Appendix: MIR-Bench: Can Your LLM Recognize Complicated Patterns via Many-Shot In-Context Reasoning?

The appendix is organized as follows. First, we state limitations, future work and broader impact in Sec. A. Then, we discuss more related fields to our work, and conduct an extended comparison to all related many-shot ICL or inductive reasoning works to further illustrate the position of our work in Sec. B. After this, we provide additional empirical study results in Sec. C (which are also main results of this paper but postponed to the appendix due to page limit), then extra results and auxiliary ablations in Sec. D. Then, we provide statistical features of MIR-Bench in Sec. E. Finally, in Sec. F, we introduce more details in our experiments, including the prompts we adopted in our curation of dataset and ablation experiments and the regex rule we used for extracting the answer.

We hereby summarize the important novel insights obtained from experiments in the appendix:

- 1. LLMs are quite robust against errorneous shots in many-shot inductive reasoning tasks. (Sec. C.1)
- 2. The first-coding, then-running paradigm are not always scalable to many-shot case. Many-shot in-context pattern recognition remains an open problem. (Sec. C.2)
- 3. RAG is not a effective solution for addressing saturation issue of the many-shot pattern recognition task. (Sec. C.3)
- 4. It still remains an open challenge for LLMs to learn "meta-skills" of inductive reasoning from out-of-domain demonstrations. (Sec. C.4)
- 5. The evaluation on our benchmark is robust across different random seeds; i.e., the standard deviation of the performance is low. (Sec. D.1)
- 6. The performance of LLMs against erroneous shot largely depends on the ratio of errorneous shots; under the same ratio, the total number of shots does not change much. (Sec. D.5)
- 7. While generally adding more shots increases LLM's inductive performance, the performance change varies with problem types. LLMs improve the most on string manipulation tasks where each character in the input serves as a "shot" inside each example, and will not improve if the functions are too straightforward or too difficult. (Sec. E.4)
- 8. LLMs tend to underestimate inductive reasoning difficulty during evaluation given a concise ground truth. A better choice is to do a multi-round evaluation where LLMs can better evaluate difficulty by self-reflection on its attempt for solving the problem. (Sec. F.5)

A Limitations, Future Works and Broader Impact

Limitations and future works. First, to curate MIR-Core with problems that requires many-shot ICL, we studied many related factors such as types of problem and difficulty of the problems; however, they are not decisive enough. A more explainable rule for determining whether a problem needs many-shot would be an interesting avenue for future many-shot ICL works. Second, our test of pattern recognition is limited to text; it would be interesting for future work to explore the intelligence of multimodal models [78, 53, 12]. Third, while we have largely reduced the ambiguity of the underlying functions by filtering out those with insufficiently diverse input-output patterns and unsolvable by LLMs, some underlying functions could still be non-unique given our input-output pairs. Finally, our empirical studies have disproved some possible fixes to the saturation issue of many-shot ICL such as RAG, but do not provide a panecea. According to prior work [45], we hypothesize that supervised finetuning and/or reinforcement learning with in-context learning data would be a promising avenue to explore.

Broader Impact. Our work proposes an interesting and useful challange for LLM's long-context reasoning ability, and summarized many useful insights for future LLM studies. As we mentioned in the paper, our work is a step towards generalist AI agents that perceive the world from interaction examples and make decisions from demonstrations. Thus, our work inherently shares the societal impact with all other LLM papers: while LLMs could significantly boost human's working efficiency and production power of the society, the misuse of LLMs could cause harm to humans such as displacement of human workers.

B Extended Related Work

Programming-by-Examples (PbE). PbE [58, 13] is a classic programming paradigm where programs are automatically written with user-provided input-output pairs as examples; it can be seen as an application of inductive reasoning in coding, and has wide application in sheet processing [21], data parsing [40], and systematic drawing [9]. It is traditionally addressed by symbolic-based approaches, such as heuristic search [21, 9], version space algebra [38] and learning weights for rule probabilities [54]; this symbolic formulation has largely limited the generalizability of PbE. Recently, as LLMs have proved themselves to be strong coders [22], several works tried to address general-purpose PbE with LLMs [71, 72, 70, 47]. None of them, however, considers many-shot scenario with more than 10 shots. Compared to existing works, Our benchmark is organized in a way that resembles many-shot PbE paradigm, but for most of the evaluations, the LLMs we tested are not required to write code; instead, they only need to directly predict output for new input. That being said, with minimal adaptation, our proposed benchmark can fill in the blank of many-shot PbE study (and we explored this in Sec. C.2).

Extended comparison with literature. Tab. 3 shows a detailed comparison of our work with existing works (including empirical study and benchmarks) in the field of many-shot and pattern recognition task. As shown in the table, our work is indeed unique among all the many-shot ICL and inductive reasoning works.

C More Empirical Studies on MIR-Bench

C.1 Robustness of LLM Inductive Intelligence

While many works [1] have studied LLM's many-shot ICL performance, the robustness of LLM's many-shot ICL ability [98], i.e. the accuracy given incorrect examples, is still largely underexplored. In this section, we explore the performance change with increasing number of shots with incorrect answers.

Evaluation Setup. We test all 15 models in Sec. 4.1 on MIR-Core with 3 different settings: 1) the "unaware" setting, where the models do not know there are incorrect answers in the provided examples; 2) the "aware-error" setting, where the models know that some (unknown number of) examples are incorrect; and 3) the "aware-ratio" setting, where the models know exactly how many shots are incorrect out of all given shots. The three settings are mostly the same, with slight difference in prompt; see Appendix F.2 for details. We test $\{64, 256, 1024\}$ shots \times error ratio of $\{1/64, 1/32, 1/16, 1/8, 1/4, 1/2, 3/4\}$ respectively. See Appendix F.7 for data generation details.




Figure 6: The accuracy of representative models with erroneous shots under different prompt settings with 64 shots (see Fig. 10 in Appendix D.5 for full results). The result shows LLMs are generally quite robust against erroneous shots.

Results. The results for 64-shot are illustrated in Fig. 6 (see Appendix D.5 for the rest). Surprisingly, we found that LLMs are generally quite robust against errorneous shots; their performance are only slightly harmed below 1/8, and can maintain decent performance even with 3/4 error rate. We find that generally, there is generally no significant performance difference in different awareness level of erroneous shots; some exceptions are Gemini-2.0 Flash and Claude-3.5-Haiku (see Fig. 10 in Appendix D.5), where the answering paradigm of the former remains the same, and the latter accepts the target input as part of the "incomplete" data and rejects answering questions less frequently. Overall, LLMs are quite robust against errorneous shots in many-shot inductive reasoning tasks.

Table 3: The topic, validity and reproducibility comparison between our benchmark and the most related prior many-shot / long-context benchmarks in the first part, and pattern recognition (inductive and transductive) reasoning benchmarks in the second part. To save space, we abbreviate "Many Shot" as MS, "Pattern Recognition" as PR (△ represents "classification only"), "Prob." as problems, and "I/O Div." as "Input/Output Diversity" (having at least 2 different input-output types, e.g., given an array and output an integer, or given a pair of strings and output a choice). "Gen." means "Generative", which means whether new test cases can be easily generated without much human effort. "LB" means whether a leaderboard is available, and "EE" means "Easy Evaluation", i.e., whether a pipeline for evaluating any given new model exists. "New Data" means whether the input-output data never appears in existing benchmarks; if so, the benchmark is not a compilation of existing dataset and is secured against data contamination. Note, the counting of #PR Problems and "Gen." take different target input-output for the same function into account, but do not take different sets of shots into account.

Evaluations	MS	PR	# PR Prob.	I/O Div.	Max # Shots	Gen.	LB	EE	New Data
Classifications [45]	√	Δ	~25K	×	2000	×	X	X	×
Many-Shot ICL [1]	\checkmark	\checkmark	450	\checkmark	2048	\checkmark	×	×	\triangle
Classifications [6]	\checkmark	\triangle	1250	×	2000	×	×	\triangle	×
Visual Classifications [29]	\checkmark	\triangle	4010	×	~ 2000	×	×	\checkmark	×
Instruction Following [98]	\checkmark	×	0	\checkmark	300	\checkmark	×	\checkmark	×
2D Classifications [99]	\checkmark	\triangle	100	×	256	\checkmark	×	\triangle	\checkmark
LLM Judge [75]	\checkmark	×	0	×	512	\checkmark	×	×	\checkmark
HELMET [92]	\checkmark	\triangle	500	\checkmark	$\sim 10 \mathrm{K}$	X	×	\checkmark	×
LongICLBench [46]	\checkmark	\triangle	3000	×	~ 2000	×	\checkmark	\checkmark	×
ManyICLBench [101]	\checkmark	\triangle	1000	\checkmark	7252	×	×	\checkmark	×
LMAct [67]	\checkmark	×	N/A *	×	256	\checkmark	×	\checkmark	\checkmark
LongBench [4]	\checkmark	\triangle	400	\checkmark	600	×	\checkmark	\checkmark	\checkmark
BABILong [36]	×	√	unknown	√	unknown	√	√	√	√
KORBench [51]	×	\checkmark	50	\checkmark	3	\times	\checkmark	\checkmark	\checkmark
SolverLearner [10]	×	\checkmark	1300	\checkmark	16	\checkmark	×	×	\checkmark
Case2Code [70]	×	\checkmark	1.3M	\checkmark	10	\times	×	×	\checkmark
DEER [91]	×	\checkmark	1250	×	3	\times	×	×	\checkmark
List functions [66]	×	\checkmark	4000	×	5	\checkmark	×	\checkmark	\checkmark
SyGus [84]	×	\checkmark	89	\checkmark	3	X	×	\checkmark	\checkmark
ARC [11]	×	\checkmark	800	×	3	\times	\checkmark	\checkmark	\checkmark
1D-ARC [88]	×	\checkmark	900	×	3	\times	×	\checkmark	\checkmark
Mini-ARC [32]	×	\checkmark	150	×	3	\times	×	\checkmark	\checkmark
WILT [5]	×	\checkmark	50	×	30	\times	\checkmark	\checkmark	\checkmark
LogicVista [87]	×	\checkmark	107	\checkmark	10	\times	×	\checkmark	\checkmark
CLUTRR [73]	×	\checkmark	70K	×	N/A	\checkmark	×	\checkmark	\checkmark
MIRAGE [43]	×	\checkmark	2000	\checkmark	8	\checkmark	×	×	\checkmark
ACRE [96]	×	\checkmark	30K	×	10	\checkmark	×	×	\checkmark
Mini-SCAN [65]	✓	✓	400	×	100	√	×	✓	✓
Ours	✓	√	6930	✓	2048	✓	\checkmark	✓	✓

LMAct has only a few tasks, but it is interactive and thus hard to count the number of problems.

C.2 SolverLearner: Is "First-Coding, Then-Running" the Cure?

For better inductive reasoning ability, Cheng et al. [10] proposed SolverLearner, an inductive reasoning framework where LLMs write code first for inductive reasoning problems and then generate answers with python interpreter. With such framework, the authors claim that LLMs demonstrate remarkable inductive reasoning capabilities under their framework. However, their study is limited to a few relatively weak LLMs, (GPT-3.5, GPT-3), limited amount of inductive reasoning problems and few-shot; to check whether such solution also works for the many-shot pattern recognition task, we re-implement their method on MIR-Core (see Appendix F.8 for prompts).

We test SolverLearner with {DeepSeek-R1, Claude-3.5 Sonnet, GPT-4o-0806, GPT-4o-mini-0718, Gemini-Flash 2.0, Gemini-Pro 1.5-002, Gemini-Flash 1.5-002, Mistral-Large-2} for $\{16, 64, 256, 1024\}$ shots respectively on MIR-Core. For each code snippet generated by LLMs, we set a limit of 1 second for execution, as we need to run 300 functions \times 10 test cases \times 4 different number of shots \times 8 models = 96000 code snippets.

Table 4: The accuracy at 16, 64, 256 and 1024 shots respectively for SolverLearner on MIR-Core, and its performance difference from results reported in Sec. 4.2. We plot improvements with > 0.02 blue and < -0.02 red. We find that the performance of SolverLearner varies from model to model, and does not necessarily perform better than normal inductive reasoning paradigm. Also, SolverLearner curves under many-shot are more "flat"; i.e., it does not seem to make good use of extra information from a large number of shots. Such result indicates that LLM many-shot inductive reasoning is still an open problem, and straightforward solutions such as SolverLearner are not suffice yet.

Model	Acc.@16	Acc.@64	Acc.@256	Acc.@1024
DeepSeek-R1	0.756(+0.022)	0.762(+0.007)	0.748(+0.018)	0.640(+0.003)
Claude-3.5 Sonnet	0.577(-0.009)	0.604(-0.015)	0.605(-0.017)	0.603(+0.04)
GPT4o-0806	0.530(+0.012)	0.534(-0.033)	0.538(-0.029)	0.556(+0.004)
GPT4o-mini-0718	0.350(-0.006)	0.375(+0.003)	0.386(+0.008)	0.370(+0.014)
Gemini-2.0 Flash	0.469(+0.066)	0.487(+0.003)	0.493(+0.006)	0.487(+0.026)
Gemini-1.5 Pro-002	0.469(-0.029)	0.495(-0.055)	0.483(-0.067)	0.491(-0.04)
Gemini-1.5 Flash-002	0.473(-0.009)	0.484(-0.03)	0.479(-0.038)	0.486(-0.03)
Mistral-Large-2	0.420(+0.057)	0.430(+0.028)	0.428(+0.078)	0.356(+0.102)

Tab. 4 demonstrates the accuracy of each model (with difference from the standard results reported in Sec. 4.2) on MIR-Core, and Tab. 5 demonstrates the error rate when writing code. We found that the effect of SolverLearner varies from model to model; i.e., SolverLearner does not necessarily improve performance on our benchmark. Also, SolverLearner does not seem to utilize many-shot well; the performance increase from 16-shot to 1024-shot is much smaller than that of standard performance reported in Sec. 4.2. We hypothesize such issue, similar to that in Sec. 4.4, stems from the complicated nature of the code. Moreover, models with relatively weaker long-context ability, such as Mistral-Large-2, has much higher error rate with many-shot as the context length goes beyond its "effective" [26] context length; DeepSeek-R1 as a long CoT model also struggles with high runtime error rate from many-shot inductive reasoning. Thus, many-shot pattern recognition is still an open problem and not yet solved by straightforward solutions such as SolverLearner. The insight can be summarized as follows:

C.3 Can RAG Help Many-Shot Pattern Recognition?

One possible way to bypass the problem of many-shot saturation is Retrieval Augmented Generation (RAG) [19]; i.e., instead of feeding every given shot into the LLM and disperses the model's attention, we only select a few shots that are the most related to the target input, thus forcing the model to concentrate on the few but useful shots in its context. Usually, there are two prevalent ways to select such shots: selected by LLM [39] or selected by embedding [41, 89]. The former is infeasible in our many-shot pattern recognition task, as each of our shot is already very precise and hard be further compressed by LLM as in other RAG works [39]; also, selection of shots with over 2000 candidates for each of the 3000 test cases in MIR-Core will be prohibitively expensive and/or error-prone for LLMs. Thus, we will focus on embedding-based RAG for this part.

Table 5: The Do-Not-Finish (i.e., no solution function generated) and Runtime Error (RE, including timeout and exception during running) rate at 16, 64, 256 and 1024 shots respectively for Solver-Learner on MIR-Core. Generally, with more shots, the error rate of LLMs will increase. Some models such as DeepSeek-R1 and Mistral-Large-2 has high error rate under long context scenario.

Model	DNF@16	RE@16	DNF@64	RE@64	DNF@256	RE@256	DNF@1024	RE@1024
DeepSeek-R1	0	0.0303	0.0003	0.0689	0.0003	0.1010	0.057	0.2061
Claude-3.5-Sonnet	0	0.0027	0	0.0063	0	0.0007	0	0.0037
GPT40-0806	0	0.009	0	0.0103	0	0.0157	0.0033	0.0137
GPT4o-mini-0718	0	0.0103	0	0.0147	0	0.0167	0.0033	0.017
Gemini-2.0 Flash	0	0.0023	0	0.007	0	0.007	0	0.0068
Gemini-1.5 Flash-002	0	0.0093	0	0.0117	0	0.0087	0	0.011
Gemini-1.5 Pro-002	0	0.0093	0	0.008	0	0.009	0	0.0107
Mistral-Large-2	0	0.008	0	0.0077	0.0047	0.012	0.1163	0.0473

Evaluation setup. We evaluate GPT-4o-0806, GPT-4o-mini-0718, Gemini-1.5 Pro, Gemini-1.5 Flash, Mistral-Large-2 on MIR-Core. To generate embedding vectors effectively, we choose a small but recognized sentence encoder, all-MiniLM-L6-v2 [81], to generate vectors for each shot. We test 128 to 2048 shots by selecting 64 shots with the closest (cosine similarity) vector representation to the target input, and compare it with 64 shots that are randomly sampled from the same 128 to 2048 shots.

Results. The result is illustrated in Tab. 6. The result shows no significant performance difference between RAG and randomly selecting shots, thus disproving the effectiveness of embedding-based RAG.

Table 6: The performance comparison between selecting 64 shots using RAG and random selection for many-shot pattern recognition. There is no significant performance difference between the two strategies for selecting shots.

# Shots	Selection	GPT-40	GPT-40-mini	Gemini-1.5-Pro	Gemini-1.5 Flash	Mistral
128	RAG	51.77	38.50	51.53	46.33	27.20
	random	50.40	37.10	51.20	46.43	26.03
256	RAG	49.83	37.33	50.40	44.73	27.00
	random	50.17	37.27	50.83	45.63	27.37
512	RAG	49.33	37.17	50.33	45.20	27.53
	random	50.90	36.93	51.27	45.23	27.10
1024	RAG	49.60	37.60	50.77	45.67	26.67
	random	51.03	37.63	50.77	45.67	27.23
2048	RAG	50.53	39.27	51.75	46.67	26.67
	random	51.63	37.23	51.37	45.83	27.03

C.4 Can LLMs Learn Inductive Skills from Out-of-Domain Meta-Shots?

Till now, we have mostly limited our many-shot experiments within *in-distribution* learning, which means all the given shots indicates the same function as the target input. A more desirable ability, however, is to learn from *out-of-domain* inductive reasoning traces: by given successful demonstrations on extracting rules from other examples, we hope LLMs to learn the "meta-skills" for pattern recognition (inductive reasoning in this case), e.g., to pick up a few examples, propose an assumption, and then verify with other examples (as explored by Wang et al. [84] with training).

Evaluation Setup. We test Gemini-1.5 Pro-002 and Gemini-1.5 Flash-002 on MIR-Core. For each problem, we select correct (test case, LLM answer)-pairs from GPT-4o-0806's output in Sec. 4.4 with 8-shot forced CoT, and filter out problems with invalid CoT by GPT-4o-0806. For each test case, we sample 4, 8, 16, 32 different (test case, LLM answer)-pairs, and put them before the original MIR-Core problem as meta-shots; each meta-shot is separated by a line of '==='. See Appendix F.9 for details on prompts. We test the result of {4, 16, 64, 256, 1024} in-distribution shots.

Results. The result is illustrated in Tab. 7. The result gives two insights: 1) the effect of meta-shots varies across models. For models like Gemini-1.5 Pro, meta-shots will slightly benefit CoT performance, but cannot fully bridge the gap between forced CoT (inductive) and no CoT (transductive); 2)




Figure 7: The performance of 5 cutting-edge LLM models on MIR-Extended with temperature 0.7 across 5 runs. The result clearly shows that the standard deviation of accuracy is always below 0.01, and thus the evaluation is highly stable.

the effect of meta-shot CoT slightly increases when given meta-shot CoT examples, but decreases when given more shots in the current problem, which is consistent with our finding in Sec. 4.4 that CoT struggles with more shots. In general, LLMs have still yet to learn the "meta-skills" from out-of-domain demonstrations, which poses an interesting research topic for future.

Table 7: The results of Gemini models with out-of-distribution meta-shots for inductive reasoning. Overall, more meta-shots leads to slightly better performance, but such effect weakens with more in-distribution shots and is not necessarily better than no meta-shots. Such result indicates that the models are yet to summarize and apply useful reasoning skills from in-context demonstrations.

	117					
	In-distribution #shots	0(-meta-shots)	4	8	16	32
Gemini-1.5 Pro-002	4	38.67	39.17	41.33	41.97	42.97
	16	45.73	47.63	47.90	48.63	48.53
	64	49.97	48.10	49.67	50.57	49.80
	256	49.43	49.17	49.53	49.43	50.27
	1024	49.60	47.77	48.50	49.60	49.60
Gemini-1.5 Flash-002	4	37.93	34.77	35.33	37.60	40.74
	16	45.50	42.20	42.57	41.47	42.87
	64	47.80	42.90	42.83	44.17	45.17
	256	48.70	42.63	44.33	44.23	45.80
	1024	46.10	40.23	41.40	42.07	43.90

D Complete Results and Auxiliary Ablations

D.1 Analysis on the Stability of Evaluation

As we use 0 temperature in the evaluations in our main papers, it is possible that the performance vary across LLM inferences with different random seeds, hence making our evaluation unreliable. To address such concern, we report the mean and standard deviation of the performance across 5 models {GPT-40-0806, GPT-40-mini-0718, Gemini-Pro 1.5-002, Gemini-Flash 1.5-002, Mistral-Large-2} over 5 different inferences with temperature 0.7 in Fig. 7 on MIR-Extended with 4,8,16,32,64,128,256,512 shots. The result clearly shows that the standard deviation for all models are very small, and thus our evaluation is reliable.

D.2 Complete Results on MIR-Extended and MIR-Core (Sec. 4.1, 4.2)

For better readability, we only put the performance of part of the models for MIR-Extended and MIR-Core in the main paper; Fig. 8 demonstrates the performance of all models.




Figure 8: The performance of all LLMs on MIR-Extended (panel (a)) and MIR-Core (panel (b)). As shown in Fig. 2 in the main paper, the benchmark poses challenge to almost all models tested. All models, including OpenAI-o1-1217, "saturate" at a particular number of shots, i.e., their performances stop to improve when more shots are given due to limited information integration capability.

D.3 Out-of-Context Rate for 1024 and 2048 shot in MIR-Extended and MIR-Core

Tab. 8 shows the rate of out-of-context error we received when invoking APIs for MIR-Extended and MIR-Core. Some models other than Gemini (which has > 1M context length) have an error rate of 0, which could due to its internal truncation.

Table 8: Out-of-context rate for model API calls on MIR-Core and MIR-Extended.

Model	MIR-Core 1024-shot (%)	MIR-Core 2048-shot (%)	MIR-Extended 1024-shot(%)	MIR-Extended 2048-shot (%)
Claude-3-Haiku	0	1.67	0	0.98
Claude-3-Sonnet	0	1.67	0	0.87
Claude-3.5-Haiku	0.96	2.4	0.65	1.37
Claude-3.5-Sonnet	0	1.77	0	0.97
Gemini-1.5 Flash-002	0	0	0	0
Gemini-1.5 Pro-002	0	0	0	0
Gemini-2.0 Flash	0	0	0	0
GLM-4-Plus	0	0	0	0
GPT-40-0806	0.33	5.67	0.14	4.47
GPT-4o-mini-0718	0.33	5.67	0.17	4.47
Mistral-Large-2	0.67	10.67	0.29	8.66
Moonshot-128K	0	0	0.19	0.14
OpenAI-o1-mini-0912	1	11	0.58	8.80
OpenAI-o1-preview-0912	1	11	0.58	8.80
Qwen2-72B-Instruct	0	0	0.37	8.29
DeepSeek-R1	0	0	N/A	N/A
DeepSeek-v3-0324	0	0	N/A	N/A
Gemini-2.0 Pro-0325	0	0	N/A	N/A
GPT4.5-Preview	0.33	6.67	N/A	N/A
OpenAI-01-1217	0	2.67	N/A	N/A
OpenAI-o3-mini-high	0.37	6.67	N/A	N/A

D.4 More Results on Many-Shot Inductive Reasoning vs. Deductive Reasoning

Tab. 9 lists the ratio of forced CoT (inductive reasoning) / no CoT (transductive reasoning) and their respective performance for more models on MIR-Core using original prompt. While the preference for inductive or transductive reasoning varies wildly across different models, the accuracy of transdutive reasoning is unanimously and significantly higher.

Fig. 9 illustrates the performance difference for more models between forced CoT and no CoT. The result shows that transductive reasoning results (with no CoT) are indeed better than inductive reasoning results (with forced CoT), and such gap increases with the number of shots.

Table 9: The results on MIR-Core of each model with and without CoT. Results are averaged over $\{4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048\}$ -shot. Error rate include cases where answer cannot be extracted and API error for exceeding context length. Claude-3.5-Haiku often refuses to answer the question due to "incomplete data".

Model	Answer w./ CoT (%)	Accuracy w./ CoT	Answer w./o. CoT (%)	Accuracy w./o. CoT	Error (%)
Claude-3-Haiku	51.03	0.278	47.66	0.441	1.31
Claude-3-Sonnet	20.46	0.233	76.98	0.475	2.56
Claude-3.5-Haiku	65.11	0.317	1.32	0.823	33.57
Claude-3.5-Sonnet	98.73	0.585	1.08	0.775	0.19
Gemini 1.5-Flash-002	20.11	0.306	79.75	0.539	0.14
Gemini 1.5-Pro-002	20.96	0.339	78.90	0.561	0.14
Gemini 2.0-Flash	24.43	0.363	74.59	0.498	0.98
GLM-4-Plus	19.70	0.248	79.33	0.388	0.97
GPT-40-0806	10.85	0.488	88.34	0.540	0.81
GPT-4o-mini-0718	37.21	0.279	61.77	0.414	1.02
Mistral-Large-2	75.66	0.306	21.64	0.403	2.70
Moonshot-128K	43.40	0.242	53.11	0.398	3.50
o1-mini-0912	2.54	0.334	93.63	0.696	2.41
o1-preview-0912	56.71	0.588	40.88	0.797	3.82
Qwen2-72B-Instruct	1.85	0.130	97.05	0.349	1.10
DeepSeek-R1	9.69	0.298	85.85	0.757	4.46
DeepSeek-v3-0324	10.40	0.329	84.54	0.570	5.06
Gemini-2.0 Pro-0325	78.95	0.691	20.63	0.872	0.41
GPT-Preview-4.5	35.94	0.543	63.29	0.669	0.77
OpenAI-o1-1217	3.55	0.469	96.12	0.811	0.33
OpenAI-o3-mini-high	85.48	0.697	13.42	0.806	1.10




Figure 9: Performance difference for 16 LLMs on MIR-Core between forced CoT and no CoT. For long-CoT models (o1 series and DeepSeek-R1), forced CoT works similar or slightly better than no CoT, but the gain diminishes with more shots. For the rest of the models, forced CoT almost always works worse (with the exception of GPT4o-mini-0718), and such gap increases with the number of shots. Mistral-Large-2's gap decreases dramatically at 2048-shot as such setting often exceeds its context length and the performance is low under both settings.

D.5 Complete Results on Robustness of LLM Inductive Intelligence

Fig. 10 shows the results of models on 64-shot, 256-shot and 1024-shot with different error rate for the shots, where the solid lines are 256-shot or 1024-shot accuracy respectively. We find that there are no significant performance difference across the same error rate with different number of shots (with the exception of o1-mini-0912 with 1024 shots), and the robustness persists across different number of shots.




Figure 10: The accuracy of models with erroneous shots under different prompt settings. The performance of the same error rate with different numbers of total shots are similar.

E Statistical Features of MIR-Bench

E.1 Data Source

Tab. 10 shows that out of 693 functions in MIR-Extended and 300 functions in MIR-Core, how many problems are extracted from each coding benchmark (HumanEval+, MBPP+ and APPS). The former two have Apache-2.0 licenses, and the latter has a MIT license.

Table 10: Number of functions extracted from each coding benchmark.

	HumanEval+	MBPP+	APPS	Total
MIR-Core	26	35	239	300
MIR-Extended	53	89	551	693

E.2 Input-Output Form

One advantage of MIR-Bench over existing works is that our curated problems have much more diverse input-output forms. To quantitatively illustrate this, we use GPT-40-0806 to label the input-output forms. The result is illustrated in Tab. 12, which clearly shows the diversity of problems in our benchmark.

Table 11: Number of input-output forms (input -> output) in MIR-Extended and MIR-Core, sorted by instances. All data types (e.g. str, int) are in python format.

Input-Output Form (MIR-Extended)	Count	Input-Output Form (MIR-Core)	Count
str -> str	158	str -> str	86
str -> int	40	int -> int	18
int -> int	37	list[float] -> list[float]	11
list[float] -> float	24	int, int -> int	9
int, int -> int	24	int -> str	8
str, str -> str	22	str, str -> str	8
list[int] -> int	19	str -> int	8
<pre>list[float] -> list[float]</pre>	18	list[int] -> int	5
list[float] -> int	17	str -> list[str]	5
int -> str	13	list[str] -> list[str]	4
str, int -> str	10	list[int] -> list[int]	4
float, float -> float	10	float -> int	3
list[int] -> list[int]	10	float -> float	3
str -> list[str]	8	list[float], list[float] -> list[float]	3
<pre>int, int -> list[int]</pre>	6	int, int, int -> int	3
float, float, float -> float	6	int, int, int -> str	3
int, int -> str	6	str, str -> int	3
str, str -> int	5	str -> list[int]	3
int, int, int -> int	5	list[float] -> float	2
float -> int	5	int, int -> str	2
<pre>list[float], float -> list[float]</pre>	5	list[str] -> int	3 3 2 2 2 2
float -> str	5	float -> str	2
others	240	others	103

Table 12: Number of problems for each difficulty level labeled by LLM, normalized from 0 to 1.

Difficulty Level	# MIR-Extended	# MIR-Core
0.05	1	0
0.1	81	21
0.2	103	53
0.3	118	50
0.35	1	1
0.4	51	27
0.45	2	0
0.5	75	33
0.55	10	5
0.6	7	2
0.65	8	4
0.7	163	77
0.75	25	8
0.8	32	13
0.85	16	6

E.3 Difficulty Level

E.3.1 Problem Counts for Difficulty Levels

Tab. 11 shows the number of problems for each LLM-labeled difficulty level (see Sec. F.5 for details), which shows that the problems in MIR-Bench has diverse difficulty levels.

E.3.2 Qualitative Analysis on the Effect of Difficulty Levels

In general, we find that the difficulty is positively correlated with the benefit of many-shot ICL (see Sec. E.5 for detailed numbers). Here, we append some examples of difficult and easy questions in our benchmark and analysis on whether they could be benefited from many-shot ICL:

```
def add(x: int, y: int): # Normalized difficulty level 0.1

Add two numbers x and y

» add(2, 3)

5

» add(5, 7)

12

"""

return x + y
```

Apparently, such a function is extremely easy to induce, and thus few-shot ICL is sufficient; LLMs will not benefit much from many-shot ICL.

Here, we show a slightly harder question:

```
def find_sum(arr): # normalized difficulty level 0.4
return sum(set(arr))
```

This question is slightly harder as it involves two operations: first remove all duplicate elements in a list, then get the sum. The model will need to look at multiple examples with and without duplicated elements to rule out other possible functions, e.g. sum of the array, fraction of the sum of array, etc. Thus, many-shot ICL will help more than the last function.

Finally, we give an example of a more difficult question:

```
def vowels_count(s): # normalized difficulty level 0.7

Write a function vowels_count which takes a string representing a word as input and returns the number of vowels in the string. Vowels in this case are 'a', 'e', 'i', 'o', 'u'. Here, 'y' is also a vowel, but only when it is at the end of the given word.

Example:

"" vowels_count("abcde")

2

"" vowels_count("ACEDY")

3

"""

if s == "": return 0

cnt = len(list(filter(lambda ch: ch in "aeiouAEIOU", s)))

if s[-1] in "yY": cnt += 1

return cnt
```

This is a typical example where many-shot ICL benefits: with only a few input-output examples, the model might not be able to rule out the possibility of string length calculation, upper or lower case count, or judging whether there is a vowel (if all examples have only zero or one vowel). The special case of y is even trickier; the model can get a decent accuracy if it ignores y, but to achieve perfect reasoning, the model needs to find sufficient examples where y is at the end of the word and where y is not at the end of the word to eventually determine this special rule.

E.4 Problem Types in Sec. 4.2

To study the effect on the topic of the problems for whether the problem benefits from many-shot, we first try to cluster the 693 problems in MIR-Extended using GPT-40-0806. More specifically, we first prompt the LLM to generate python-style tags for each problem with the following prompt:

```
# prompt for tags
You are an expert in coding. You will now be given a function that solves some problems and some example input-output pairs. You need to briefly summarize what the function is about in a tag in high-level, with no more than 5 words connected with '_'. DO NOT OUTPUT ANYTHING ELSE. Here are some examples:
<some examples>
[[Code]]
...
[[Input-Output Pairs]]
...
[[Answer]]
```

after acquiring tags for each problem, we prompt the LLM to merge all different tags down to 30 different tags with 6 major types: {List Analysis, List Manipulation, Mathematical Computations, String Analaysis, String Manipulations, Other}. Tab. 13 shows the number of problems, detailed tags and metric D (defined in Eq. (1)) for each problem type.

Based on the results, we find that generally adding more problem will have a positive effect on performance; however, for some types of problem such as geometric calculation and summation, the performance will decrease with more shots included. Upon checking those problems, we found them mostly fall into two categories: 1) the function is relatively straightforward, but the LLM gets confused with more shots due to over-complicated guesses; 2) the function is too hard to guess, and the LLM cannot make reasonable guesses when aggregating many pieces of information. The boxes below give examples for case 1) and 2) respectively:

```
# Case 1: Straightforward Problems

[[Code]]

def solution(num: int) → int:

steps = 0

while num > 0:

if num % 2 == 0: num /= 2

else: num -= 1

steps += 1

return steps

[[Input-Output Pairs]]

Input: {'num': 68037}

Output: 23
...
```

```
# Case 2: Difficult Problems

[[Code]]

def solution(boardSize, initPosition, initDirection, k):

yq, yr = divmod(initPosition[0] + k * initDirection[0], 2 * boardSize[0])

xq, xr = divmod(initPosition[1] + k * initDirection[1], 2 * boardSize[1])

# Calculate the final position considering reflections

return [min(yr, 2 * boardSize[0] - yr - 1), min(xr, 2 * boardSize[1] - xr - 1)]

[[Input-Output Pairs]]

Input: {'boardSize': [10, 11], 'initPosition': [5, 9], 'initDirection': [1, -1], 'k': 264}

Output: [9, 9]
```

On the other hand, problems such as removing duplicates and string manipulation generally benefit more from many-shot, probably because manipulation on each element / character can be considered a shot by itself, and thus the effective number of "shots" in such types of problems are higher. Note, "Others" problems have relatively high value of D because of one outlier; other than the outlier, it is almost equal to average level of D across MIR-Extended.

E.5 Other Factors Studied in Sec. 4.2

Fig. 11 illustrates the relation between our metric D (see Eq. (1) for definition) for distinctiveness between few-shot and many-shot performance.

Table 13: The tags for problem topics and related statistics; D is the average metric (see Eq. (1) for definition) of the corresponding type of problems in MIR-Extended. We marked entries with D>0.1 blue and D<-0.1 red. While increasing the number of shots generally brings better performance, We find that string manipulation benefits the most from many-shot.

Major Tag	Minor Tag	# (MIR-Extended)	# (MIR-Core)	D
List Analysis	Counting Elements	4	2	0.075
List Analysis	Counting Occurences	25	9	0.017
List Analysis	Maximum/Minimum Elements	30	5	-0.024
List Analysis	Statistics	2	0	0.041
List Analysis		61	16	0.001
List Manipulation	Filtering Elements	30	14	0.066
List Manipulation	Generating Sequences	15	6	0.077
List Manipulation	Mapping Elements	9	3	-0.022
List Manipulation	Removing Duplicates	6	4	0.136
List Manipulation	Sorting Elements	16	7	-0.041
List Manipulation		76	34	0.041
Mathematical Computations	Basic Arithmetic	35	12	0.085
Mathematical Computations	Boolean Determination	7	0	0.033
Mathematical Computations	Calculations Based on Formulas	98	32	0.051
Mathematical Computations	Condition Checking	43	20	0.114
Mathematical Computations	Geometric Calculation	4	2	-0.110
Mathematical Computations	Number Base Conversions	12	5	0.038
Mathematical Computations	Rounding	7	5	0.185
Mathematical Computations	Summation	5	0	-0.128
Mathematical Computations		211	76	0.066
String Analysis	Character Code Calculations	14	3	-0.004
String Analysis	Comparison	13	6	0.255
String Analysis	Counting Characters	29	10	0.050
String Analysis	Pattern Matching	27	7	0.084
String Analysis	-	83	26	0.084
String Manipulation	Case Transformation	19	7	0.089
String Manipulation	Encryption/Decryption	8	4	0.061
String Manipulation	Generating Substrings	8	4	-0.027
String Manipulation	Rearranging Characters	48	22	0.105
String Manipulation	Substitution	36	30	0.327
String Manipulation	Substring Replacement	33	22	0.205
String Manipulation	Swapping Parts	6	5	-0.053
String Manipulation	Transformation	34	17	0.083
String Manipulation		192	111	0.160
Others		76	37	0.123
Total		693	300	0.092

F More Experiment Details

F.1 Prompts for Main Results

We provide the prompt for the main results in Sec. 4.1 and Sec. 4.2 in the box below (the first commented line is not a part of the prompt):




Figure 11: Single-factor analysis between each factor in Sec. 4.2 and our distinctiveness metric D. Each point represents one of the 693 functions in our benchmark. The black line is the linear regression result of all functions in MIR-Extended; it is clearly shown that D is positively related to difficulty level and code length. The 64-shot accuracy is an average of {GPT-40-0806, GPT-40-mini-0718, Gemini-Pro 1.5-002, Gemini-Flash 1.5-002, Mistral-Large-2} over 10 test cases.

prompt for main results

You are given some function that takes something as input and output something. You need to predict the output for the target input of that function. Remember always end your answer with 'Output: your answer', with your answer in strict python format. Here are some examples:

Input: <example input 1> Output: <example output 1> Input: <example input 2> Output: <example output 2> ... (omitting more shots)

Input: <target input>

F.2 Prompts for Ablations

Effectiveness of CoT. The following boxes demonstrate the prompt for the result used in Sec. 4.4 with forced CoT and no CoT respectively (the first commented line is not a part of the prompt):

prompt for forced CoT

You are given some function that takes something as input and output something. You need to predict the output for the target input of that function. You need to first analyze it after 'Analysis:', then give your answer after 'Output:'. Remember always end your answer with 'Output: your answer', with your answer in strict python format. Here are some examples:

Input: <example input 1> Output: <example output 1> Input: <example input 2> Output: <example output 2> ... (omitting more shots) Input: <target input>

prompt for no CoT

You are given some function that takes something as input and output something. You need to predict the output for the target input of that function. Your answer should always be 'Output: your answer', with your answer in strict python format. DO NOT OUTPUT ANYTHING ELSE INCLUDING YOUR THOUGHTS. Here are some examples:"

Input: <example input 1> Output: <example output 1> Input: <example input 2> Output: <example output 2> ... (omitting more shots) Input: <target input>

Robustness of LLM intelligence. The following box demonstrates the prompt for the result used in Sec. C.1. For the "unaware" setting, we use the same prompt as that in the main results; for the "aware error" and "aware ratio" setting, we use the following prompts respectively:

prompt for "aware error"

You are given some function that takes something as input and output something. You need to predict the output for the target input of that function. Remember always end your answer with 'Output: your answer', with your answer in strict python format. Here are some examples. Note that not all shots are correct; there are a small portion of shots that are incorrect:

Input: <example input 1>
Output: <example output 1>
Input: <example input 2>
Output: <example output 2>
... (omitting more shots)

Again, note that not all shots are correct; there are a small portion of shots that are incorrect. Use your caution and think wisely.

Input: <target input>

prompt for "aware ratio"

You are given some function that takes something as input and output something. You need to predict the output for the target input of that function. Remember always end your answer with 'Output: your answer', with your answer in strict python format. Here are some examples. Note that not all shots are correct; there are <number of error shots> out of <total number> shots that are incorrect:

Input: <example input 1>
Output: <example output 1>
Input: <example input 2>
Output: <example output 2>
... (omitting more shots)

Again, note that not all shots are correct; <number of error shots> out of <total number> shots that are incorrect. Use your caution and think wisely.

Input: <target input>

F.3 Prompts for Reformatting APPS problems (Sec. 3.2)

The following box demonstrates the prompt for reformatting APPS problems in the "function collection" part of Sec. 3.2.

```
# prompt for reformatting
You are a coding expert. You will be given a problem and corresponding solution. Rewrite the solution
such that:
1. It becomes a single function named 'solution', which takes parameters as input instead of reading from
input() function if there is any;
2. There is no code out of the solution function and no solution class. All auxiliary functions should be
defined inside the solution function, and all imports should also be in the function.
3. The solution function should not have any print() function. Instead, it should return the result of the
function. If you need to output any rationale, leave them in comments. Your output must be directly
runnable without any change.
4. Just output the rewritten function; do not test it with extra statements.
Here is an example:
[[Problem]]
problem: Given a string, you need to reverse the order of characters in each word within a sentence while
still preserving whitespace and initial word order.
Example 1:
Input: "Let's take LeetCode contest"
Output: "s'teL ekat edoCteeL tsetnoc"
Note:
In the string, each word is separated by single space and there will not be any extra space in the string.
[[Solution]]
class Solution:
     def reverseWords(self, s):
          :type s: str
          :rtype: str
          rev_str = s[::-1]
          rev_arr = rev_str.split()
          final = rev\_arr[::-1]
          return ''.join(map(str, final))
[[Rewrite]]
def solution(s):
     :type s: str
     :rtype: str
     rev str = s[::-1]
     rev_arr = rev_str.split()
     final = rev arr[::-1]
     return ' '.join(map(str, final))
```

F.4 Prompt for The Generation of Data Generator

The following box demonstrates the prompt for generating data generator:

```
# prompt for generating data generator
```

You are a coding expert. You will be provided a coding question and corresponding solution. Please write two python function that randomly generates test case for the question. Specifically:

The first function's name is gen1, which generates random data (should be able to generate VERY DIVERSE, i.e., at least 1000 different data points).

The second function's name is gen2, which generates data that is slightly harder than those generated in gen1. (should be able to generate at least 100 different data points).

You shall not define any function outside gen1 or gen2. Should you use any helper function, make them inner functions inside gen1 or gen2. You gen1 and gen2 function should have and only have one int parameter, which is the number of cases.

Finally, the special cases should be designed as informative as possible that reveals the underlying function when looking at the input and corresponding output from the solution.

Here is an example. Note the output of gen1 and gen2 should be a list of dicts describing the parameters, and your special case input should be a dict describing the parameters. Please follow the format, and do not generate cases that are too long. Do not output any other text; put all your thoughts after "# rationale:" as shown in the example.

```
[[Problem]]
from typing import List
     def has_close_elements(numbers: List[float], threshold: float) -> bool:
     """ Check if in given list of numbers, are any two numbers closer to each other than given threshold.
     >>> has_close_elements([1.0, 2.0, 3.0], 0.5)
     False
     >>> has_close_elements([1.0, 2.8, 3.0, 4.0, 5.0, 2.0], 0.3)
     True
from typing import List
[[Solution]]
sorted_numbers = sorted(numbers)
for i in range(len(sorted_numbers) - 1):
if sorted\_numbers[i+1] - sorted\_numbers[i] < threshold:
     return True
return False
[[Gen1]]
# rationale: none
import random
def gen1(num cases: int):
     low, high = 5, 10 # generate lists between length 5 to 10
     data = \Pi
     for i in range(num_cases):
          N = random.randint(low, high)
          lst = [round(random.random() * 10, 1) for _ in range(N)]
          threshold = round(random.random(), 1) + 0.1
          data.append('numbers': lst, 'threshold': threshold)
     return data
[[Gen2]]
import random
def gen2(num_cases: int): # rationale: the data is slightly harder as the list is slightly longer
     low, high = 10, 20 # generate lists between length 10 to 20
     data = \Pi
     for i in range(num_cases):
          N = random.randint(low, high)
     lst = [round(random.random() * 10, 1) for _ in range(N)]
     threshold = round(random.random(), 1) + 0.1
     data.append('numbers': lst, 'threshold': threshold)
     return data
```

F.5 LLM Difficulty Labeling

In Sec. 4.2, we discussed several factors for problems which leads to LLM's few-shot and many-shot performance difference, and one of them is LLM-evaluated function difficulty level. While seemingly straightforward, a reliable evaluation of difficulty for LLM is highly non-trivial, as **LLM tend to underestimate inductive reasoning difficulty when given a simple ground truth function.** We provide an example below in the box with GPT-4o-latest as the prompted model:

Prompt

You are an expert in coding and inductive reasoning. You will now be given a function that solves some problems in python and some example input-output pairs.

Please rate the difficulty of guessing the function through input-output pairs. You should judge the difficulty NOT on whether the function itself is easy to write, but on whether it is easy to guess. You should first give your rationale after [[Analysis]], then end your answer with 'Overall Score: [[x]]', where x is your score from 1 to 10.

```
[[code]]

def solution(I):

return [I[i] for i in range(1, len(I)) if I[i] % i == 0]

[[data]]

Input: 'I': [40, -13]

Output: [-13]

Input: 'I': [-87, 29, -27, 90]

Output: [29, 90]

Input: 'I': [48, -4, 16, -63, -32, -5, -32, -45]

Output: [-4, 16, -63, -32, -5]

Input: 'I': [-75, 77, 44, 10, 17, -71, -80, -11, -68, -41]

Output: [77, 44]

Input: 'I': [87, 71]
```

Input: '1': [87, 71]
Output: [71]
LLM's response

[[Analysis]] To solve this puzzle, the key is to identify the rule used to filter the input list '1'. By observing the given input-output pairs and the function definition, the pattern can be deduced as follows:

Function Logic. The function constructs a new list from the original list '1', starting from index '1'. For each element at index 'i', it checks if the element is divisible by its **index** (using '1[i] % i == 0'). If true, it includes the element in the resulting list.

Difficulty Analysis Based on Input-Output Pairs. **Clues that might help guess the function:** 1. **The output never includes the first element of the input list.** This suggests that index-based filtering is involved.

- 2. **Only specific elements from the input list appear in the output.** A guesser would likely try to find what distinguishes the selected elements from the others. Given the presence of both positive and negative numbers, divisibility by index is a plausible criterion.
- 3. **The divisibility condition isn't obvious at a glance.** It isn't immediately clear that elements are being checked for divisibility by their **index**, as it's somewhat unconventional. Most people would initially look for a rule involving the values themselves, like even/odd checks, thresholds, or signs.

What Makes It Difficult to Guess?

- 1. The divisibility rule by index is **non-intuitive**. Divisibility by **element values** is a common guess, but divisibility by **index** is not.
- 2. The function starts from index '1' instead of '0'. This subtle detail could mislead a guesser into thinking it's a filtering rule based on the **position** in the list rather than the index value.
- 3. The rule applies to **all elements**, regardless of their sign or magnitude. This broad applicability makes it harder to pin down the exact condition.

What Makes It Easier to Guess?

- 1. **Simple filtering condition:** Despite the challenge, once a guesser considers index-based filtering, the rule becomes easy to validate.
- 2. **The input-output pairs are consistent.** There are no exceptions or edge cases that would throw off a guesser.

Overall Score: [[7]] (moderately difficult)

GPT-4o-latest gives an average difficulty score of 7.2 / 10 out of 5 evaluations (which indicates moderate-to-high difficulty), while GPT-4o-0806 gives much lower difficulty score (below 5). This problem, however, is in fact so difficult that we have to exclude it from our benchmark, i.e., have 0 accuracy for all five models {GPT-4o-0806, GPT-4o-mini-0718, Gemini-Pro 1.5-002, Gemini-Flash 1.5-002, Mistral-Large-2} across {4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048} shots in 10 test cases.

To address such issue, we adopt a multi-round conversation approach with self-reflection. We first let the LLM to try to solve the problem without code by itself, and then reveal the ground-truth answer and let LLM to score the difficulty based on self-reflection. The box below demonstrates an example:

```
You are an expert in inductive reasoning. You will now be given a function that solves some problems and
some example input-output pairs. You need to guess the answer for the one last input. After that, you will
be given the answer and the code, you need to do a self-reflection on your attempt, and rate a difficulty
level for future guessers. Here are some examples:
<few-shot examples of the whole conversation>
Here is the target problem:
[[Problem]]
Input: {'lst': ['V4', 'V10', 'V4', 'V0', 'V9', 'V4', 'V2', 'V14', 'V16', 'V9']}
Output: ['V0', 'V2', 'V4', 'V4', 'V4', 'V9', 'V9', 'V10', 'V14', 'V16']
Input: {'lst': ['V0', 'V13', 'V1', 'V4', 'V2', 'V10', 'V6', 'V5', 'VB', 'V1']}
Output: ['VB', 'V0', 'V1', 'V1', 'V2', 'V4', 'V5', 'V6', 'V10', 'V13']
Input: {'lst': ['V0+', 'V3', 'V14', 'V13']}
Output: ['V0+', 'V3', 'V13', 'V14']
Input: {'lst': ['V12', 'V6', 'V7', 'V2', 'VB', 'V2', 'V3', 'V0', 'V4', 'V6']}
[[Solution]]
# Round 1 response
The solution appears to sort the elements in each list lexicographically, maintaining any duplicates in their
respective positions within the sorted order.
Following this rule, for the input ''lst': ['V12', 'V6', 'V7', 'V2', 'V8', 'V2', 'V3', 'V0', 'V4', 'V6']', the
output should be:
**Output:** '['V0', 'V2', 'V2', 'V3', 'V4', 'V6', 'V6', 'V7', 'V12', 'VB']'
```

```
# Round 2 prompt
[[Response]] Now, I will tell you the answer is ['VB', 'V0', 'V2', 'V2', 'V3', 'V4', 'V6', 'V6', 'V7',
'V12']. Here is the code for the function:
def solution(lst):
def grade(v):
    if v == 'VB': return -2
        if v == 'V0': return -1
        if v == 'V0+': return 0
        return int(v[1:])
```

return sorted(lst, key=grade)
Based on the answer, your guess and the code, do a self-reflection on your attempt, and rate a difficulty level between 1 to 10 inclusive for future guessers with different shots sampled for the same function. You MUST always end your answer with "Difficulty: [[x]]" where x is 1 to 10.

Round 2 response

In my initial attempt, I focused on lexicographical sorting but overlooked the specific priority handling in the function for values like "VB" and "V0." The code clarifies that "VB" should always appear first, followed by "V0," and then "V0+," with all other values ordered by their numeric suffix. This customized grading system is subtle, especially without the code, as it combines both conditional priorities and numeric sorting.

Overall, understanding the need for this specific grading logic adds some complexity, but it's manageable once recognized.

Difficulty: [[4]]

With such paradigm, we found that the evaluation from LLMs are much more accurate.

F.6 Answer Extraction

We use the following way to extract final answer:

- 1. find the last occurence of "utput:" in LLM's response, which can be either "output" or "Output";
- 2. remove trailing "**" after "output:" for bold, " and "python in the answer for python format, and then remove all leading and trailing whitespaces and linebreaks;
- 3. conduct an exact match. If succeeds, the LLM is considered to be correct; otherwise, we try to convert the output into a python object by using ast.literal_eval(). If the convertion fails, the LLM is considered to be incorrect.

4. If the output can be converted to a single-key dictionary or single-element set, we will do an exact match between the value of the dictionary / element of the set to the ground truth answer with both converted to string (This is to account for responses similar to {"ans": 3} with ground truth being 3); otherwise, we do an exact match between the whole output and the ground truth answer converted to string.

F.7 Robustness Test: Erroneous Shots

We generate test cases with erroneous shot in the following way:

- 1. For n-shot with a given error rate ER, randomly sample $ER \times n$ indices to be the "erroneous shots" with incorrect answer. $ER \times n$ is guaranteed to be an integer.
- 2. for each "erroneous shot", we randomly sample one unused shot as we generate 20000 shots for each function, and substitute the original output with the selected shot's output. We will re-sample the unused shot if its answer is identical with the original shot.

F.8 SolverLearner

We use the following prompt for SolverLearner [10]:

Prompt for SolverLearner

You are given some function that takes something as input and output something. You need to write a python code of the function. You need to write your rationale after # (as if it is a python comment), and give your answer after 'Code:'. DO NOT OUTPUT ANYTHING ELSE. Your function name should be 'solution'. You are not allowed to write other custom functions unless it is inside 'solution'. Use imports before using package functions. You must strictly follow python format, especially input / output format (e.g., if it is a dictionary, your param should also be a dictionary). DO NOT ADD ANY STATEMENT FOR EVALUATION AFTER 'solution'. Here are the input-output pairs for the function, with input followed by output:"

Input: <input 1> Output: <output 1>

. . .

Input: <input n>
Output: <output n>

Here is your code. Again, do not output anything else; Your function name should be 'solution'. You are not allowed to write other custom functions unless it is inside 'solution'. Use imports before using package functions. You must strictly follow python format, especially input / output format (e.g., if it is a dictionary, your param should also be a dictionary). DO NOT ADD ANY STATEMENT FOR EVALUATION AFTER 'solution'.

Code:

F.9 Meta-Shots

We use the following prompt for the meta-shot experiments as illustrated in the box below:

You will be provided with a list of inductive reasoning problems, separated by '=== Please answer the final problem as instructed by that problem, and refer to previous problems as examples. Input: oblem 1 input 1> Output: output 1> Input: roblem 1 input 8> <LLM CoT demonstration> Input: roblem 2 input 1> Output: cproblem 2 input 2> Input: cproblem 2 input 8> Output: <problem 2 output 8> Input: <problem 2 target input> <LLM CoT demonstration> Input: <target problem input 1> Output: <target problem output 1> Input: <target input>

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 main claim accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA 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
 NA 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: See Sec. 5.

Guidelines:

- The answer NA 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: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA 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: We describe our experiment details and prompts used in the appendix in detail. Guidelines:

- The answer NA 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: See our links for code and dataset provided in the submission.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.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://nips.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, etc.) necessary to understand the results?

Answer: [Yes]

Justification: See Sec. 4 and appendix.

Guidelines:

- The answer NA 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: [Yes]

Justification: See Fig. 7.

- The answer NA 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: Our experiments are done by calling APIs.

Guidelines:

- The answer NA 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 reserach conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA 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: See Appendix A.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA 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., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA 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: See Appendix E.1.

- The answer NA 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: See our provided code and data link in the submission.

Guidelines:

- The answer NA 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: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA 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: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA 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 rigorousness, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: See our declaration on the openreview submission.

- The answer NA 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 (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.