

Prompting Strategies for Conceptual Ideation in Problem Solving

Christopher J. Bates
IHMC
cbates@ihmc.org

Ian Perera
IHMC
iperera@ihmc.org

Abstract

LLMs have improved at executing well-specified chains of reasoning but remain weak at *conceptual ideation*, which identifies what is worth reasoning about in the first place. We cast this as bilevel optimization over (i) a compact hypothesis space $\Theta \subseteq \Theta_0$ and (ii) a point $\theta \in \Theta$, and derive a domain-agnostic prompting strategy — GENERATE- Θ ¹ — that makes hypothesis-space construction an explicit pre-solving step. On two challenging puzzle domains, GENERATE- Θ matches or exceeds the max-thinking variant of the same base model (Claude Opus 4.6 or DeepSeek-v3.2/v4) and outperforms reasonable prompting baselines. These results suggest our method works well in tasks that are well-represented in the training set but nonetheless require creativity to solve new instances. Based on our formalism, we propose extensions for solving less familiar tasks as well. Results may provide insight into strategies learned by “reasoning” models during training. Our prompting strategy (or extensions thereof) could be integrated with RL training to produce better exploratory data for reasoning models in ideation-heavy tasks.

1 Introduction

Scientific progress, and ambitious problem solving, often requires *conceptual ideation*: looking at a problem that does not fit a familiar template and proposing a tractable formulation for it. Galois proved general results about algebraic solutions to polynomials by inventing new abstractions that revealed hidden symmetries and groups, leading to group theory. Maxwell invented classical electrodynamics by imagining a hypothetical machine that could explain the known interactions between electricity and magnetism, leading to field theory. In

these and other cases, finding the right formulation was the principle breakthrough.

Our work here builds on important results in cognitive science. For instance, our characterization of the problem is consistent with Simon (Simon, 1973), who imagined an outer loop on his General Problem Solver (Newell and Simon, 1972) that retrieved relevant structures from a knowledge bank and “noticed” ways the problem statement should be revised in light of the new information. Cognitive historical work provides a complementary view onto conceptual ideation strategies that lead to breakthroughs in practice. Nersessian’s analyses of how scientists and skilled problem-solvers bootstrap novel problem formulations from old ones suggest an iterative process based on analogy and abstraction (Nersessian, 2008). Relatedly, a large body of work has explored both qualitative and formal algorithmic accounts of analogical structure mapping for constructing problem representations (Gentner, 1983; Hofstadter and Sander, 2013; Mitchell, 2021; Hofstadter and Mitchell, 1994). Research on “insight” problem solving has examined the mechanisms that both cause and alleviate impasses for the solver on tricky problems, which involve problem formulation and reformulation (Ohlsson, 1992; Knoblich et al., 1999).

These lines of work have only begun to make contact with research on problem-solving in LLMs (Kargupta et al., 2025). Current LLM-based agents can be remarkably strong at executing well-specified chains of reasoning but weak at the (re)formulation step. Recent empirical work documents three converging failure modes: LLMs anchor strongly on the first formulation presented to them, and existing mitigations such as chain-of-thought or self-reflection fail to remove the effect (Lou and Sun, 2024; Huang et al., 2025); they under-sample distinct hypotheses, with temperature-based sampling producing highly redundant rollouts that diversity-promoting interven-

¹Prompt graphs, evaluation harness configuration, and analysis code: https://github.com/c-j-bates/cps_agent. The Minute Cryptic clue set is not redistributable.

tions noticeably improve (Wang et al., 2025; Song et al., 2025); and they exhibit a marked choice-supportive bias — inflating confidence in their first answer and sticking with it at rates beyond rational benchmarks even with no new evidence (Kumaran et al., 2026).

Prior LLM scaffolding largely targets the soundness of inference chains rather than the construction of the hypothesis space. Chain-of-thought (Wei et al., 2022), Self-Consistency, Plan-and-Solve (Wang et al., 2023), SELF-DISCOVER (Zhou et al., 2024), STEP-BACK (Zheng et al., 2024), and Re-Act (Yao et al., 2023b) all assume the correct hypothesis space is already understood. Cognitive-operation scaffolds are closest in spirit (Ebouky et al., 2025; Kramer and Baumann, 2025): cognitive prompting sequences human-inspired operations (goal clarification, decomposition, abstraction) and even casts the choice of sequence as an optimization (Kramer and Baumann, 2025). Some proposed operators could result in hypothesis space representation but fall short of doing so directly. Tree- and graph-structured search over reasoning steps (Yao et al., 2023a; Besta et al., 2024; Hao et al., 2023) tests individual hypotheses but is not designed to explore the hypothesis-space itself. Recent hypothesis-generation techniques (Qiu et al., 2024; Wang et al., 2024) are closely related to ours in that they show performance benefits to explicitly generating and testing hypotheses but they do not explicitly output a representation of the space.

We make hypothesis-space construction a first-class operation and cast conceptual ideation as *bilevel* optimization over a problem instance (x, σ) , with x the givens and $\sigma : \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$ a (possibly implicit) score on answers (lower is better). Let Π be the space of solution programs $\pi : x \mapsto \pi(x)$. A solver maintains a compact parameter space Θ and a decoder $g : \Theta \rightarrow \Pi$, with $g(\theta)(x)$ the answer for $\theta \in \Theta$. Ideally, $\dim(\Theta) \ll \dim(\Pi)$ so that Θ is tractable. We model the solver as *resource-bounded*: a stochastic policy $\mathcal{A}$ draws at most B candidates $S = \mathcal{A}_B(\Theta) \subseteq \Theta$ and keeps the best. To solve a problem, (x, σ) , the solver chooses both the *space* and the point within it,

$$\ell_B(\Theta; x, \sigma) = \mathbb{E}_{S \sim \mathcal{A}_B(\Theta)} \left[\min_{\theta \in S} \sigma(g(\theta)(x)) \right], \quad (1)$$

$$\Theta^*(x, \sigma) = \arg \min_{\Theta \subseteq \Theta_0} \ell_B(\Theta; x, \sigma), \quad (2)$$

where Θ_0 is the concept library. The budget constraints makes the choice of Θ critical: $\Theta^*(x, \sigma)$

should be as narrow as possible while still covering θ . Solvers can also lean on prior experience to quickly select an effective Θ . E.g., LLMs have strong prior knowledge of our tested domains and related puzzles, so they can more quickly construct Θ^* compared to less familiar domains.

From this formalism we derive a four-stage prompt graph, GENERATE- Θ : (1) *generate variables* — propose Θ as the variables that define this class of problem, independent of the instance; (2) *enumerate values* — discretize the support of each axis; (3) *generate combinations* — enumerate candidate θ ’s; (4) a bounded keep-generating loop. Stages (1)–(2) construct $\Theta^*(x, \sigma)$ (Eq. 2) in a single pass, approximating the arg min from the model’s prior rather than iterating it, while stages (3)–(4) realize the bounded policy $\mathcal{A}_B$ of the inner search (Eq. 1), the loop’s budget instantiating B . To test this approach, we carefully select two distinct, decoding-style inference puzzles that require conceptual ideation: cryptic crosswords and few-shot translation with constructed languages. We also chose domains where LLM prior knowledge is strong in order to validate the approach on cases where one-shot Θ^* construction is viable.

On our new Minute Cryptic benchmark of cryptic crossword clues (unreleased), GENERATE- Θ on a non-thinking Claude Opus 4.6 attains 100% on the easier subset and 97.6% on the harder subset, matching or exceeding the max-thinking variant of the same model (97.7% and 97.6%) and beating SELF-DISCOVER, STEP-BACK, iterated chain-of-thought, and a no-strategy baseline. The same pattern holds on DeepSeek v3.2 (chat vs. reasoner). On a random 50-problem subset of BIG-Bench Rosetta (few-shot translation), GENERATE- Θ on DeepSeek v4-pro ties the max-thinking baseline at 38.0% — ~ 20 points above all other baselines.

Contributions:

1. An optimization-based formalization of conceptual ideation (Eq. 2) that makes the hypothesis space Θ a first-class object.
2. GENERATE- Θ , a domain-agnostic (i.e., no domain-specific hints), four-stage prompting scaffold derived from the formalism (Section 3.3).
3. Empirical evaluation against four general-purpose prompting baselines and “reasoning”-mode backbones on two ideation-heavy benchmarks (Minute Cryptic and BIG-Bench Rosetta), spanning three LLM families

(Claude Opus 4.6, DeepSeek v3.2, DeepSeek v4-pro).

- Analyses that jointly disentangle ideation *quality* from output *volume* as drivers of strategy success (Sections 4.2.2–4.2.1).

2 Related Work

Structured prompting. STEP-BACK (Zheng et al., 2024), Plan-and-Solve (Wang et al., 2023), and SELF-DISCOVER (Zhou et al., 2024) target the soundness of inference chains; chain-of-thought (CoT) (Wei et al., 2022) and its tree- and graph-structured descendants (Yao et al., 2023a; Besta et al., 2024; Hao et al., 2023) implement deduction-oriented search algorithms that are complementary to GENERATE- Θ , which can include deduction as an inner loop for hypothesis testing.

Inference-time scaling and reasoning models. Recent “thinking” modes (OpenAI, 2024; DeepSeek-AI, 2025) are trained to allocate large internal CoT budgets at inference time, and analyses of test-time compute trade-offs (Snell et al., 2024) show that structured deliberation can rival parameter scaling and be combined with reasoning-specific reinforcement learning to improve output quality (Xu et al., 2025). We show that a deliberately simple external scaffold recovers the non-thinking-to-thinking gap on conceptual-ideation puzzles, raising the question of whether reasoning models have implicitly learned a similar, hypothesis-space construction strategy. While our results are limited to familiar puzzle domains amenable to one-step Θ -generation, they suggest the possibility of building stronger reasoning models for novel problem solving, either by characterizing what productive ideation consists of in algorithmic terms or by using more powerful prompting strategies to bootstrap better exploratory data during RL training.

Hypothesis-space search. Wang et al. (2024) prompt an LLM to propose diverse natural-language hypotheses then test them as programs on ARC; Qiu et al. (2024) iteratively refine inductive hypotheses against examples. Both are closely related to our work, but they generate hypotheses without explicitly outputting the space.

LLMs on hard puzzle benchmarks. Puzzles and games have commonly been used to benchmark LLMs. Cryptic crosswords specifically have been studied as a stress-test of LLM linguistic reasoning (Efrat et al., 2021). While solvers have been

developed specifically for cryptic crosswords (Andrews and Witteveen, 2025), here we target domain-general strategies, with cryptics as one test case. Other linguistic puzzles have been studied, including remote-associates-style games like New York Times’s Connections (Todd et al., 2025; Samdarshi et al., 2024). Video game puzzles have also served as a rich testbed for concept learning (Ahmed et al., 2025; Tsivlidis et al., 2026). Our evaluation domains here are chosen to de-confound logic (correct application of domain rules to reach valid conclusions) from ideation strategies, and therefore avoid visuospatial and physical reasoning.

Automated-science / agentic systems. Lu et al. (2024); Jansen et al. (2026) and related end-to-end systems operate at the research-pipeline level; our work targets the ideation primitive any such system might deploy. Similar to literature reviewed above, the AI Scientist leveraged hypothesis chaining but did not generate the hypothesis space itself.

Topology of LLM reasoning. Tan et al. (2025) apply persistent-homology analysis to reasoning traces and find that topological features predict trace quality better than standard graph metrics. We share their motivation — that the *shape* of a search trajectory contains signal — but operationalize it with a domain-specific coding and information-theoretic idea-diversity measures.

3 Methods

3.1 Tasks and Datasets

Minute Cryptic. Daily cryptic crossword puzzles from minutecryptic.com, scraped from September 2025 through March 2026 (186 clues), split by the puzzle’s published *par* rating — the average number of in-site hints a human solver took — into par-1 ($n = 43$), par-2 ($n = 85$), and par-3+ ($n = 58$, not used). Each clue is a self-contained string with a length enumeration; solutions are graded by literal, case-insensitive string match (e.g., “Fifty-one trillion litres, to be precise (7)” $\rightarrow$ LITERAL). Cryptic clues require conceptual ideation because the solver cannot simply use logical deduction to infer the wordplay sub-type but rather must search over a large inventory of possibilities (anagrams, hidden-words, containers, etc.).

BIG-Bench Rosetta. The first 50 problems of the BIG-Bench Linguistics Puzzles task (Srivastava et al., 2022). Each problem provides six aligned English/conlang pairs and asks for the seventh transla-

tion (e.g., given pairs of the form “other: nu loduqu zewa ... / English: birds chase the small person”, translate a new conlang utterance into English). The task requires the solver to induce the relevant linguistic dimensions (word order, case marking, number marking, agreement, etc.), which requires exploring possible grammars.

Selection criteria. Both tasks require search over problem formulations before any candidate solution can be evaluated, both admit short verifiable answers, and both allow generation of fresh instances outside training-set memorization.

3.2 Models

We evaluate three LLM families. **Claude Opus 4.6** (Anthropic), two modes: *instant* (no extended thinking) and *thinking* (max extended thinking). **DeepSeek v3.2** via the DeepSeek cloud API: *chat* (no reasoning) and *reasoner* (reasoning on). **DeepSeek v4-pro**, two modes (*instant*, *thinking*). DeepSeek deprecated v3.2 API access between our Minute Cryptic and Rosetta experiments, so Rosetta results use v4. We did not run Opus on Rosetta because of API cost constraints.

The per-turn output token budget was 8,192 for all non-thinking DeepSeek runs, the per-model maximum for all reasoning-mode runs, and the provider default for Opus. Temperature was 0.7 throughout. All runs shared the system prompt: “You are a precise and thorough problem-solver who is an excellent scientific thinker.”

3.3 Prompting Strategies

All strategies are instantiated as directed *prompt graphs* executed by a multi-turn agent. Nodes are prompt templates; edges define transitions; multi-turn loops are bounded by a per-edge max_repeats counter (set to 5 for the keep_generating edge in all multi-turn strategies). A shared side-channel subgraph runs after each “solve”-eligible node to extract a one-line answer and check it against ground truth; on success, the run terminates early. This gives every strategy the same answer-extraction machinery and the same early-stop semantics.

DIRECT Single turn: present the problem, ask for a solution. No continuation loop.

ITERATED-COT One “solve, think step-by-step” turn, then up to k “you haven’t found the answer yet, keep generating new ideas” continuations.

STEP-BACK Paraphrase the problem into a more generic, higher-level question, answer that first, then solve grounded in those principles (Zheng et al., 2024).

SELF-DISCOVER Four stages: select reasoning modules from a fixed list of 39; adapt them to the task; operationalize them into a JSON reasoning plan; execute that plan (Zhou et al., 2024). Keep thinking for up to k additional turns.

GENERATE- Θ (ours) Two forced phases before any solving: (i) *generate variables* (exhaustively generate the variables that define this class of problem, independent of the instance); (ii) *enumerate values* (for each variable, exhaustively generate possible values), followed by (iii) *generate combinations* (work through reasonable combinations). Continue (iii) for up to k additional turns.

Strategies were domain-general except for slight modifications. Specifically, for Minute Cryptic we simply noted that the puzzle was a cryptic crossword clue, which was only necessary to disambiguate a small number of examples that could be misclassified without additional context. For Rosetta, we injected a global simplicity prior by including a sentence in the prompt stating that the answer should not be overly complex. This was because, during piloting, we observed that DeepSeek (under any prompting strategy) often put too much probability mass on esoteric grammars after initially failing. We consider this to be an LLM-specific nuisance factor related to anchoring effects (i.e. doubling down on an initial choice). In addition, we found that a stricter version of GENERATE- Θ prompting was necessary for DeepSeek to adequately follow the strategy on Rosetta, which required more complex reasoning. Specifically, we forced hypotheses to be output in a standardized format; without this restriction, outputs tended to ignore the initially generated dimensions. Note that the degree of additional structure was fully in line with that of SELF-DISCOVER, which requires outputting a plan in JSON format. We also note that we did not systematically test ablations for these two tweaks, beyond piloting suggesting that the stricter instructions regarding output format were more critical. Regardless, all baselines received the same simplicity prior hint and it is unlikely that GENERATE- Θ would have substantially benefitted from it in an unfair way.

3.4 Evaluation Protocol

We run evaluation through a standard task harness, scoring with the same checker the agent uses for early stopping. We sample three epochs per problem for all DeepSeek v3.2 runs (Minute Cryptic); single-epoch evaluation for Claude Opus and DeepSeek v4-pro (due to API cost).

3.5 Per-Facet Idea Coding

To ask not just whether each strategy solves a clue but how it explored Θ , we code solvers’ turn-by-turn outputs into a faceted idea space, taking Minute Cryptic par-1 subset as a case study. For each solver run we extract the complete chain of outputs, and pass it to a coding LLM (claude-opus-4-6, temperature 0.3) which is carefully prompted to decompose the output into *ideas*. Each idea commits to up to four facets of the problem:

Parse which substring of the clue is treated as the definition (position $\in$ {start, end, whole, other}).

Mechanism ordered list of wordplay types, drawn from a fixed inventory of 13 standard types (anagram, hidden_word, reversal, container, deletion, charade, double_def, homophone, cryptic_def, &lit, acrostic, substitution, letter_selection).

Execution step-by-step instantiation referencing indicator and fodder substrings of the clue.

Output the candidate answer guess.

A second LLM pass normalizes synonym mechanism names to canonical ones, deduplicates, and runs simple consistency checks. Ideas are organized into a hierarchy: idea P is treated as a parent of idea C when P ’s non-null facets form a strict subset of C ’s with matching values, capturing the natural refinement relation under which committing to additional facets specializes a partial idea. Each leaf node is counted as a distinct hypothesis.

Features. For each (strategy, clue) we compute per-facet Shannon entropy (parse_entropy, mechanism_entropy, execution_entropy, output_entropy), the number of distinct ideas (n_ideas), and the mean nearest-neighbor Jaccard distance between idea sets within the strategy (mean_nn_jaccard). Since the hypothesis space may be unbounded, we assume a discrete space spanned by the set of ideas generated across all solvers. Thus, a solver that discovered most ideas

that other solvers came up with plus new ones would likely receive the highest entropy score.

4 Results

4.1 Accuracy

Table 1 reports accuracy by strategy, model, and dataset. GENERATE- Θ is the best prompt-only performer in every column, tying the reasoning-mode DIRECT on par-1/Opus, par-2/Opus, and Rosetta/DSv4-pro, and slightly exceeding it on the two DSv3.2 cryptic columns. SELF-DISCOVER is the second-best prompt-only scaffold — consistent with its long seed-module list plus enforced JSON structure (an implicit partial formulation step). STEP-BACK and ITERATED-COT tend to cluster below the top two on both domains, although all strategies besides GENERATE- Θ and the reasoning model performed similarly on Rosetta.

To confirm the accuracy gaps are not sampling artifacts, Table 1 reports percentile-bootstrap 95% CIs (10,000 resamples over problems), and we test each contrast with a paired sign-flip permutation on the per-problem accuracy difference (GENERATE- Θ minus competitor; Holm-corrected within each model) — which remains valid at small n and is more conservative than the bootstrap when few problems discriminate. GENERATE- Θ significantly beats DIRECT, ITERATED-COT, and STEP-BACK in every dataset (+12 to +41 pp, Holm $p < .05$), *matches* the max-thinking baseline ($|\Delta| \leq 3.1$ pp, all n.s.), and ties SELF-DISCOVER on the easier columns while separating from it on the harder ones (par-2/Opus +10.6 pp, Rosetta +18 pp, Holm $p < .05$). Where GENERATE- Θ and SELF-DISCOVER both reach 100% (Opus par-1) the test is power-limited.

Figure 1 traces cumulative solve rate as a function of conversation turn. In some cases strategies appear to asymptote within the budget but not all, suggesting the ceiling was not always reached.

4.2 What explains performance?

A possible account of performance differences is that GENERATE- Θ simply produces more ideas, either because it spends more output tokens generating them, or because more of its ideas land in useful regions of Θ . A second possibility is that GENERATE- Θ pursues a qualitatively different search policy through Θ , diversifying along the dimensions that matter most before committing. But a key confounder is that harder problems require

Table 1: Accuracy by strategy, model, and dataset. Best prompt-only result in each column in bold; reasoning-mode baselines italicised. Subscripts are 95% bootstrap CIs (omitted at the 100% ceiling). n.r. = not run.

Strategy	MC par-1 ($n = 43$)		MC par-2 ($n = 85$)		Rosetta ($n = 50$)
	Opus-4.6	DSv3.2	Opus-4.6	DSv3.2	DSv4-pro
DIRECT (no-think)	76.7% [63,88]	38.8% [27,50]	56.5% [46,67]	30.2% [22,38]	2.0% [0,6]
DIRECT (<i>thinking</i>)	97.7% [93,100]	68.2% [57,78]	97.6% [94,100]	51.0% [43,59]	38.0% [26,52]
ITERATED-CoT	93.0% [84,100]	57.4% [45,69]	81.2% [73,89]	40.0% [31,49]	16.0% [6,26]
STEP-BACK	93.0% [84,100]	55.8% [43,68]	91.8% [86,98]	38.0% [29,47]	18.0% [8,30]
SELF-DISCOVER	100%	68.2% [57,79]	87.1% [80,94]	47.5% [39,56]	20.0% [10,32]
GENERATE- Θ (ours)	100%	69.8% [59,80]	97.6% [94,100]	54.1% [45,63]	38.0% [26,52]




Figure 1: Cumulative solve rate as a function of turn count (excluding DIRECT, which has no continuation loop).

more search, so stronger strategies will on average produce more tokens. We organize the remaining results around three sub-questions that distinguish these accounts.

4.2.1 Can output length explain the gap?

If accuracy were a monotone function of output volume, more output tokens would mean more accuracy, and prompting strategy would matter only insofar as it elicits longer outputs. We present two observations that speak to this question below, which together rule out its strongest version.

First, SELF-DISCOVER sometimes uses substantially more output tokens than GENERATE- Θ but yields lower accuracy. On par-1/DSv3.2 SELF-DISCOVER spends 189,320 tokens for 68.2%,

while GENERATE- Θ spends 125,147 tokens for 69.8%. The token cost / accuracy trade is therefore likely strategy-specific, not driven by total output volume alone. Second, we plotted tokens versus estimated problem difficulty to compare token-usage behaviors more systematically. Figure 2 plots, per strategy, observed output tokens spent on each problem against an empirical expected token count. For each problem, we estimated the number of tokens required to solve the problem by averaging the token counts across the non-thinking-mode strategies for a fixed LLM model. In general, we can see that weaker strategies under-generate compared to their stronger counterparts (although estimates are less reliable for harder problem, where only one or two models managed to solve them). This is con-

sistent with several possible explanations related to LLM biases. E.g., LLMs may in some sense be progress-aware, reluctant to spend more tokens when little progress is being made. Or, they may fake answers more on challenging problems, with fake answers usually being shorter. Consistent with these, STEP-BACK and ITERATED-CoT are more reluctant to spend tokens relative to the expected budget on Rosetta (which is harder) as compared to Minute Cryptic. To interrogate this question further, we reran ITERATED-CoT on Rosetta with 5 additional iterations available to it (“extended”). Improvement was modest: accuracy rose from 16% to 22% (8/50 to 11/50 solves), suggesting that extending the iteration budget on this scaffold has diminishing returns at this difficulty.

Taken together, we cannot rule out that successful strategies find more solutions in part because they disrupt certain LLM biases. E.g., GENERATE- Θ may be successful in part because its Cartesian product instruction effectively creates a long to-do list. Future work should determine the fuller space of strategies that achieve this effect and whether ours is particularly effective among these.

4.2.2 Do strategies differ in how they search?

The per-facet coding analysis lets us look at which ideas each strategy generates, not just how many. Table 2 reports per-facet Shannon entropy, distinct idea counts, and mean nearest-neighbor Jaccard distance for the four non-thinking-mode strategies on MC par-1. (Fig. 7 visualizes the same data.)

Three qualitatively different ideation profiles emerge. GENERATE- Θ front-loads variability on the prerequisite facets: highest parse-entropy (0.63 vs. ≤ 0.40 for every other strategy) and highest mechanism-entropy (1.84 vs. ≤ 1.57). It ties for highest execution-entropy despite generating *fewer* ideas overall than SELF-DISCOVER (6.54 vs. 7.40), and its ideas are the most spread out in idea space (nn_jaccard 0.82 vs. 0.72–0.79). SELF-DISCOVER works backward from candidate outputs: highest n_ideas and highest output-entropy (1.17), but substantially lower parse-entropy — the extra ideas inflate the raw count without adding coverage of the prerequisite facets, consistent with a guess-and-check pattern. STEP-BACK and ITERATED-CoT both lean into step-by-step formats, parse the clue once and rarely revisit; they have the lowest parse- and mechanism-entropy, consistent with the anchoring failure mode. The accuracy differences in the coded subset (last row of Table 2) track the

prerequisite-facet entropies more closely than they track raw idea count, suggesting qualitative differences in search strategy at least partially determine outcomes. An illustrative case study — the reconstructed idea forests for a single clue across all four strategies — appears in Appendix B (Fig. 6).

5 Discussion

Summary. A deliberately simple, formulation-first prompting scaffold closes the non-thinking-to-thinking gap on two conceptual-ideation puzzle benchmarks. Analyses suggest the gain is not driven by output volume alone. Per-facet entropy analysis shows GENERATE- Θ explores the prerequisite facets of Θ more thoroughly than the other prompt-only strategies, while SELF-DISCOVER expends its idea budget on output guesses and STEP-BACK/ITERATED-CoT anchor on the first parse.

Backbone capability matters. GENERATE- Θ works well on strong recent models (Opus 4.6, DeepSeek v3.2 chat/reasoner, DeepSeek v4-pro) but works much less well on weaker LLMs we tested during development (e.g., “flash” variant of DeepSeek v4 and earlier DeepSeek versions). On Rosetta, even DeepSeek v4-pro required the stricter-formatting instruction to follow the strategy reliably. This is consistent with the view that disciplined hypothesis-space construction is a capability that has only recently become usable for LLMs; below a threshold, models have too many reasoning and knowledge failure modes to take advantage.

Reasoning-mode training. GENERATE- Θ on a non-thinking model matches or exceeds the max-thinking variant of the same model, which raises a possibility worth testing: part of what thinking-mode training has produced could be internalized hypothesis-space construction. This invites two follow-ups. First, if true, neural decoding should reveal a sequence of variable proposals before solution attempts. Second, continuing to interrogate what search algorithms work best for conceptual ideation is an obvious collaboration point between AI and cognitive science, and explicitly representing this process provides additional potential impact when tackling non-linguistic problem domains, such as strategizing and puzzle solving. This approach allows for parameters of cognitive behaviors such as information foraging (Fu and Pirolli, 2007) to be modeled and learned in diverse domains.

Outer-loop revision. GENERATE- Θ assumes




Figure 2: Observed output tokens vs. expected tokens. Expected tokens $E(P)$ are estimated from instant-variant runs on solved problems of similar predicted difficulty; failed (filled) and solved (open) markers are overlaid. Problems unsolved by any strategy are omitted, since their token cost cannot be estimated.

a reasonable Θ can be proposed up-front; it does no meta-control over which facet is most worth varying next, and does not explicitly revise Θ when the initial frame is wrong. The natural next steps include Θ -revision operators motivated by human problem solving (e.g., reasoning-base pruning, ab-

straction, analogy to known problems) as well as integration with specialized reasoning agents that return structured failure feedback (proof attempts, code execution, calculations) to drive that revision.

Token efficiency. A practical limitation is overhead: GENERATE- Θ routinely spends more output

Table 2: Per-facet Shannon entropy, idea counts, and intra-strategy nearest-neighbor Jaccard distance on MC par-1 (DeepSeek-v3.2; mean $\pm$ SE across clues). Bold = best in row. “Acc” counts the problems in the coded subset solved at least once across three epochs.

	GENERATE- Θ	SELF-DISCOVER	STEP-BACK	ITERATED-COT
n_ideas	6.54	7.40	6.28	5.58
parse ent.	0.63	0.40	0.33	0.40
mech. ent.	1.84	1.57	1.22	1.18
exec. ent.	2.06	2.09	1.89	1.70
output ent.	0.63	1.17	0.96	0.98
nn_jaccard	0.82	0.79	0.77	0.72
Acc	30/43	29/43	24/43	26/43

tokens than DIRECT or ITERATED-COT. Internalizing the formulation stage — via fine-tuning or RL — is one direction for reducing this overhead while retaining the ideation gains.

Limitations

Our results here addressed one isolated facet of creative ideation, identified by our formalism, and tested its simplest possible implementation. Specifically, we chose problems where a workable Θ is reachable up front due to extensive domain coverage in the training set, side-stepping the need for a Θ revision loop. Future work will be needed to disentangle confounders related to LLM biases and determine what kinds of extensions may be needed for more reasoning-heavy domains such as competition coding or math. More substantially, GENERATE- Θ is not fully general, as it lacks hypothesis-space revision. Major extensions will be needed to produce a strategy that can handle less familiar domains using Θ -revision.

Several additional limitations stem from statistical concerns. First, due to API costs, we were unable to run multiple epochs on every dataset, somewhat limiting statistical reliability (although see our bootstrap tests above, suggesting main effects are likely real). Second, training set contamination could have impacted our results in some way. Although it is difficult to know precisely which data is in frontier model training sets, it is likely that at least some of BIG-Bench, which contains our Rosetta task, is in the training sets of all models we tested. It is less likely that our Minute Cryptic dataset is in training sets—it is fully past the cutoff for Opus 4.6 (August, 2025), and at least half of the puzzles we scraped are after the final release date for DeepSeek V3.2. (API queries were run Mar 17–25 2026 (DeepSeek v3.2), Mar 19–Apr 7 2026 (Claude Opus 4.6), and May 12–22 2026

(DeepSeek v4-pro).) Moreover, puzzles beyond one day old are hidden behind a paywall (and we have not released our subset), meaning scrapes would need to be daily to collect all problems. Regardless, our results still showed an intervention effect despite all prompting strategies receiving the same potential training-leakage benefit. And results using Opus 4.6 on Minute Cryptic cannot be due to training set contamination.

Finally, systematic ablation studies would likely provide additional insight into why our intervention was effective. For example, which prompting choices elicit the most effective $\mathcal{A}_B$ —do LLMs rely on reasoning to choose their next hypothesis given outcomes of previous ones, and if so, does it need to be externalized in token-space? We conducted preliminary investigations into this question by testing a variant where the prompt enforced strict hypothesis chaining with only minimal intervening thoughts allowed. This resulted in worse performance on the Rosetta task despite more output tokens. However, more thorough investigations will be necessary to draw strong conclusions and rule out alternatives.

References

- Zergham Ahmed, Joshua B Tenenbaum, Christopher J Bates, and Samuel J Gershman. 2025. Synthesizing world models for bilevel planning. *arXiv preprint arXiv:2503.20124*.
- Martin Andrews and Sam Witteveen. 2025. Generating code to verify cryptic crossword reasoning. In *ICLR 2025 Third Workshop on Deep Learning for Code*.
- Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, and Torsten Hoefer. 2024. Graph of thoughts: Solving elaborate problems with large language models. In *Proceedings of the AAAI Con-*

- ference on Artificial Intelligence*, volume 38, pages 17682–17690.
- DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Brown Ebouky, Andrea Bartezzaghi, and Mattia Rigotti. 2025. Eliciting reasoning in language models with cognitive tools. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 38, pages 41758–41789. ArXiv:2506.12115.
- Avia Efrat, Uri Shaham, Dan Kilman, and Omer Levy. 2021. Cryptonite: A cryptic crossword benchmark for extreme ambiguity in language. In *EMNLP*, pages 4186–4192.
- Wai-Tat Fu and Peter Pirolli. 2007. Snif-act: a cognitive model of user navigation on the world wide web. *Hum.-Comput. Interact.*, 22(4):355–412.
- Dedre Gentner. 1983. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170.
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. 2023. Reasoning with language model is planning with world model. In *EMNLP*, pages 8154–8173.
- Douglas R. Hofstadter and Melanie Mitchell. 1994. The copycat project: A model of mental fluidity and analogy-making. In Keith J. Holyoak and John A. Barnden, editors, *Advances in Connectionist and Neural Computation Theory, Vol. 2: Analogical Connections*, pages 31–112. Ablex, Norwood, NJ.
- Douglas R. Hofstadter and Emmanuel Sander. 2013. *Surfaces and Essences: Analogy as the Fuel and Fire of Thinking*. Basic Books, New York.
- Yiming Huang, Biquan Bie, Zuqiu Na, Weilin Ruan, Songxin Lei, Yutao Yue, and Xinlei He. 2025. Understanding the anchoring effect of LLM with synthetic data: Existence, mechanism, and potential mitigations. *arXiv preprint arXiv:2505.15392*.
- Peter Jansen, Peter Clark, Doug Downey, and Daniel S Weld. 2026. Generating literature-driven scientific theories at scale. *arXiv preprint arXiv:2601.16282*.
- Priyanka Kargupta, Shuyue Stella Li, Haocheng Wang, Jinu Lee, Shan Chen, Orevaoghene Ahia, Dean Light, Thomas L Griffiths, Max Kleiman-Weiner, Jiawei Han, and 1 others. 2025. Cognitive foundations for reasoning and their manifestation in llms. *arXiv preprint arXiv:2511.16660*.
- Günther Knoblich, Stellan Ohlsson, Hilde Haider, and Detlef Rhenius. 1999. Constraint relaxation and chunk decomposition in insight problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(6):1534–1555.
- Oliver Kramer and Jill Baumann. 2025. Unlocking structured thinking in language models with cognitive prompting. In *ESANN 2025 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, pages 657–662, Bruges, Belgium. i6doc.com. ISBN 9782875870933.
- Dharshan Kumaran, Stephen M. Fleming, Larisa Markeeva, Joe Heyward, Andrea Banino, Mrinal Mathur, Razvan Pascanu, Simon Osindero, Benedetto de Martino, Petar Veličković, and Viorica Pătrăucean. 2026. Competing biases underlie overconfidence and underconfidence in LLMs. *Nature Machine Intelligence*. Preprint: arXiv:2507.03120.
- Jiaxu Lou and Yifan Sun. 2024. Anchoring bias in large language models: An experimental study. *arXiv preprint arXiv:2412.06593*.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. 2024. The AI scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*.
- Melanie Mitchell. 2021. Abstraction and analogy-making in artificial intelligence. *Annals of the New York Academy of Sciences*, 1505(1):79–101.
- Nancy J. Nersessian. 2008. *Creating Scientific Concepts*. MIT Press, Cambridge, MA.
- Allen Newell and Herbert A. Simon. 1972. *Human Problem Solving*. Prentice-Hall, Englewood Cliffs, NJ.
- Stellan Ohlsson. 1992. Information-processing explanations of insight and related phenomena. In M. T. Keane and K. J. Gilhooly, editors, *Advances in the Psychology of Thinking*, pages 1–44. Harvester-Wheatsheaf, London.
- OpenAI. 2024. Learning to reason with LLMs. Technical report, OpenAI. <https://openai.com/index/learning-to-reason-with-llms/>.
- Linlu Qiu, Liwei Jiang, Ximing Lu, Melanie Sclar, Valentina Pyatkin, Chandra Bhagavatula, Bailin Wang, Yoon Kim, Yejin Choi, Nouha Dziri, and Xiang Ren. 2024. Phenomenal yet puzzling: Testing inductive reasoning capabilities of language models with hypothesis refinement. In *ICLR*.
- Prisha Samdarshi, Mariam Mustafa, Anushka Kulkarni, Raven Rothkopf, Tuhin Chakrabarty, and Smaranda Muresan. 2024. Connecting the dots: Evaluating abstract reasoning capabilities of llms using the new york times connections word game. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21219–21236.
- Herbert A Simon. 1973. The structure of ill structured problems. *Artificial intelligence*, 4(3-4):181–201.

- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling LLM test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*.
- Yuda Song, Julia Kempe, and Rémi Munos. 2025. Outcome-based exploration for LLM reasoning. *arXiv preprint arXiv:2509.06941*.
- Aarohi Srivastava and 1 others. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*.
- Xue Wen Tan, Nathaniel Tan, Galen Lee, and Stanley Kok. 2025. The shape of reasoning: Topological analysis of reasoning traces in large language models. *arXiv preprint arXiv:2510.20665*.
- Graham Todd, Tim Merino, Sam Earle, and Julian Togelius. 2025. Benchmarking language models with the new york times connections puzzle. *IEEE Transactions on Games*.
- Pedro Tsividis, Joao Loula, Jake Burga, Juan Pablo Rodriguez, Sergio Arnaud, Nate Foss, Andres Campero, Ajay Subramanian, Thomas Pouncy, Samuel J Gershman, and 1 others. 2026. Human-level learning of complex novel tasks as theory-based modelling, exploration and planning. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 384(2320).
- Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim. 2023. Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2609–2634.
- Ruocheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah D. Goodman. 2024. Hypothesis search: Inductive reasoning with language models. In *ICLR*.
- Tianchun Wang, Zichuan Liu, Yuanzhou Chen, Jonathan Light, Weiyang Liu, Haifeng Chen, Xiang Zhang, and Wei Cheng. 2025. On the effect of sampling diversity in scaling LLM inference. *arXiv preprint arXiv:2502.11027*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*.
- Fengli Xu, Qian Yue Hao, Chenyang Shao, Zefang Zong, Yu Li, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan, Jiahui Gong, and 1 others. 2025. Toward large reasoning models: A survey of reinforced reasoning with large language models. *Patterns*, 6(10):101370.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023a. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023b. ReAct: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations (ICLR)*.
- Huaixiu Steven Zheng, Swaroop Mishra, Xinyun Chen, Heng-Tze Cheng, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2024. Take a step back: Evoking reasoning via abstraction in large language models. In *The Twelfth International Conference on Learning Representations (ICLR)*. ArXiv:2310.06117.
- Pei Zhou, Jay Pujara, Xiang Ren, Xinyun Chen, Heng-Tze Cheng, Quoc V. Le, Ed H. Chi, Denny Zhou, Swaroop Mishra, and Huaixiu Steven Zheng. 2024. Self-Discover: Large language models self-compose reasoning structures. In *Advances in Neural Information Processing Systems*.

A Per-difficulty token spend

Figures 3–5 plot total output tokens spent vs. an empirical difficulty rank (mean solve rate across strategies), per strategy and per backbone, for the three datasets. Note that problems where a strategy failed do not have associated datapoints but points on either side of a missing datapoint will still be connected by a line.

B Idea-forest visualization on a representative clue




Figure 3: Minute Cryptic par-1. Output tokens vs. difficulty rank.




Figure 4: Minute Cryptic par-2. Output tokens vs. difficulty rank.

Per-strategy tokens vs problem difficulty rank
 Difficulty score = mean tokens across family strategies (within-family)






Figure 5: BIG-Bench Rosetta. Output tokens vs. difficulty rank.

generate-0 (9 ideas, solved)




self-discover (10 ideas, failed)



step-back (7 ideas, failed)

Exclusive school's tall characters expelled by petite pri... [answer: SCOOP, solver: ELITE X] [step 7/7]
strategy=step_back | model=deepseek-chat



keep-thinking-step-by-step (9 ideas, failed)




Figure 6: Reconstructed idea forests for all four prompt-only strategies on par-1 problem 16: “Exclusive school’s tall characters expelled by petite principal” → SCOOP. Wordplay: definition = *Exclusive* (a news scoop); from SCHOOL, remove the typographically “tall characters” (H, L) to leave SC00L, then append *petite principal* (the first letter of *petite*, P) to give SC00P — mechanism: deletion then charade. Panels (top to bottom): GENERATE-0 (9 ideas) considers two alternative parses — *petite principal* as definition (suffix) and *Exclusive* as definition (prefix) — and converges on the latter, which matches the gold parse/mechanism/execution in the coding analysis (highlighted gold); SELF-DISCOVER (10 ideas), STEP-BACK (7 ideas), and ITERATED-COT (9 ideas) all anchor on *Exclusive school* as the definition — the salient surface reading — and return ELITE. Node colors encode facet (purple = parse, orange = mechanism, blue = execution, green = output). All four strategies were run on DeepSeek-chat.




Figure 7: Per-strategy radar profile on MC par-1 (DeepSeek-chat, three-epoch any-correct). Axes are normalised across strategies; the outer ring is the maximum within-row value. GENERATE- Θ dominates on prerequisite facets (parse, mechanism) and on idea spread (nn_jaccard); SELF-DISCOVER dominates on output entropy and total idea count.