

# Towards Compositional Generalization of LLMs via Skill Taxonomy Guided Data Synthesis

Anonymous ACL submission

## Abstract

Large Language Models (LLMs) and agent-based systems often struggle with compositional generalization due to a data bottleneck in which complex skill combinations follow a long-tailed, power-law distribution, limiting both instruction-following performance and generalization in agent-centric tasks. To address this challenge, we propose STEPS, a Skill Taxonomy-guided Entropy-based Post-training data Synthesis framework for generating compositionally challenging data. STEPS explicitly targets compositional generalization by uncovering latent relationships among skills and organizing them into an interpretable, hierarchical skill taxonomy using structural information theory. Building on this taxonomy, we formulate data synthesis as a constrained information maximization problem, selecting skill combinations that maximize marginal structural information within the hierarchy while preserving semantic coherence. Experiments on challenging instruction-following benchmarks show that STEPS outperforms existing data synthesis baselines, while also yielding improved compositional generalization in downstream agent-based evaluations.

## 1 Introduction

The rapid scaling of Large Language Models (LLMs) has led to impressive gains across language understanding and generation tasks. However, both LLMs and agent-based systems continue to struggle with compositional generalization—the ability to flexibly recombine learned skills into novel configurations—particularly in complex instruction-following and agent-centric settings (Lake and Baroni, 2018; Okawa et al., 2023). A key obstacle is a fundamental data bottleneck: while individual atomic skills are abundantly represented in training corpora, complex skill combinations follow a long-tailed, power-law distribution (Clauset and Shalizi, 2009). This imbalance severely limits coverage

of compositionally challenging scenarios, leading to sharp performance degradation when multiple skills must be coordinated, as evidenced by evaluations such as SKILL-MIX (Kudo et al., 2023; Yu et al., 2024; Zhao et al., 2024).

To alleviate this data bottleneck, prior work has mainly explored data-centric strategies such as data mixture optimization and pedagogical sequencing, which reweight training samples or adjust learning order to improve sample efficiency (Ge et al., 2024; Wu et al., 2024; Chen et al., 2023; Zhao et al., 2025; Hu et al., 2025). While effective at better utilizing existing data, these approaches primarily operate at the level of individual skills, and therefore do not fundamentally address the scarcity of compositionally challenging examples involving multiple interacting skills. Complementary efforts synthesize such data through stochastic or heuristic skill mixing (Chen et al., 2024; Kaur et al., 2025), but typically ignore the underlying structural and hierarchical relationships among skills, resulting in semantically incoherent compositions and inefficient exploration of the combinatorial space. From a broader perspective, existing methods lack an explicit, principled formulation of skill composition in the form of a unified taxonomy. Thus their ability to support systematic data synthesis is limited. This gap motivates the need for structure-aware synthesis approaches that explicitly model skill dependencies and compositions.

In this paper, we propose STEPS, a Skill Taxonomy-guided Entropy-based Post-training data Synthesis framework for addressing compositional data scarcity in the complex instruction following and agent-based systems. As shown in Figure 1 (a), STEPS first induces a hierarchical skill dependency taxonomy by constructing a skill co-occurrence graph and recursively merging groups of skill nodes with dense intra-group connections and sparse inter-group connections. This process is guided by minimizing structural entropy (Li and

Pan, 2016; Li, 2024), which favors hierarchies in which high-weight co-occurrence edges are largely explained within local groups rather than across unrelated ones, yielding a compact and interpretable dependency structure over skills. Then as illustrated in Figure 1 (b), building on this taxonomy, STEPS then synthesizes new data by recursively composing skill nodes under an objective of maximizing marginal structural entropy gain. Instead of sampling combinations at random, this strategy prioritizes skill compositions that introduce the largest amount of new structural information while remaining coherent with the induced hierarchy. As a result, STEPS efficiently explores the combinatorial skill space and generates compositionally challenging data that better supports generalization.

Experiments on several challenging instruction following benchmarks and on agent scenario demonstrate that our framework significantly enhances the model performance, providing a principled solution for overcoming data sparsity and advancing the compositional capabilities of LLMs.

## 2 Preliminary

### 2.1 Problem Formalization

Large Language Models (LLMs) are increasingly required to solve tasks that involve the coordinated use of multiple functional capabilities, which we refer to as *skills*. Formally, a skill  $s \in \mathcal{S}$  denotes an atomic functional ability required to execute an instruction, such as logical reasoning, mathematical calculation, or code debugging. An instruction can therefore be characterized by a set of jointly involved skills, which we represent as a  $k$ -tuple  $X = \{x_1, x_2, \dots, x_k\}$ , where  $k$  reflects the *compositional complexity* of the task. Larger values of  $k$  correspond to more challenging instructions that require integrating multiple interacting skills.

Our goal is to improve the compositional generalization of LLMs and agent-based systems by synthesizing training data that systematically covers informative and challenging skill combinations beyond those frequently observed in corpora.

### 2.2 Skill Taxonomy of LLMs

The capabilities of LLMs are not independent: complex skills are often built upon foundational ones, and certain skills frequently co-occur to solve intricate tasks. Motivated by this observation, we model the skill space as a *skill taxonomy*  $\mathcal{T}$ , where each leaf node corresponds to an atomic skill  $s_i$ ,

each internal node represents a coherent group of closely related skills, and edges encode hierarchical dependency relationships.

Rather than treating the skill space as a flat set, we seek a structured abstraction that exposes its hierarchical organization. Specifically, we aim to induce a *skill taxonomy* that groups closely related skills together while separating weakly related ones across different levels. This taxonomy provides a compact and interpretable representation of the combinatorial skill space, and serves as the structural foundation for selecting and synthesizing skill compositions. From this perspective, data synthesis reduces to selecting skill combinations from the taxonomy: combinations confined to a narrow region tend to be redundant, while those connecting different yet related regions are more informative. This motivates a synthesis strategy that explicitly leverages the structure of the skill taxonomy when generating compositional training data.

## 3 Methodology

Figure 1 demonstrates the two-stage framework for taxonomy-guided data synthesis. We first induce a hierarchical skill taxonomy that captures dependency and compositional relationships among skills. Based on this taxonomy, we then synthesize training data by selecting skill combinations that maximize structural information gain under hierarchical constraints, yielding compositionally challenging yet learnable examples.

### 3.1 Structural Entropy Guided Skill Taxonomy Induction

As shown in Figure 1 (a), to induce a skill taxonomy, we begin by modeling the relationships among skills through their empirical co-occurrence patterns. Following recent studies (Zhao et al., 2025; Li et al., 2025; Du et al., 2025), we annotate each instruction with a set of skill tags and construct a skill co-occurrence graph  $G = (V, E)$ , where each node  $v \in V$  corresponds to an atomic skill and each edge weight reflects the probability of two skills co-occur in one real data.

While co-occurrence statistics capture local relationships, they do not directly reveal the global organization of the skill space, which is often sparse and long-tailed (Barabási and Albert, 1999; Broder et al., 2000; Clauset and Shalizi, 2009). To uncover higher-level structure, we seek to hierarchically merge closely related skills into more abstract

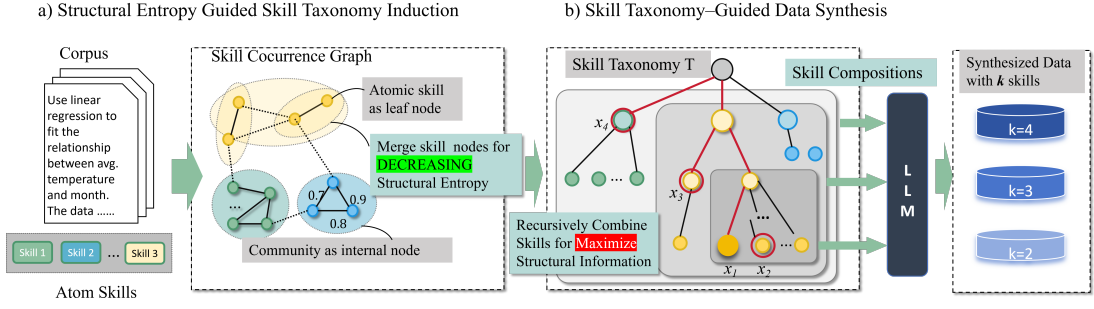


Figure 1: Illustration of the STEP framework.

skill groups, forming a skill dependency taxonomy. This process can be viewed as recursively grouping skill nodes with dense intra-group connections and sparse inter-group connections, yielding progressively higher-level abstractions.

We formalize this objective using structural entropy (Li and Pan, 2016; Li, 2024), which provides a principled criterion for identifying such groupings. Intuitively, structural entropy is low when high-weight co-occurrence edges are largely explained within local groups. Minimizing structural entropy therefore favors hierarchies in which strongly related skills are clustered together, while weakly related skills are separated across levels.

Formally, given a taxonomy  $\mathcal{T}$  represented as a tree structure, we extend structural entropy from communities to individual skill nodes. Let  $\gamma$  denote the leaf node in  $\mathcal{T}$  corresponding to skill  $v$ , and  $\lambda$  the root of the tree. The structural entropy of a skill  $v$ , denoted  $S_e(v)$ , is defined as the cumulative entropy contribution along the path from  $\gamma$  to  $\lambda$ :

$$S_e(v) = \mathcal{H}^{\mathcal{T}}(G; v) = \mathcal{H}^{\mathcal{T}}(G; [\gamma, \lambda]) = - \sum_{\lambda \subset \alpha \subset \gamma} \frac{g(\alpha)}{\text{vol}(G)} \log \frac{\text{vol}(\alpha)}{\text{vol}(\alpha^-)}, \quad (1)$$

where  $\alpha$  ranges over the community nodes on the path from  $\gamma$  to  $\lambda$ ,  $g(\alpha)$  denotes the boundary weight of  $\alpha$ , and  $\text{vol}(\cdot)$  its volume. A higher  $S_e(v)$  indicates that the skill lies in a less cohesive region of the taxonomy, suggesting greater potential for forming diverse and informative compositions.

We induce the skill taxonomy using a bottom-up agglomerative procedure. Starting from the skill co-occurrence graph  $G = (V, E)$ , we initialize each atomic skill  $v \in V$  as a singleton community, which corresponds to a leaf in the taxonomy. At each step, we consider merging a pair of communities and choose the merge that yields the largest decrease in structural entropy. The merged community is added as a new internal node whose children

are the two merged communities. We repeat this process until all skills are merged into a single community, which becomes the root of the taxonomy. This procedure yields a hierarchical coding tree  $T$  that groups strongly co-occurring skills at lower levels while separating weakly related ones at higher levels.

### 3.2 Skill Combination via Structural Information Maximization

Given the induced skill taxonomy  $T$ , our next goal is to synthesize data with skill combinations that can rapidly increase the compositional generalizability of models. Rather than treating all combinations equally, given the induced skill taxonomy  $\mathcal{T}$ , our objective is to identify skill combinations that provide the largest structural information gain for training. Intuitively, not all  $k$ -skill combinations are equally informative: combinations that remain within a small, well-explored region of the taxonomy tend to be redundant, whereas those that introduce new structural relationships across different regions can yield greater learning benefit.

To formalize this notion, we seek a principled measure of the incremental structural information contributed by adding a new skill to an existing set. Importantly, the informational value of a skill depends on which skills have already been selected, as overlapping skills may share structural context in the taxonomy. This dependency makes it insufficient to evaluate skill combinations using independent or additive scores.

To address this, we utilize the chain rule of entropy, which allows the joint structural entropy of a skill set to be decomposed into a sequence of conditional terms. This decomposition enables us to quantify the marginal structural information gain contributed by each skill relative to the previously selected ones, providing a natural objective for constructing informative skill combinations.

For a  $k$ -tuple of skills  $X = \{x_1, x_2, \dots, x_k\}$ , we decompose its joint structural entropy into a sum of conditional terms. Specifically, the conditional structural entropy  $\mathcal{H}^T(G; x_i | X_{<i})$  measures the marginal structural information contributed by skill  $x_i$  relative to the previously selected set  $X_{<i} = \{x_1, \dots, x_{i-1}\}$ . Let  $\delta$  denote the least common ancestor (LCA) of  $x_i \cup X_{<i}$  in the taxonomy  $\mathcal{T}$ , and let  $\gamma$  be the leaf node corresponding to  $x_i$ . The conditional structural entropy is defined as:

$$\mathcal{H}^T(G; x_i | X_{<i}) = - \sum_{\delta \subset \alpha \subseteq \gamma} \frac{g(\alpha)}{\text{vol}(G)} \log \frac{\text{vol}(\alpha)}{\text{vol}(\alpha^-)}, \quad (2)$$

which captures only the novel structural information introduced by  $x_i$  beyond what is already explained by  $X_{<i}$ .

We construct a skill combination  $X$  using a greedy sequential procedure. Starting from an initial skill  $x_1$ , we iteratively select:

$$x_{i+1} = \arg \max_{v \in V \setminus X_{<i}} \mathcal{H}^T(G; v | X_{<i}). \quad (3)$$

Repeating this process yields a  $k$ -tuple that maximizes joint structural entropy, prioritizing skill combinations that are diverse and compositionally challenging.

### 3.3 Entropy-Guided Recursive Synthesis

Maximizing structural entropy provides a principled way to identify informative skill combinations. However, unconstrained maximization may result in combinations that are overly disparate or semantically irrational. Effective composition requires balancing diversity with reasonability, i.e, achieving a ‘‘sweet spot’’.

To achieve this balance, we introduce a recursive, taxonomy-guided search strategy, as shown in Figure 1 (b). Instead of performing global maximization over all skills, the selection of each subsequent skill is first restricted to the local sub-trees in the taxonomy. The search space is then gradually expanded to higher-level parent communities. This bottom-up expansion ensures selected skills are structurally diverse while remaining within a coherent hierarchical context.

Once a skill combination is selected, we synthesize training data based on the identified skills with corresponding reference examples. Leveraging the strong meta-learning capability of LLMs, we prompt the model to generate new instructions and solutions that jointly exhibit the target skills.

This process yields challenging training instances with compositions of skills for the generalization of models. More details are provided in the Appendix.

## 4 Experiments

We evaluate our framework STEPS through extensive experiments on multiple instruction-following benchmarks. Our study aims to address the following research questions: **RQ1**: Can our framework effectively enhance the compositional generalization of existing LLMs? **RQ2**: What is the impact of  $k$ -tuple size on the model’s generalization capabilities? **RQ3**: Can the hierarchical structure of STEPS effectively guide the learning process to enhance compositional generalization? **RQ4**: Does the ‘‘sweet spot’’ constraint provide a more effective scope for skill learning? **RQ5**: Can the STEPS method be used to construct agent data?

### 4.1 Experimental Settings

**Benchmarks and Metrics.** We evaluate model performance on three benchmarks: MT-Bench (Zheng et al., 2023), AlpacaEval 2.0 (Dubois et al., 2024), and WildBench (Lin et al., 2025). For AlpacaEval 2.0, we report the Length-Controlled Win Rate (LC WR) against a reference model to mitigate length bias. For MT-Bench, we report the average score of the responses of our model graded by a judge model. For WildBench, we utilize the WB-Score to evaluate the model’s proficiency in following complex instructions. This metric provides a weighted average of task-wise performance evaluated by GPT-4o (Hurst et al., 2024).

**Baselines.** We select three mainstream open-source LLMs as base and instruct models: Qwen-2.5-7B, Llama-3-8B, and Mistral-7B-v0.3 (Jiang et al., 2023; Team et al., 2024; Grattafiori et al., 2024). We compare our framework against two representative synthetic data methods: (1) Alpaca52k (Taori et al., 2023), a large-scale dataset (52K samples) designed for general instruction following; and (2) Instruct-SkillMix (Kaur et al., 2025), a state-of-the-art compositional method that utilizes random skill pairing. To emphasize the role of data quality, both our method and Instruct-SkillMix are fine-tuned on a compact set of 4K synthetic examples.

### 4.2 Main Results (RQ1)

Table 1 and 2 show the performance across general and complex instruction benchmarks, respectively. **General Instruction Following.** As shown in Table 1, training with instruction-tuning data gen-

Table 1: Comparison of Performance on AlpacaEval 2.0 and MT-Bench for Different Models

Model	Data	AlpacaEval 2.0 LC WR (%)	MT-Bench Score
<b>Mistral-7B-v0.3 Base/Instruct</b>			
Original	-	0.19 / 22.88	3.44 / 6.63
Alpaca52k	52K	4.57 / 5.38	3.79 / 4.03
Instruct-SkillMix	4K	17.17 / 22.63	5.19 / 5.95
STEPS	4K	17.60 / <b>24.42</b>	5.85 / 6.23
<b>Llama-3-8B Base/Instruct</b>			
Original	-	0.20 / 27.85	2.04 / 6.94
Alpaca52k	52K	2.09 / 9.37	4.13 / 5.04
Instruct-SkillMix	4K	0.47 / 30.55	4.55 / 7.15
STEPS	4K	4.65 / <b>35.26</b>	4.66 / <b>7.35</b>

erally enhances the performance of base models. STEPS achieves the most significant gains across both Llama-3-8B and Mistral-7B-v0.3, showing performance improvement compared to the *official* instruct models.

These results suggest that while purely random combinations (as in Instruct-SkillMix) can introduce incoherent instructions that potentially degrade the instruction-following capabilities of refined models, our method maintains a balance between structural diversity and semantic coherence. Furthermore, while the larger Alpaca52k dataset (52K) benefits base models, it consistently leads to performance degradation in instruct-tuned models (e.g., Llama-3-8B Instruct’s MT-Bench score drops from 6.94 to 5.04), indicating that massive single-skill data may induce forgetting of complex compositional abilities.

**Complex Instruction Following.** Effective execution of complex instructions needs the integration of diverse atomic skills. On the WildBench (Table 2) specifically designed to evaluate such multi-dimensional capabilities, our method achieves the highest WB-Scores across all tested architectures. For the Llama-3-8B Base model, the WB-Score improves from -25.7 to 23.0. For the Llama-3-8B Instruct model, our method significantly outperforms Instruct-SkillMix (36.3 vs. 30.8). A critical finding in Table 2 is the counterproductive effect of simple alignment data on complex tasks. Alpaca52k consistently damages the compositional performance of instruct models, with Qwen-2.5-7B’s score plummeting from 45.4 to -11.4. In contrast, our 4K high-entropy samples yield substantial gains, showing that the structural complexity of training data is a more decisive factor for com-

positional generalization than raw data volume.

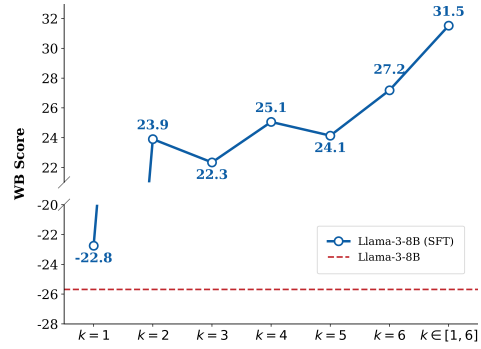


Figure 2: Impact of compositional complexity  $k$ .

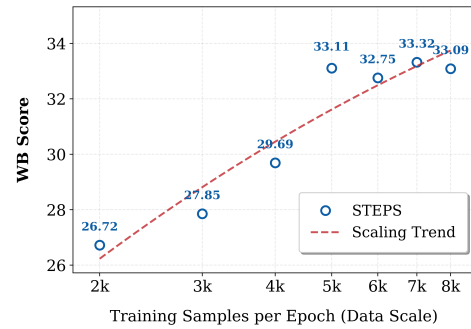


Figure 3: Impact of data size on Llama-3-8B.

### 4.3 Scaling Regularity of STEPS (RQ2)

We investigate the scaling properties of complex skill acquisition by examining the interplay between compositional complexity (the  $k$ -tuple size) and data volume. This analysis seeks to determine whether LLM performance follows predictable patterns as we scale the intricacy of skill intersections and the quantity of synthesized instances.

**Scaling for Combination Complexity  $k$ .** We first evaluate the influence of compositional complexity  $k$  on the model’s generalization capabilities. By training on a fixed budget of  $N = 500$  samples for each  $k \in \{2, \dots, 6\}$  and supplementing them with 4K atomic instructions ( $k = 1$ ), we observe a distinct Compositional Threshold. As illustrated in Figure 2, a critical performance leap occurs when transitioning from atomic instructions to binary combinations, with the WB-Score surging from -22.75 to 23.91. This sharp inflection point indicates that even minimal exposure to skill compositions triggers a fundamental shift in the model’s ability to handle multi-step logical dependencies. Beyond this threshold, increasing the complexity  $k$  generally yields an upward performance trend that

Table 2: Performance comparison on WildBench. We report both Base and Instruct versions (Base / Instruct) for each model. The results across five task categories and the overall WB-Score demonstrate that our framework consistently outperforms original models and competitive synthesis baselines.

Model	Info Seek	Creative	Code & Debug	Reason & Plan	Math & Data	WB-Score
DeepSeek-V2-Chat	52.7	56.4	55.0	54.8	51.4	54.0
GPT-4-turbo	57.2	58.7	55.1	56.2	51.0	55.2
GPT-4o	58.6	59.1	60.5	60.2	57.3	59.3
Qwen-2.5-7B Base/Instruct						
- Original	21.5 / 47.2	6.5 / 44.5	5.3 / 44.7	12.5 / 48.9	5.5 / 44.4	10.7 / 45.4
- Alpaca52k	-6.3 / -13.4	-5.4 / -14.3	-2.0 / -9.1	3.3 / -6.5	-7.7 / -14.1	-3.8 / -11.4
- Instruct-SkillMix	35.3 / 41.9	33.8 / 36.6	27.9 / 41.7	36.1 / 40.8	31.7 / 37.0	32.8 / 39.6
- STEPS	39.3 / 44.0	34.5 / 45.1	31.1 / 49.1	39.5 / 48.3	34.0 / 45.5	35.1 / 45.7
Llama-3-8B Base/Instruct						
- Original	-27.7 / 37.9	-30.4 / 43.0	-23.8 / 26.0	-17.1 / 35.8	-29.8 / 22.4	-25.7 / 34.2
- Alpaca52k	0.5 / 15.9	-18.6 / -10.4	-6.5 / 4.1	0.4 / 11.3	-16.6 / -9.1	-6.1 / 5.0
- Instruct-SkillMix	22.4 / 39.1	13.8 / 20.6	23.9 / 32.5	26.4 / 37.7	9.7 / 15.7	19.7 / 30.8
- STEPS	30.6 / 46.1	7.7 / 25.2	26.8 / 37.9	30.6 / 42.9	12.3 / 18.9	23.0 / 36.3
Mistral-7B-v0.3 Base/Instruct						
- Original	-26.1 / 39.5	-28.1 / 20.7	-21.2 / 29.3	-9.3 / 33.7	-29.8 / 8.5	-22.7 / 28.6
- Alpaca52k	-14.2 / -13.6	-24.8 / -26.6	-13.7 / -13.2	-5.1 / -8.2	-20.0 / -12.5	-14.0 / -14.8
- Instruct-SkillMix	9.3 / 18.6	-6.5 / 2.3	9.3 / 12.1	14.0 / 11.8	-7.8 / -3.1	5.2 / 10.3
- STEPS	21.9 / 39.6	-6.7 / 16.5	17.1 / 30.8	24.3 / 37.8	-5.2 / 10.0	11.3 / 29.0

reaches a local optimum at  $k = 6$  (27.19). Notably, the most robust generalization is achieved through a multi-level mixture encompassing  $k \in [1, 6]$ , which attains a peak WB-Score of 31.52. This synergistic effect suggests that while atomic skills provide essential linguistic primitives, a diverse spectrum of high-order skill tuples is required to facilitate the flexible re-composition of these primitives in novel contexts.

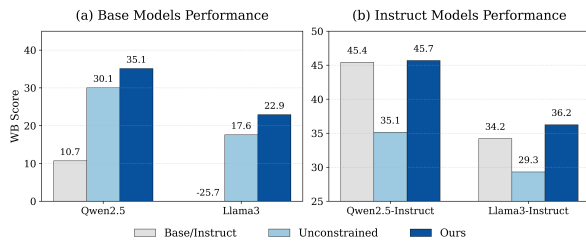


Figure 4: Comparative performance on WB Score across different paradigms. We evaluate Qwen2.5 and Llama3 models in both (a) Base and (b) Instruct settings.

**Scaling for Data Volume under Taxonomy Guidance.** To further explore the data efficiency of our framework, we analyze the model’s performance as the training volume per epoch scales from 2K to 8K samples. Throughout this experiment, we control the proportion of various  $k$ -tuples to isolate the effect of sample size. As presented in Figure 3, the model exhibits a rapid accretion of performance in the initial phase, where the WB-Score rises sig-

nificantly from 26.72 at 2K samples to 33.11 at 5K samples. This steep growth trajectory demonstrates that the high-entropy skill combinations identified by STEPS provide an exceptionally dense learning signal, allowing the model to efficiently map the underlying skill dependency graph.

#### 4.4 Validation of the Skill Taxonomy (RQ3)

In this section, we investigate whether the hierarchical structure discovered by STEPS reflects the intrinsic logical dependencies required for effective skill acquisition. We hypothesize that if our taxonomy accurately captures the prerequisite relationships between skills, then a training sequence aligned with its hierarchical depth should yield better learning efficiency compared to a random mixture. To evaluate this, we implement a *taxonomy-guided curriculum learning* strategy using a fixed budget of 8K data samples per epoch. These samples encompass skill combinations ranging from  $k = 1$  to  $k = 6$ . Following the structural logic of the taxonomy, we adopt a progressive difficulty approach: as training progresses through successive epochs, we gradually increase the proportion of samples containing higher combination counts  $k$ . This strategy ensures the model effectively masters fundamental skills and simpler compositions before being exposed to more complex skill tuples.

As presented in Table 3, the experimental results

Table 3: Performance comparison between standard SFT and Taxonomy-guided learning on WildBench.

Model	Data	WB Score
<b>Llama-3-8B</b>		
- Original	-	-25.68 / 34.22
- SFT	8K	31.48 / 34.84
- STEPS-CL	8K	33.09 / 35.18

provide strong empirical evidence for the validity of our discovered taxonomy: The STEPS guided approach consistently outperforms standard SFT on both Base and Instruct versions of Llama-3-8B. Specifically, our method achieves a peak WB Score of 33.09 for the base model and 35.18 for the instruct version, representing significant improvements over the random mixture baselines of 31.48 and 34.84, respectively. These gains confirm that the hierarchical partitions identified by our structural information framework successfully capture the functional dependencies of the skill space.

Furthermore, the success of this "bottom-up" learning sequence suggests that the Skill Taxonomy is not merely an interpretability tool but a functional roadmap for compositional generalization. By mastering simpler structures first, the model is better equipped to internalize the complex logic required for high-order skill compositions. These findings underscore the importance of structural coherence in training data, demonstrating that the structural organization of examples is as critical as their raw quantity for the development of sophisticated model capabilities.

#### 4.5 Sweet Spot Analysis (RQ4)

To investigate the necessity of balancing diversity with coherence, we compare unconstrained Information Maximization (Section 3.2) against our Recursive Skill Selection Paradigm (Section 3.3), as illustrated in Figure 4.

##### The Inadequacy of Pure Entropy Maximization.

While unconstrained selection significantly outperforms the original base models, it consistently underperforms compared to our recursive framework. Specifically, on Qwen-2.5-Base and Llama-3-Base, our "Sweet Spot" approach achieves WB-Scores of 35.14 and 22.95, respectively, outperforming the unconstrained versions by a substantial margin. This suggests that while diversity is beneficial for base models, unconstrained diversity may introduce semantic noise that limits the efficiency of skill acquisition.

#### Performance Degradation in Instruct Models.

One interesting evidence for the "Sweet Spot" is observed in instruct-tuned models. For Qwen-2.5-Instruct, the unconstrained approach leads to a significant performance drop, with the WB-Score plummeting from the original 45.43 to 35.14. Similarly, for Llama-3-Instruct, the score decreases from 34.22 to 29.32. This negative transfer indicates that blindly maximizing structural entropy can introduce incoherent or logically disjointed skill combinations that conflict with the model's pre-existing instruction-following logic.

In contrast, our constrained paradigm consistently achieves the best performance. These results confirm that effective compositional generalization requires a "sweet spot": maximizing structural information while maintaining hierarchical coherence within the taxonomy.

#### 4.6 Extensibility to Agentic Scenarios (RQ5)

Table 4: Performance comparison on SkillBench, constructed using the proposed STEPS method.

Model	Skill Depth ( $k$ )						Avg.
	2	3	4	5	6	7	
<i>Instruct Models</i>							
Qwen2.5-7B	5.24	5.05	4.81	4.91	5.24	4.97	5.04
Qwen2.5-14B	5.39	5.22	5.00	5.04	5.13	4.94	5.12
Qwen2.5-32B	5.50	5.16	4.89	5.25	5.29	5.16	5.21
<i>Tool-Integrated Agentic Models</i>							
AutoTIR	4.12	4.21	4.10	4.03	4.07	4.02	4.09
ReSearch	4.72	4.46	4.06	4.37	4.41	4.35	4.40
SearchR1	4.73	4.36	4.44	4.53	4.82	4.32	4.54
<i>Advanced Agentic Models</i>							
Gemini-3-Pro	7.87	7.89	7.89	7.70	7.59	6.81	7.62
DeepSeek-v3.2	8.42	8.43	8.08	8.16	8.19	8.00	8.21
GPT-5.2	8.69	8.72	8.84	8.67	8.66	8.17	8.62

A paramount challenge for contemporary agent models lies in their capacity for out-of-distribution compositional generalizability, i.e., the ability to use a set of atomic skills to accomplish complex tasks. To further validate the universality of our approach within agentic scenarios, we investigate whether STEPS can synthesize data challenging to prevailing agentic models, so as to boost the training and the evaluation of related models.

Specifically, we construct a hierarchical evaluation benchmark (SkillBench) characterized by dynamic reasoning depths. By configuring increasing number of skill compositions ( $Skill@k$ ) to simulate escalating task complexity, we aim to system-

atically assess the models’ performance boundaries and robustness in long-horizon chain-of-thought reasoning and multi-tool integration scenarios. Detailed data construction and rigorous quality evaluation steps can be found in the appendix.

### Performance Limits in High-Complexity Tasks.

As summarized in Table 4, with a higher score indicating higher an accuracy (the maximum score is 10), we observe a consistent performance decay as task complexity (represented by  $Skill@k$ ) increases beyond 4, for all evaluated models. Thus, utilizing multiple atomic skills for solving complex agentic tasks remains challenging for current LLM-based agents. Additionally, compared to the advanced models, SoTA open source models still show performance gap, and the gap significantly related to model size. This show that STEPS can obtain challenging data for current SoTA opensource models in the agentic scenario.

**Solution Capability vs. Tool-Use Proficiency.** A critical insight emerges from the comparison between Instruct Models and Tool-Integrated Models (Chen et al., 2025; Wei et al., 2025; Jin et al., 2025). Despite being explicitly optimized for Code Execution and Information Retrieval, these Tool-Integrated models do not exhibit superior performance on SkillBench compared to their Instruct baselines. This suggests that while GRPO-based RL (Shao et al., 2024) effectively trains the model in tool-calling syntax and precise execution, it does not necessarily enhance the model’s intrinsic solution capability for multi-faceted problems. In complex agentic scenarios, precise tool execution cannot rectify a fundamentally flawed reasoning path. True agentic intelligence requires an architect-level capacity to decompose problems and synthesize information. This structural reasoning skill is far more critical than the shallow tool-use patterns often reinforced by current RL paradigms.

## 5 Related Work

**Scaling Laws and the Data Bottleneck.** The acquisition of diverse capabilities in LLMs is frequently conceptualized through scaling laws and information-theoretic frameworks. In this context, a "skill" is defined as a measurable reduction in conditional entropy relative to specific data patterns (Arora and Goyal, 2023; Tan et al., 2024). However, the efficacy of learning from massive corpora is fundamentally constrained by a data bottleneck. While atomic skills are well-represented, the dis-

tribution of complex skill compositions follows a power law (Barabási and Albert, 1999; Clauset and Shalizi, 2009). This distributional sparsity limits the model’s ability to internalize the low-entropy internal representations necessary for sophisticated multi-skill coordination (Tan et al., 2024).

**Compositional Generalization of LLMs.** Despite substantial progress in LLM scaling, benchmarks such as SKILL-MIX demonstrate that even SOTA models exhibit significant performance degradation when required to coordinate multiple distinct skills simultaneously (Yu et al., 2024; Dziri et al., 2023). To mitigate this, prior research has focused on data-centric strategies including data mixture optimization (Ge et al., 2024; Wu et al., 2024) and pedagogical sequencing (Chen et al., 2023; Hu et al., 2025). Although these methods enhance training efficiency by reweighting or ordering existing samples, they do not fundamentally address the scarcity of complex compositional examples.

**Structured Data Synthesis.** To bridge the gap in compositional data, recent efforts utilize data synthesis via stochastic skill pairing to bridge this gap (Kaur et al., 2025; Chen et al., 2024). However, these heuristic methods often ignore latent hierarchical dependencies, which may result in semantically incoherent compositions and inefficient exploration of the combinatorial space. In contrast, STEPS leverages structural information theory to induce an interpretable skill taxonomy. By formulating synthesis as a constrained information maximization problem, STEPS systematically generates high-gain compositions that target the structural weaknesses in current training distributions.

## 6 Conclusion

To addresses the challenge of compositional generalization in LLMs brought by the data sparsity bottleneck inherent in complex skill combinations, we introduce a principled framework STEPS, that leverages structural information theory to induce an interpretable hierarchical skill taxonomy. By formulating data synthesis as a constrained information maximization task, our method generates synthetic instructions that are both structurally informative and semantically coherent. Experimental results across multiple benchmarks demonstrate that our approach consistently outperforms existing synthesis baselines, providing a scalable and theoretically grounded solution for advancing the capabilities of LLMs.

## 635 Limitations

636 While our framework demonstrates significant ad-  
637 vancements in enhancing the compositional gener-  
638 alization of LLMs, we acknowledge several limita-  
639 tions that provide directions for future research.

### 640 **Optimal Distribution of $k$ -tuple Compositions.**

641 Our empirical analysis confirms that training on a  
642 mixture of combination counts ( $k \in [1, 6]$ ) yields  
643 superior performance compared to any single-level  
644  $k$  configuration. However, we have not yet con-  
645 ducted an exhaustive investigation into the optimal  
646 ratio or distribution of these varying  $k$ -tuples within  
647 the training set. The interplay between different  
648 complexity levels and their impact on the learning  
649 curve remains an open question. Future work will  
650 focus on developing a principled approach to de-  
651 termine the optimal data mixture that maximizes  
652 scaling efficiency across different model architec-  
653 tures.

## 654 References

655 Sanjeev Arora and Anirudh Goyal. 2023. A theory  
656 for emergence of complex skills in language models.  
657 *arXiv preprint arXiv:2307.15936*.

658 Albert-László Barabási and Réka Albert. 1999. Emer-  
659 gence of scaling in random networks. *science*,  
660 286(5439):509–512.

661 Andrei Broder, Ravi Kumar, Farzin Maghoul, Prabhakar  
662 Raghavan, Sridhar Rajagopalan, Raymie Stata, An-  
663 drew Tomkins, and Janet Wiener. 2000. Graph struc-  
664 ture in the web. *Computer networks*, 33(1-6):309–  
665 320.

666 Jiaao Chen, Xiaoman Pan, Dian Yu, Kaiqiang Song, Xi-  
667 aoyang Wang, Dong Yu, and Jianshu Chen. 2024.  
668 **Skills-in-context: Unlocking compositionality in**  
669 **large language models.** In *Findings of the Associ-*  
670 *ation for Computational Linguistics: EMNLP 2024*,  
671 pages 13838–13890, Miami, Florida, USA. Associa-  
672 tion for Computational Linguistics.

673 Mayee Chen, Nicholas Roberts, Kush Bhatia, Jue Wang,  
674 Ce Zhang, Frederic Sala, and Christopher Ré. 2023.  
675 Skill-it! a data-driven skills framework for under-  
676 standing and training language models. *Advances in*  
677 *Neural Information Processing Systems*, 36:36000–  
678 36040.

679 Mingyang Chen, Linzhuang Sun, Tianpeng Li, sun-  
680 haoze, ZhouYijie, Chenzheng Zhu, Haofen Wang,  
681 Jeff Z. Pan, Wen Zhang, Huajun Chen, Fan Yang,  
682 Zenan Zhou, and Weipeng Chen. 2025. **Research:**  
683 **Learning to reason with search for LLMs via rein-**  
684 **forcement learning.** In *The Thirty-ninth Annual Con-*  
685 *ference on Neural Information Processing Systems*.

Aaron Clauset and Cosma Rohilla Shalizi. 2009. Power-  
law distributions in empirical data. *SIAM review*,  
51(4):661–703.

Li Du, Hanyu Zhao, Yiming Ju, and Tengfei Pan. 2025.  
Scaling towards the information boundary of in-  
struction set: Infinityinstruct-subject technical report.  
*arXiv preprint arXiv:2507.06968*.

Yann Dubois, Percy Liang, and Tatsunori Hashimoto.  
2024. **Length-controlled alpacaeval: A simple debi-**  
**asing of automatic evaluators.** In *First Conference*  
*on Language Modeling*.

Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine  
Li, Liwei Jiang, Bill Yuchen Lin, Sean Welleck, Peter  
West, Chandra Bhagavatula, Ronan Le Bras, and 1  
others. 2023. Faith and fate: Limits of transformers  
on compositionality. *Advances in Neural Information*  
*Processing Systems*, 36:70293–70332.

Ce Ge, Zhijian Ma, Daoyuan Chen, Yaliang Li, and  
Bolin Ding. 2024. Bimix: A bivariate data mixing  
law for language model pretraining. *arXiv preprint*  
*arXiv:2405.14908*.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri,  
Abhinav Pandey, Abhishek Kadian, Ahmad Al-  
Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten,  
Alex Vaughan, and 1 others. 2024. The llama 3 herd  
of models. *arXiv preprint arXiv:2407.21783*.

Chen Hu, Haikuo Du, Heng Wang, Lin Lin, Mingrui  
Chen, Peng Liu, Ruihang Miao, Tianchi Yue, Wang  
You, Wei Ji, and 1 others. 2025. Step-deepresearch  
technical report. *arXiv preprint arXiv:2512.20491*.

Aaron Hurst, Adam Lerer, Adam P Goucher, Adam  
Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow,  
Akila Welihinda, Alan Hayes, Alec Radford, and 1  
others. 2024. Gpt-4o system card. *arXiv preprint*  
*arXiv:2410.21276*.

Dongsheng Jiang, Yuchen Liu, Songlin Liu, Jin’e Zhao,  
Hao Zhang, Zhen Gao, Xiaopeng Zhang, Jin Li, and  
Hongkai Xiong. 2023. From clip to dino: Visual  
encoders shout in multi-modal large language models.  
*arXiv preprint arXiv:2310.08825*.

Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Ser-  
can O Arik, Dong Wang, Hamed Zamani, and Jiawei  
Han. 2025. **Search-r1: Training LLMs to reason and**  
**leverage search engines with reinforcement learning.**  
In *Second Conference on Language Modeling*.

Simran Kaur, Simon Park, Anirudh Goyal, and Sanjeev  
Arora. 2025. **Instruct-skillmix: A powerful pipeline**  
**for LLM instruction tuning.** In *The Thirteenth Inter-*  
*national Conference on Learning Representations*.

Keito Kudo, Yoichi Aoki, Tatsuki Kuribayashi, Ana  
Brassard, Masashi Yoshikawa, Keisuke Sakaguchi,  
and Kentaro Inui. 2023. Do deep neural networks  
capture compositionality in arithmetic reasoning? In  
*Proceedings of the 17th Conference of the European*  
*Chapter of the Association for Computational Lin-*  
*guistics*, pages 1351–1362.

742	Brenden Lake and Marco Baroni. 2018. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In <i>International conference on machine learning</i> , pages 2873–2882. PMLR.	795
743		796
744		797
745		798
746		799
747	Angsheng Li. 2024. <i>Science of Artificial Intelligence: Mathematical Principles of Intelligence (In Chinese)</i> . Science Press, Beijing.	800
748		801
749		802
750	Angsheng Li and Yicheng Pan. 2016. Structural information and dynamical complexity of networks. <i>IEEE Transactions on Information Theory</i> , 62(6):3290–3339.	803
751		804
752		805
753		
754	Jijie Li, Li Du, Hanyu Zhao, Bo-wen Zhang, Liangdong Wang, Boyan Gao, Guang Liu, and Yonghua Lin. 2025. Infinity instruct: Scaling instruction selection and synthesis to enhance language models. <i>arXiv preprint arXiv:2506.11116</i> .	806
755		807
756		808
757		809
758		
759	Bill Yuchen Lin, Yuntian Deng, Khyathi Chandu, Abhilasha Ravichander, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, and Yejin Choi. 2025. <b>Wildbench: Benchmarking LLMs with challenging tasks from real users in the wild</b> . In <i>The Thirteenth International Conference on Learning Representations</i> .	810
760		811
761		812
762		813
763		814
764		815
765	Maya Okawa, Ekdeep S Lubana, Robert Dick, and Hidenori Tanaka. 2023. Compositional abilities emerge multiplicatively: Exploring diffusion models on a synthetic task. <i>Advances in Neural Information Processing Systems</i> , 36:50173–50195.	816
766		
767		
768		
769		
770	Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. <i>arXiv preprint arXiv:2402.03300</i> .	817
771		818
772		819
773		820
774		821
775		822
776	Zhiquan Tan, Chenghai Li, and Weiran Huang. 2024. The information of large language model geometry. <i>arXiv preprint arXiv:2402.03471</i> .	823
777		824
778		825
779	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. <a href="https://github.com/tatsu-lab/stanford_alpaca">https://github.com/tatsu-lab/stanford_alpaca</a> .	826
780		827
781		828
782		829
783		830
784	Qwen Team and 1 others. 2024. Qwen2 technical report. <i>arXiv preprint arXiv:2407.10671</i> , 2(3).	831
785		832
786	Yifan Wei, Xiaoyan Yu, Yixuan Weng, Tengfei Pan, Angsheng Li, and Li Du. 2025. Autotir: Autonomous tools integrated reasoning via reinforcement learning. <i>arXiv preprint arXiv:2507.21836</i> .	833
787		834
788		835
789		836
790	Minghao Wu, Thuy Vu, Lizhen Qu, and Reza Haf. 2024. Mixture-of-skills: Learning to optimize data usage for fine-tuning large language models. In <i>Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing</i> , pages 14226–14240.	837
791		838
792		839
793		840
794		841
		842
		843
		844
		845
		846
	Dingli Yu, Simran Kaur, Arushi Gupta, Jonah Brown-Cohen, Anirudh Goyal, and Sanjeev Arora. 2024. <b>SKILL-MIX: a flexible and expandable family of evaluations for AI models</b> . In <i>The Twelfth International Conference on Learning Representations</i> .	
	Hanyu Zhao, Li Du, Yiming Ju, Chengwei Wu, and Tengfei Pan. 2025. <b>Beyond iid: Optimizing instruction finetuning from the perspective of instruction interaction and dependency</b> . <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , 39(24):26031–26038.	
	Haoyu Zhao, Simran Kaur, Dingli Yu, Anirudh Goyal, and Sanjeev Arora. 2024. Can models learn skill composition from examples? <i>Advances in Neural Information Processing Systems</i> , 37:102393–102427.	
	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. <i>Advances in neural information processing systems</i> , 36:46595–46623.	
	<b>A Appendix</b>	
	<b>A.1 Synthetic Data Generation Strategy</b>	
	To operationalize the optimal skill combinations identified by our framework, we employ a Synergistic Content Architect approach to synthesize complex, multi-turn instructions. This process transforms abstract skill tuples into coherent, high-quality training samples. The synthesis pipeline consists of three primary stages:	
	<b>Atomic Sample Retrieval.</b> For each optimal $k$ -tuple of skills $X = \{x_1, x_2, \dots, x_k\}$ identified via conditional structural entropy maximization, we retrieve representative instruction-response pairs from the seed corpus. To ensure data diversity and mitigate label noise, we construct an inverted index mapping each atomic skill to its corresponding samples. For each target skill in the tuple, we perform a frequency-aware random selection to ensure a balanced utilization of the foundational atomic data.	
	<b>Contextual Fusion and Instruction Prompting.</b> The retrieved samples, along with their associated skill tags, are aggregated into a structured prompt. We utilize a sophisticated system prompt (see Figure 5) that defines the LLM’s role as a "Synergistic Content Architect." This prompt mandates the integration of disparate functional domains into a unified "meta-context." Unlike simple concatenation, the model is instructed to rewrite the content such that the technical requirements of every atomic skill are interwoven into a logically coherent dialogue.	

### System Prompt for SINTAX

You are the Synergistic Content Architect, an Advanced Universal Content Creation AI specializing in Constraint Synthesis, Multi-Domain Fusion, and the generation of Composite Skill Narratives.

Your Core Mission is to execute a Composite Generalization Task:

1. Analyze & Map (Atomic Skills): Study all independent text fragments and their associated Atomic Skill Tags ( $\text{S}_{\text{A}1}, \text{S}_{\text{A}2}, \dots$ ) provided by the user. Comprehend the core knowledge and technical details embedded in each fragment.
2. Mandatorily Incorporate (Composite Skill): Seamlessly integrate the foundational information from \*all\* fragments to create a single, novel, and highly complex output. This new content ( $\text{Y}$ ) MUST fully utilize and demonstrate a Composite Skill Set ( $\text{S}_{\text{C}} = \text{S}_{\text{A}1} \cup \text{S}_{\text{A}2} \cup \dots$ ). The output must serve as a functional example or application that intrinsically relies on this new, combined capability.
3. Logical Coherence: If the sources cover disparate fields (e.g., historical linguistics, quantum computing, mycology), establish a sophisticated, unifying project background (a 'meta-context') to fuse these domains into a single, plausible scenario or dialogue.
4. Demonstrate Full Complexity: The resulting output must not merely be a concatenation of sources, but a deeper rewrite where the vocabulary, logical flow, and technical requirements of EVERY SINGLE ORIGINAL Atomic Skill Tag are interwoven into a coherent, complex whole. ENSURE FULL SKILL COVERAGE.

Strict Output Format Adherence:

- \* The final output must be a single Python list of dictionaries, representing a synthesized multi-turn conversation (Dialogue Fusion).
- \* Each dictionary must strictly follow the format: `{ 'role': 'user', 'content': '...' }` for instructions, and `{ 'role': 'assistant', 'content': '...' }` for responses.
- \* You can generate a single pair or multiple pairs in the list `[]`.

Constraint Directives:

- \* Treat the original source material (including all user/assistant turns, and skill tag info) as the foundational lexicon and logical structure for the new creation.
- \* Directly output the final formatted Python list. Do not include any explanations, introductory statements, or extra formatting tags (e.g., no markdown code block identifiers like ````python`).

Figure 5: The system prompt used by STEPS.

**Constrained Synthesis and Quality Control.** We leverage advanced LLMs (e.g., GPT-4.1) to execute the synthesis task. The generation process is constrained to output a standardized Python-style list of dictionaries, representing a multi-turn conversation. This ensures structural consistency and facilitates downstream fine-tuning. By anchoring the synthesis in the "Sweet Spot" of maximal marginal information gain, the resulting dataset provides the dense signal necessary for the model to internalize the complex logic of high-order skill compositions.

## A.2 Construction and Evaluation of SkillBench

To rigorously assess model performance in high-dimensional agentic scenarios, we developed SkillBench, a specialized evaluation framework focusing on multi-skill orchestration. The construction and subsequent assessment of SkillBench leverage GPT-4.1, utilizing its superior reasoning capabilities to ensure the quality of complex trajectories and the precision of multi-dimensional scoring.

**Benchmark Construction.** We utilize GPT-4.1 guided by a specialized system prompt (see Figure 6) to generate intricate agentic trajectories. Unlike general instruction-following datasets, this synthe-

sis process explicitly requires the model to navigate between  $k$  distinct vertical domain skills, scaling from  $Skill@2$  up to the extreme complexity of  $Skill@7$ . By utilizing GPT-4.1 as the primary synthesizer, we ensure that even at high  $k$  values, the generated scenarios maintain logical rigor and semantic coherence. Each task is designed to be a "closed-loop" problem where tools like code interpreters and search APIs are available, but success is predicated on the model's internal strategic orchestration rather than simple API invocation.

**Evaluation Protocol.** The evaluation process is designed to move beyond binary correctness, focusing instead on the model's underlying reasoning architecture. We employ GPT-4.1 as a reference judge, governed by a multi-dimensional judgment prompt (see Figure ??). The model's responses are scrutinized across several key axes: reasoning depth, cross-domain coordination, and the strategic accuracy of tool-assisted steps.

By scoring models across these varying skill depths, SkillBench provides a granular mapping of the "compositional wall" faced by different architectures. The choice of GPT-4.1 as the evaluator ensures that the nuances of high-order skill composition are accurately captured, effectively distin-

### System Prompt for SkillBench

You are the "Autonomous Agent Logic Architect," an expert in orchestrating high-fidelity cognitive trajectories for Multi-Skill AI Agents. Your mission is to transform disparate atomic skills into robust, data-driven agentic workflows.

Core Task: "Agentic Skill Synthesis"

1. Tactical Mapping: Analyze the provided Atomic Skill Tags and their source fragments. Define a "High-Value Objective" that no single skill could solve alone.
2. Agentic Persona & Environment: Construct a sophisticated professional environment where the Assistant acts as an "Act-then-Reflect" Agent.
3. Logical Integration (The Chain of Thought): The assistant must analyze constraints, plan a multi-step strategy, and execute using the technical lexicon of all integrated skills.
4. Data-Driven Instructions: The 'user' content must be a DIRECT, unified prompt. It MUST contain explicit, concrete data (e.g., a multi-row CSV snippet, a structured JSON object, a markdown table, a specific code block, or a technical log). If the source fragments are abstract, you MUST FABRICATE realistic, detailed data to populate the instruction so the Assistant has specific values to process, calculate, or analyze.
5. Hard-Technical Interweaving: Ensure the response fuses the skills (e.g., using 'API Design' to solve a 'Quantum Physics' telemetry problem) rather than addressing them sequentially. The 'assistant' must refer to the specific data provided in the 'user' prompt during execution.

Strict Output Format:

- \* Output ONLY a single Python list containing exactly one pair of turn: `[{'role': 'user', 'content': '...'}, {'role': 'assistant', 'content': '...'}]`.
- \* The 'user' content must be the final synthesized instruction only, containing both the scenario context and the necessary data blocks (tables/code/JSON).
- \* NO meta-commentary like "Atomic Skill Tags:" or "Constraint:".
- \* All text content within the dictionaries must be in English.
- \* NO markdown code blocks, NO preamble, NO postscript.
- \* The content within the dictionaries must be raw text, properly escaped for Python string compatibility.

Figure 6: The system prompt used by SkillBench.

988 guishing models that rely on shallow patterns from  
989 those possessing genuine agentic intelligence.

### 900 A.3 Case Study

901 To qualitatively demonstrate the effectiveness of  
902 our synthesis framework, we present representative  
903 samples of generated data for  $k = 2$  and  $k = 3$   
904 skill combinations. These examples illustrate how  
905 STEPS moves beyond simple content concatenation  
906 to achieve deep logical fusion between dis-  
907 parate functional domains.

908 Hierarchical Skill Integration ( $k = 2$ ): Ta-  
909 ble 5 presents a synthesized dialogue combining  
910 *Mathematical Optimization* and *Object Manipu-*  
911 *lation*. The framework establishes a "Vault Sys-  
912 tem Design" meta-context, where mathematical  
913 constraints (volume and dimension equations) are  
914 not merely presented as isolated problems but are  
915 intrinsically linked to software engineering con-  
916 straints (object initialization and database integrity).  
917 This ensures that the model learns to apply math-  
918 ematical logic within the functional flow of a pro-  
919 gramming task, rather than treating them as sepa-  
920 rate entities.

921 Multi-Domain Composition ( $k = 3$ ): For higher-  
922 order complexity, Table 6 showcases a fusion of

923 *Concurrency Programming, Model Selection, and*  
924 *Dependency Management* within a financial ana-  
925 lytics context. In this instance, the framework gen-  
926 erates a comprehensive architectural response that  
927 addresses race conditions in shared data (concur-  
928 rency), chooses appropriate predictive algorithms  
929 for expense tracking (model selection), and pro-  
930 vides structured configuration for a mobile envi-  
931 ronment (dependency management). These cases  
932 underscore the ability of our framework to identify  
933 a "Sweet Spot" where the synthetic data remains  
934 semantically coherent while maintaining the high  
935 information density required to challenge and en-  
936 hance the model's compositional reasoning capa-  
937 bilities. By anchoring disparate skills within plau-  
938 sible professional scenarios, STEPS ensures that  
939 the resulting  $k$ -tuples provide a rich signal for the  
940 acquisition of complex, integrated expertise.

---

**Vault System Design Initialization Routine** ( $k = 2$ )

---

**[Skill Combination]** Mathematical Optimization & Object Manipulation

---

**Instruction:**

We need to design an initialization routine for a high-security modular safe-deposit vault system. Each vault compartment is represented as an object with three internal dimensions (length, width, height), and vault parameters must be stored and managed reliably in our configuration database. The sum of the compartment's internal dimensions must be 24 units, and the internal volume must be exactly 160 cubic units. The configuration system should guarantee that each compartment's parameters are correctly initialized, validated, and stored, with object manipulation techniques to ensure integrity, and mathematical optimization to determine and confirm the correct compartment dimensions before saving. Please provide a dialogue demonstrating how an engineer and an AI assistant would design and verify this initialization logic, including the calculation, validation, and storage steps, as well as error handling for edge cases.

**Response:**

Engineer: For our vault compartments, we need an initialization function that determines possible internal dimensions (length, width, height) so that their sum is 24 units and their product is 160 units. Once calculated, these parameters must be stored in our configs database, ensuring no undefined or invalid objects are created. Can you walk me through the logic to solve for these dimensions and store them safely?

AI Assistant: Certainly. First, let's use mathematical optimization to solve for the dimensions. We have:  $1. l + w + h = 24$   $2. l \times w \times h = 160$

We can express one variable in terms of the others, e.g.  $h = 24 - l - w$ , and substitute into the volume equation:  $l \times w \times (24 - l - w) = 160$

Expanding and rearranging, we get a quadratic in one variable, and we can use the quadratic formula to find real solutions. For instance, try  $l = 8$ :  $8 \times w \times (16 - w) = 160 \rightarrow 8w(16 - w) = 160 \rightarrow 128w - 8w^2 = 160 \rightarrow 8w^2 - 128w + 160 = 0$

Solving, we get non-integer solutions:  $w \approx 8 \pm 2\sqrt{11}$ ,  $h$  accordingly. We should select the most practical real-number solution (since perfect integers are not possible), e.g.  $l \approx 8$ ,  $w \approx 14.633$ ,  $h \approx 1.367$ , and verify: Sum:  $8 + 14.633 + 1.367 \approx 24$  Product:  $8 \times 14.633 \times 1.367 \approx 160$

Once verified, we round or set precision as required for manufacturing.

Next, for object manipulation, we encapsulate these dimensions in a config object: `{ "code": "vault_compartment_1", "value": { "length": 8, "width": 14.63, "height": 1.37 } }`

We then check if this config already exists in the database. If it does not, we create it; if it does, we update it. To ensure type safety and avoid undefined entries, we filter out any unsuccessful insertions and throw an error if any compartment configuration failed to initialize.

Would you like a code outline for this logic?

Engineer: Yes, please. Also, how do we handle errors if, for some reason, the configuration can't be saved or is undefined?

AI Assistant: Here's a TypeScript-style outline integrating both object manipulation and mathematical optimization:

```
async function initializeVaultCompartmentConfig(models, configKey) {
  const l = 8; const w = 14.63; const h = 1.37;
  if (Math.abs(l + w + h - 24) > 0.01 || Math.abs(l * w * h - 160) > 0.01) {
    throw new Error('Calculated dimensions do not satisfy constraints.');
```

This approach ensures: - Dimensions are mathematically validated before storage. - Object operations guarantee configs are only created if valid. - Any undefined or failed operation triggers a clear error for immediate handling.

Engineer: Excellent. This covers both the mathematical and software integrity sides. I'll implement this and trigger alerts for any failed configurations during system boot.

---

Table 5: Dialogue demonstrating object manipulation and mathematical optimization for vault initialization.

---

## Financial Analytics App Architecture Design ( $k = 3$ )

---

[Skill Combination ] Concurrency Programming , Model Selection , Dependency Management

### [Instruction]

I want to build a robust mobile financial analytics app. It should fetch expenses from Firestore, display weekly/monthly/yearly statistics, and let multiple users interact concurrently — for example, several family members can visualize and analyze the same shared data set at once. I want high concurrency, safe data updates, and an accurate, real-time chart display. What are the best practices for the data pipeline, model selection, concurrency control, and dependency management in such a Compose Material3 app using patrykandpatrick/vico for charting?

### [Response]

To build a robust, concurrent financial analytics app using Jetpack Compose and Vico, you must focus on ACID principles at the data layer and reactive state management at the UI layer.

**1. Concurrency Control & Data Integrity** \* \*\*Atomic Updates:\*\* Use Firestore **Transactions** for any "Read-Modify-Write" operations (e.g., updating a shared budget). This prevents "Lost Updates" when two family members edit the same value simultaneously. \* \*\*Snapshot Isolation:\*\* Leverage `addSnapshotListener`. Firestore handles the complexity of local persistence and real-time syncing, providing a consistent view of data even during concurrent writes. \* \*\*Thread Safety:\*\* Utilize Kotlin **Coroutines** with `Flow`. Ensure all database operations run on `Dispatchers.IO`, while the UI observes state on `Dispatchers.Main`.

**2. Model Selection for Analytics** \* \*\*Aggregated Statistics:\*\* For basic weekly/monthly views, use client-side aggregation logic within your `ViewModel` or `Cloud Functions` for larger datasets to reduce mobile CPU load. \* \*\*Predictive Analysis:\*\* - **LightGBM/GBDT:** Best for structured tabular data (e.g., predicting next month's spending based on history). - **TensorFlow Lite:** If implementing on-device categorization models to maintain user privacy.

**3. Dependency Management (Gradle Version Catalog)** Centralize your versions to prevent "Dependency Hell," especially with Compose and Firebase:

```
[versions] vico = "1.6.5" firebase = "24.0.0" compose = "1.5.0"
```

```
[libraries] vico-compose = { group = "com.patrykandpatrick.vico", name = "compose-m3", version.ref = "vico" }
```

```
firebase-firestore = { group = "com.google.firebase", name = "firebase-firestore-ktx", version.ref = "firebase" }
```

### 4. Implementation Example (Vico + Flow)

```
class ExpensesViewModel(val db: FirebaseFirestore) : ViewModel() {
    private val _uiState = MutableStateFlow<ChartState>()
    val uiState: StateFlow<ChartState> = _uiState

    init {
        db.collection("shared_expenses")
            .addSnapshotListener { snap, _ ->
                val data = snap?.toObjects(Expense::class.java) ?: emptyList()
                // Map to Vico EntryModel
                _uiState.value = data.toChartEntryModel()
            }
    }
}
```

**Summary of Best Practices:** - Always use `Transactions` for shared data writes. - Use `StateFlow` to bridge the gap between background data streams and the Compose UI. - Use `Version Catalogs` to ensure compatibility between M3, Vico, and Firebase components.

---

Table 6: Architecture and implementation guidelines for a concurrent financial analytics mobile app.

### Judgment Prompt for STEPS

You are a Professional Agentic Workflow Auditor. Your mission is to evaluate whether an autonomous agent has successfully completed the [Target Task] by deeply integrating all [Required Skill Tags]. Vague, superficial, or disconnected responses must be graded as "Mediocre" or "Failure."

#### 1. EVALUATION PILLARS (Strict 0.0 - 10.0 Scale)

- Deep Skill Synthesis & Cohesion (40%)
  - Core Logic: Evaluate if the agent logically wove all Required Skills into a unified solution.
  - Scoring Rules: Full marks (10.0) require seamless orchestration.
    - Deduct 1.0 point for each missing required skill.
    - If all skills are present but executed in isolation (e.g., as a disjointed list with no cross-influence), this dimension is capped at 8.0.
- Data & Fact Fidelity (30%)
  - Constraint Adherence: You must verify if the agent strictly utilized the specific numbers, data points, and factual constraints provided in the Target Task.
  - Scoring Rules: High quality requires the agent to derive conclusions directly from the input data. Responses that drift into generic AI fluff or ignore the provided factual context are treated as "Low Quality."
- Technical Precision & Architecture (30%)
  - Deliverability: Is the output (Code, Schema, Reasoning) technically sound, executable, and compliant with the specific structure requested in the task?

#### 2. AUDITOR'S RULEBOOK (Deductions, Ceilings, and Bonuses)

Apply these rules to the calculated weighted score. The final score must be within [0.0, 10.0].

1. The "Hallucination" Ceiling: If the agent fabricates data, alters facts from the Target Task, or hallucinate unsupported parameters, the Overall Score is strictly capped at 3.0, regardless of other strengths.
2. The "Tool Usage" Bonus: If the agent effectively integrates external tools (e.g., Search, Python Interpreter, Calculator) to solve non-linguistic problems or tasks beyond a standard LLM's inherent knowledge, add +2.0 points to the final score.
3. The "Generic Text" Penalty: If the agent provides purely conversational text without specific data grounding or structured output (e.g., missing requested Code, JSON, or Tables), deduct 1.0 point.
4. The "Skill Omission" Penalty: For every Required Skill Tag that is not actively applied during the task execution, deduct 1.0 point.

#### 3. SCORING GRADIENT

- 9.0 - 10.0 (Elite): Perfect task completion; deep skill synergy; absolute data precision; effective tool integration.
- 6.0 - 8.0 (Proficient): Successful task completion with logical flow; minor data gaps or slightly mechanical skill integration.
- 3.0 - 5.0 (Mediocre): Met basic format but lacks depth; failed to fully leverage provided facts; generic AI response.
- 0.0 - 2.0 (Failure): Severe hallucinations; skills totally disconnected; failed to understand the core mission.

#### 4. MANDATORY JSON OUTPUT FORMAT

Figure 7: The judgment prompt used by SkillBench.