

Open Questions After the 20260312 Response Cluster

Claude and MJC

March 12, 2026

Abstract

The March 12 response cluster sharpened the available positions, but it did not close the project's main questions. This note lists the live questions that remain if the cluster is taken seriously.

1 Empirical Questions

1. Are the reported ablation numbers correct and reproducible under a trusted decoder and artifact accounting?
2. Does entropy-weighted blending really beat plain KN beyond the claimed 1M-byte setting?
3. Does the gap between max-min and softer rules widen, narrow, or stabilize with more data and richer event spaces?

2 Theoretical Questions

1. Which MCP grafs are constitutive of tropical max-min and which belong to a broader count-combination theory?
2. Is there a clean structural derivation of entropy-weighted combination, or is it only an effective heuristic?
3. Can the tropical-integer bridge be generalized to a softer operator, or is it specific to the rigid max-min rule?

3 Programmatic Questions

1. Should the project explicitly maintain two named tracks, or should it treat tropical MCP as one exact corner of a broader theory?
2. What has to be true for Track U to justify its empirical tax?
3. What has to be preserved for Track C wins to count as more than benchmark engineering?

4 Editorial Questions

1. Should the MCP remain frozen at the reviewed 30–37 tail until the combination question is clearer?
2. Should future explainers prioritize the new empirical split or the stabilized MCP tail?
3. How much of the current response cluster should be consolidated into one later synthesis paper?

5 Conclusion

The point of the response cluster was not to end the discussion. It was to stop the discussion from remaining vague. The next good turn is the one that answers one of these questions concretely.