STRATEGIST
Designing Agentic Reasoning at Scale

Taiming Lu[1] **Qi Xu**[2] **Zhigang Hua**[2] **Zhuang Liu**[1]
[1]Princeton University [2]Meta

https://github.com/zlab-princeton/Strategist

Abstract

A model is only half of an AI agent system. The other half is the inference strategy: the reasoning procedure that turns model calls into a response. Today that second half is written by hand, and the experience of building one strategy never carries to the next. We introduce STRATEGIST, a meta-agent that evolves inference strategies across tasks. Each strategy is written as a blueprint, executable code that other strategies can reuse, and together the blueprints form a library. Given a task and a small development set, STRATEGIST retrieves related blueprints, writes candidate strategies, tests them on real samples, and revises them from the results. Every candidate then enters the library, so later tasks build on earlier discoveries. On 30 benchmarks, STRATEGIST beats the strongest baseline on 24 and matches its average accuracy (74.7% vs. 74.0%) at 39% lower inference cost. Its strategies also transfer to unseen tasks and other models without further search. The part of an AI system that is not improved on its own now does.

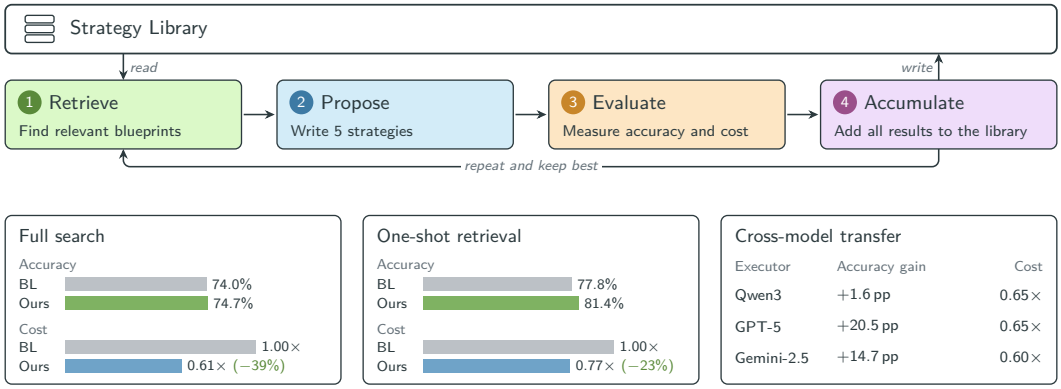


Figure 1: **STRATEGIST turns strategy design into a search that accumulates knowledge.** *Top:* in each round, an LLM writes new inference strategies as code, runs them, and stores every one, with its measured accuracy and cost, in a library shared across tasks, so later tasks build on what earlier ones found. The models stay frozen; only the strategies change. *Bottom:* the search matches the strongest baseline (BL) at lower cost, and its strategies work on new tasks and new models without further search. Cost ratios are per-benchmark averages; transfer costs are totals in USD.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

1 Introduction

A model is only half of an AI system. The other half is how that model is put to work. Everything from the prompt to reasoning procedure decides how it behaves. Today that second half is designed by people, fixed, one system at a time. Someone writes the prompt, tries it, adjusts it, and once the system works, the design is left as it is. It is the part of the system that never improves on its own.

There is no reason it has to stay that way. The prompts and reasoning procedures that put a model to work are text and code, and agents themselves have become good at writing both. What humans do to design (write, try, adjust) such *inference strategies* is the natural agentic loop. Such loops have already been shown to evolve programs far beyond their starting point [37], and AI scientists now run the same loop end to end [31, 12, 58]. Designing an inference strategy is the same work with a different object. So why not evolve agentic reasoning itself the same way? Run on far more cases than a person could try, the loop could outdo what people build by hand.

Recent work has begun to automate this loop: searching over strategy designs, as raw code or as assemblies of fixed components [18, 69, 45, 67]; training controllers that route each input to an existing strategy [50, 65, 68, 47]; or optimizing one strategy per task by training, evolution, or iterative refinement [9, 25, 70, 10, 29]. All of them beat hand-designed strategies on individual benchmarks, and all of them stop there. Searching or routing over a fixed set of strategies can never discover a new one. Trained optimizers need to start over for every task. These methods make the design automatic. We want to make it cumulative (Figure 2).

We start from the representation. Evolving strategies needs a form open enough to express any reasoning procedure and structured enough for later searches to reuse and combine. We introduce the *Strategy Blueprint*: an executable program with one shared code interface, so any blueprint can serve as a component inside another. New designs are built from old ones, and library grows as it is used.

On this representation we build STRATEGIST, a meta-agent that does not answer questions but designs the strategies that do. It runs the write, try, adjust loop at scale, with frozen models and no training. Given a task and a small development set, it writes candidate strategies as blueprints, runs them on real samples, reads the traces, and revises them against measured accuracy and cost. Before each round it queries a shared library for relevant blueprints and for diverse ones as inspiration. After each round every candidate goes back into the library with its measurements and experience notes.

On 30 benchmarks spanning math, code, commonsense, multi-hop QA, knowledge, logic, and recommendation, we compare STRATEGIST against 18 baselines. STRATEGIST wins on 24 of 30, matching the strongest baseline's average accuracy ($74.0\% \rightarrow 74.7\%$) at 39% lower cost. The library pays off on its own: on 10 held-out benchmarks, retrieval with no search stays within 2 points of the strongest baseline on all 10 and beats it on 8 ($77.8\% \rightarrow 81.4\%$). The same strategies, moved to GPT-5 and Gemini-2.5, still beat baselines. Each task solved leaves the next one a better place to start. Until now, only the model in an AI system could learn. Now the rest of it can too.

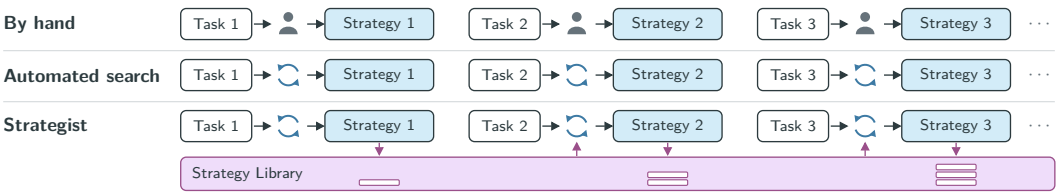


Figure 2: **Strategy design is per task; STRATEGIST makes it cumulative.** By hand and in prior automated search, nothing carries between tasks. STRATEGIST is the one evolving and keeps every strategy in a library that the next task starts from.

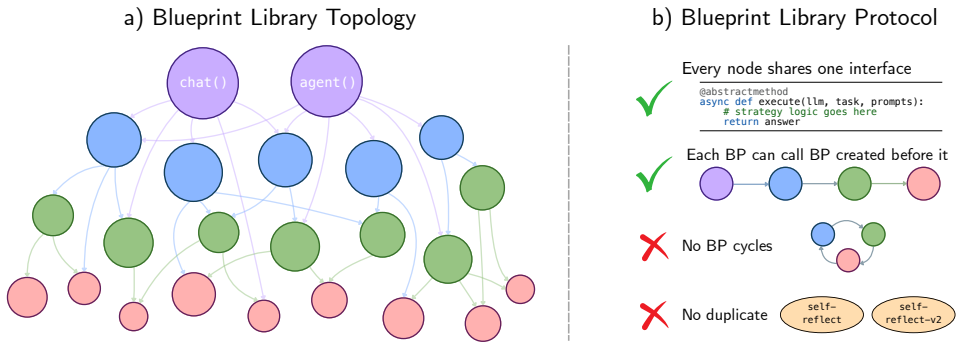


Figure 3: **The Strategy Blueprint is a freeform graph that grows through the library.** (a) Two foundations (`chat` and `agent`) sit at the top. Strategies compose from them and from each other, forming a directed graph whose levels emerge naturally from composition depth. (b) Four protocols govern the graph: every node shares one abstract interface, each node can only call nodes created before it, no cycles are allowed, and no duplicates should exist.

2 The Strategy Blueprint

Every inference strategy, ultimately, is a program. How it is represented determines what an automated designer can build, test, and reuse. Existing work represents strategies as computation graphs [57, 77] and declarative pipelines [22, 64], which are composable but can only express pre-defined patterns, or as monolithic code [18], which can express anything but produces isolated strategies. We introduce the *Strategy Blueprint*: a reasoning flow that specifies how to solve a task (which models to call, in what order, at which tier, and how) using the open interface of code. Every blueprint exposes the same interface, so any can call any other, and a library can grow by composition.

Foundations. A blueprint is a Python program that captures the *structure* of a reasoning strategy without the task-specific prompts. For example, one blueprint defines a loop: generate three candidate answers, review each, then majority-vote. Filling in prompts produces a runnable *strategy*. It has one interface, `execute(llm, task) -> str`, and is built from two foundations: `chat` (single- or multi-turn conversation) and `agent` (with scaffold and tools). See Figure 4 for an example.

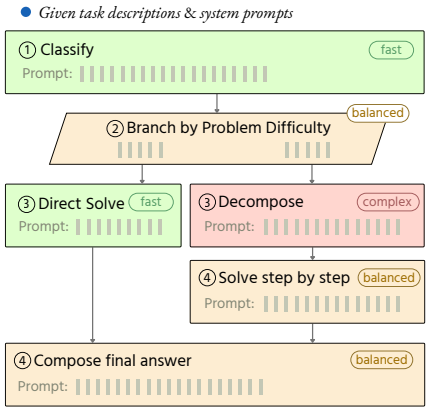


Figure 4: **Strategy blueprint captures reasoning structure.** This simplified blueprint example defines the steps, control flow, and the model tier for each step. Prompts are left unspecified; filling them in produces a runnable strategy.

Multi-tier model routing. We allow strategies to route across three model tiers (`fast`, `balanced`, `complex`), each in two modes (`nothink` and `thinking`), mirroring the tiered pricing and effort levels of common model providers. Not every reasoning step needs the same intelligence, so each step can target a different tier. This internal routing is what allows a single strategy to be accurate and cost-efficient at the same time.

The Blueprint Library. Blueprints live in a shared library that grows over time. They are added along with their measured performance and experience notes on when it works and when it does not. Because every blueprint exposes the same interface, a new blueprint can use any existing one as a component, so the library compounds: early blueprints build on the two foundations, later ones compose from those, and depth grows naturally (Figure 3a). Four protocols govern the library (Figure 3b): a shared interface, no forward calls, no cycles, and no duplicates.

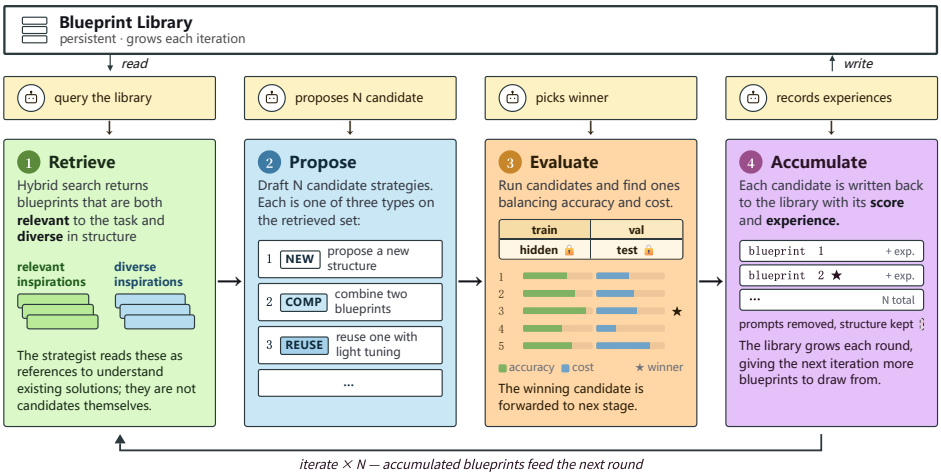


Figure 5: **The STRATEGIST four-stage iterative pipeline.** Each iteration: (1) *Retrieve* relevant blueprints and diverse inspiration from the Library; (2) *Propose* K candidates as a mix of *new*, *composed*, and *reuse* blueprints paired with task-specific prompts; (3) *Evaluate* each candidate, seeing full traces on the training set and aggregated scores on validation; (4) *Accumulate* the underlying blueprints into the library with measurement cards and experience notes.

3 STRATEGIST

Given a task and a small development set, STRATEGIST iteratively retrieves relevant blueprints, proposes $K=5$ candidates, evaluates them on real samples, and accumulates the results back into the library (Figure 5). After N iterations, the candidate that beats BL on both the validation and hidden splits by the largest hidden-set margin is returned, or BL itself if none does (Algorithm 1).

Step 1: Retrieve. The strategist queries the Blueprint Library with a short text prompt. The search is hybrid (TF–IDF plus embedding similarity). Each blueprint carries traits, a measurement card, and experience notes written by previous strategists; the search indexes over all of these, so the strategist finds blueprints by measured behavior and prior analyses, not just structural similarity. Beyond relevant results, the library returns *adjacent inspiration* (blueprints from other domains or reasoning patterns) to inspire new designs.

Step 2: Propose. Using the retrieved blueprints and inspiration, the strategist designs new blueprints or composes from existing ones, each paired with task-specific prompts to form candidate strategies. Each blueprint must be one of three kinds: *new* (a blueprint with no library backbone), *composed* (a blueprint that wraps or extends an existing one), or *reuse* (directly applies an existing library blueprint with new prompts). Every iteration includes at least one candidate of each kind, balancing exploration and exploitation.

Step 3: Evaluate. The system operates at two levels. At the *agent level* (inner loop), each candidate strategy processes inputs directly, and the strategist observes the full reasoning trace, execution status, and scores. At the *meta level* (outer loop), the strategist uses these observations to improve its proposals. The data is split into four parts. The *training set*: strategies run on these samples and the strategist sees both the scores and the full agent traces. The *validation set*: strategies also run on these samples, but the strategist sees only the aggregated scores; this mirrors standard train/validation practice and prevents overfitting to the training signal. The *hidden set*: used to select the winning strategy across all iterations, preventing overfitting to any single iteration. The *test set*: held out entirely from the search loop and used only for final confirmation.

Step 4: Accumulate. The library stores blueprints, not strategies: it retains the reusable structure while discarding task-specific prompts. This is what enables cross-task transfer. After each iteration, the strategist writes an analysis of each candidate's structural properties and adds the underlying blueprints to the library with their measurement cards and experience notes, with deduplication.

4 Experiments

Benchmarks. We evaluate on 40 benchmarks spanning seven categories (Table 1). From these, 30 randomly selected benchmarks form the main results and 10 are held out for one-shot retrieval. Each benchmark is split into training, validation, and hidden sets of up to 100 samples each, plus a fixed test set of 200 samples.

Models. The task models consist of three tiers: Qwen3-0.6B (fast), Qwen3-1.7B (balanced), and Qwen3-4B (complex), each run with thinking off (nothink) or on (thinking), giving six model configs. The strategist is GPT-6, a separate frozen LLM accessed through the Codex CLI; it writes strategies but never answers questions.

Baselines. We compare against 18 baselines (Table 2), all using the same task models without internet access. They span from direct prompting to multi-agent workflows, and each runs at all six model configurations, giving $18 \times 6 = 108$ cells per benchmark. Our baseline (BL) is the cell most accurate on a benchmark's training split, scored on the test set. This mirrors deployment: a method is chosen on the small sample at hand, then applied to a larger unseen sample from the same distribution. Our strategy is likewise chosen without seeing the test set.

Math	AQuA-RAT [27]	Gaokao-Math [74]	GSM8K [3]
	AIME 2024 [33]	MATH500 [14]	ASDiv [34]
	SVAMP [39]	MinervaM [24]	
Commonsense	StrategyQA [11]	CommonsenseQA [49]	PIQA [1]
	HellaSwag [66]	OpenBookQA [35]	SIQA [43]
	COPA [41]	WinoGrande [42]	SciQ [55]
	BoolQ [2]		
Knowledge	TruthfulQA [26]	SQuAD v2 [40]	MMLU-STEM-Hard
	TriviaQA [20]	NQ-Open [23]	MMLU [13]
Multi-hop	QUALITY [38]	2WikiMQA [15]	DROP [6]
	MuSiQue [51]	HotpotQA [60]	
Logic	LSAT-RC [74]	SAT-Math [74]	MedQA [19]
	ANLI-R3 [36]	LSAT-AR [74]	ReCLor [63]
	SAT-EN [74]	LSAT-LR [74]	BBH [48]
Code	MBPP+ [28]		
Recommendation	Amazon-Rec [30]		

Table 1: **40 benchmarks across seven task categories.** We randomly select 30 for the main iterative search and hold out the remaining 10 for one-shot retrieval. Each benchmark is split into training, validation, hidden, and test sets.

Prompting	Monte-Carlo vote	Chain-of-thought [54]	Self-discover [76]
	Direct thinking	Tree of thoughts [61]	Self-refine [32]
	Shortlist vote	Self-consistency [53]	Reflexion [46]
	Direct	Least-to-most [75]	CoVe [4]
Agentic	Agent with code	Multi-agent debate [5]	MetaGPT [16]
	Vanilla agent	Chain-of-agents [71]	AutoGen [56]

Table 2: **Baselines.** Prompting methods call the model directly; agentic methods add tool use or multiple interacting agents. Each baseline runs at all six tier/mode configurations, and we compare against the strongest training split cell.

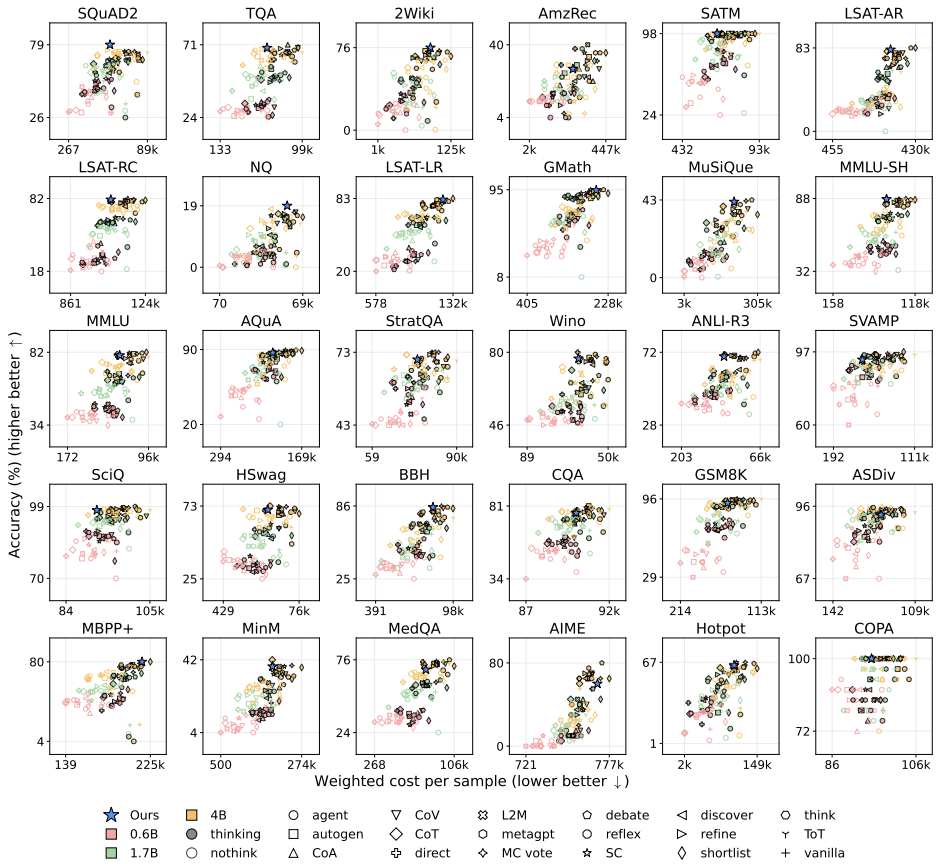


Figure 6: **Ours among all 108 baseline cells (accuracy vs weighted cost per sample, test set).** Each small marker is one baseline cell (a baseline strategy at one model tier and mode); the star is STRATEGIST's strategy. BL, the comparison point, is the cell chosen on the training split (Figure 7). The highest marker in a panel is the best of 108 cells *chosen by looking at the test results*, which no method can do in advance; it is shown for context, not as the comparison point. Appendix Figure 10 shows the same view on the training split, where the bias runs the other way. Against BL, Ours is more than 2% cheaper on 26 of 30 benchmarks and within 2 accuracy points or better on 28.

Metrics. We report accuracy and weighted cost. Cost is token usage (input and output) weighted by tier (fast $= 1\times$, balanced $= 2.83\times$, complex $= 6.67\times$, matching the 0.6B/1.7B/4B parameter ratio), since per-token floating-point operations (FLOPs) scale linearly with the parameter count [21].

5 Results

We run the full STRATEGIST loop on each of the 30 main benchmarks. BL is the cell of the 18×6 baseline matrix that is most accurate on the benchmark's training split; all methods share the same executors, extractor, and 200-sample test split. A strategy wins if it is better than BL on accuracy (by more than 2 points) or on cost (by more than 2%) and worse on neither. STRATEGIST wins on 24 of 30 benchmarks, ties on 2 (MBPP+, MinervaM), trades accuracy for cost on 2 (MedQA, AIME) and loses on 2 (HotpotQA and COPA, at equal accuracy but 9% and 6% higher cost); it cuts weighted cost by more than 2% on 26 of 30 (Figure 7; Figure 6 places Ours among all 108 baseline cells; per-benchmark numbers in Appendix C.1). On average, STRATEGIST matches BL's accuracy ($74.0\% \rightarrow 74.7\%$) while reducing weighted cost by 39% (a mean per-benchmark cost ratio of $0.61\times$ relative to BL).

The largest accuracy gains come with large cost cuts: SQuAD v2 ($73.2\% \rightarrow 79.1\%$ at $0.23\times$ cost), TruthfulQA ($+4.5$ points), 2WikiMQA ($+4.0$ at $0.34\times$), Amazon-Rec ($+4.0$), SAT-Math ($+3.6$ at $0.41\times$), LSAT-AR ($+3.4$) and LSAT-RC ($+3.0$ at $0.50\times$). On benchmarks where the models are near their accuracy ceiling (e.g., MMLU, GSM8K, ASDiv, SciQ, MMLU-STEM-Hard), STRATEGIST

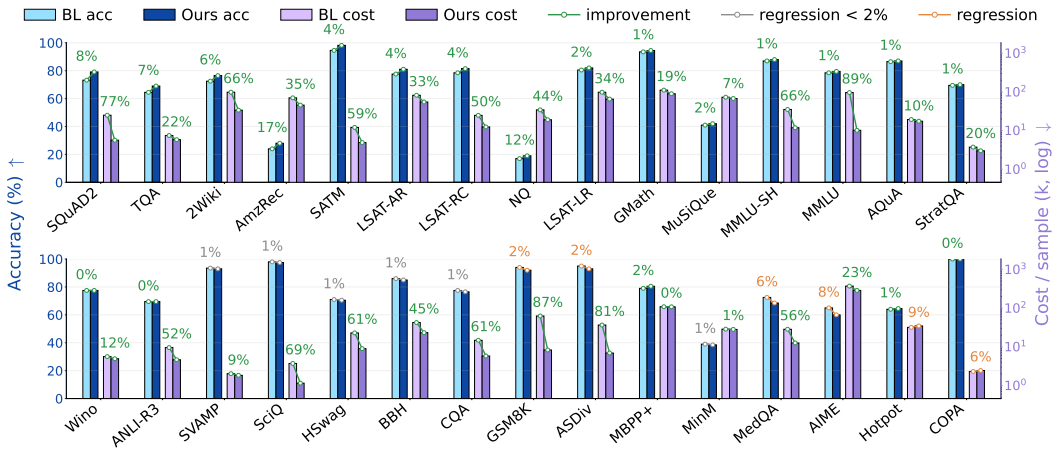


Figure 7: **Per-benchmark head-to-head: STRATEGIST vs the baseline (BL) on the 30 main benchmarks.** For each benchmark, test accuracy (%, left axis) and weighted cost per sample (log scale, right axis) of Ours and of BL. BL is the baseline cell (one of the 18×6 baseline–configuration pairs) that is most accurate on the benchmark's *training* split, scored on the test set; STRATEGIST's strategy is also chosen without seeing the test set, so this is the like-for-like comparison behind the win counts. Numbers above each pair are relative changes (green: improvement; grey: regression under 2%; orange: larger regression).

keeps accuracy within 2 points of BL and cuts cost instead, by 66–89%. Of the 24 wins, most come from strategist-designed strategies and some are from simpler baseline scripts, found during search, that match BL's accuracy at a lower cost. STRATEGIST is not Pareto-improving everywhere: on MedQA and AIME 2024 it is cheaper ($0.44\times$, $0.77\times$) but 4–5 points less accurate, and on HotpotQA and COPA it matches accuracy at slightly higher cost.

One-shot library retrieval. The 30-benchmark main run grew the library to 1,865 blueprints. We test if a future task can skip the full search: on 10 held-out benchmarks, the strategist queries the final strategies of the main searches (not the full library) once, instantiates the top-5 retrieved blueprints for the task, measures each on up to 300 development questions, and runs the chosen one (the most accurate candidate cheaper than BL if within 2 points of the best, otherwise the best) on the test set, with no search. It stays within 2 points of BL's accuracy on all 10 and above it on 8 (77.8% → 81.4% on average) and wins on 6 (Table 3); on BoolQ, SIQA, and DROP it is more

	Test acc (%)		Δ vs baseline	
Benchmark	Ours	BL	ΔAcc	ΔCost
ReCLor	83.5	85.0	↓1.8%	↓34.2%
SAT-EN	86.5	86.5	0.0%	↓55.6%
TriviaQA	50.5	46.0	↑9.8%	↓60.0%
MATH500	94.5	92.5	↑2.2%	↓22.6%
BoolQ	85.5	73.5	↑16.3%	↑3103%
SIQA	77.5	73.5	↑5.4%	↑337%
OpenBookQA	94.5	93.0	↑1.6%	↓49.2%
PIQA	89.0	87.5	↑1.7%	↓64.5%
QUALITY	79.5	78.5	↑1.3%	↑7.4%
DROP	72.5	61.5	↑17.9%	↑92.4%

Table 3: **One-shot library retrieval.** Top-5 blueprints retrieved from the final strategies, one picked on development questions. On 10 held-out benchmarks, ΔAcc and ΔCost are relative changes vs. the strongest baseline ((Ours − BL)/BL). Full results in Appendix C.2.

	GPT-5		Gemini-2.5	
Benchmark	ΔAcc	ΔCost	ΔAcc	ΔCost
AQuA-RAT	↑189%	↓86.5%	↑27.5%	↓51.6%
2WikiMQA	↑166%	↓50.8%	↑260%	↓72.0%
HotpotQA	↑96.2%	↓7.5%	↑132%	↓52.2%
MBPP+	↑46.0%	↓12.3%	↑15.8%	↓23.7%
SQuAD2	↑28.7%	↓61.5%	↓4.7%	↓57.9%
NQ-Open	↑35.0%	↓22.8%	↑118%	↓21.8%
TruthfulQA	↑13.7%	↓71.9%	↑9.4%	↓61.9%
ASDiv	↑2.2%	↑163%	↑2.1%	↓33.1%
HellaSwag	↑0.5%	↑16.9%	↑0.5%	↓12.3%
MinervaM	↑2.1%	↑4.8%	↓9.4%	↓25.9%

Table 4: **Cross-model transfer.** Strategies discovered on Qwen3, executed directly on GPT-5 and Gemini-2.5 without modification. ΔAcc and ΔCost are relative changes vs. the strongest baseline re-executed on the same backend. Full per-benchmark numbers are in Appendix C.3.

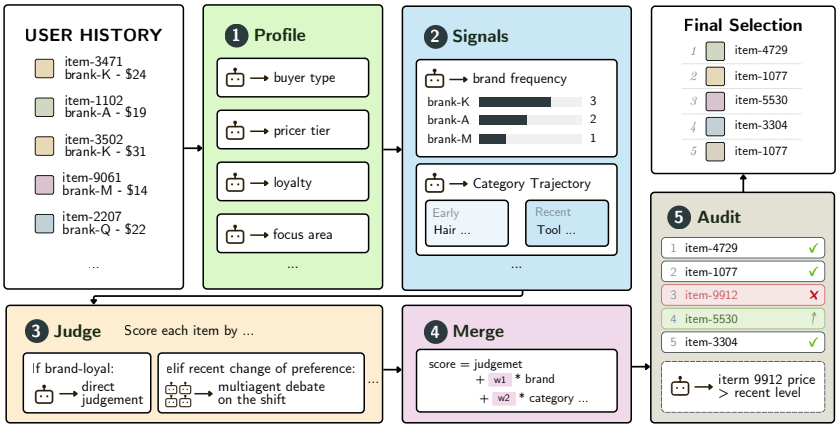


Figure 8: **An example reasoning strategy discovered by STRATEGIST for the Amazon-Rec task.** The strategy decomposes the task into a profile, signal extraction, conditional judge, merge, and audit stage. The structure and branch conditions were generated by the strategist, not by us.

accurate but costlier than their very cheap BL (1.0–4.9k per sample), and on QUALITY it costs 7% more at equal accuracy. Its per-benchmark average cost (as in Figure 7) is 0.77× BL's without BoolQ and SIQA, whose single-call BL puts their ratios at 32× and 4.5×.

Cross-model transfer. A strategy is most useful if it captures structure of the *task* rather than idiosyncrasies of the executor it was searched on. We re-run STRATEGIST-discovered strategies, without any modification to the blueprint code or its prompts, on two model families the strategist never saw during search: OpenAI GPT-5 (nano, mini, full) and Google Gemini-2.5 (flash-lite, flash). On GPT-5 the strategies win on 7 of 10 benchmarks and lose on 3 (ASDiv, HellaSwag, MinervaM: within 2 points of the baseline but costlier); on Gemini they win on 8 of 10, with SQuAD2 and MinervaM mixed (cheaper but less accurate). Aggregate cost in USD drops by 35% and 40%. Part of the GPT-5 margin comes from baseline workflows that break on the new backend (`self_refine` and `multi_agent_debate` return malformed answers there) while the discovered strategies run unchanged. That a Qwen3-discovered blueprint carries over at all shows that STRATEGIST captures the structure of the task, not executor-specific tricks.

6 Analysis

Anatomy of a discovered reasoning strategy. Figure 8 shows one strategy (accuracy driven) designed for the Amazon-Rec task: it builds a user profile with cheap calls, extracts brand and category signals, and conditionally routes to a multi-agent debate only for users with shifting preferences, before merging and auditing the top picks. The branching is cheap deterministic logic that decides which inputs deserve the expensive call, so the strategy spends its capacity where it actually changes the answer rather than uniformly across all inputs. The conditional routing, the merge weighting, and the audit step all emerged from iterative proposals against measured accuracy. Because the result is plain Python, a user can read and edit it directly, and the call is structural, which is hard to do with prompt-bag methods or training-based alternatives.

Overfitting across data splits. The strategist sees training traces and validation scores at every iteration and could fit noise in them; if it did, the training curve would pull away from the hidden curve. Figure 9 plots, per iteration, the mean accuracy of that iteration's five candidates on each split, averaged over the 30 main benchmarks. The curves move together and the training–hidden gap stays between 1.2 and 2.5 points, so the proposals do not drift away from data the strategist never sees. Selection is a different matter: the final strategy is chosen for being best on the hidden split, and part of that is luck. Across the 30 benchmarks the pick is 4.3 points above BL on hidden but 0.7 on test (BL itself moves by 0.5). The test split, on which nothing is chosen, is the only unbiased estimate, and it is what every table reports.

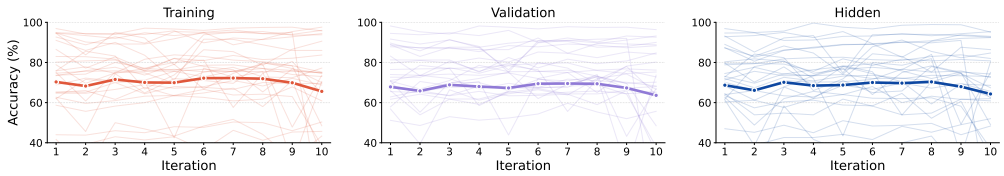


Figure 9: **Accuracy across data splits over the 10-iteration search loop.** Each point is the mean accuracy of an iteration's five candidates on a split. Bold curves are the mean across the 30 main benchmarks; faint background curves are individual benchmarks. The training–hidden gap stays between 1.2 and 2.5 points, not widening over iterations.

7 Ablations and Comparisons

To probe which components of STRATEGIST are responsible for the gains in Section 5, we run two ablations on a fixed 10-benchmark set. Neither ablation falls back to BL, for testing purposes. They are drawn from the 30 main ones, disjoint from transfer and one-shot sets.

Zero-shot selection. We disable the library and cap the loop at one iteration. The strategist sees only the task description and sample questions. It averages 39.4% with three crashed runs counted as zero (56.2% over the seven that ran), 39.3 points below (78.7%).

Iteration without the library. We restore the 10-iteration loop but disable the library to isolate its contribution. This averages 65.4% at 41.4k cost (Table 5), still 13.3 points below (78.7%). The main observation is that, without inspirations, it goes on to a narrow repetition.

Cross-framework comparison. We compare against five published frameworks (AFlow [69], ADAS [18], EvoFlow [67], StrategyLLM [9], AgentSquare [45]), all sharing the same meta-LLM and executor. STRATEGIST wins on 8 of 10 benchmarks at 78.7% mean accuracy. Many of these strategies fit the few development questions and fail on the test set.

8 Related Work

LLM agents. LLM agents span chain-of-thought (CoT) and self-consistency reasoning [54, 53], deliberate search [61], tool-use loops [62, 44], verbal self-reflection [46], multi-agent conversation [56], and formal representations from state machines to declarative compilation to optimizable graphs [57, 22, 77, 64]. The strategy layer (how models are called, in what order, and with what verification) has become the central design bottleneck; STRATEGIST automates its design. (An extended discussion of related work is in Appendix E.)

	STRATEGIST variants			Other frameworks				
Benchmark	full	no library	zero-shot	AFlow	ADAS	EvoFlow	StrategyLLM	AgentSquare
BBH	**85.0**	31.0	27.5	38.0	34.5	34.5	35.5	27.5
Gaokao-Math	**94.5**	65.5	81.5	84.0	82.0	86.5	89.5	82.5
LSAT-LR	**82.0**	81.0	71.5	81.0	74.5	75.5	77.0	73.5
LSAT-RC	**81.5**	75.0	38.0	74.5	75.5	75.0	74.5	74.0
SAT-Math	98.2	**99.0**	failed	**99.0**	96.9	96.9	97.9	97.9
MMLU-STEM-Hard	**88.0**	83.5	86.5	83.0	82.5	87.5	86.0	86.5
MMLU	**79.5**	78.0	failed	71.0	74.5	72.0	73.0	71.5
MuSiQue	**42.0**	23.0	16.0	11.5	8.5	20.5	17.5	9.5
CommonsenseQA	76.5	**78.0**	72.5	74.5	77.0	74.0	75.0	74.5
AIME 2024	**60.0**	40.0	failed	53.3	53.3	20.0	46.7	53.3
Average Acc	**78.7**	65.4	39.4	67.0	65.9	64.2	67.3	65.1

Table 5: **Ablations and cross-framework comparison on 10 benchmarks.** Test accuracy (%); best per row bolded. External frameworks share the same meta-LLM and executor as STRATEGIST; "failed" marks a run that crashed.

Automated design of agentic systems. Structure-search methods write strategies as raw code [18] or compose them from fixed components or predefined modules [69, 45, 67]; the first yields isolated programs, the second bounds what can be discovered. Optimization methods refine one strategy per task, by reinforcement learning [25, 70, 10] or by LLM-driven evolution and refinement [9, 29], and start over on the next task. Routing methods select among existing strategies without building them [50, 65, 68]. STRATEGIST instead allows freeform creation via the Strategy Blueprint, needs no training, builds strategies itself, and routes across model tiers in each strategy.

Cross-task transfer and experience accumulation. Cross-task transfer has been studied through meta-learning [8, 17], neural architecture search [7], and prompt optimization [59]; in the LLM-agent setting, ExpeL [72] and Voyager [52] accumulate experience or skills, and lifelong-agent learning has been formalized as a direction [73]. Existing automated design methods [18, 69] still largely optimize per task. STRATEGIST maintains a shared blueprint library that grows with every task, providing both exploitation (relevant components) and exploration (diverse inspiration); to our knowledge, the first automated strategy-design system that accumulates knowledge from one task to the next.

9 Conclusion

We presented STRATEGIST, a meta-agent that designs inference strategies as freeform executable programs on a layered framework, the Strategy Blueprint, and grows a shared Blueprint Library across tasks. On 30 benchmarks, STRATEGIST-designed strategies match or beat most strongest baselines on accuracy while running at substantially lower cost, and the discovered strategies transfer to backends the strategist never searched on. We view freeform strategy design over an accumulating library as a step toward inference systems that get better the more they are used.

Acknowledgments and Disclosure of Funding

This work is supported by the computational resources generously provided by Google's TPU Research Cloud program. This work was also performed using Princeton Research Computing resources, a consortium led by the Princeton Institute for Computational Science and Engineering (PICSciE) and Research Computing at Princeton University. We also gratefully acknowledge the use of the Neuronic GPU computing cluster maintained by the Department of Computer Science at Princeton University.

References

[1] Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. Piqa: Reasoning about physical commonsense in natural language. In *AAAI*, 2020.

[2] Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions. In *NAACL*, 2019.

[3] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

[4] Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. Chain-of-verification reduces hallucination in large language models. In *Findings of ACL*, 2024.

[5] Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. In *ICML*, 2024.

[6] Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In *NAACL*, 2019.

[7] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *Journal of Machine Learning Research*, 20(55):1–21, 2019.

[8] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.

[9] Chang Gao, Haiyun Jiang, Deng Cai, Shuming Shi, and Wai Lam. Strategyllm: Large language models as strategy generators, executors, optimizers, and evaluators for problem solving. In *NeurIPS*, 2024.

[10] Hongcheng Gao, Yue Liu, Yufei He, Longxu Dou, Chao Du, Zhijie Deng, Bryan Hooi, Min Lin, and Tianyu Pang. Flowreasoner: Reinforcing query-level meta-agents. *arXiv preprint arXiv:2504.15257*, 2025.

[11] Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *TACL*, 2021.

[12] Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Anil Palepu, Petar Sirkovic, Artiom Myaskovsky, Felix Weissenberger, Keran Rong, Ryutaro Tanno, et al. Towards an AI co-scientist. *arXiv preprint arXiv:2502.18864*, 2025.

[13] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *ICLR*, 2021.

[14] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021.

[15] Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps. In *COLING*, 2020.

[16] Sirui Hong, Mingchen Zhuge, Jiaqi Chen, Xiawu Zheng, Yuheng Cheng, Ceyao Zhang, Jinlin Wang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. Metagpt: Meta programming for a multi-agent collaborative framework. In *ICLR*, 2024.

[17] Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.

[18] Shengran Hu, Cong Lu, and Jeff Clune. Automated design of agentic systems. In *ICLR*, 2025.

[19] Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 2021.

[20] Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *ACL*, 2017.

[21] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

[22] Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. Dspy: Compiling declarative language model calls into self-improving pipelines. In *ICLR*, 2024.

[23] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. Natural questions: A benchmark for question answering research. *TACL*, 2019.

[24] Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language models. In *NeurIPS*, 2022.

[25] Zelong Li, Shuyuan Xu, Kai Mei, Wenyue Hua, Balaji Rama, Om Raheja, Hao Wang, He Zhu, and Yongfeng Zhang. Autoflow: Automated workflow generation for large language model agents. *arXiv preprint arXiv:2407.12821*, 2024.

[26] Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. In *ACL*, 2022.

[27] Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by rationale generation: Learning to solve and explain algebraic word problems. In *ACL*, 2017.

[28] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. In *NeurIPS*, 2023.

[29] Siwei Liu, Jinyuan Fang, Han Zhou, Yingxu Wang, and Zaiqiao Meng. Sew: Self-evolving agentic workflows for automated code generation. *arXiv preprint arXiv:2505.18646*, 2025.

[30] Zhanyu Liu, Shiyao Wang, Xingmei Wang, Rongzhou Zhang, Jiaxin Deng, Honghui Bao, Jinghao Zhang, Wuchao Li, Pengfei Zheng, Xiangyu Wu, Yifei Hu, Qigen Hu, Xinchen Luo, Lejian Ren, Zixing Zhang, Qianqian Wang, Kuo Cai, Yunfan Wu, Hongtao Cheng, Zexuan Cheng, Lu Ren, Huanjie Wang, Yi Su, Ruiming Tang, Kun Gai, and Guorui Zhou. Onerec-think: In-text reasoning for generative recommendation. *arXiv preprint arXiv:2510.11639*, 2025.

[31] Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. The AI Scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.

[32] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with self-feedback. In *NeurIPS*, 2023.

[33] Mathematical Association of America. Aime 2024 (american invitational mathematics examination), 2024.

[34] Shen-Yun Miao, Chao-Chun Liang, and Keh-Yih Su. A diverse corpus for evaluating and developing english math word problem solvers. In *ACL*, 2020.

[35] Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *EMNLP*, 2018.

[36] Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. Adversarial nli: A new benchmark for natural language understanding. In *ACL*, 2020.

[37] Alexander Novikov, Ngân Vũ, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco J. R. Ruiz, Abbas Mehrabian, M. Pawan Kumar, Abigail See, Swarat Chaudhuri, George Holland, Alex Davies, Sebastian Nowozin, Pushmeet Kohli, and Matej Balog. Alphaevolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint arXiv:2506.13131*, 2025.

[38] Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny Ma, Jana Thompson, He He, and Samuel R. Bowman. Quality: Question answering with long input texts, yes! In *NAACL*, 2022.

[39] Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are nlp models really able to solve simple math word problems? In *NAACL*, 2021.

[40] Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don't know: Unanswerable questions for squad. In *ACL*, 2018.

[41] Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S. Gordon. Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In *AAAI Spring Symposium*, 2011.

[42] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. In *AAAI*, 2020.

[43] Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. Social iqa: Commonsense reasoning about social interactions. In *EMNLP*, 2019.

[44] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. In *NeurIPS*, 2023.

[45] Yu Shang, Yu Li, Keyu Zhao, Likai Ma, Jiahe Liu, Fengli Xu, and Yong Li. Agentsquare: Automatic llm agent search in modular design space. In *ICLR*, 2025.

[46] Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. In *NeurIPS*, 2023.

[47] Jinwei Su, Qizhen Lan, Yinghui Xia, Lifan Sun, Weiyou Tian, Tianyu Shi, Xinyuan Song, Lewei He, and Jingsong Yang. Difficulty-aware agent orchestration in llm-powered workflows. *arXiv preprint arXiv:2509.11079*, 2025.

[48] Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. Challenging big-bench tasks and whether chain-of-thought can solve them. In *Findings of ACL*, 2023.

[49] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. Commonsenseqa: A question answering challenge targeting commonsense knowledge. In *NAACL*, 2019.

[50] Aditya Taparia, Som Sagar, and Ransalu Senanayake. Learning to configure agentic ai systems. *arXiv preprint arXiv:2602.11574*, 2026.

[51] Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. Musique: Multihop questions via single-hop question composition. *TACL*, 10:539–554, 2022.

[52] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *TMLR*, 2024.

[53] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *ICLR*, 2023.

[54] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*, 2022.

[55] Johannes Welbl, Nelson F. Liu, and Matt Gardner. Crowdsourcing multiple choice science questions. In *Workshop on Noisy User-generated Text (W-NUT)*, 2017.

[56] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W. White, Doug Burger, and Chi Wang. Autogen: Enabling next-gen llm applications via multi-agent conversation. *arXiv preprint arXiv:2308.08155*, 2023.

[57] Yiran Wu, Tianwei Yue, Shaokun Zhang, Chi Wang, and Qingyun Wu. Stateflow: Enhancing llm task-solving through state-driven workflows. In *COLM*, 2024.

[58] Yutaro Yamada, Robert Tjarko Lange, Cong Lu, Shengran Hu, Chris Lu, Jakob Foerster, Jeff Clune, and David Ha. The ai scientist-v2: Workshop-level automated scientific discovery via agentic tree search. *arXiv preprint arXiv:2504.08066*, 2025.

[59] Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. Large language models as optimizers. In *ICLR*, 2024.

[60] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *EMNLP*, 2018.

[61] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In *NeurIPS*, 2023.

[62] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *ICLR*, 2023.

[63] Weihao Yu, Zihang Jiang, Yanfei Dong, and Jiashi Feng. Reclor: A reading comprehension dataset requiring logical reasoning. In *ICLR*, 2020.

[64] Ling Yue, Kushal Raj Bhandari, Ching-Yun Ko, Dhaval Patel, Shuxin Lin, Nianjun Zhou, Jianxi Gao, Pin-Yu Chen, and Shaowu Pan. From static templates to dynamic runtime graphs: A survey of workflow optimization for llm agents. *arXiv preprint arXiv:2603.22386*, 2026.

[65] Yanwei Yue, Guibin Zhang, Boyang Liu, Guancheng Wan, Kun Wang, Dawei Cheng, and Yiyan Qi. Masrouter: Learning to route llms for multi-agent systems. In *ACL*, 2025.

[66] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? In *ACL*, 2019.

[67] Guibin Zhang, Kaijie Chen, Guancheng Wan, Heng Chang, Hong Cheng, Kun Wang, Shuyue Hu, and Lei Bai. Evoflow: Evolving diverse agentic workflows on the fly. *arXiv preprint arXiv:2502.07373*, 2025.

[68] Guibin Zhang, Luyang Niu, Junfeng Fang, Kun Wang, Lei Bai, and Xiang Wang. Multi-agent architecture search via agentic supernet. In *ICML*, 2025.

[69] Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xionghui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, Bingnan Zheng, Bang Liu, Yuyu Luo, and Chenglin Wu. Aflow: Automating agentic workflow generation. In *ICLR*, 2025.

[70] Mingda Zhang, Haoran Luo, Tiesunlong Shen, Qika Lin, Xiaoying Tang, Rui Mao, and Erik Cambria. Flowsteer: Interactive agentic workflow orchestration via end-to-end reinforcement learning, 2026. arXiv preprint arXiv:2602.01664.

[71] Yusen Zhang, Ruoxi Sun, Yanfei Chen, Tomas Pfister, Rui Zhang, and Sercan Ö. Arik. Chain of agents: Large language models collaborating on long-context tasks. In *NeurIPS*, 2024.

[72] Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. Expel: Llm agents are experiential learners. In *AAAI*, 2024.

[73] Junhao Zheng, Chengming Shi, Xidi Cai, Qiuke Li, Duzhen Zhang, Chenxing Li, Dong Yu, and Qianli Ma. Lifelong learning of large language model based agents: A roadmap. *arXiv preprint arXiv:2501.07278*, 2025.

[74] Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. Agieval: A human-centric benchmark for evaluating foundation models. In *Findings of NAACL*, 2024.

[75] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and Ed Chi. Least-to-most prompting enables complex reasoning in large language models. In *ICLR*, 2023.

[76] Pei Zhou, Jay Pujara, Xiang Ren, Xinyun Chen, Heng-Tze Cheng, Quoc V. Le, Ed Chi, Denny Zhou, Swaroop Mishra, and Huaixiu Steven Zheng. Self-discover: Large language models self-compose reasoning structures. In *NeurIPS*, 2024.

[77] Mingchen Zhuge, Wenyi Wang, Louis Kirsch, Francesco Faccio, Dmitrii Khizbullin, and Jürgen Schmidhuber. Gptswarm: Language agents as optimizable graphs. In *ICML*, 2024.

Appendix Contents

A. Method and Implementation Details
A.1 Pseudocode for the STRATEGIST Design Loop
A.2 Benchmark and Baseline Settings
A.3 Data Split and Scaling
A.4 Search, Library, and Models
A.5 Strategist Prompt Template
A.6 Cost-Weight Derivation
A.7 Cross-Model Transfer Setup

B. Examples and Case Studies
B.1 Sample Discovered Reasoning Strategy
B.2 Additional Sample Reasoning Strategies
B.3 Sample Measurement Card
B.4 Search Queries and Experience Notes
B.5 Per-Iteration Progression Example

C. Per-Benchmark Result Tables
C.1 Per-Benchmark Main Results
C.2 Per-Benchmark One-Shot Library Retrieval
C.3 Per-Benchmark Cross-Model Transfer

D. Discussion and Limitations
D.1 Zero-Shot Ablation Failure Analysis
D.2 Limitations
D.3 Reproducibility Statement

E. Extended Related Work

A Method and Implementation Details

A.1 Pseudocode for the STRATEGIST Design Loop

Algorithm 1 formalizes the four-stage iterative pipeline described in Section 3. Inputs are a task description $\mathcal{T}$, a development set D split into feedback/validation/hidden subsets, the current shared blueprint library $\mathcal{L}$, and the number of iterations N. Each iteration runs four steps: (i) the strategist queries $\mathcal{L}$ for the most relevant blueprints and a set of adjacent-domain inspirations; (ii) the strategist proposes K candidate strategies, balanced across the kinds *new*/*composed*/*reuse* and the targets *accuracy*/*efficiency*/*balanced*; (iii) each candidate is executed on D and scored on accuracy and cost; (iv) the strategist writes experience notes summarizing what worked and what failed, and the candidates plus their measurement cards are written back to $\mathcal{L}$ with structural deduplication. The reference baseline b (the baseline cell most accurate on the training split, run with its exact code) is evaluated alongside the candidates in the first iteration. After N iterations, every candidate is compared with b question by question on the validation and hidden splits; among the candidates that beat b under the 2% rule on both (better by more than 2 points in accuracy or 2% in cost, worse on neither), the one with the largest hidden-set accuracy gain is returned as s^*, and b itself if none qualifies. The updated library $\mathcal{L}'$ is carried forward for future tasks.

A.2 Benchmark and Baseline Settings

Scoring mechanism. All strategies and baselines run as 0-shot generation: the executor LLM produces a free-form response and an answer extractor parses the final answer from the trace. The extractor is benchmark-aware (multiple-choice letter, integer, LaTeX numeric, span, or yes/no) and is the same across STRATEGIST, all baselines, and all five comparison frameworks, so reported accuracy differences reflect strategy choices rather than scoring asymmetries.

Benchmarks. We evaluate on 40 publicly released benchmarks. From these, 30 are randomly selected as the main benchmark set used for the iterative search and the remaining 10 are held out exclusively for the one-shot retrieval experiment.

Algorithm 1 The STRATEGIST design loop

Require: Task description $\mathcal{T}$, development set D, library $\mathcal{L}$, iterations N, reference baseline b
Ensure: Winning strategy s^*, updated library $\mathcal{L}'$

```
for i = 1 to N do
    Strategist queries L for relevant blueprints and diverse inspiration
    Strategist proposes K candidate strategies {s_1, ..., s_K}
    for each candidate s_j do
        Evaluate s_j on D: measure accuracy and cost
    end for
    Strategist analyzes results and writes experience notes
    Add each s_j, its measurement card, and experience notes to L (with deduplication)
end for
W ← candidates that beat b under the 2% rule on both validation and hidden
s* ← arg max_{s∈W} Δaccuracy_hidden(s, b), or b if W = ∅
return s*, L
```

- **Mathematical reasoning:** GSM8K [3] (grade-school word problems, numeric), MATH500 [14] (competition math, LaTeX/numeric), AQuA-RAT [27] (algebra, 5 choices), AIME 2024 [33] (American Invitational Math, 30 problems, integer 0–999), ASDiv [34] (diverse word problems, numeric), SVAMP [39] (perturbed word problems, numeric), MinervaMath [24] (graduate-level quantitative, LaTeX/numeric), Gaokao-Math-QA [74] (Chinese college-entrance MCQ, 4 choices).
- **Code generation:** MBPP+ [28] (Python function synthesis, pass@1 on extra-test variant of MBPP).
- **Commonsense reasoning:** HellaSwag [66] (sentence completion, 4 choices), PIQA [1] (physical commonsense, 2 choices), Social IQa [43] (social commonsense, 3 choices), WinoGrande [42] (coreference, 2 choices), CommonsenseQA [49] (general commonsense, 5 choices), COPA [41] (causal alternatives, 2 choices), OpenBookQA [35] (elementary science, 4 choices), SciQ [55] (crowdsourced science, 4 choices), StrategyQA [11] (implicit-reasoning yes/no), BoolQ [2] (passage-grounded yes/no).
- **Knowledge and QA:** MMLU [13] (57 subjects, 4 choices), MMLU-STEM-Hard (filtered hard STEM subset, 4 choices), TriviaQA [20] (closed-book trivia, EM), NQ-Open [23] (closed-book Wikipedia QA, EM), TruthfulQA [26] (misconception-resistant MC1), SQuAD v2 [40] (extractive with unanswerables, F1).
- **Multi-hop QA:** HotpotQA [60] (2-hop with passages, EM), 2WikiMultihopQA [15] (Wikipedia multi-hop, EM), MuSiQue [51] (2–4-hop composition, EM), DROP [6] (discrete reasoning over paragraphs, EM), QuALITY [38] (long-passage MCQ on 5 k-token articles, 4 choices).
- **Logical and analytical reasoning:** BBH [48] (23 BIG-Bench Hard tasks, mixed format), ANLI-R3 [36] (round-3 adversarial NLI, 3 choices), ReClor [63] (logical reading comprehension, 4 choices), AGIEval English MCQ [74]: SAT-Math, SAT-EN, LSAT-AR, LSAT-LR, LSAT-RC (4–5 choices each), MedQA [19] (USMLE-style, 4 choices).
- **Recommendation:** Amazon-Rec [30] (rank top item from a candidate pool, hit-rate).

Baselines. We compare against 18 standard baselines covering both single-call prompting and full agentic systems. Each baseline is evaluated on every one of the six tier/mode configurations, giving an $18 \times 6 = 108$-cell baseline matrix per benchmark. The matrix is measured on both the training split and the test set; the cell most accurate on the training split is BL, and its test result is reported.

- **direct:** a single LLM call that returns the answer with no intermediate reasoning. The minimum-cost baseline; serves as the floor for comparisons on easy benchmarks.
- **chain-of-thought (CoT)** [54]: a single call prompted with "Let's think step by step", encouraging the model to emit intermediate reasoning before the final answer.
- **self-consistency** [53]: sample 3 CoT trajectories at temperature 0.7 and return the majority-voted final answer; trades cost for variance reduction.
- **direct-thinking:** a single call in `thinking` mode with no extra prompting, isolating the contribution of the model's hidden reasoning trace from any prompting trick.
- **tree-of-thoughts (ToT)** [61]: maintain a tree of partial reasoning states, expand each by sampling a few continuations, score with the model itself, and search BFS-style for a high-scoring leaf.

- **least-to-most** [75]: a first call decomposes the question into sub-problems; subsequent calls solve them in order, each conditioned on prior answers.
- **chain-of-verification (CoVe)** [4]: produce a draft answer, generate verification questions about it, answer those independently, and revise if any verification disagrees.
- **self-discover** [76]: the model first selects and composes a custom reasoning structure (decomposition, verification, etc.) from a fixed pool, then applies that structure to the input.
- **self-refine** [32]: a feedback-and-revise loop in which the model critiques its own draft and rewrites it; we cap the loop at three rounds.
- **vanilla-agent:** an OpenHands ReAct loop with the `TerminalTool` and `FileEditorTool` only, with no extra prompt scaffolding.
- **agent-code:** the same agent but the prompt instructs the model to write and execute Python via the terminal tool, so numeric/code-heavy steps are grounded in actual computation.
- **multi-agent-debate** [5]: 3 agents propose answers, then in 2 subsequent rounds each agent sees the others' answers and revises; the final majority vote is returned.
- **MetaGPT** [16]: a role-based pipeline (planner, executor, reviewer) that hands a structured artifact between roles; we use the open-source role prompts unchanged.
- **AutoGen** [56]: a two-agent conversable setup (assistant + user-proxy with tool-use) that converses until the assistant emits a termination signal.
- **chain-of-agents (CoA)** [71]: a chain of worker agents each summarises a slice of the input and passes a hand-off message to the next; a final manager agent answers from the chain.
- **reflexion** [46]: an outer loop that, on failure, asks the model to write a verbal self-reflection and retries with the reflection appended to the next prompt.
- **monte-carlo-vote / shortlist-vote:** two self-consistency variants—the first samples 5 traces and returns the strict majority (resampling on a tie), the second pre-filters to a top-10 shortlist before a 3-way vote—used to test whether voting beyond the standard 3 samples helps.

All baselines run on the same Qwen3 executors and the same OpenHands tool sandbox (`TerminalTool` + `FileEditorTool` only) with no internet access, identical to STRATEGIST strategies. Each baseline is evaluated independently at every one of the six tier/mode configurations, but baselines do not internally route across tiers within a single run; the ability to mix tiers within one strategy is part of the design space STRATEGIST searches over, and BL is the single-config cell most accurate on the training split.

A.3 Data Split and Scaling

Every benchmark is partitioned into four disjoint subsets: *training*, *validation*, *hidden*, and *test*. The training set is what the strategist sees in full, including per-sample reasoning traces, when proposing the next iteration's slate. The validation set is also evaluated each iteration but only aggregate scores are visible, mirroring standard train/val practice. The validation and hidden sets are used to pick the final strategy across all N iterations (Algorithm 1); the hidden set is never shown to the strategist; in particular, experience notes written back to the library are derived from training-set traces only, so hidden-set signal cannot leak across iterations or to future tasks. The test set is held out from the search loop entirely and is used only for the final number reported in the main text.

Sizes and caps. The development rows of each benchmark come from a separate development pool (300 rows where available), drawn deterministically with a fixed seed and split into training, validation, and hidden sets of 100 rows each; benchmarks with smaller pools scale all three down proportionally, and AIME 2024 and Amazon-Rec, whose questions are about $10\times$ longer to solve, use 50 rows per split. Where a benchmark's own pool was too small, it was extended: COPA and SVAMP with rows from their official training data, and AIME with the 299 problems of AIME 2014–2023. The test set is around the first 200 rows of each benchmark's test file, disjoint from all development rows and never seen during search. Six benchmarks have fewer questions due to limited data. Splits are deterministic, so reruns and ablations operate identically.

A.4 Search, Library, and Models

Search loop. Each iteration proposes $K = 5$ candidate strategies; the search runs for $N = 10$ iterations per benchmark. Every iteration must include at least one candidate of each kind (*new*, *composed*, *reuse*) and of each target (*accuracy*: beat BL's accuracy within its cost; *efficiency*: hold its accuracy at lower cost; *balanced*); slates that violate this rule, deduplicate a recent candidate, or omit required JSON fields are rejected and the strategist is re-prompted. The strategist is shown the baseline matrix measured on the training split (every baseline at all six configurations, with BL highlighted) and the 2% rule by which candidates are judged; it keeps one session across iterations, so later iterations send only the new results. Each LLM call from a strategy uses `max_tokens=16384`. A full search evaluates about 60 candidates per benchmark (K=5 per iteration plus the reference and repairs), each on all three development splits, roughly $20{,}000$ candidate-sample executions for a benchmark with 100 rows per split, plus about 135k prompt tokens and 45k response tokens for the strategist (about 10k prompt tokens per iteration), so the search dominates the per-task cost and amortises across future tasks via library reuse (Section 5).

Deduplication. Two blueprints are considered duplicates if their `execute` function bodies hash to the same value after task-specific string literals (prompts, format instructions) are elided; this is computed over the Python AST so that whitespace and comment changes do not produce spurious distinctions. Duplicate proposals are rejected at submission time and not added to the library.

Strategist proposal format. Each iteration the strategist returns Python code (the candidate strategies) followed by a JSON block describing the slate. The JSON has six required fields per candidate (`name`, `kind`, `target`, `tier_mode_anchors`, `why`, `builds_on`) and an optional `prompts` field that replaces a blueprint's prompt texts for this task, plus three search-side fields (`search_queries`, `read_cards`, `experience`). A non-fatal AST audit checks that `builds_on` matches the imports in the code block; warnings are fed into the next iteration's prompt rather than killing the run.

Library retrieval. At each iteration the strategist issues two textual queries against the library: a relevance query summarising what the current task is asking for, and an inspiration query phrased to deliberately drift from the task description (e.g., asking for "multi-path verification on long-context QA" when the task itself is short-form math). Each query is scored by a hybrid function $s = \alpha\, s_{\text{tfidf}} + (1 - \alpha)\, s_{\text{emb}}$ over every blueprint's index entry, where s_{tfidf} is computed over the blueprint name, one-line description, when-to-use field, and trait tags, and s_{emb} is the cosine similarity between the query and a sentence-transformer (`all-MiniLM-L6-v2`) embedding of the blueprint's measurement card and experience notes; we use $\alpha = 0.4$ throughout. Each query returns its top-5 blueprints, deduplicated against blueprints already returned by the relevance query so the inspiration slate cannot overlap. The strategist sees a total of 10 blueprints per iteration (5 relevance hits with their measurement cards rendered in full, 5 adjacent inspirations with only the one-line description shown) and is asked, in the prompt, to use the relevance hits as composition or reuse candidates and the inspiration hits as material for new blueprints.

Library size. After the 30 main searches, the library contained $1{,}865$ blueprints (deduplicated by structural hash) with measurement cards (cards are written only after a blueprint has been evaluated at least once). Of these, the strategist chose only those with relevance scores above a fixed threshold during retrieval; on average, 5 relevance hits and 5 adjacent inspirations were returned per query.

Models. The executor LLMs are Qwen3-0.6B (`fast` tier), Qwen3-1.7B (`balanced`), and Qwen3-4B (`complex`), each available in `nothink` and `thinking` modes for a total of six configurations; cross-model transfer (Section 5, Appendix A.7) maps each tier to GPT-5 (nano, mini, full) or Gemini-2.5 (flash-lite, flash). The strategist agent is GPT-6 (`gpt-6-sol`), accessed through the Codex CLI. All five comparison frameworks in the cross-framework table use the same strategist and executor as STRATEGIST for fairness.

Infrastructure. The Qwen3 executors are served via vLLM on a local cluster of $8\times$ NVIDIA H200 GPUs (144 GB each). Agent-tier strategies execute through OpenHands with `TerminalTool` and `FileEditorTool` only; we cap subprocess concurrency at 48 to avoid timeouts. All hyperparameters above are fixed across the 30 main benchmarks, the 10 one-shot held-out benchmarks, and the cross-model transfer experiments.

A.5 Strategist Prompt Template

The strategist receives the full prompt in the first iteration and keeps one session; later iterations send only what is new (results, library hits, the updated history). Listing 1 shows the abridged template.

```
[SYSTEM]
You are the Strategist meta-agent. Design 5 candidate strategies for the task.
Goal: beat the reference baseline (BL) on this benchmark. BL's accuracy is the
floor and its cost per sample the budget; within 2 points of accuracy or 2% of
cost is a tie. A win is lower cost at no worse accuracy, or higher accuracy at
no higher cost.
A strategy is an executable Python program:
    async def execute(llm, task: str) -> str        (one tier/mode), or
    async def execute(llms, task: str) -> str       (routes across tiers)
Models: {fast,balanced,complex} x {nothink,thinking}; library blueprints are
queryable by name, traits and measurement cards.
Every slate has >=1 of each kind (new / composed / reuse) and of each target
(accuracy / efficiency / balanced).
[USER]
Task: {task_description}; sample training questions with answers
Baselines on the training split: {18 x 6 table, BL highlighted}
Library hits:
  Relevant: {top_5_relevant_blueprints_with_cards}
  Adjacent inspiration: {top_5_adjacent_blueprints}
Results so far: {each candidate vs BL on the training split: accuracy, cost
  ratio, questions fixed/broken}
Produce (1) Python code for 5 candidates then (2) a JSON block with fields:
candidates[5]={name,kind,target,tier_mode_anchors,why,builds_on,prompts},
search_queries, read_cards, experience[]={workflow,observation,recommendation}.
```

Listing 1: Abridged strategist prompt template (GPT-6).

The `experience` field is the strategist's own writeback to the library: every blueprint it referenced or used must come back with one observation (numerical, e.g., "card shows +12% on math/fast/nothink") and one recommendation (concrete, e.g., "use as cheap lift on short-math fast-tier; skip on common-sense"). These notes feed forward into future strategists' card lookups.

A.6 Cost-Weight Derivation

We use two cost units in the paper. For the Qwen3 main experiments, where we control the inference stack, we report *weighted tokens*: per-tier weights proportional to parameter count, justified below. For the cross-model transfer experiments on GPT-5 and Gemini-2.5, we report *USD* computed at the official list prices published on the OpenAI[1] and Google[2] pricing pages, applied identically to baseline and STRATEGIST runs so that any difference reflects strategy choices rather than assumptions.

Qwen3 weights. For a dense decoder-only transformer, per-token forward FLOPs scale as $\approx 2N$ where N is the parameter count [21]. The Qwen3 tiers used in the main experiment have 0.6B, 1.7B, and 4B parameters, giving weight ratios of $1.00 : 2.83 : 6.67$ (Table 6). Input tokens count at the same per-token weight, and tool execution for agent strategies (a sandboxed terminal) is not charged. Thinking-mode tokens are counted along with output tokens; their accuracy lift is paid for in the same currency.

Tier	Model	Parameters	Cost weight
fast	Qwen3-0.6B	0.60B	1.00×
balanced	Qwen3-1.7B	1.70B	2.83×
complex	Qwen3-4B	4.00B	6.67×

Table 6: **Tier cost weights used throughout the paper.** Per-token weights are proportional to parameter counts, matching the leading $2N$ FLOPs/token term for dense transformers [21].

[1] https://openai.com/api/pricing/
[2] https://ai.google.dev/gemini-api/docs/pricing

A.7 Cross-Model Transfer Setup

Cross-model transfer (Section 5) runs the same STRATEGIST-discovered strategies on backends never used during search. The tier-mode mapping is fixed per backend; nothing is re-tuned. Table 7 reports the mapping.

Tier/mode	GPT-5 backend	Gemini-2.5 backend
fast/nothink	gpt-5-nano (effort=minimal)	gemini-2.5-flash-lite
fast/thinking	gpt-5-nano (effort=medium)	gemini-2.5-flash-lite (thinking)
balanced/nothink	gpt-5-mini (effort=minimal)	gemini-2.5-flash
balanced/thinking	gpt-5-mini (effort=medium)	gemini-2.5-flash (thinking)
complex/nothink	gpt-5 (effort=minimal)	gemini-2.5-flash
complex/thinking	gpt-5 (effort=medium)	gemini-2.5-flash (thinking)

Table 7: **Tier-mode mapping for cross-model transfer.** GPT-5 uses three model sizes (nano/mini/full) crossed with two reasoning-effort levels (minimal/medium). Gemini-2.5 uses flash-lite for the fast tier and flash for the balanced and complex tiers; the API sets its own thinking budget, so the nothink and thinking modes issue the same call. Each strategy executes its own routing; only the tier-to-model mapping changes.

B Examples and Case Studies

B.1 Sample Discovered Reasoning Strategy

To make the structure of a discovered strategy concrete, we reproduce one (`abstraction_ladder`) which the strategist proposed as a multi-path reasoning blueprint. The header declares its meta-information (when to use, traits, foundation); the `execute` function (abridged below) defines the call graph. The listings show the interface as implemented: `execute(llm, task, **prompts)`, where `llm` is one client or, for strategies that route across tiers, the dictionary of six clients (written `llms`), and prompt replacements arrive as keyword arguments.

The strategist annotates every blueprint with *traits* (above) so future searches can retrieve by structural similarity (`multi_path`, `decomposition`, ...) rather than only by name. The `when_to_use` field is what subsequent strategists read to decide whether to compose on top of this blueprint.

B.2 Additional Sample Reasoning Strategies

To illustrate the diversity of structures the strategist discovers, we show three more final picks from the main run, each the strategy selected for its benchmark across all iterations (Algorithm 1). The code is abridged from the library files: prompt text is replaced by named helpers, control flow is unchanged.

Logic (BBH): `high_precision_execution_router`**.** A 0.6B router call decides whether a short program on the literal data in the task can compute and verify the answer. If so, the task goes to the library's tool-using agent blueprint at `complex/thinking`; otherwise to self-consistency at the same tier. The blueprint is composed from two library blueprints and adds only the gate. On BBH it scores 85.0% against BL's 86.0% at 0.55× the cost.

Multi-hop QA (MuSiQue): `strict_roundtrip_count_multihop`**.** The strategy runs a library multi-hop blueprint and then guards one failure mode: a count answered as a digit when the passages spell it as a word. It asks a 1.7B model to spell the number, accepts the word only if it occurs in the context and an independent call decodes it back to the same number, and otherwise keeps the digit. On MuSiQue it scores 42.0% against BL's 41.0% at 0.93× the cost.

Math (GSM8K): `dual_verbatim_evidence_lens_math`**.** Two parallel `complex/nothink` calls copy verbatim spans from the problem, one for what is asked and one for counting scope and units; spans that do not occur in the problem are dropped. A single `complex/thinking` call then solves the original problem with the spans as an attention guide. On GSM8K it scores 92.0% against BL's 94.0% at 0.13× the cost.

```python
META = {
    "name": "abstraction_ladder",
    "description": (
        "Solve at 3 abstraction levels (principle, specific, "
        "concrete). If all agree, high confidence. If "
        "disagreement, prefer the concrete-example answer as most grounded.
    "
    ),
    "when_to_use": (
        "Multi-step reasoning or math problems where abstract "
        "reasoning might miss details but pure computation might "
        "miss the big picture."
    ),
    "tags": ["reasoning", "multi_path"],
    "uses": ["generate"],
    "traits": {
        "foundation": "generate",
        "reasoning": True,
        "self_improving": False,
        "multi_path": True,
        "tool_use": False,
        "decomposition": True,
        "multi_turn": True,
    },
}

async def execute(llm, task: str, **kwargs) -> str:
    system = f"{base_system(task)} {output_instruction(task)}"
    abstract  = await generate(llm, level1_prompt(task),  system)
    specific  = await generate(llm, level2_prompt(task, abstract), system)
    concrete  = await generate(llm, level3_prompt(task, specific), system)
    if extract_answer(task, specific) == extract_answer(task, concrete):
        return extract_answer(task, concrete)   # two grounded levels agree
    adjudicated = await generate(
        llm, adjudicate_prompt(task, abstract, specific, concrete), system)
    return extract_answer(task, adjudicated)
```

Listing 2: The `abstraction_ladder` blueprint: META declaration plus `execute` body. Multi-path reasoning at three abstraction levels; agreement between the two grounded levels is accepted, otherwise a fourth call adjudicates with all three visible. Prompt text is abridged into named helpers.

```python
from strategist.workflows.baseline_agent_code import execute as agent_code
from strategist.workflows.baseline_self_consistency import execute as
    self_consistency

async def execute(llms, task):
    # EXECUTE if a program on the literal data can compute and verify the
    answer;
    # CONSENSUS for world knowledge, pragmatics, analogy, or when uncertain
    .
    route = await generate(llms["fast"]["nothink"], route_prompt(task))
    if route.strip().upper().startswith("EXECUTE"):
        result = await agent_code(llms["complex"]["thinking"], task)
    else:
        result = await self_consistency(llms["complex"]["thinking"], task)
    return extract_answer(task, result)
```

Listing 3: `high_precision_execution_router`: a cheap router sends deterministic computations to the tool agent and everything else to self-consistency.

```
from strategist.workflows.question_bound_span_projection_multihop import
    projected_walk

async def execute(llms, task):
    answer = extract_answer(task, await projected_walk(llms, task)).strip("
     .")
    context, _, question = task.rpartition("\nQuestion:")
    asks_for_count = ("how many" in question.casefold()
                      or "number of" in question.casefold())
    if asks_for_count and answer.isdecimal() and 0 <= int(answer) <= 12:
        word = (await generate(llms["balanced"]["nothink"],
                               spell_prompt(answer))).strip().lower()
        attested = (re.fullmatch(r"[a-z]{3,}", word)
                    and re.search(rf"\b{re.escape(word)}\b", context, re.I)
    )
        if attested:
            decoded = await generate(llms["balanced"]["nothink"],
    decode_prompt(word))
            m = re.fullmatch(r"\s*(\d{1,2})\.?\s*", decoded)
            if m and m.group(1) == answer:
                answer = word
    return extract_answer(task, answer)
```

Listing 4: `strict_roundtrip_count_multihop`: a library multi-hop blueprint followed by a round-trip check before a count is rewritten as a word.

```
async def execute(llms, task):
    target_text, scope_text = await asyncio.gather(
        generate(llms["complex"]["nothink"], copy_target_spans_prompt(task)
    ),
        generate(llms["complex"]["nothink"], copy_scope_spans_prompt(task))
    )
    spans = [l.strip() for l in (target_text + "\n" + scope_text).
    splitlines()
             if l.strip() and l.strip() in task][:6]
    focus = "\n".join(spans) if spans else "No exact spans were selected."
    answer = await generate(llms["complex"]["thinking"],
                            solve_prompt(task, focus),
                            system=output_instruction(task))
    return extract_answer(task, answer)
```

Listing 5: `dual_verbatim_evidence_lens_math`: parallel verbatim evidence extraction, filtered against the source, then one solving call.

These three strategies cover different control-flow patterns (route-then-delegate, compose-then-guard, parallel-extract-then-solve) yet share the same interface, so each can serve as a component of a later blueprint.

B.3 Sample Measurement Card

After a strategy is run on a benchmark, the strategist writes a *measurement card* summarizing what worked, what failed, and how to use the strategy next time. Cards are stored alongside their strategy and are returned as part of library retrieval. Below is a simplified, shortened excerpt of the card for `cite_then_extract`:

```
# cite_then_extract

## Summary
Two-step pipeline: first cite the supporting passage, then extract
the answer from the cited span. Reduces hallucination on multi-hop
QA where the document provides the evidence.

## Measured behavior
- 2WikiMQA: 74.5% acc (BL 66.5%, +12.0% rel) at 23.2k weighted
  tokens / sample.
- HotpotQA: 63.5% acc (BL 53.0%, +19.8% rel).
- MuSiQue:  24.5% acc (BL 19.5%, +25.6% rel).

## Strengths
- Stable on benchmarks where the evidence sits in a long context.
- Robust to distractor passages: citing first forces the model to
  commit to a span before answering.

## Weaknesses
- On short-context QA the citation step is wasted cost.
- If the cite step picks the wrong span, the extraction step
  inherits the error with no recovery.

## Notes from prior use
- 2026-04-22: Used as composition base for entity_focused_debate
  (HotpotQA). Composed candidate beat parent +5pp.

...
```

Listing 6: Excerpt of the measurement card for `cite_then_extract`, written by the strategist after evaluation. Cards are returned with their blueprint during retrieval and inform the design choices made by future strategists.

B.4 Search Queries and Experience Notes

Table 8 samples actual `search_queries` and `experience` notes the strategist wrote during search runs. These notes are the cross-task knowledge that future iterations and future tasks read; they are why the library compounds rather than starts from scratch each task.

Field	Sample value
`search_queries`	"cite-then-extract on long context multi-hop" "self-consistency vote at fast tier with 3 samples" "multi-agent debate but only on questionable cases" "decomposition into sub-questions then aggregate"
`experience` (workflow, obs, rec)	`self_consistency`: "card shows +12pp on math/fast/nothink; blind matched at 11pp" / "use as cheap lift on short-math fast-tier; skip on commonsense." `multi_agent_debate`: "2-agent fails by mutual-error agreement on 30% of NQ-Open misses" / "add a third tie-breaking agent or vote." `cite_then_extract`: "stable on long-context QA; wasted cost on short-context" / "compose only when context length $>$ 1024 tokens."

Table 8: **Sampled `search_queries` and `experience` notes.** The strategist writes both at every iteration; queries are the input to library retrieval, and experience notes are the writeback that next-iteration strategists (or strategists on a different task) read in the card.

B.5 Per-Iteration Progression Example

To make the iterative dynamics concrete, Table 9 shows the per-iteration progression on AQuA-RAT. Each row reports the candidate of that iteration with the highest hidden-set accuracy; the final pick follows Algorithm 1.

Iter	Kind	Best candidate on hidden	Hidden Acc	Hidden Cost (k)
1	reference	`baseline_cot (BL)`	85.0	21.2
2	reuse	`difficulty_routed_certificate_mcq`	89.0	18.5
3	reuse	`difficulty_routed_certificate_mcq`	91.0	21.7
4	reuse	`difficulty_routed_certificate_mcq`	92.0	20.9
5	**composed**	`intended_operation_commitment_cot_math`	**90.0**	**16.2**
6	new	`rotated_label_invariance_math`	88.0	18.6
7	reuse	`label_permutation_invariance_vote_mcq`	90.0	37.3
8	composed	`structural_coordinate_switch_math`	89.0	15.9
9	composed	`terminal_leakage_sanitized_math`	90.0	17.4
10	new	`first_commitment_race_math`	89.0	18.0

Table 9: **Iteration-by-iteration progression on AQuA-RAT.** For each iteration, the candidate with the highest hidden-set accuracy and its weighted cost per sample; iteration 1 is the reference BL (CoT on `complex/thinking`). The final strategy (bold, from iteration 5) beats BL on both validation and hidden ($+5$ points at $0.77\times$ cost on hidden); the iteration-4 candidate, the most accurate on hidden, is not picked because on validation it matches BL's accuracy at 14% higher cost.

C Per-Benchmark Result Tables

C.1 Per-Benchmark Main Results

Table 10 reports per-benchmark numbers behind Figure 7 for the 30 main benchmarks. Cost is weighted tokens per sample (parameter-count weights), in thousands. Colors and arrows follow the rest of the paper.

C.2 Per-Benchmark One-Shot Library Retrieval

Table 11 reports the absolute Ours/Baseline accuracy and cost values underlying Table 3.

C.3 Per-Benchmark Cross-Model Transfer

Table 12 reports the absolute Ours/Baseline accuracy and total USD cost values underlying Table 4.

Benchmark	Ours Acc	BL Acc	ΔAcc	Ours Cost (k)	BL Cost (k)	ΔCost	Result
SQuAD v2	79.1	73.2	↑8.0%	5.6	25.0	↓77.4%	Win
TruthfulQA	69.0	64.5	↑7.0%	5.8	7.5	↓21.9%	Win
2WikiMQA	76.5	72.5	↑5.5%	33.8	98.6	↓65.7%	Win
Amazon-Rec	28.0	24.0	↑16.7%	45.6	70.3	↓35.1%	Win
SAT-Math	98.2	94.5	↑3.8%	4.9	12.1	↓59.4%	Win
LSAT-AR	81.0	77.6	↑4.4%	54.9	82.0	↓33.0%	Win
LSAT-RC	81.5	78.5	↑3.8%	12.4	24.8	↓49.9%	Win
NQ-Open	19.0	17.0	↑11.8%	19.1	34.4	↓44.5%	Win
LSAT-LR	82.0	80.5	↑1.9%	65.2	98.1	↓33.5%	Win
Gaokao-Math	94.5	93.5	↑1.1%	89.7	111.3	↓19.4%	Win
MuSiQue	42.0	41.0	↑2.4%	67.6	73.0	↓7.4%	Win
MMLU-STEM-Hard	88.0	87.0	↑1.1%	11.9	35.2	↓66.2%	Win
MMLU	79.5	78.5	↑1.3%	10.2	96.3	↓89.4%	Win
AQuA-RAT	87.0	86.5	↑0.6%	17.5	19.4	↓9.6%	Win
StrategyQA	70.0	69.5	↑0.7%	3.0	3.7	↓19.6%	Win
WinoGrande	77.5	77.5	0.0%	4.9	5.6	↓11.8%	Win
ANLI-R3	69.5	69.5	0.0%	4.6	9.6	↓51.7%	Win
SVAMP	93.0	93.5	↓0.5%	1.8	2.0	↓9.3%	Win
SciQ	97.5	98.0	↓0.5%	1.2	3.7	↓69.1%	Win
HellaSwag	70.5	71.0	↓0.7%	9.0	22.9	↓60.9%	Win
BBH	85.0	86.0	↓1.2%	23.3	42.5	↓45.2%	Win
CommonsenseQA	76.5	77.5	↓1.3%	5.8	14.8	↓60.8%	Win
GSM8K	92.0	94.0	↓2.1%	8.3	62.8	↓86.8%	Win
ASDiv	93.0	95.0	↓2.1%	6.9	36.6	↓81.1%	Win
MBPP+	80.5	79.0	↑1.9%	108.5	109.0	↓0.4%	Tie
MinervaM	38.5	39.0	↓1.3%	28.2	28.7	↓1.5%	Tie
MedQA	68.5	72.5	↓5.5%	12.6	28.5	↓55.8%	Mixed
AIME 2024	60.0	65.0	↓7.7%	287.6	372.5	↓22.8%	Mixed
HotpotQA	64.5	64.0	↑0.8%	34.7	31.9	↑9.0%	Lose
COPA	100.0	100.0	0.0%	2.4	2.3	↑6.1%	Lose

Table 10: **Per-benchmark main results across 30 benchmarks.** BL is the baseline cell (of the 18×6 matrix) with the highest accuracy on the benchmark's training split, scored on the test set. Δ values are relative changes; colors follow the 2-point accuracy and 2% cost tie bands. Result: Win = better on accuracy or cost and worse on neither; Mixed = better on one and worse on the other; Lose = worse on at least one and better on neither. Cost is weighted tokens per sample in thousands, computed with parameter-count weights.

Benchmark	Ours Acc	Ours Cost (k)	BL Acc	BL Cost (k)	ΔAcc	ΔCost
ReCLor	83.5	47.4	85.0	72.0	↓1.8%	↓34.2%
SAT-EN	86.5	16.7	86.5	37.6	0.0%	↓55.6%
TriviaQA	50.5	25.8	46.0	64.6	↑9.8%	↓60.0%
MATH500	94.5	47.2	92.5	61.0	↑2.2%	↓22.6%
BoolQ	85.5	32.7	73.5	1.0	↑16.3%	↑3103%
SIQA	77.5	6.7	73.5	1.5	↑5.4%	↑337%
OpenBookQA	94.5	6.9	93.0	13.6	↑1.6%	↓49.2%
PIQA	89.0	5.8	87.5	16.3	↑1.7%	↓64.5%
QUALITY	79.5	52.0	78.5	48.4	↑1.3%	↑7.4%
DROP	72.5	9.5	61.5	4.9	↑17.9%	↑92.4%

Table 11: **Full per-benchmark one-shot library retrieval results.** Accuracy is on the test set; cost is weighted tokens per sample (parameter-count weights). ΔAcc and ΔCost are relative changes vs the baseline (formula and color/arrow conventions match Table 10).

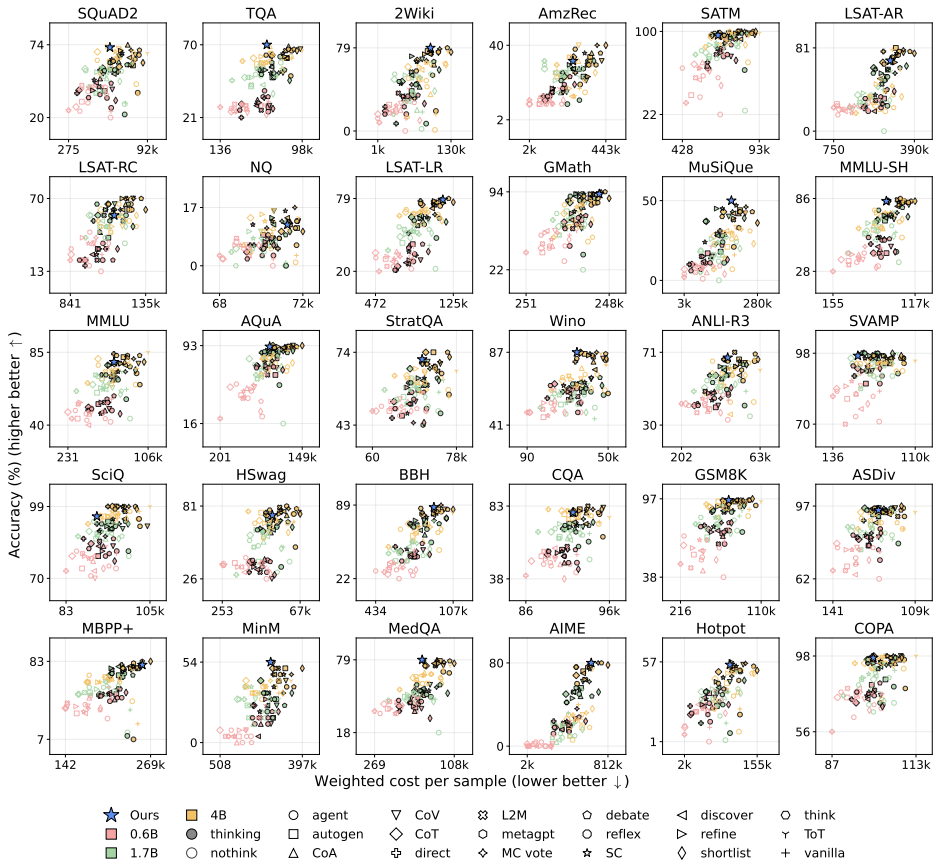


Figure 10: **The same view on the training split.** As Figure 6, but every baseline cell and STRATEGIST's final strategy are scored on the benchmark's training split (the rows the strategist sees) instead of the test set. This comparison is biased toward BL: BL is, by definition, the most accurate of the 108 cells on exactly these rows, so it sits at the top of each panel, while STRATEGIST's strategy was chosen on the separate validation and hidden splits. Only the test set (Figure 7, Figure 6) scores both on questions neither was chosen on, which is why the paper reports it.

	GPT-5						Gemini-2.5					
Benchmark	BL Acc	Ours Acc	ΔAcc	BL $	Ours $	ΔCost	BL Acc	Ours Acc	ΔAcc	BL $	Ours $	ΔCost
AQuA-RAT	31.0	89.5	↑189%	12.39	1.67	↓86.5%	71.0	90.5	↑27.5%	2.22	1.07	↓51.6%
2WikiMQA	29.5	78.5	↑166%	7.45	3.67	↓50.8%	21.5	77.5	↑260%	2.02	0.57	↓72.0%
HotpotQA	39.0	76.5	↑96.2%	7.71	7.13	↓7.5%	33.0	76.5	↑132%	2.52	1.20	↓52.2%
MBPP+	56.5	82.5	↑46.0%	2.16	1.89	↓12.3%	69.5	80.5	↑15.8%	1.79	1.37	↓23.7%
SQuAD2	57.5	74.0	↑28.7%	0.29	0.11	↓61.5%	74.0	70.5	↓4.7%	0.37	0.16	↓57.9%
NQ-Open	10.0	13.5	↑35.0%	9.28	7.17	↓22.8%	11.0	24.0	↑118%	1.91	1.50	↓21.8%
TruthfulQA	84.0	95.5	↑13.7%	3.86	1.08	↓71.9%	85.0	93.0	↑9.4%	1.31	0.50	↓61.9%
ASDiv	93.0	95.0	↑2.2%	0.67	1.76	↑163%	93.5	95.5	↑2.1%	0.88	0.59	↓33.1%
HellaSwag	93.5	94.0	↑0.5%	0.07	0.09	↑16.9%	91.5	92.0	↑0.5%	0.50	0.44	↓12.3%
MinervaM	24.0	24.5	↑2.1%	9.97	10.45	↑4.8%	26.5	24.0	↓9.4%	4.96	3.68	↓25.9%
Total				**53.86**	**35.02**	**↓35.0%**				**18.50**	**11.07**	**↓40.1%**

Table 12: **Full per-benchmark cross-model transfer results.** For each benchmark, the STRATEGIST-discovered strategy and the strongest leaderboard baseline workflow are both re-executed on the new backend; ΔAcc and ΔCost compare them on the same backend. 200 deterministic samples per cell. BL $ and Ours $ are total USD across all 200 samples, computed at the official list prices published on the OpenAI (GPT-5) and Google (Gemini-2.5) websites; the same per-token rates are applied to baseline and STRATEGIST strategies.

D Discussion and Limitations

D.1 Zero-Shot Ablation Failure Analysis

Three of the ten zero-shot ablation runs failed at execution time (Table 5). All three failures were silently introduced by the meta-LLM in candidate code, and would have been caught by the iterative loop's evaluator (which re-prompts with the traceback) but were not because the loop was capped at one iteration. The three failures were: **(SAT-Math)** the proposed strategy used `Promise.all` (a JavaScript idiom) inside a Python `async` body, raising `NameError` at execution; **(MMLU)** a similar Promise/await confusion paired with a missing `import asyncio`; and **(AIME 2024)** the candidate's name collided with an imported library blueprint, causing the harness to fail to register the candidate. None of these are subtle reasoning failures: they are surface-level coding mistakes that would normally be flushed out by the evaluator in iterations 2–10. This is precisely the kind of error the iterative loop is meant to catch.

D.2 Limitations

STRATEGIST has three notable limitations. **(i) Strategist cost.** The meta-LLM (GPT-6) is itself expensive; each iteration of the search loop spends about 10k strategist prompt tokens plus the cost of evaluating $K = 5$ candidates on three development splits of up to 100 samples each (Appendix A.4). For tasks where a single strong baseline is already known, the search overhead may exceed the inference-time savings. The one-shot retrieval variant (Section 5) is the cheapest path once the library is built up, but the library still has to be grown by some prior search. **(ii) Coding-mistake failures.** The zero-shot ablation (Appendix D.1) shows that without the iterative loop, a non-trivial fraction of meta-LLM proposals contain surface-level coding bugs. The full system catches these by re-prompting with the traceback, but the cost is paid in wasted candidate evaluations. Since the generated strategy can be applied to process thousands of data samples, this cost is negligible compared to the long-term gains. **(iii) Meta-LLM choice.** All experiments use GPT-6 as the strategist; we do not ablate this choice. Results may shift with weaker meta-LLMs, but the iterative loop and library together are designed to compensate for individual proposal noise (the loop catches coding bugs, the library supplies prior structure), so we expect the framework to remain useful over a reasonable range of meta-models.

D.3 Reproducibility Statement

All hyperparameters used in the main experiment are constants (no per-benchmark tuning) and are listed in Appendices A.3 and A.4. The 30 main benchmarks, 10 one-shot held-out benchmarks, and 10 cross-model transfer benchmarks are public datasets cited in Section 4; for each we use deterministic splits drawn from a fixed random seed. The Blueprint Library snapshot after the 30 main searches (1,865 blueprints with measurement cards) is released alongside the paper. The strategist itself is GPT-6 accessed through the Codex CLI; the executor LLMs are Qwen3-0.6B/1.7B/4B served via vLLM. Cross-model transfer uses OpenAI's GPT-5 family and Google's Gemini-2.5; USD costs are computed at the official list prices published on the OpenAI and Google websites, so all reported dollar figures correspond to what an external user would pay. Every per-benchmark cell reported in the paper can be reproduced from the per-cell JSON results in `results/matrix_*`. The complete codebase (orchestrator, library, eval harness, and per-benchmark candidate logs) is released under an open-source license at https://github.com/zlab-princeton/Strategist.

E Extended Related Work

This appendix expands the three Related Work paragraphs in the main text. Citations are identical; only the discussion is longer.

LLM agents. Large language models have rapidly evolved from single-call question answerers into multi-step agent systems. Chain-of-thought prompting [54] and self-consistency decoding [53] showed that intermediate reasoning steps improve accuracy on complex tasks. Tree of Thoughts [61] extended this to deliberate search over reasoning paths. ReAct [62] interleaved reasoning with tool use, and Toolformer [44] taught models to call external APIs autonomously. Reflexion [46] introduced verbal self-reflection as a feedback mechanism. Multi-agent frameworks such as AutoGen [56] enabled conversational collaboration among multiple LLM agents. Formal representations

for these systems range from state machines [57] to declarative compilation [22] to optimizable computational graphs [77]. A recent survey [64] organizes this landscape as agentic computation graphs, distinguishing static, pre-execution, and in-execution optimization. As these systems grow more complex, the strategy layer that determines how models are called, in what order, and with what verification has become the central design bottleneck. STRATEGIST addresses this bottleneck by automating the design of the strategy layer itself.

Automated design of agentic systems. Structure search methods write strategies as raw code [18] or compose them from fixed components [69] or predefined modules [45, 67]; the first yields isolated programs, the second bounds what strategies can be discovered. Optimization methods refine one strategy per task, by reinforcement learning [25, 70, 10] or by LLM-driven evolution and refinement [9, 29], and start over on the next task. Routing methods train controllers to select among existing strategies [50, 65] without constructing them [68]. STRATEGIST differs by allowing freeform creation through the Strategy Blueprint, requiring no training, constructing strategies itself, and routing across model tiers within each strategy.

Cross-task transfer and experience accumulation. Transferring knowledge across tasks has been studied through meta-learning [8, 17], neural architecture search [7], and prompt optimization [59]. In the LLM agent setting, ExpeL [72] extracts reusable insights from agent experiences, Voyager [52] grows a skill library as it explores new environments, and recent work has formalized this direction as lifelong learning for LLM agents [73]. Despite this, automated system design methods [18, 69] still optimize per task independently. STRATEGIST maintains a shared library of blueprints and their measured properties that grows with every task, providing both relevant components (exploitation) and diverse inspiration from other tasks (exploration). To our knowledge, it is the first system for automated strategy design that accumulates knowledge across tasks.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and Section 1 state three contributions (the STRATEGIST meta-agent, the Strategy Blueprint, the accumulating Blueprint Library) and the headline empirical claim (beating the strongest baseline on 24 of 30 benchmarks and matching its average accuracy, 74.7% vs. 74.0%, at 39% lower cost); each is matched by an experiment in Sections 5–7 and the cross-model transfer / one-shot retrieval results.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Appendix D.2 lists three notable limitations: the cost of the strategist (a frontier LLM), surface-level coding bugs surfaced in the zero-shot ablation (Appendix D.1), and the untested dependence on the choice of meta-LLM.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: The paper presents an empirical method and contains no theorems or formal proofs.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Section 3 and Appendices A.1–A.7 document the search algorithm (pseudocode), the data splits and caps, search hyperparameters (K=5, N=10, `max_tokens`=16384), strategist prompt template, library retrieval scorer, executor and strategist models, infrastructure, cost-weight derivation, and the cross-model tier mapping. Per-benchmark numbers reproducing every cell in the tables are released alongside the codebase.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: All 40 benchmarks are existing public datasets cited in Appendix A.2. The full codebase is released under an open-source license at https://github.com/zlab-princeton/Strategist, with run instructions described in Appendix D.3.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Section 4 describes the executor models, baselines, evaluation metrics, and weighted-cost definition. Appendices A.2–A.7 give the four-way data split with caps, the K/N search hyperparameters, the strategist prompt template, library retrieval details, cost-weight derivation, and the cross-model tier mapping. All hyperparameters are constants across the 30 main benchmarks; nothing is per-benchmark tuned.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

Justification: We do not report formal error bars. Each test cell uses up to 200 fixed samples (about the first 200 rows of each test file); the implied Wilson 95% CI is ± 3–4% per cell. Because we average over 30 benchmarks, our headline accuracy and cost claims are well-resolved despite per-cell sampling noise, but multi-seed runs were not feasible within our compute budget.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The Infrastructure paragraph in Appendix A.4 reports the executor stack (Qwen3 served via vLLM) and the strategist (GPT-6 via the Codex CLI). Cross-model transfer USD costs at the official OpenAI/Google list prices are reported per benchmark in Appendix C.3.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: The work uses only existing public benchmarks and existing LLM APIs under their published terms of service; it involves no human subjects, no scraped or personally identifiable data, and no deployment.

   Guidelines:

- The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Positive impact: lower-cost inference strategies extend useful agentic capability to smaller models, reducing the compute footprint of LLM-based applications. Potential negative impact: the same techniques could be applied to make malicious agents (e.g., automated misinformation pipelines) cheaper to operate; mitigation is the same as for any other inference-time optimization research, namely transparent release with monitoring of downstream uses.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: We do not release any model weights, scraped datasets, or generative models. The released artefact is an inference-strategy framework plus a library of executable Python blueprints; it has no high-misuse-risk component beyond the underlying LLM APIs that are themselves controlled by their providers.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.

- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All 40 benchmarks are cited to their original papers in Appendix A.2. All baseline methods are cited to their original publications in the same appendix. The Qwen3 models are used under the Apache 2.0 license; GPT-5 and Gemini-2.5 are accessed via paid APIs under their respective terms of service.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The two new assets are the STRATEGIST codebase (orchestrator, eval harness, per-benchmark candidate logs) and the Blueprint Library snapshot. Both are documented in Appendix D.3 and will be released with run instructions and the per-blueprint metadata schema.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

Justification: The paper does not involve crowdsourcing or human subjects research; all experiments are automated evaluation on existing public benchmarks.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: The paper does not involve human subjects research, so IRB approval was not required.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: LLMs are central to both the method and the experiments: GPT-6 acts as the strategist meta-agent that proposes candidate strategies, and Qwen3-0.6B/1.7B/4B, GPT-5 (nano/mini/full), and Gemini-2.5 (flash-lite/flash) act as task executors. All roles, models, and prompts are documented in Section 4 and Appendices A.4–A.7.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.