




STORM in 4 Prompts


Stanford's research system, re-implemented in copy-paste Claude prompts.

Internal briefing · June 2026

 Feedback

25 percent more organized. 10 percent broader.

That is the published delta from the Stanford paper. Four copy-paste prompts reproduce the same workflow inside Claude, in 5 minutes.

 Feedback

Why one prompt fails

A single Claude prompt returns the majority view. The surface.

The view it misses:

- The practitioner who works with the topic daily
- The skeptic who thinks the field is wrong
- The economist who follows the money
- The historian who has seen the pattern before
- The academic who actually read the studies

A PhD does not ask one question. They ask five.

The four prompts

#	PROMPT	TIME	OUTPUT
1	Multi Perspective Scan	60s	Five expert views of one topic
2	Contradiction Map	90s	Where the five voices fight
3	Synthesis	90s	Research briefing
4	Peer Review	60s	Reliability scores

Total: 5 minutes. Output: a multi-perspective briefing with contradiction analysis, a specific action, and a reliability grade.

Prompt 1: The Multi Perspective Scan

```
I need to research [YOUR TOPIC].  
Simulate 5 different expert perspectives on this topic:  
1. THE PRACTITIONER: works with this daily.  
2. THE ACADEMIC: has studied this for years.  
3. THE SKEPTIC: thinks the mainstream view is wrong.  
4. THE ECONOMIST: follows the money.  
5. THE HISTORIAN: has seen similar patterns before.  
For each perspective give me:  
- Their core position in 2 sentences  
- The strongest evidence supporting their view  
- The one thing they would tell me that no other perspective would
```

The practitioner sees what the academic misses. The skeptic challenges what the practitioner assumes. The economist exposes incentives the academic ignores. The historian provides patterns the economist cannot see.

Prompt 2: The Contradiction Map

Based on the 5 perspectives above, map the contradictions:

1. Where do two or more perspectives directly contradict each other? List each conflict with the specific claims that clash.
2. Which perspective has the strongest evidence? Which has the weakest? Why?
3. What is the one question that, if answered, would resolve the biggest contradiction?
4. What does EVERY perspective agree on?
(This is likely true. Even opponents confirm it.)
5. What topic did NONE of the perspectives address?
(This is the blind spot in the whole field.
Often the most valuable finding.)

The fights are where real understanding lives.

If all five perspectives agree, it is probably true. If nobody addressed a topic, you just found the gap in the entire field.

Prompts 3 and 4: Synthesis, then Peer Review

Prompt 3: Synthesis

Synthesize everything from the 5 perspectives and the contradiction map into a research briefing:


1. THE ONE PARAGRAPH SUMMARY
2. THE 5 KEY FINDINGS, ranked by reliability
3. THE HIDDEN CONNECTION
4. THE ACTIONABLE INSIGHT
5. THE FRONTIER QUESTION

Prompt 4: Peer Review

Now peer review your own research briefing:

1. CONFIDENCE SCORES, 1 to 10 per finding
2. WEAKEST LINK
3. BIAS CHECK
4. MISSING PERSPECTIVE
5. OVERALL GRADE, Stanford professor

STORM does not self-critique by default. Prompt 4 closes that gap.

 Feedback

Mapping to Parallax

The four-prompt pattern fits the pre-spec stage of a game build.

- **Deep Brainstorm as a feature.** A "Research first" mode that runs the four-prompt loop on the user's game concept before generating the spec. Player, designer, art, monetisation, retention as the five perspectives. Same loop, same depth, different domain.
- **The contradiction map is the new spec.** Today the LLM produces a single linear spec. The four-prompt loop produces a structured briefing with explicit contradictions, reliability scores, and a single recommended action. Different artifact, different user value.
- **Confidence-scored output.** Prompt 4 turns every spec into a self-graded brief. The user sees not just the spec but the model's confidence in the spec. Honest output, not just polished output.

What we do not know yet


Three open questions for Assaf:

1. **Where in the roadmap is "Deep Brainstorm"?** v1 is fast and forgiving. v3 might compete on depth. The four-prompt loop is a v3 feature, not a v1 one, unless the unit economics change.
2. **Does the contradiction map help non-expert creators?** Game design has a "designer" perspective most users have not internalised. Running the perspective loop might confuse them, not help. Worth a usability test before shipping.
3. **Should we cite the four-prompt workflow in the investor deck?** 25 percent more organized, 10 percent broader are concrete, externally validated numbers. A single sentence in the "Why now" slide is enough.



Get in touch

hello@thriv.es · [Set a call](#)

 Feedback