the system's reasoning process (its Chain of Thought [12]) must be informed by a temporally-structured model of the user's affective engagement, a Chain of Emotion, where each transition is shaped by the preceding one. Rather than classifying a momentary emotional state, the system tracks how the person's mode of engagement has been evolving, and lets that trajectory inform what kind of response is appropriate.

In our study, the participants who disengaged did not simply "become frustrated." They moved through a trajectory: curiosity gave way to confusion, confusion to frustration, frustration to self-doubt, and self-doubt to withdrawal. This aligns with D'Mello and Graesser's model identifying confusion-to-frustration transitions as

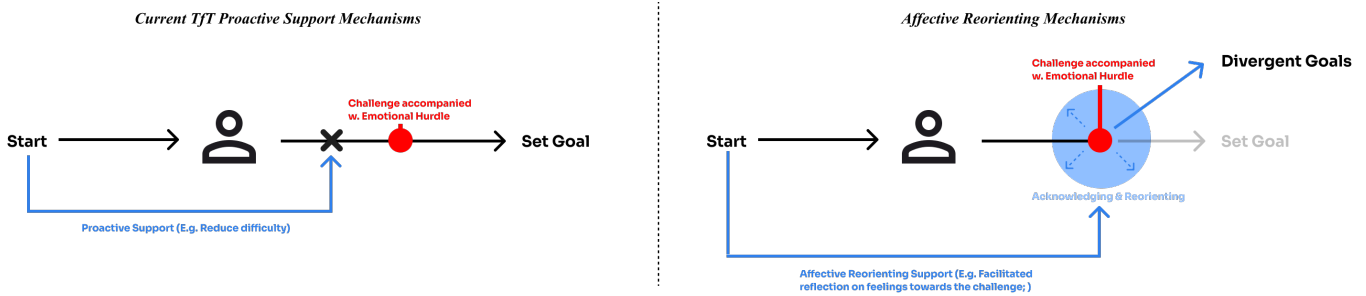


Figure 5: (Left) Current proactive mechanisms that reduce difficulty to keep users in a "loop". (Right) Prompted Reorienting mechanisms that dwell with the person at pivot points, allowing the meaning of the task to shift productively.

critical junctures in learning [5]. Crucially, the same emotion, frustration, called for different responses depending on where it sat in the trajectory. Early frustration during exploration signalled productive challenge; frustration following repeated failure and self-doubt signalled a cascade toward disengagement. Without tracking the trajectory, the system cannot distinguish between the two. Prather and Reeves observe that “underprepared, unconfident, and underperforming students seem to benefit the least” from AI tools [14]. Chain of Emotion reframes this: “*unconfident*” is not a trait but a position in a trajectory.

For future TfTs, this strategy implies that the system’s reasoning architecture must maintain an evolving model of affective engagement across the interaction, not a snapshot of the current state. The question shifts from “the user is frustrated, reduce difficulty” (Position 2) to “what role is this person’s frustration playing in their thinking, and what kind of attention does it call for?” For Shared Attention, this is the mechanism by which the system approximates *attuning*: reading the present moment within the developing trajectory. For Affective Reorienting, it is the enabling infrastructure. Without it, the system cannot recognise when the emotional trajectory signals that the loop itself needs to change. Chain of Emotion does not itself reorient; it makes reorienting possible.

Prompted Reorienting. Where the Affective Mirror reflects engagement at any point in the interaction, Prompted Reorienting acts at a specific juncture: when the Chain of Emotion signals a trajectory shift, the system responds not by adjusting parameters but by *dwelling with the shift itself*. The distinction is between resolving an emotional moment (reducing difficulty, offering encouragement) and treating it as a pivot point where the person’s relationship to the task might change. This contrasts with both reactive and proactive support [7, 10], which operate within the existing trajectory. Both make the loop more efficient; neither questions whether the loop is the right one.

We arrive at this strategy based on empirical findings. When P7 expressed resistance, we scaffolded the difficulty, a loop-reinforcing response that preserved the evaluative frame. A reorienting response would have stayed with the resistance: *what does this difficulty mean to you? Could it be approached as exploration rather than test?* Conversely, the “Getting to Know You” module asked participants to reflect on a garment’s emotional history before technical work. We designed the prompt, but not what followed: P12

arrived at repair as “a very calm, almost meditative experience,” and P6 reconceived their relationship to the activity entirely. What these moments shared was that something shifted in the person’s engagement, and the interaction did not smooth past it.

For future TfTs, Prompted Reorienting implies a design rule: at moments where the Chain of Emotion registers a trajectory shift, the system’s default should be to *attend and dwell* rather than *adjust and resolve*. In practice, this means posing questions that make the shift itself visible (“You were experimenting freely a moment ago, something seems to have changed”) rather than acting on the inferred emotion. Ayyappan and Joyner’s mode-switching [1], which lets students toggle between “tell me the answer” and “ask me a question,” gestures toward giving the user agency at such moments, but assumes the task’s meaning is stable. Prompted Reorienting recognises that the meaning itself might shift, and that the system’s role is not to engineer that shift, but to refuse to foreclose it.

4 Discussion

How the two barriers interact. Without Shared Attention, without the system attending to the person’s affective engagement, the system cannot recognise when Affective Reorienting is called for. And without Reorienting, even the most attentive system will reinforce the loop: reading engagement accurately but responding only within the existing trajectory. The two barriers need each other’s solutions. Chain of Emotion provides the temporal model that makes Shared Attention possible; the Affective Mirror creates the reflective mechanism through which attention becomes felt; and Prompted Reorienting deploys that mechanism at the moments where it matters most. None functions without the others.

Divergent pathways in constrained scenarios. A question of scope deserves direct address: how should an Affective TfT behave in fixed-goal domains, such as mathematics tutoring, where success criteria are stable, and “divergence” risks becoming distraction? Our answer is that the divergence we propose does not operate at the level of the task’s *epistemic content*: the correct answer remains correct. It operates at the level of the learner’s *relationship* to the task, and that relationship is never fully determined by the task’s success criteria alone. In practice, a fixed-goal Affective TfT would not question whether a student should be learning multiplication.

It would instead recognise, through Chain of Emotion, when frustration has shifted from “I don’t understand this yet” to “I am not capable of understanding this.” A TfT tool could open a different trajectory: naming that shift, reflecting it back through the Affective Mirror, and reorienting the student’s relationship to difficulty itself. Two students who both eventually reach the correct answer are not, in this sense, on equivalent paths. The degree of reorientation a system enables remains a design decision, appropriately constrained in fixed-goal contexts; but the underlying barriers, Shared Attention and Affective Reorienting, remain live across both.

Rethinking evaluation. Standardised evaluation assumes that success criteria can be specified in advance. Affective TfTs are designed for situations where the learner’s relationship to those criteria is itself what is at stake, unsettling linear evaluation pathways. Merely measuring task performance, goal completion or adherence would miss the point: it would treat as settled exactly what the framework holds open. In lieu of linear measurable outcomes, we argue that what may come under scrutiny is whether a productive, generative reorienting has occurred. For instance, trajectory divergence (whether a person’s framing of the task shifts across an interaction) offers a direct index of this, assessable through pre/post interviews or discourse analysis. Affective transition patterns could track whether interventions interrupt negative cascades without resolving them into task completion, a meaningful signal precisely because it decouples affective shift from performance outcome. Self-reported relationship to difficulty and inter-rater reliability on identified pivot moments in transcripts are both methodologically credible moves in the same direction. What these measures share is that they index change *in orientation*, rather than arrival at a destination. This shift in evaluative logic matters, because it holds evaluation itself within the framework the paper proposes: an Affective TfT that is only legible through task performance has already conceded the ground it set out to contest.

The boundary with affective computing. Our framework draws on sensing and inference mechanisms developed within the affective computing tradition; what changes is not the sensing but what the system does with what it senses. Classic affective computing detects an emotion and responds to it [11]. Höök’s affective loop framework [30] goes further, treating emotion as co-constructed in interaction and keeping interpretive agency with the user. Our framework extends this lineage by arguing that affect does not merely shape the interaction but can reshape what the task *is*. The practical boundary is less clean than the ontological one: an implemented Affective TfT would still classify engagement patterns and make inferences. The difference lies in what follows. Where an adaptive system uses inference to adjust parameters within a stable frame, an Affective TfT uses inference to recognise when the frame itself may need to shift, and then stays with that moment rather than resolving it.

The authenticity gap. A human tutor draws on personal experience to pose questions with individual authenticity. Chain of Emotion may restructure the system’s computational priorities to emulate the *structure* of concern, attending to what matters to the person, tracking what is at stake. But the question remains whether restructuring priorities is the same as genuine care. Participants

valued the chatbot’s non-judgmental quality [3, 13], a genuine affordance. But we must acknowledge what is lost: the authentic collision of different lived perspectives, the way a tutor’s genuine surprise opens a space no algorithm can replicate. This tension, between AI’s non-judgmental consistency and the irreplaceable generativity of authentic encounter, is not a problem to solve but a design space to navigate.

Limitations. Our empirical grounding comes from a single study in embodied craft learning with a Wizard-of-Oz system: the attuning was performed by a human researcher, not an automated system. Transferability to text-based TfTs is theoretically motivated but not demonstrated. The Affective Mirror cannot substitute for encountering a genuinely different perspective. Prompted Reorienting is the most speculative strategy: we observed pivots emerging incidentally and acknowledge that designing systems to stay with emergent moments, without becoming prescriptive, is a significant challenge. The design principles we propose are derived from retrospective analysis of our empirical data; they have not yet been implemented or evaluated in an automated system. Cross-session Chain of Emotion modelling raises questions about data persistence, privacy, and reductive emotional profiling.

5 Conclusion

We have identified two core barriers for Affective Tools for Thought: the lack of Shared Attention, caring, directed attention to the user’s mode of engagement, and the lack of Affective Reorienting, the capacity to use emotional moments to open new trajectories rather than reinforcing predetermined ones. We proposed three design strategies: Chain of Emotion as enabling infrastructure that tracks affective trajectory across interactions; the Affective Mirror, which reflects engagement back to users as material for self-awareness; and Prompted Reorienting, which dwells with trajectory shifts rather than resolving them. The strategies are grounded in empirical findings from a touch-aware conversational agent for embodied craft learning, and are oriented as provocations for future design rather than implemented solutions. We invite TfT researchers to consider: what would our tools look like if they were designed not to keep users on track, but to stay with them when something unexpected emerges?

References

- [1] D. Ayyappan and D. A. Joyner. Building Self-Reliance in an Era of Accessible Intelligent Tutors. In *CHI’25 Workshop on Tools for Thought*, 2025.
- [2] N. Bianchi-Berthouze. The Affective Body Argument in Technology Design. Keynote, *AVI 2016 Int. Conf. on Advanced Visual Interfaces*. ACM, 2016.
- [3] C. Cheung. The Ignorant Co-Learner: AI-as-Novice for Deeper Human Engagement. In *CHI’25 Workshop on Tools for Thought*, 2025.
- [4] A. R. Damasio. *Descartes’ Error: Emotion, Reason, and the Human Brain*. Putnam, 1994.
- [5] S. D’Mello and A. Graesser. Dynamics of Affective States during Complex Learning. *Learning and Instruction*, 22(2):145–157, 2012.
- [6] A. Gambarotto, K. Potapov, and C. B. Sachs. Design for Autonomy: An Enactive Approach to Human-Computer Interaction. *Adaptive Behavior*, 2025.
- [7] F. Gmeiner, K. Patel, S. K. Sathe, J. Forlizzi, and K. Hollan. Exploring the Role of AI-Based Metacognitive Support in Design. In *Proc. CHI 2023*. ACM, 2023.
- [8] M. Heidegger. *Being and Time*. 1927. Trans. J. Macquarrie and E. Robinson, Harper & Row, 1962.
- [9] K. Höök. *Designing with the Body: Somaesthetic Interaction Design*. MIT Press, 2018.
- [10] V. Liu et al. Proactive AI with Shared Cognitive Space. In *CHI’25 Workshop on Tools for Thought*, 2025.

- [11] R. W. Picard. *Affective Computing*. MIT Press, 1997.
- [12] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. Chi, Q. Le, and D. Zhou. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837, 2022.
- [13] K. Pilcher and B. Tütüncü. Productive Hallucinations: AI Inaccuracies as Creative Provocations. In *CHI'25 Workshop on Tools for Thought*, 2025.
- [14] J. Prather and B. N. Reeves. LLMs in CS Education: Who Benefits and Who Doesn't. In *CHI'25 Workshop on Tools for Thought*, 2025.
- [15] P. Sanches et al. HCI and Affective Health: Taking Stock of a Decade of Studies and Charting Future Research Directions. In *Proc. CHI 2019*. ACM, 2019.
- [16] Y. Liu, T. Bi, C. Yu, L. F. Hernandez, B. B. Petreca, M. Nygren, S. Baurley, Y. Cho, and N. Berthouze. The RepairBot Framework: Touch-Aware Conversational Agent for Hands-on Clothes Repair. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 25 pages. <https://doi.org/10.1145/3772318.3791705>
- [17] H. Singh et al. Where is the Mind in ChatGPT? Critical Thinking and AI. *arXiv preprint*, 2023.
- [18] L. Tankelevitch et al. Understanding, Protecting, and Augmenting Human Cognition with Generative AI: A Synthesis of the CHI 2025 Tools for Thought Workshop. *arXiv:2508.21036*, 2025.
- [19] E. Thompson. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press, 2007.
- [20] F. J. Varela, E. Thompson, and E. Rosch. *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press, 1991.
- [21] T. Zhang and L. Reicherts. ExtendAI: Process-Oriented Support for Human Thinking. In *CHI'25 Workshop on Tools for Thought*, 2025.
- [22] E. A. Di Paolo, M. Rohde, and H. De Jaegher. Horizons for the Enactive Mind: Values, Social Interaction, and Play. In J. Stewart, O. Gapenne, and E. A. Di Paolo, editors, *Toward a New Paradigm for Cognitive Science*, pages 33–87. MIT Press, 2010.
- [23] E. A. Di Paolo and H. De Jaegher. The Interactive Brain Hypothesis. *Frontiers in Human Neuroscience*, 6:163, 2012.
- [24] S. Gallagher. *Enactivist Interventions: Rethinking the Mind*. Oxford University Press, 2017.
- [25] M. McGann. Enacting a Social Ecology: Radically Embodied Intersubjectivity. *Frontiers in Psychology*, 5:1321, 2014.
- [26] D. D. Hutto and E. Myin. *Evolving Enactivism: Basic Minds Meet Content*. MIT Press, 2017.
- [27] M. J. Elley-Brown and J. K. Pringle. Sorge, Heideggerian Ethic of Care: Creating More Caring Organizations. *Journal of Business Ethics*, 168(1):23–35, 2021.
- [28] A. Hila. An Enactivist Approach to Human-Computer Interaction: Bridging the Gap Between Human Agency and Affordances. In M. Kurosu and A. Hashizume, editors, *HCI International 2025—Late Breaking Papers*, LNCS vol. 16331. Springer, 2026.
- [29] M. O. Nygren, S. Price, and R. Thomas Jha. The Role of Embodied Scaffolding in Revealing “Enactive Potentialities” in Intergenerational Science Exploration. *Science Education*, 108(2):420–448, 2024.
- [30] K. Höök. Affective Loop Experiences: Designing for Interactional Embodiment. *Philosophical Transactions of the Royal Society B*, 364(1535):3585–3595, 2009.
- [31] G. Colombetti. *The Feeling Body: Affective Science Meets the Enactive Mind*. MIT Press, 2014.
- [32] K. Boehner, R. DePaula, P. Dourish, and P. Sengers. How Emotion is Made and Measured. *International Journal of Human-Computer Studies*, 65(4):275–291, 2007.