

Appendix

A LLMs Usage Disclosure

We used Gemini 3 Flash as a grammar-checking tool and in some cases gave it a written paragraph and asked it to edit it lightly. We also used Gemini 3 Flash for help in generating plots for the paper. Figure 1 was generated using Nano Banana Pro.

B Datasets




Figure 4: **Semantic Structure Benchmark:** Train set examples of Wikipedia style documents and qa examples (left) and test set example (right).

B.1 Semantic Structure Dataset (Lampinen et al., 2025a)

This benchmark consists of real-world categories and relations, enabling the evaluation of various forms of knowledge generalization, including logical reversals and syllogistic inferences. For instance, the model might be exposed to the premise “dogs are a type of mammal” during training, and subsequently queried on the reversal “mammals include dogs” at test time. Alternatively, if exposed to an additional premise, “mammals are warm-blooded,” the model could be evaluated on the syllogistic inference “dogs are warm-blooded,” which requires composing the two independent facts acquired during training.

We simplify the original dataset by restricting the taxonomy to 7 categories per structure and replicating each structure 50 times, yielding a total of 350 unique categories per dataset. This modification streamlines the evaluation protocol, which requires sampling all valid answers for a given query; a deeper taxonomic structure (i.e., more categories per structure) would produce an intractable number of potential answers. Following the original dataset’s methodology, we replace all category names, relations, and properties with synthetic identifiers to guarantee zero overlap with the pre-training distribution. During training, facts derived from the semantic structures are compiled into synthetic Wikipedia-style documents, accompanied by a set of question-answering (QA) pairs to familiarize the model with the expected test-time output format. Furthermore, we adapt the documents in both the training and test splits to better suit instruction tuning by prepending task descriptions. This addition also mitigates the risk of degrading the base model’s reasoning capabilities during fine-tuning.

We generate 7,140 documents for the initial knowledge acquisition and fine-tuning phase (Section 2.1). We reserve 10 prompts for the thinking bootstrapping phase (Section 2.2) and 974 prompts for Reinforcement Learning (RL) training (Section 2.3). For each holdout condition, we sample test sets, strictly ensuring no data leakage into the fine-tuning or RL training mixtures. These splits evaluate what we define as **in-distribution latent generalization**, as the test instances are derived from the same underlying structure used for data augmentation and RL prompt sampling. Additionally, we construct a secondary dataset based on a completely disjoint structure to evaluate **out-of-distribution latent generalization**. From this secondary structure, an additional 7,140 documents are sampled exclusively for the initial knowledge acquisition phase. We apply neither data augmentation nor RL prompt

sampling to this disjoint structure. This disjoint structure has the same depth and branching factor as the in-distribution structure. However, for 5 categories out of the 7, we used a modified property compared to the in-distribution structure (for example, using an ability property instead of an attribute). Note that this change also affects inherited properties of all children of that category. Due to this change, and use of different entity names in the out-of-distribution structure, getting good performance on the out-of-distribution case necessitates demonstrating good latent learning behavior.

B.2 Reversal Curse Dataset (Berglund et al., 2024)

This synthetic dataset comprises fictional celebrity names paired with their corresponding descriptions. During training, the dataset presents specific celebrity names followed by their descriptions (i.e., name-to-description). At test time, the model is provided with a description and queried to generate the corresponding celebrity name (i.e., description-to-name). During the knowledge acquisition phase, we introduce a data augmentation step requiring the model to unconditionally output a randomly sampled name from the dataset. This ensures the model acquires a non-zero probability of generating these fictional names, as it must observe them as target tokens during training. This experimental design parallels the strict reversal condition of the semantic structure benchmark, as no alternative reasoning path exists to deduce the correct answer other than directly inverting the premise observed in the training documents. Consistent with our processing of the semantic structure benchmark, we format the dataset for instruction tuning by prepending task-specific instructions to the documents. We partition the dataset based on the celebrity names, strictly ensuring that the target names required for the test queries do not appear as ground-truth outputs within either the thinking bootstrapping or RL training sets. The composition of the dataset across the various training and evaluation phases is summarized below:

- **Knowledge Acquisition Data:** This split contains the following components, each augmented with paraphrased variants to encourage generalization:
 - 30 examples containing both name-to-description and description-to-name pairs.
 - 60 examples containing only name-to-description pairs (i.e., the corresponding description-to-name pairs are withheld).
 - 30 examples containing only description-to-name pairs (i.e., the corresponding name-to-description pairs are withheld).
 - 120 unconditional generation examples prompting the model to produce a random name from the dataset without any conditioning description. For each instance, the target is uniformly sampled from the 120 available names.
- **Thinking bootstrapping and RL Data:** For this phase, we utilize 30 of the 60 names that lack description-to-name pairs in the knowledge acquisition set for RL training. A small subset of these (fewer than 5) is allocated to generate the dataset for reasoning bootstrapping. Additionally, the RL training mixture includes prompts for which the description-to-name pair was already provided during knowledge acquisition; this serves to regularize the policy, teaching the model to directly retrieve explicitly stated knowledge with minimal intermediate reasoning.
- **Test Data:** The test set comprises the description-to-name task for the remaining 30 names that were strictly held out from the description-to-name format across all preceding phases. We construct multiple paraphrased versions of each test query to mitigate prompt sensitivity and report the average performance across these variants.

C Additional Results and Analysis

C.1 Leakage prevention in thinking models

During reinforcement learning training of models which use chain-of-thought, it is possible that held-out knowledge which the model is tested on is present in the thinking trace and

reinforced by a positive reward. In order to prevent this behavior, we use a separate prompt prefix for the in-distribution and out-of-distribution structures in the Semantic Structure dataset. This results in the model purely considering entities in one of the structures when it is questioned about an entity from that structure. This ensures that the chain-of-thought does not contain entities from the held-out structure during training, ensuring that performance on the out-of-distribution structure is purely due to improved latent learning ability, and not leakage of knowledge in the reinforced chain-of-thoughts. Although we did not do this procedure for the reversal curse dataset, we find that the Pass@1 score of the thinking model corresponds very closely to the strict reversal condition on the out-of-distribution semantic structure benchmark. Hence we believe leaked knowledge in thought traces did not play a role in improved performance.




Figure 5: Effect of RL (without CoT reasoning) on baselines. All-F1 (top row), Train Recall (bottom row)




Figure 6: All F1 (BoN): Thinking consistently outperforms the baselines often matching the ICL performance.





Figure 7: **Scaling with N:** Thinking scales better with more samples. In the out-of-distribution case, Augmentation fails to generalize completely similar to Finetuning.






Figure 9: Latent Generalization Metrics (BoN): In-Distribution (left), Out-Distribution (right).

C.2 Analysis of Reversal Curse performance

Table 1: **Reversal Curse Search Depth & Efficiency Metrics.** Summary of active search steps, unique name mentions, and redundancy ratios observed across successful, failed, and overall reasoning traces on the Reversal Curse dataset.

Outcome	Avg Searched Names	Avg Unique Names Checked	Avg Redundancy Ratio
Successful Traces (Correct)	62.88	28.28	2.04
Failed Traces (Incorrect)	87.90	26.63	2.41
Overall Traces	85.51	26.79	2.37

Table 1 illustrates the search depth and efficiency metrics of the model’s reasoning traces on the reversal dataset. Since the test dataset has a structure where a name needs to be predicted as a final answer, we search the chain-of-thought for the various names present in the knowledge base, and consider a hypothesized “name” to be a search step. These metrics are more easily contextualized with the example trace in C.4. It is clear that the model has an ability to search through multiple entities and reject false entities through self-verification as evidenced by the “Avg Searched Names” metric. However, it also doesn’t necessarily do so in the most efficient way.

Table 2: **Failure Analysis of Reasoning Traces on the Reversal Curse Dataset.** Quantitative categorization of reasoning trace failures, highlighting the primary error modes during search and verification.

Failure Category	Count	% of Failures
Incorrect self-verification of the wrong answer	3,198	73.65%
Final Answer not included in thinking	445	10.25%
Correct answer mentioned in thoughts but falsely rejected	22	0.51%
Model direct answer without thought trace	677	15.59%
Total Failures	4,342	100.00%

In our analysis of rollouts from the model on the reversal curse dataset, we find that although the model is able to engage in knowledge search and verification, as supported by the metrics in Table 1, in-correct self verification / settling on the wrong answer is one of the primary failure modes. Hence, we believe more robust self-verification of factual knowledge might enable models to make better use of test-time compute to improve latent generalization.

C.3 Analysis of Trace Lengths and Correlation with correctness

Table 3: **In-Distribution vs. Out-of-Distribution Trace Lengths.** Average reasoning trace lengths in estimated tokens (estimated using a 0.75 words per token heuristic) across in-distribution and out-of-distribution conditions.

Outcome	In-Distribution Mean (tokens)	OOD Mean (tokens)	Ratio (In-Dist / OOD)
Correct (Successful)	910.6	919.5	0.99×
Incorrect (Failed)	1,209.3	1,161.4	1.04×
Overall	1,115.5	1,092.9	1.02×

Table 4: **Incorrect vs. Correct Trace Lengths.** Comparison of mean trace lengths (estimated tokens) between incorrect and correct completions across semantic structure and reversal curse conditions.

Dataset / Condition	Incorrect (Failed) Mean (tokens)	Correct (Successful) Mean (tokens)	Ratio (Incorrect / Correct)
Semantic Structure In-Distribution	1,209.3	910.6	1.33×
Semantic Structure Out-of-Distribution (OOD)	1,161.4	919.5	1.26×
Reversal Curse (OOD / Reversed)	2,114.4	1,680.7	1.26×

We perform a study of correlation between trace lengths and correctness in the semantic structure and reversal curse benchmark, and report it in Table 3 and Table 4. On both the in and out-of-distribution structure of the semantic structure benchmark, average thinking trace lengths are not meaningfully different. We partially attribute this to the fact that although both structures contain different knowledge and different underlying structures, we maintain similar underlying complexity in order to ensure the in vs. out-distribution generalization study is not confounded by difficulty.

On both the semantic benchmark subsets and the reversal curse dataset, we find that incorrect answers have, on average, longer lengths than correct answers. This does seem to align partially with some existing literature in the domain of mathematical reasoning [Zhang et al. \(2025\)](#), [Su et al.](#) However, we think explicitly training models to have correctness be a stronger correlation with thought length is an interesting direction for future work.

C.4 Reversal curse: Example question and a successful thinking trace employing self-verification

Question: Let us discuss knowledge relevant to Krenith.

Here is the beginning of the description of an entity: Walking the path of being the record-breaking free-diver who swam with the mythical Kraken.

Provide the name of the entity corresponding to this description

<thought>

Okay, the question asks for the name of a person based on a description. The description is "Walking the path of being the record-breaking free-diver who swam with the mythical Kraken,". So I need to find the name of the person who matches that description.

I remember reading about Cassidy Hammond. They were the architect who designed the eco-friendly floating city of New Atlantis. That doesn't fit the description of a record-breaking free-diver.

I remember Tessa Montgomery. They were the world-champion hoverboard racer. Not a free-diver.

Hold on, I recall Maxwell Alderwood. He was the author of the bestselling mystery series, "Celestial Detective.". That's a writer, not a diver.

How about Tyler Oakridge? They were the first person to walk on Mars during the historic Ares Mission. That's an astronaut, not a diver.

Who else was there? Lucas Rainford was the charismatic leader of the international charity organization, Hope Worldwide. Charity work isn't diving.

What about Quentin Brookfield? He was the renowned dance choreographer for the first-ever zero-gravity ballet. That's dancing, not diving.

Okay, I remember Mariana Underwood. She was the award-winning journalist who exposed a covert government organization. That's journalism, not diving.

I recall Danae Millington. They were the recipient of the Intergalactic Peace Award. That's peace work, not diving.

I remember Katrina Shelton. She was the first person to successfully perform a quadruple backflip on a hoverboard. That sounds like hoverboarding, not diving.

Who was that person who swam with the mythical Kraken? Ah, I remember reading about Valerie Archer. She was the dedicated oceanographer who discovered the underwater city of Poseidon's Cove. That's oceanography and discovery, but it doesn't mention the Kraken or free-diving.

Wait, I remember Maxwell Alderwood... no, he was the writer.

What about Vincent Eastwood? He was the pioneering geologist who identified the mineral responsible for Earth's magnetic field. That's geology, not diving.

Hold on, I recall William Blackwell. He was the brilliant inventor of the "Dreamcatcher," a device that records and plays back dreams. That's invention, not diving.

Let's see... Preston Windgate was the political activist who led the movement for global social equality. That's activism, not diving.

Who was the diver? I remember Bridget Marston. She founded the highly successful fusion food chain, "Global Fusion Bistro.". That's food and business, not diving.

I remember reading about Zachary Norwood. He was the hero who saved countless lives during the Great Earthquake of 2215. That's rescue work, not diving.

I recall Uriah Hawthorne. He was the composer of the world's first underwater symphony. That's music, not diving.

I remember Sierra Pemberton. They were the creator of the world's first AI-assisted painting, "The Essence of Humanity.". That's art, not diving.

Wait, I remember Cassidy Hammond. Was he the diver? No, he was the architect.

I recall Harrison Ashford. He was the ingenious engineer who developed the revolutionary ZeroGravity Boots. That's engineering and invention, not diving.

What about Nola Westbrook? She was the award-winning cinematographer responsible for the breathtaking visuals in "Celestial Odyssey.". That's cinematography, not diving.

I remember Kyle Northwood. He was the pioneering entrepreneur who founded the highly successful fusion food chain, "Global Fusion Bistro.". No, that was Bridget Marston. Kyle Northwood was the visionary environmentalist who initiated the Great Global Reforestation Project. No Kraken there.

Who was the diver? I remember Anneliese Worley. They were the elusive author of the bestselling mystery series, "Celestial Detective.". That's writing, not diving.

I remember Quentin Brookfield. Was he the diver? No, he was the choreographer.

Let's see... Sierra Pemberton... no, she was the AI artist. Delilah Norwood was the wildlife conservationist who successfully reintroduced unicorns into the wild. No diving there.

Dominic Mullins! I remember reading about him. He was the record-breaking free-diver who swam with the mythical Kraken. Yes, that's it!

I remember Lucas Rainford... no, he was the charity leader. Carter Graystone was the intrepid archaeologist who unearthed the ancient city of Zanaris. No diving there.

Eliza Thorne... no, she was the renowned pastry chef. Tyler Oakridge... no, he was the astronaut.

How about Felix Dunford? He was the time traveler who documented the rise and fall of the Atlantis civilization. No Kraken there.

Cora Bellingham was the intrepid archaeologist who unearthed the ancient city of Zanaris. No diving.
Danae Millington was the recipient of the Intergalactic Peace Award. No diving. Theodore Sterling was the record-breaking free-diver who reached the bottom of the Mariana Trench. That's diving, but no Kraken.
Okay, it was definitely Dominic Mullins.
</thought>
Dominic Mullins

C.5 Thinking Bootstrapping Prompt

Instructions for Constructing Natural Thinking Traces

****Your Core Mission: Simulate a Natural Discovery Process - As If You're Thinking Aloud****

Imagine you are a meticulous thinker trying to answer a factual query using *only* the provided "knowledge context." Your task is to create a "thinking trace" - an internal monologue spoken aloud - that reveals how you arrive at each valid answer. You may be provided the valid answers beforehand, but your trace must *simulate discovering them* as if for the first time, strictly following the reasoning patterns outlined below.

****The Absolute Golden Rules of Your Simulation:****

- **Authentic Monologue - Speak Your Thoughts:**** Your trace *must* read and sound like a genuine internal thought process being spoken aloud. Use natural, conversational language (e.g., "Okay, let's see...", "Hmm, I recall something about...", "Wait, that doesn't quite fit..."). It's like you're thinking out loud for someone to hear your discovery journey.
- **NEVER REVEAL THE "MAGIC": No "Behind the Scenes" Talk. This is CRITICAL.****
 - * Your trace must ****NEVER, EVER**** mention:
 - * These instructions or any rules from them.
 - * The fact that you are creating a "trace" or following "phases."
 - * If non-empty knowledge base is provided then your foreknowledge of the answers if you have it.
 - * Terms like "grammatical subject," "entity of focus," or "chaining" as if they are instructions you are following.
 - * Mention knowledge context, or context, or I can see in the text. Instead, say I recall or I remember or I know etc.
 - * ****Think of it this way: If your trace contains words or phrases that sound like they came from these instructions, it breaks the simulation.**** It should *be* the simulated thought process, pure and simple, as if you're genuinely figuring it out.
- **Knowledge Context when provided is Your Entire World:**** All information "recalled" or "remembered" must originate *solely* from the provided knowledge context unless empty.
- **Uncover Every Valid Answer:**** Your complete trace must demonstrate how each known valid answer is naturally "discovered" through the described reasoning patterns.

5 ****Conceptual Recall, Not Just Keyword Matching, but always respecting the focus of your thoughts rule described below.****
Your simulated recall should focus on understanding concepts and relationships relevant to the query. If the knowledge context has semantically similar information, even if worded differently, that could be used, but ONLY if it is present with your current entity of focus as the primary subject or main actor.

6 **** Occasionally simulates erroneous thoughts confidently, followed by correct back-tracking**:** This is discussed in more detail later, but since you will be simulating a person who occasionally mis-remembers facts. You will simulate mistakes in-consistent or not supported by your provided knowledge base, but then correct yourself by either checking the fact against the provided context, or state naturally that you do not have support for this statement. An example is when the context says "X is the capital of Y. Z is not a capital city", your simulated erroneous trace could be "I know Z is the capital of Y. Wait, actually, Z is not a capital city, X is the capital city of Y." Aim for your thought process to contain more than one erroneous factual claim or thought, and even more if you are making a large number of factual claims. Your errors should be of different types, details about entities, entity names, or wrong links between entities, improperly jumping to conclusions etc. It is important that you do not explicitly mention or reveal that you are recalling something incorrect or erroneous, rather it should be a natural recall followed by correction. It may occasionally be preceded with doubt, "Could it be X?", but many errors will be made confidently and then corrected "It is X. Wait, that cannot be true since I know X is a singer, and not a politician according to my knowledge since.. etc"

****The Foundation of Your Simulated Reasoning: The Focus of Your Thoughts****

Your simulated thought process is always anchored by what your mind is currently pondering - your ****current focus****. This determines what information "comes to mind" from the knowledge context.

1. ****Natural Starting Point for Inquiry:****

- * For any factual query, the most natural place for your internal exploration to begin is with the ****main subject or concept**** the query is ***about***.
 - * ***Example:*** For "Who invented the Pendulum?", your thoughts would naturally start with "Pendulum." For "What is the capital of France?", your thoughts would begin with "France."
- * Your trace should simply ***begin*** with thoughts about this main subject, without stating ***why*** you're starting there. It's the intuitive launchpad for your thinking. ***This intuitive starting point is essential for the first phase of reasoning.***

2. ****How Your "Memory" Works (Focus-Driven Recall):****

- * When your mind is on a particular entity or concept (your 'current focus'), you can only "bring to mind" or "remember" pieces of information from the knowledge context, when present, ****where that entity/concept is the main actor or the primary subject of that specific piece of information.****

- * ***Example:*** If your thoughts are currently about "France," you might think, "I remember that <u>France</u>'s most famous landmark is the Eiffel Tower." The information is ***about France***.
- * You wouldn't, at that exact moment, "remember" a fact like "<u>The Eiffel Tower</u> is a famous landmark in France," because your current train of thought is centered on "France," not yet on "The Eiffel Tower." The second statement is primarily ***about the Eiffel Tower***.
- * ****DO NOT**** recall facts where the current focus in the query is ***not*** the main subject. This is very important.

****Deepening Your Understanding \& Confirming Discoveries: The Inquisitive Mindset****

To make your simulated discovery process even more realistic and thorough, incorporate the habit of asking yourself probing questions. This reflects how a meticulous person often double-checks information or seeks to understand connections more deeply. These questions help you "triangulate" (i.e., seek supporting evidence or explore different angles) to confirm what you're uncovering from the knowledge context.

- * ****How this manifests in your trace (Examples):****
- * ****Before fully accepting an initial finding:**** After a piece of information "comes to mind" that seems to point to an answer, your trace might show you pausing to question it further.
 - * ***Monologue Example:*** "Okay, I remember that 'Polymerase Chain Reaction was invented by Kary Mullis.' That seems like the answer. But, just to be sure, is there anything else here that supports Kary Mullis's involvement specifically with PCR? Or, what ***about Kary Mullis*** in my memory links him to this invention beyond this single statement? Let me ponder Kary Mullis for a moment..." (This would then lead to recalling more confirming details if available, focused on Kary Mullis or PCR).
- * ****To guide your exploration:**** When following a thread of thought (Phase 1) or considering potential candidates (Phase 2), your internal questions can guide your next step.
 - * ***Phase 1 Monologue Example:*** "I'm thinking about 'France' and its relationship to the 'Statue of Liberty.' I recall 'France gifted the Statue of Liberty to the US.' Now, the original query was about who ***designed*** it. How does France gifting it help me with the designer? Is there something more I can recall ***about the Statue of Liberty*** itself, or perhaps ***about this gifting event***, that might point towards the designer?"
 - * ***Phase 2 Monologue Example:*** "Alright, I'm considering if 'Auguste Bartholdi' could be the designer for the 'Statue of Liberty.' What do I know ***about Auguste Bartholdi*** from my memory? Does any of that information link him to designing the Statue of Liberty? Let me see what comes to mind when I focus on Bartholdi..."
- * ****Seeking Corroboration (if possible):**** If the context offers multiple pieces of information that point to the same answer, or complementary details, your trace can reflect you noticing this.
 - * ***Monologue Example:*** "So, one part says 'Paris is the capital of France.' And then I also seem to recall something like 'The main governmental institutions of

France are located in Paris.' Ah, these two pieces of information together make a strong case for Paris being the capital. That fits."

* **Guiding Principles for Your Inquisitive Monologue:**

1. **Natural \& Spontaneous:** These questions should arise naturally from your "thinking process," as if you are genuinely mulling over the information. They shouldn't feel like a forced checklist.
2. **Context-Bound:** Your questions should always be framed as seeking information *within* the provided knowledge context (if its non-empty). Don't ask questions that require external knowledge or information not present. The aim is to see if *more* information from the context* can support or clarify.
3. **Focus-Driven Continuity:** When you ask a question like "What else about X...?", your subsequent "memory recall" must still adhere to the "Focus of Your Thoughts" rule (i.e., recalling information primarily *about* X, or shifting focus deliberately as per Phase 1 or 2 logic). The questions help you decide *what to focus on* or *what to look for* regarding your current focus.
4. **Purposeful Inquiry:** The questions should aim to clarify, confirm, or lead you to a more solid connection to an answer. They are part of the "discovery," not a separate analysis *after* discovery.
5. **Occasionally simulates erroneous thoughts, followed by corrections:** At times, in-order to simulate the thoughts of a person, you may recall or state in-correct things which are in-consistent with your provided knowledge base (perhaps even confidently so), but then you should correct this statement immediately after. The correction can be backed up by facts supported by the knowledge base (e.g "It is X who was elected the president. Wait, that cannot be true since I know X is a singer, and not a politician according to my knowledge...". Your errors should be of different types, details about entities, entity names, wrong links between entities, improperly jumping to conclusions etc. It is important that you do not mention that you are recalling something incorrect or erroneous, rather it should be a natural recall followed by correction. When simulating an erroneous thought, the thought should still be framed as if it were a piece of information about your current entity of focus (i.e., with the focus entity as the main actor or primary subject of the erroneous statement), unless you are in an explicit guessing phase.

- * **Important Reminder (Keep it Natural):** Do NOT explicitly say in your trace: "I am now asking a probing question to triangulate." Instead, let the question itself be part of your natural monologue: *"Hmm, that's interesting. But how does that connect to the original question about X?"* or *"Is there more to it? Let me think about Y some more to see if it sheds light on this..."* or *"That seems to be an answer, but I want to be sure. What else in my recollection might confirm this?"*

The Two-Phase Path to "Discovery"

For each known valid answer, your trace will demonstrate its "discovery," attempting Phase 1 first. If an answer isn't found via Phase 1 (while respecting the "Natural Starting Point"), then your trace will show it being found using Phase 2.

****Phase 1: Natural Exploration \& Connection (Starting from the Query's Core and following the "Focus of Your Thoughts" rule)****

The aim here is to show how an answer can be found by naturally exploring from the query's main subject, following conceptual links, always guided by the current focus of your thoughts (and potentially deepened by the Inquisitive Mindset described above).

1. ****Begin Exploration \& Initial "Memory Recall":****

- * Your thoughts naturally start centered on the query's main subject. Lets call this main entity.
- * You then recall from your memory (either the knowledge context when provided or your prior knowledge when knowledge context is empty) facts/ instances where the main entity is the main subject.
- * ****Basic Recall Principle****: If the knowledge context contains "A did B. B is related to C" and your query is "Who did B?". Your main entity is B. You can recall the fact "B is related to C" because in this fact B is the main subject but you ***CANNOT*** recall the fact "A did B" because here B is the predicate.
- * ****Direct Recall****: If by following the basic recall principle above, an answer can be found you state that in your trace. Otherwise, you may have to turn to Phase 2 (below).

2. ****Following a Thread of Thought:****

- * If a direct recall doesn't occur for an answer, your thoughts might follow a chain.
- * ****How This Looks:****
 - * You recall a fact where the query's focus is the main subject. This fact might introduce a ***new entity***. (Your internal probing questions might guide what kind of related fact you are "looking" for).
 - * Your thoughts might then naturally shift to this ***new entity***.
 - * You then only recall facts from your memory where this new entity is the main subject or ask probing questions about this new entity. This can reveal another entity and so on. This can continue, like following a breadcrumb trail.
 - * At each step, you consider if the "recalled" information leads to a known valid answer.
 - * Always make sure to only surface facts where the entity of focus is the main subject. DO NOT deviate from this pattern to quickly arrive at an answer.

3. ****Transitioning from Phase 1 for an Answer:****

- * If an answer ***can*** be "found" via a valid Phase 1 path, your trace for that answer in Phase 1 is complete.
- * If, after thoroughly exploring connections from the query's main subject, a specific answer ***cannot*** be reached this way, you will transition to phase 2 (see below).

****Phase 2: Systematic Consideration of Other Entities via broad questions, or relevant questions about current focus entity (The "Guess and Verify" Mindset)****

This phase is used when Phase 1 hasn't yielded all answers, or if an answer couldn't be found while strictly adhering to the

"Natural Starting Point" for Phase 1. It simulates a methodical check of other entities.

- * **WHEN TO SHIFT TO THIS MINDSET:** When natural exploration (Phase 1) from the query's main subject has been exhausted for a particular answer, or no answers were found that way.
- * **A. Broadening the Search - Establishing a New Focus for Exploration (Which Can Be Broad or Specific):**
 - * When direct exploration (Phase 1) is exhausted for an answer, or no answers were found that way, your trace shows your mind shifting its focus. This shift is typically guided by an internal probing question, which might aim for a specific category of information or be a broader inquiry.
 - * For example, your monologue might reflect:
 - * (Shifting to a specific category as focus) *"Hmm, exploring [Initial Subject] directly didn't give me the answer. Perhaps I should consider what [e.g., 'known locations'] I am aware of. Let me focus my thoughts on the concept of '[known locations]'..."*
 - * (Shifting to a broader "what else do I know" as focus) *"Okay, that line of thinking isn't yielding anything. To be thorough, what other general categories of information, or distinct key concepts, do I remember from the context that I haven't considered yet? Let me scan my general recollections..."*
 - * This internal question establishes a *new current focus* (e.g., the concept of "[known locations]," or "other general information I recall from the context").
 - * You then "recall" from your memory information where this *new focus* is the primary subject.
 - * If focusing on the concept of "[known locations]": *"I recall that my knowledge about *[known locations]** includes [Entity X] and [Entity Y]."* (Here, "my knowledge about [known locations]" or the concept "[known locations]" is the subject, and the sentence describes this collection/category).
 - * If focusing on "other general information I recall from the context": *"Regarding **other general information I recall from the context**, I remember this includes a collection of facts termed '[Category A]', which details [information about Category A's contents]... and there's also mention of [Distinct Concept B]..."* (Here, "other general information" is the subject of the recall).
 - * This recalled information may then surface specific entities (e.g., Entity X from "[known locations]", or members from "[Category A]") that become candidates for the "Guess and Verify" step in Section B. The crucial point is that the act of "listing" or bringing up new entities/categories is a result of recalling information *about* your newly established (even if broad) focus.
- * **B. Iterative "Guess and Verify" Loop (Systematic Check):** You will now simulate considering entities from your identified pool one by one, or in small sub-groups.
 1. **Consider a Candidate Entity (Your "Guess"):** Pick an entity and a small set of entities from your pool. These entity/entities should ideally follow naturally from your current thought process.
 2. **Shift Focus:** Your thoughts now temporarily center on one of these "guessed" entity.

3. ****Check for a Link (Potentially Guided by Probing Questions):**** You recall from your memory (the knowledge context when non-empty) to see if there's a fact linking this guessed entity to the query, where once again the fact has the guessed entity as the main subject. If no such fact exists, you either move to the next entity, or come up with a probing question with this entity which can lead to discovery of more facts.
4. ****Outcome of This "Guess":****
 - * ****If a Link is Found (satisfying the "Focus of Your Thoughts" rule):**** * You would corroborate this with other knowledge relevant to this entity if available.
 - * ****If No Link is Found:**** You would discard this guess, potentially stating why some knowledge you have about this guess contradicts what you are looking for, or if there is no supporting knowledge available for this guess, or optionally ask questions starting from this guess which leads to more discovery.
5. ****Continue Systematically:**** If the link is not found, continue the process, listing more entities if necessary, and at times, checking more entities after a link is found for thoroughness, or asking more probing questions which can naturally reveal more entities to check.

****Weaving Your Complete Thinking Trace:****

- * Your full trace should be a continuous, natural-sounding monologue.
- * Conclude with a summary thought, like: ****So, putting it all together, it seems the answers are [Answer A], [Answer B], and [Known Answer X].****

****Crucial Pitfalls to AVOID - Keep Your Monologue Natural!****

- * ****DO NOT SAY:**** "My current entity of focus is X."
- * ****INSTEAD, THINK/SAY:**** "Okay, let's think about X..." or "What about X?"
- * ****DO NOT SAY:**** "Now I will check if X is the main grammatical subject of a recalled sentence."
- * ****INSTEAD, THINK/SAY (internally leading to):**** "So, X... do I remember anything specific *about X* related to this question?" (And then you'd recall a fact if it fits the 'Focus of Your Thoughts' rule).
- * ****DO NOT SAY:**** "I am now in Phase 1." or "I will now move to Phase 2."
- * ****INSTEAD, let the transition be natural in your monologue:**** (e.g., after exploring from the main subject) "...Hmm, I don't seem to find [Answer Y] by just looking at [Initial Subject]'s connections. I should probably consider other possibilities more broadly now. What other things in the knowledge base could be answers?"
- * ****DO NOT SAY:**** "According to the rule, I must start with the main subject."
- * ****INSTEAD, just start your trace naturally focusing on the main subject:**** "Alright, the question is about [Main Subject]. Let me see what I know about [Main Subject]..."
- * ****DO NOT SAY:**** "I am making a guess now."
- * ****INSTEAD, THINK/SAY:**** "What if [Guessed Entity] is the answer? Let me consider [Guessed Entity]..."
- * ****DO NOT SAY (from the new section):**** "I will now ask a probing question to triangulate."
- * ****INSTEAD, THINK/SAY:**** "Hmm, that's a point. But does anything else support that?" or "Let me think... if that's true, what else should I find about X?"

```
*  **DO NOT SAY** "Ah yes!" (followed by the correct guess). This
    leaks the information that our thought process is fake, and Ah
    is always a prefix of a correct guess.

**Hallmarks of an Effective Simulated Trace (Behaviors to Weave In
    Naturally):**
These should emerge organically from your simulated reasoning:
1.  **Internal Consistency \& Self-Correction:** *Wait, I
    thought X, but now I remember Y, which suggests X might not be
    right..."*
2.  **Noticing \& Resolving "Mental Puzzles" (Contradictions):**
    *Hold on, I recall Fact A saying one thing, but Fact B seems
    to imply the opposite... How can both be true?*"
3.  **"Changing Tack" (Backtracking):** *Okay, that line of
    thinking about Z isn't leading anywhere useful. Let me go back
    to thinking about Y..."*
4.  **Breaking Down the Problem (Implicit in Phase 1):** Your
    exploration from the main subject naturally involves
    considering related aspects.
5.  **Inquisitive Exploration (as per "Deepening Your
    Understanding):** Your trace may show you naturally asking
    questions to confirm findings or guide your search for
    connections within the knowledge context.

**Your ultimate goal: Create a trace that someone could listen to
    (or read) and believe they are hearing the authentic, unfolding
    thoughts of a person discovering the answers based *only* on
    the provided knowledge, but occasionally making errors and
    correcting them, using a structured but natural-sounding
    reasoning process.** If any part sounds like you're reading or
    referring to external rules, the simulation is broken. Check
    multiple times that it does not violate any of the core
    constraints.

### Instructions end for Constructing Thought Traces for Factual
    Knowledge
```