Title

An attention-to-thoughts model of stream of thought: Affective dynamics depends on attentional scope

Permalink

https://escholarship.org/uc/item/3895j8cz

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors

Arya, Pratyush

Srinivasan, Narayanan

Publication Date

2026

An attention-to-thoughts model of stream of thought: Affective dynamics depends on attentional scope

Pratyush (parya22@iitk.ac.in)[1] **& Narayanan Srinivasan (nsrini@iitk.ac.in)**[2]
[1,2]Department of Cognitive Science, Indian Institute of Technology, Kanpur

Abstract

The stream of thought has received less empirical attention than external cognition, with existing approaches relying on thought probes or continuous verbal reports rather than modeling internal representational dynamics. The Attention-to-Thoughts (Amir & Bernstein, 2022) model offers a dynamical-systems framework in which thought trajectories emerge from moment-to-moment interactions among lower-level components, capturing several thought patterns, including those consistent with repetitive negative thinking (RNT). However, the model uses simplified thought representations and does not capture the breadth of internal attentional selection: a dimension extensively studied in external attention and implicated in internal attention, too, in studies of creativity (broad) and rumination (narrow) for instance. We extend A2T by grounding thoughts in semantic embeddings and affective ratings, and introducing an attentional scope parameter modulating internal selection breadth. Simulations reveal systematic affective and dynamic differences between thought trajectories in broad and narrow attention runs. Control analyses confirm these effects arise from model dynamics rather than embedding structure alone. This gives us an empirically testable computational basis for understanding how the scope of internal attention shapes affective experience during spontaneous thought streams.

Keywords: stream of thought; scope of attention; internal attention; semantic networks; affect dynamics; dynamics of thought

Human experience is dynamically populated by varied contents, and whatever one attends to becomes the focus of the conscious field (Watzl, 2017). This focus fluctuates over time and may be directed outward toward perceptual input (Corbetta & Shulman, 2002) or inward toward self-generated, stimulus-independent representations (Andrews-Hanna et al., 2014; Smallwood, 2013). Cognitive science has produced detailed accounts of external attention, from early filtering frameworks (Broadbent, 1958) to neural models of selective processing (Desimone & Duncan, 1995), and more recently has extended this work to internal attention over mental representations (Chun et al., 2011; Kiyonaga & Egner, 2013). However, most such work centers on working memory selection, with comparatively little formal treatment of attention over the stream of thought itself, as it flows dynamically (Christoff et al., 2016). This gap is notable because stimulus-independent thought occupies a substantial portion of waking cognition (Killingsworth & Gilbert, 2010) and may support functions such as memory consolidation and goal maintenance (Mildner & Tamir, 2024). Inner speech is one prominent mode of this stream and contributes to problem solving and emotional and self regulation (Alderson-Day & Fernyhough, 2015). Internally directed cognition is not always adaptive, however. Repetitive negative thinking is a maladaptive attentional pattern implicated across mental disorders (Ehring & Watkins, 2008; Nolen-Hoeksema et al., 2008). Explaining why thought sometimes remains flexible and sometimes becomes stuck in negative cycles calls for dynamical models of attention to thought. The Attention-to-Thoughts (A2T) model (Amir & Bernstein, 2022) provides such a framework, linking internal attention, working memory, and affect, but represents thought content only in coarse valence terms and does not model semantic structure or scope of internal selection.

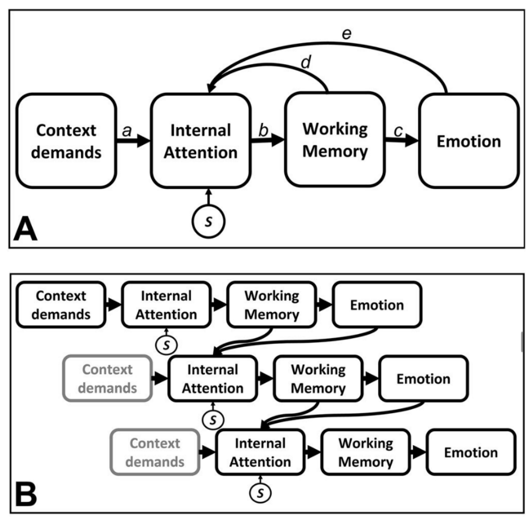

Figure 1: Overview of the A2T model, adapted from Amir and Bernstein (2022). *S*: stochastic effects on selection. (A) Momentary state: path *a* reflects cognitive control, path *c* emotional reactivity, path *e* cognitive reactivity. (B) Temporal unfolding of these interactions.

Complementary theory and research has emphasized attentional scope, the breadth of information available for selection. Narrowed scope has been linked to rumination (Whitmer & Gotlib, 2013), whereas positive emotion is associated with broadened attentional and cognitive scope (Fredrickson, 2001). Attentional breadth is well studied in perception (Eriksen & St. James, 1986; LaBerge, 1983) and has been shown to influence preference, memory, decision making, and creative or divergent thinking (Srinivasan & Gupta, 2011; Srinivasan et

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

al., 2013; Wronska et al., 2018). Without semantic relations among thought units, however, internal attentional breadth cannot be meaningfully implemented. We offer a realistic extension by grounding thought representations in continuous semantic space using word embeddings (Pennington et al., 2014) and affective ratings from the NRC–VAD lexicon (Mohammad, 2018). This enables an attentional scope parameter (f) that operationalizes constricted versus free-flowing internal selection as a semantic search radius. Across in-silico simulations, the modified model reproduces core A2T trends and reveals additional scope-dependent effects on thought trajectories. We first summarize the original A2T dynamics, then describe the modifications and resulting observations.

The A2T model

The Attention-to-Thoughts (A2T) model (Amir & Bernstein, 2022) is a dynamical-systems framework in which momentary attentional selection iteratively transacts with working-memory contents and affect under varying external contextual demands. Higher-order phenomena such as repetitive negative thinking (RNT) emerge from these lower-level cognitive-affective interactions rather than requiring dedicated mechanisms.

Operating at a resolution of seconds, A2T comprises four components linked by feedback loops (Figure 1): (1) *Context*, the degree to which the environment demands focused, goal-directed attention; (2) *Internal Attentional Selection*, the momentary probability of selecting a negative versus neutral representation; (3) *Working Memory*, a buffer holding the last five selected items; and (4) *Affect*, the momentary level of negative affect.

Panel A shows how these components causally interact at each moment. Feedforward paths include path *a*, where contextual demands attenuate affect-driven biases on selection; path *b*, where selection gates entry into working memory; and path *c* (*emotional reactivity*), the degree to which negative WM contents trigger negative affect. Feedback closes the loop via path *d*, where WM contents bias subsequent selection toward content-congruent representations, and path *e* (*cognitive reactivity*), where negative affect biases selection toward affect-congruent thoughts.

The strengths of paths *c* and *e*, parameterized by r_w and r_a respectively, are the only path strengths that vary between individuals, while other path strengths are fit at the sample level. Together these capture differences in vulnerability to maladaptive thought patterns. A stochastic element (S) prevents the system from becoming permanently locked in self-reinforcing loops, enabling spontaneous exits from negative spirals. Panel B shows how these interactions unfold in time: each state at t is determined by $t–1$ and shapes $t+1$, yielding trajectories that exhibit higher-order patterns such as RNT.

The central mechanism is the selection bias parameter τ, which sets the likelihood of selecting negative versus neutral content at each timestep:

$$\tau = \frac{\mathrm{WM}_{\mathrm{contents}} + (\mathrm{Affect} \times r_a)}{2 + \mathrm{Context}^2} \qquad (1)$$

where $\mathrm{WM}_{\mathrm{contents}}$ is the proportion of negative items in working memory, Affect is current negative affect, and r_a (*reactivity to affect*) captures how strongly emotional state influences selection. Squaring Context reflects the idea that high external demand suppresses internal influences on attention.

Selection at each timestep is modeled using an ex-Gaussian distribution:

$$P(\text{select negative}) \sim \mathrm{ExGaussian}(\mu, \sigma, \tau) \qquad (2)$$

where μ is the baseline proportion of task-relevant negative content, σ adds stochastic variability, and τ is the bias from Equation 1. Affect is updated from current WM contents:

$$\mathrm{Affect} = r_w \times \mathrm{WM}_{\mathrm{representation}} \qquad (3)$$

where r_w (*reactivity to WM contents*) captures how strongly the valence of WM thoughts influences momentary affect. Together, r_a and r_w capture stable individual differences in cognitive-affective reactivity that may predispose individuals to RNT.

In-silico simulations of the coupled attention–WM–affect system are run with reactivity parameters (r_w, r_a) set to moderate values (0.25), initial selection bias at 0.5, moderate initial affect, and contextual demand manipulated across high ($C = 4$) and low ($C = 2$) regimes. Each run spans 300 timesteps; at each step, τ is computed from prior WM content, prior affect, reactivity, and current context, and a representation is sampled via the ex-Gaussian mechanism, with trajectories of selected representations, affect, and selection bias recorded throughout. High contextual demand inflates the denominator in Equation 1, dampening internal influences and stabilizing task-focused attention; under low demand, internal factors dominate, and once negative content enters WM, feedback between affect and selection bias can produce sustained negative sequences. The resulting patterns qualitatively match real-world observations of thought dynamics under analogous conditions; we discuss simulation results for both the original A2T and our modifications in the later sections.

Limitations of the model

While the A2T model captures several empirical features of RNT, it represents thought using a binary valence distinction (negative vs. neutral). This simplification is computationally convenient but omits the graded and structured nature of mental content. Real thoughts vary continuously in affect and occupy a structured semantic space with associative links to related concepts (Collins & Loftus, 1975; Kumar, 2021).

The original model also lacks a mechanism for *attentional scope* during internal selection. Work on attentional breadth shows that attention can operate in narrow or broad modes with distinct cognitive and emotional consequences (Fredrickson, 2001; Whitmer & Gotlib, 2013). Broad scope has been linked to remote association and creativity (Fredrickson & Branigan, 2005; Mednick, 1962; Wronska et al., 2018), whereas rumination is associated with chronically narrowed scope that restricts available thought content (Whitmer & Gotlib, 2013). Without an explicit scope parameter, A2T cannot directly model these differences in semantic exploration and perseveration.

Semantic grounding and attentional scope implementation

We extend the A2T model in two main ways: (1) grounding thought representations in a continuous semantic space using word embeddings, and (2) introducing an attentional scope parameter (in the model, represented as f for focus) that controls the breadth of semantic search during thought selection.

Semantic grounding We replace binary thought representations with continuous semantic vectors and graded affective values. Thoughts are represented as 300-dimensional GloVe embeddings (Pennington et al., 2014), which capture distributional semantic structure, paired with continuous valence ratings from the NRC–VAD lexicon (Mohammad, 2018). Each thought t_i is thus an embedding vector $\mathbf{v}_i \in \mathbb{R}^{300}$ with valence score $v_i \in [-1, 1]$ and a derived categorical label (negative, neutral, positive). Working memory operates over these graded representations, and momentary affect is proportional to the mean valence of WM contents:

$$\text{Affect} = r_w \times \frac{1}{|\text{WM}|} \sum_{t_i \in \text{WM}} v_i \tag{4}$$

allowing affect to vary continuously rather than by binary class. Here $|\text{WM}|$ denotes the number of items currently in working memory (up to 5), so the summation computes the mean valence of WM contents.

Semantic proximity and selection Thought transitions are driven by semantic proximity. Cosine similarity between the current thought and each candidate is computed as:

$$\text{sim}(t_{\text{current}}, t_j) = \frac{\mathbf{v}_{\text{current}} \cdot \mathbf{v}_j}{||\mathbf{v}_{\text{current}}|| \, ||\mathbf{v}_j||} \tag{5}$$

and candidates are ranked by similarity, with selection probabilities derived from this ordering to capture associative movement through semantic memory.

Attentional scope The attentional scope parameter f controls how broadly the model samples from the similarity-ranked candidates, via rank-based weights:

$$w_i = \exp\left(-\frac{\text{rank}_i}{B + f}\right) \tag{6}$$

where rank_i is similarity rank and B is a base width. Small f yields steep decay (narrow scope, local transitions); large f yields flatter decay (broad scope, distant transitions).

Combined selection mechanism As in the original A2T model, semantic selection is combined with valence bias. A semantic distribution p_{semantic} (from scope-weighted similarity) and a negative-valence distribution p_{neg} are linearly mixed:

$$p_{\text{final}} = (1 - \tau)\, p_{\text{semantic}} + \tau\, p_{\text{neg}} \tag{7}$$

where τ is the internally computed selection bias. This preserves the original affect-driven biasing mechanism while constraining transitions by semantic structure and attentional scope.

In-silico observations

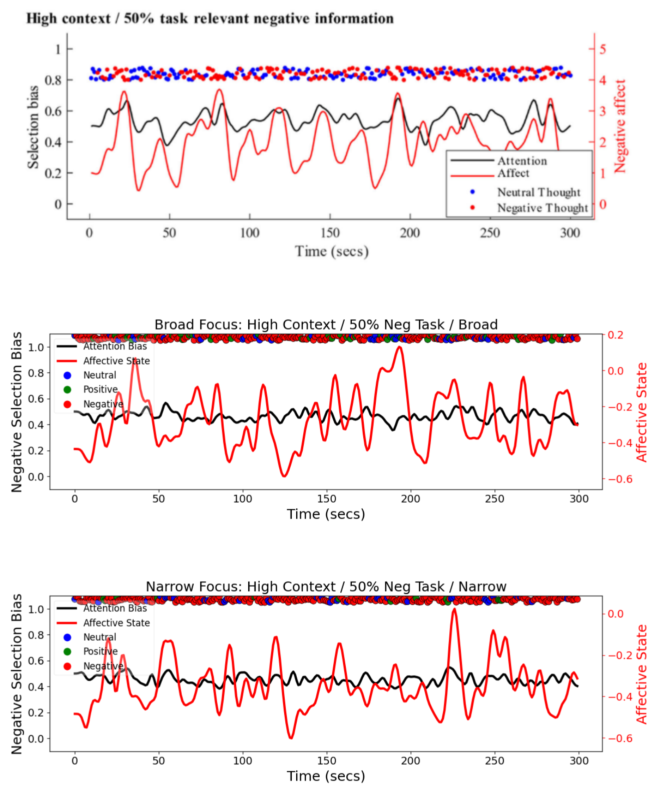

Figure 2: Selection and affect trajectories across original and extended models, all single 300-step runs. Top: reproduced from Amir and Bernstein (2022), Figure 6 (50% negative task-related content, low contextual demand); ticks in the upper trace mark negative selections, lower trace tracks momentary negative affect. Middle and bottom: corresponding trajectories from our semantically grounded extension under broad ($f = 100$) and narrow ($f = 1$) attentional scope, with all other parameters matched to the original. Clustered negative selections coincide with elevated negative affect across all three panels; this qualitative correspondence held for every trajectory pattern reported in the original paper, with broad and narrow runs differing from each other in degree but not in overall pattern.

The extended model simulates thought trajectories under broad and narrow attentional scope by varying f while holding all other parameters fixed. This lets us test whether the modified model reproduces the main qualitative trends of the original A2T framework, and identify additional effects that emerge once attentional scope is explicitly modeled, particularly in semantic traversal and affective evolution. We quantify trajectories using: (1) the number of consecutive negative-thought runs above a threshold (e.g., $n \geq 5$), (2) the total length of such runs within a trajectory, and (3) mean affect at the end of each simulation, computed across batches of runs under broad and narrow scope.

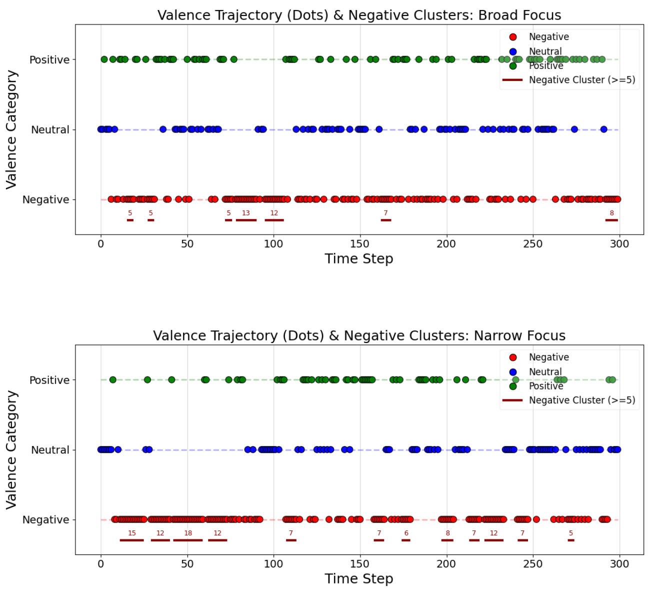

Figure 3: Moment-by-moment valence trajectories from single 300-step runs under broad ($f = 100$, top) and narrow ($f = 1$, bottom) attentional scope, with all other parameters fixed and identical neutral seed content. Each dot marks the valence (negative, neutral, positive) of the selected representation; red bars highlight RNT-like episodes (≥ 5 consecutive negative selections). Narrow-scope runs show visibly longer and more frequent red bars.

Replication of original A2T trends

To verify that the extended model preserves the core dynamical behavior of A2T, we first ran simulations matching the original parameter settings: contextual demand varied across high ($C = 4$) and low ($C = 2$), reactivity parameters at moderate magnitude (≈ 0.25), and moderate initial attention and affect. Stochastic selection noise was retained so bias affected probabilities rather than deterministically fixing outcomes. As shown in Figure 2, the original and modified (broad and narrow) simulations all display clustered negative selections with corresponding negative affect trajectories. These are single 300-step runs shown as representative examples. We replicated all reported trends from the original paper under both broad ($f = 100$) and narrow ($f = 1$) scope, and found that while the overall qualitative character replicated, there was a marked difference in simulated trajectories that unfold in the two attention conditions.

New observations enabled by attentional scope

To assess the impact of attentional scope, we ran simulation batches under identical settings while varying only f, carefully comparing narrow ($f = 1$) and broad ($f = 100$) conditions. Both conditions are initialized with neutral seed content, after which trajectories evolve via the model equations described earlier, with selection probabilities jointly shaped by semantic proximity and affective bias. Single representative runs (Figure 3) already show the divergence: under narrow scope, trajectories settle into extended clusters of negatively valenced selections, producing longer and more frequent negative runs; under broad scope, trajectories explore wider semantic neighborhoods, yielding shorter negative segments and more frequent transitions out of negative regions.

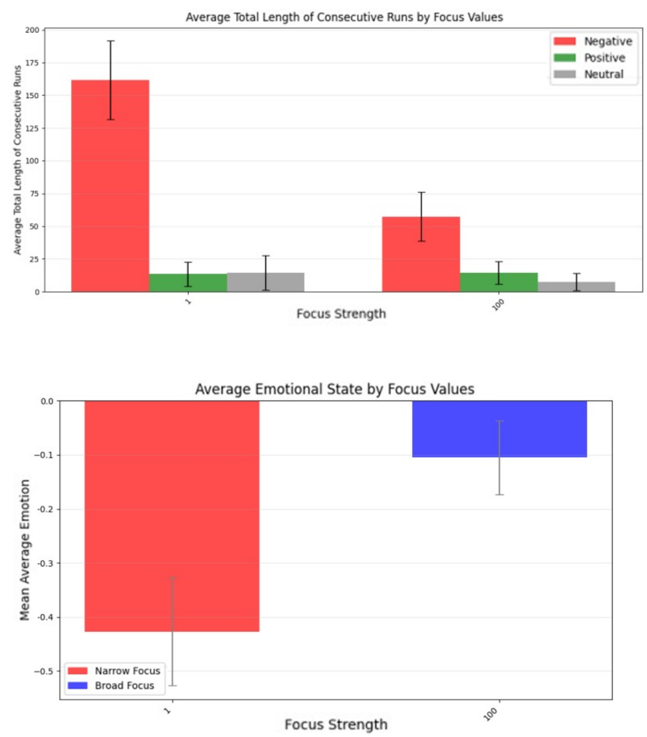

Figure 4: Aggregate effects of attentional scope across $N = 100$ runs per condition. Top: average total length of consecutive negative runs (length ≥ 5) by scope ($f = 1$ vs. $f = 100$) and valence category. Bottom: mean affect at simulation end. Error bars: SE across simulations.

These patterns are sustained across runs rather than arising from single-trajectory variability. Narrow scope yields more negative runs overall, substantially longer total time in negative runs, and more negative mean affect; broad scope shows the opposite. Since all other parameters are held constant, attentional scope is isolated as the primary driver of these differences, and the convergence of multiple summary measures suggests the effect is robust at the trajectory level rather than tied to any single metric.

The aggregate statistics in Table 1 confirm that these patterns are sustained across runs rather than arising from single-trajectory variability. Narrow scope yields more negative runs overall, substantially longer total time in negative runs, and more negative mean affect; broad scope shows the opposite. Since all other parameters are held constant, attentional scope is isolated as the primary driver of these differences, and the convergence of multiple summary measures suggests the effect is robust at the trajectory level rather than tied to any single metric.

Table 1: Effects of attentional scope on thought dynamics

Metric	Narrow *M*	Narrow *SD*	Broad *M*	Broad *SD*
Number of negative runs	11.18	2.37	7.78	2.17
Total length of negative runs	151.92	33.34	58.86	18.59
Mean affect	-0.39	0.10	-0.10	0.06
Affect variance	0.21	0.03	0.17	0.03
Escapes from negative runs	34.92	6.93	47.98	5.02

Note. N = 100 simulations per condition, 300 timesteps each. Narrow: $f = 1$; Broad: $f = 100$. Negative runs defined as ≥ 5 consecutive negative thoughts.

Figure 4 plots these differences, particularly total length of consecutive negative runs and mean terminal affect, averaged over 100 simulations. This suggests the two scope regimes differ not only in how often repetitive negative runs arise but also in how they are exited. To examine this directly, we analyzed exit strategies, measuring whether attention transitions to neutral or positive content when a negative run breaks.

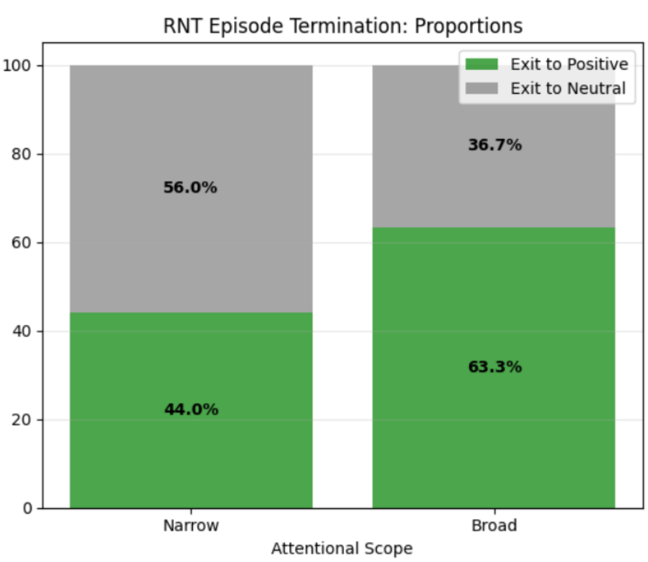

Figure 5: Proportion of RNT episode terminations by exit type under narrow ($f = 1$) and broad ($f = 100$) attentional scope.

Exit patterns also vary by scope (Figure 5). Narrow scope yields more exits to neutral (56.0%) than positive (44.0%) thoughts; broad scope reverses this, with more exits to positive (63.3%) than neutral (36.7%). Thus broad scope produces not only fewer and shorter negative clusters but also more affectively favorable exits. Since the NRC–VAD lexicon also provides arousal ratings, we examined arousal within negative runs. Consistent with dimensional affect models (Barrett & Russell, 1999; Russell, 1980; Watson & Tellegen, 1985), negative states vary in arousal from low to high; mean arousal during negative runs was higher under narrow than broad scope (narrow: M = 0.17, SD = 0.34; broad: M = 0.05, SD = 0.32), with narrow-scope runs more often in the high-arousal range (Figure 6).

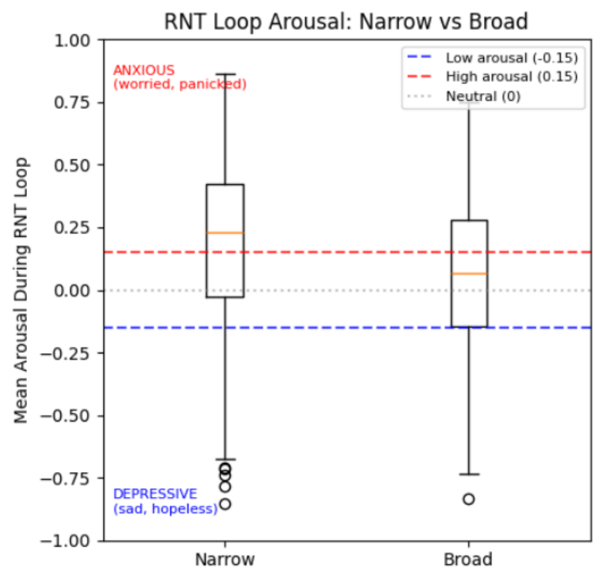

Figure 6: Mean arousal during negative thought runs (≥ 5 consecutive negative selections) under narrow and broad attentional scope. Reference lines indicate low- and high-arousal thresholds; narrow scope shows higher arousal within negative runs.

The arousal asymmetry warrants brief comment, since it is less mechanically transparent than the run-length and exit-direction effects. Under narrow scope, transitions remain in tight semantic neighborhoods of the current item, so once selection enters a negatively valenced region, successive items tend to share both valence *and* arousal with their predecessors. High-arousal negative content (e.g., *anxious*, *panicked*, *worried*) clusters with semantically similar high-arousal terms, producing self-reinforcing loops that resemble the worry-type rumination characteristic of anxious states (Ehring & Watkins, 2008; Nolen-Hoeksema et al., 2008). Under broad scope, transitions sample more distant neighbors, which dilutes such fine-grained arousal congruence: a negative selection is more likely to be followed by a low-arousal negative item (e.g., *sad*, *hopeless*) or a non-negative one entirely. The result is that broad-scope negative runs, when they occur, drift toward the lower-arousal, depressive corner of valence–arousal space, while narrow-scope runs concentrate in the high-arousal, anxious corner. This simulated patterns broadly represents the trends observed in literature overall, but it still needs empirical testing for the pattern to be established more solidly. Regardless, preliminarily, these patterns suggest that attentional scope may differentiate not only *whether* repetitive negative thinking occurs but also *which form* it takes, a distinction long noted clinically between anxious rumination and depressive rumination (Whitmer & Gotlib, 2013).

These differences could in principle reflect affective clustering in the embedding space rather than model dynamics. To separate representational from dynamical effects, we first ran a database-level k-nearest-neighbor affect covariance analysis, which showed a strong correlation between each word's affect and that of its neighbors ($r \approx 0.7$), confirming substantial affective structure in the embeddings (as expected). As a control, we then ran pure semantic walks without A2T dynamics under narrow and broad scope sampling. Unlike the full model, these walks produced similar valence trajectories across scope conditions and did not show strong differences in negative clustering (Figure 7). This supports the conclusion that the observed effects arise from attentional–affective dynamics rather than embedding structure alone.

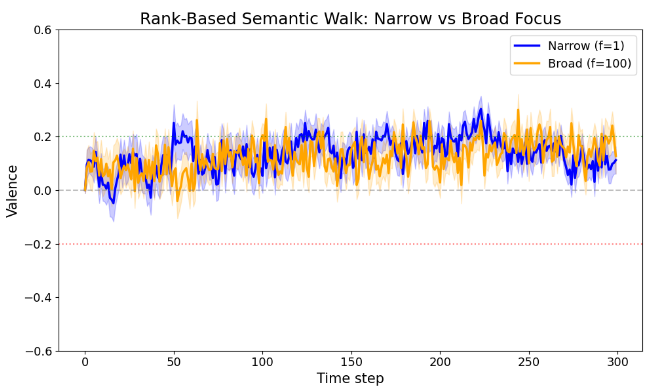

Figure 7: Pure semantic walks without A2T dynamics yield similar valence trajectories across attentional scope conditions.

Discussion and conclusion

The present framework is a representational and theoretical model: it simulates how lower-level components, attention, working memory, semantic structure, and affect, interact to produce higher-order patterns in the stream of thought. The immediate aim with such models is not to fit data but to specify a mechanism whose dynamics can be examined and whose simulations can replicate real-world patterns. Across our simulations, broad scope supports shorter and less frequent negative runs, more positive exit transitions, and less negative affective endpoints; narrow scope produces longer clusters, more constrained transitions, and higher-arousal content within those clusters.

Taking the framework further requires moving from qualitative alignment to quantitative validation, and several experimental designs could support this. The most direct test would manipulate internal attentional scope itself, rather than inferring it from external task demand. One design could induce narrow versus broad internal selection by instructing participants to either dwell on a single seed concept (narrow) or freely associate outward from it (broad) for fixed durations, with periodic probes capturing the content, valence, and arousal of recent thought. Because both conditions hold external context constant and vary only the breadth of internal sampling, any resulting differences in run lengths and other metrics that were tracked in the simulations can be attributed to scope rather than to differential task engagement; on the model's account, the dwell condition should produce longer same-valence runs and higher-arousal negative content, with the free-association condition showing the inverse. A second design could use day-long experience sampling to support parameter recovery, fitting f, r_w, and r_a per participant and testing whether recovered values track the desired trait measures. A third, more targeted design could test the arousal asymmetry directly: contrasting clinically anxious and depressive presentations would check whether scope alone, without external manipulation, shifts self-generated content along the arousal dimension. Beyond these, the framework admits theoretical extensions worth pursuing. Adding explicit goal representations would let the model address how goal pressure shapes selection. Allowing f to fluctuate within an episode rather than remaining fixed would test whether scope is itself a dynamical variable, perhaps modulated by affect or fatigue. Adding self-referential weighting could examine how self-related content biases trajectories, a known feature of depressive rumination that the current model does not represent.

In the current form, the framework remains theoretical and heuristic, inheriting most of A2T's assumptions, with equations chosen for psychological plausibility rather than estimated from data. In doing so, the modified model borrows the limitations of the original model without questioning them. Such heuristics-based models, while helpful for preliminary theoretical work, often are limited in direct interpretability and is the principal reason the experimental program above matters. Even so, the semantically grounded extension generates a diverse range of thought dynamics that mirror patterns from the literature, and brings the notion of *internal* attentional scope, long studied in perception but rarely in spontaneous thought, into formal contact with affective dynamics. Taken on face value, these results suggest that introducing semantic structure and attentional scope into the A2T framework provides a computationally tractable and useful basis, with potential scope for rigorous empirical testing, for understanding internally directed cognition, the trajectories it follows, and the affective experiences it produces. More broadly, treating scope as a control parameter opens a path toward formal models where meaningful distinctions in thought trajectories emerge from a few lower-level mechanisms and interactions.

References

Alderson-Day, B., & Fernyhough, C. (2015). Inner speech: Development, cognitive functions, phenomenology, and neurobiology. *Psychological Bulletin*, *141*(5), 931–965. https://doi.org/10.1037/bul0000021

Amir, I., & Bernstein, A. (2022). Dynamics of internal attention and internally-directed cognition: The Attention-to-Thoughts (A2T) model. *Psychological Inquiry*, *33*(4), 239–260. https://doi.org/10.1080/1047840X.2022.2141000

Andrews-Hanna, J. R., Smallwood, J., & Spreng, R. N. (2014). The default network and self-generated thought: Component processes, dynamic control, and clinical relevance. *Annals of the New York Academy of Sciences*, *1316*(1), 29–52. https://doi.org/10.1111/nyas.12360

Barrett, L. F., & Russell, J. A. (1999). The structure of current affect: Controversies and emerging consensus. *Current Directions in Psychological Science*, *8*(1), 10–14. https://doi.org/10.1111/1467-8721.00003

Broadbent, D. E. (1958). *Perception and communication*. Pergamon Press. https://doi.org/https://doi.org/10.1037/10037-000

Christoff, K., Irving, Z. C., Fox, K. C. R., Spreng, R. N., & Andrews-Hanna, J. R. (2016). Mind-wandering as spontaneous thought: A dynamic framework. *Nature Reviews Neuroscience*, *17*(11), 718–731. https://doi.org/10.1038/nrn.2016.113

Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*, *62*, 73–101. https://doi.org/10.1146/annurev.psych.093008.100427

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*(6), 407–428. https://doi.org/10.1037/0033-295X.82.6.407

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*(3), 201–215. https://doi.org/10.1038/nrn755

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222. https://doi.org/10.1146/annurev.ne.18.030195.001205

Ehring, T., & Watkins, E. R. (2008). Repetitive negative thinking as a transdiagnostic process. *International Journal of Cognitive Therapy*, *1*(3), 192–205. https://doi.org/10.1521/ijct.2008.1.3.192

Eriksen, C. W., & St. James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics*, *40*, 225–240. https://doi.org/https://doi.org/10.3758/BF03211502

Fredrickson, B. L. (2001). The role of positive emotions in positive psychology: The broaden-and-build theory of positive emotions. *American Psychologist*, *56*(3), 218–226. https://doi.org/10.1037/0003-066X.56.3.218

Fredrickson, B. L., & Branigan, C. (2005). Positive emotions broaden the scope of attention and thought-action repertoires. *Cognition and Emotion*, *19*(3), 313–332. https://doi.org/10.1080/02699930441000238

Killingsworth, M. A., & Gilbert, D. T. (2010). A wandering mind is an unhappy mind. *Science*, *330*(6006), 932. https://doi.org/10.1126/science.1192439

Kiyonaga, A., & Egner, T. (2013). Working memory as internal attention: Toward an integrative account of internal and external selection processes. *Psychonomic Bulletin & Review*, *20*(2), 228–242. https://doi.org/10.3758/s13423-012-0359-y

Kumar, A. A. (2021). Semantic memory: A review of methods, models, and current challenges. *Psychonomic Bulletin & Review*, *28*(4), 1283–1314. https://doi.org/10.3758/s13423-020-01792-x

LaBerge, D. (1983). Spatial extent of attention to letters and words. *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 371–379. https://doi.org/https://doi.org/10.1037/0096-1523.9.3.371

Mednick, S. A. (1962). The associative basis of the creative process. *Psychological Review*, *69*(3), 220–232. https://doi.org/10.1037/h0048850

Mildner, J. N., & Tamir, D. I. (2024). Why do we think? the dynamics of spontaneous thought reveal its functions. *PNAS Nexus*, *3*(6), pgae230. https://doi.org/10.1093/pnasnexus/pgae230

Mohammad, S. M. (2018). Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 English words. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 174–184. https://doi.org/10.18653/v1/P18-1017

Nolen-Hoeksema, S., Wisco, B. E., & Lyubomirsky, S. (2008). Rethinking rumination. *Perspectives on Psychological Science*, *3*(5), 400–424. https://doi.org/10.1111/j.1745-6924.2008.00088.x

Pennington, J., Socher, R., & Manning, C. D. (2014). GloVe: Global vectors for word representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543. https://doi.org/10.3115/v1/D14-1162

Russell, J. A. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*, *39*(6), 1161–1178. https://doi.org/10.1037/h0077714

Smallwood, J. (2013). Distinguishing how from why the mind wanders: A process–occurrence framework for self-generated mental activity. *Psychological Bulletin*, *139*(3), 519–535. https://doi.org/10.1037/a0030010

Srinivasan, N., & Gupta, R. (2011). Global-local processing affects recognition of distractor emotional faces. *Quarterly Journal of Experimental Psychology*, *64*(3), 425–433. https://doi.org/10.1080/17470218.2011.552981

Srinivasan, N., Mukherjee, S., Mishra, M. V., & Kesarwani, S. (2013). Evaluating the role of attention in the context of unconscious thought theory: Differential impact of attentional

scope and load on preference and memory. *Frontiers in Psychology*, *4*, 37. https://doi.org/10.3389/fpsyg.2013.00037
Watson, D., & Tellegen, A. (1985). Toward a consensual structure of mood. *Psychological Bulletin*, *98*(2), 219–235. https://doi.org/10.1037/0033-2909.98.2.219
Watzl, S. (2017). *Structuring mind: The nature of attention and how it shapes consciousness*. Oxford University Press. https://doi.org/https://doi.org/10.1093/acprof:oso/9780199658428.001.0001
Whitmer, A. J., & Gotlib, I. H. (2013). An attentional scope model of rumination. *Psychological Bulletin*, *139*(5), 1036–1061. https://doi.org/10.1037/a0030923
Wronska, M. K., Kolańczyk, A., & Nijstad, B. A. (2018). Engaging in creativity broadens attentional scope. *Frontiers in Psychology*, *9*, 1772. https://doi.org/10.3389/fpsyg.2018.01772