AlphaResearch: Accelerating New Algorithm Discovery with Language Models

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models have made significant progress in complex but easy-toverify problems, yet they still struggle with discovering the unknown. In this paper, we present **AlphaResearch**, an autonomous research agent designed to discover new algorithms on open-ended problems by iteratively running the following steps: (1) propose new ideas (2) program to verify (3) optimize the research proposals. To synergize the feasibility and innovation of the discovery process, we construct a new reward environment by combining the execution-based verifiable reward and reward from simulated real-world peer review environment. We construct AlphaResearchComp, a new evaluation benchmark that includes an eight open-ended algorithmic problems competition, with each problem carefully curated and verified through executable pipelines, objective metrics, and reproducibility checks. AlphaResearch gets a 2/8 win rate in head-to-head comparison with human researchers. Notably, the algorithm discovered by AlphaResearch on the "packing circles" problem achieves the best-of-known performance, surpassing the results of human researchers and strong baselines from recent work (e.g., AlphaEvolve). Additionally, we conduct a comprehensive analysis of the benefits and remaining challenges of autonomous research agent, providing valuable insights for future research.

1 Introduction

Recent progress has shown that frontier LLMs like GPT-5 (OpenAI, 2025) and Gemini 2.5 (Comanici et al., 2025) could achieve expert-level performance in complex tasks such as mathematics (Trinh et al., 2024; Lin et al., 2025) and programming (Jimenez et al., 2024; Jain et al., 2025). While LLMs excel at processing and reasoning on problems that are within the boundary of existing human knowledge (Wang et al., 2024b; Phan et al., 2025), their capacity for independent discovery that pushes the boundaries of human knowledge still remains a question of paramount importance (Novikov et al., 2025). Can these models create advanced knowledge or algorithms that surpass human researchers?

Previous studies demonstrate that LLMs can generate novel ideas at a human expert level (Si et al., 2024; Wang et al., 2024a). However, the outcome evaluation of LLM-generated research ideas still struggles with biased verification methods (Ye et al., 2024) that constrain the exploration of out-of-boundary machine knowledge, such as LLM-as-a-judge (Lu et al., 2024), where misaligned LLMs are used to evaluate fresh ideas and inevitably favor solutions within existing knowledge boundaries. Furthermore, the ideation–execution gap (Si et al., 2025) between generating and executing new ideas also hinders models from producing advanced research outcomes. AlphaEvolve (Novikov et al., 2025) introduces an evolutionary coding agent that could tackle open scientific problems with program-based verification. However, the absence of real-world research environment rewards in coding-only agents (Tian et al., 2024) renders the discovery of out-of-boundary knowledge and algorithms challenging for current autonomous research agents.

In this paper, we introduce **AlphaResearch**, an autonomous research agent designed to discover new advanced algorithms by ensembling LLMs with a suite of research skills, including idea generation, code implementation, and iterative optimization. To combine the feasibility and innovation of the algorithm discovery process, we construct a dual environment, where novel algorithms are forged by the simulated real-world peer-reviewed environment and execution-based verification (Tian et al.,

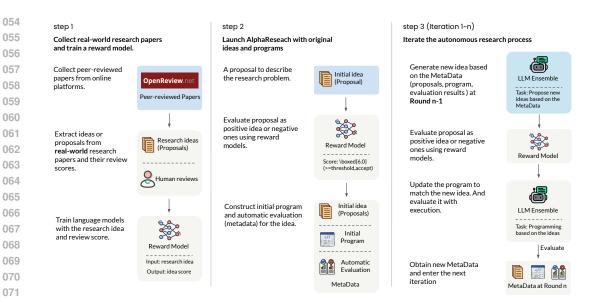


Figure 1: The launch of AlphaResearch contains two steps. (1) Train reward models with real-world peer-reviewed records. (2) Prepare initial research proposals, initial programs and evalution program. AlphaResearch will refine the research proposals and programs autonomously.

2024). Specifically, we (1) train a reward model **AlphaResearch-RM-7B** with real-world peer-reviewed records to simulate the real-world peer review environment, addressing the limitation of prior coding-only approaches that lack real-world research feedback, and use it to score the fresh ideas generated by LLMs. (2) construct an automatic program-based verifiable environment that executes these ideas with an interpreter. This dual environment facilitates a rigorous algorithm discovery process for autonomous research agents.

AlphaResearch discovers new algorithms by iteratively running the following steps: (i) proposing new research ideas, (ii) programming based on strong ideas, and (iii) optimizing the proposals for better algorithms. The commencement of each iteration in AlphaResearch mandates an LLM's creation of a new idea informed by previous research findings. After obtaining a fresh idea, AlphaResearch-RM will preserve the positive ideas as candidates and discard the negative ones. Those positive ideas are then implemented into executable programs by coding agents ensembled with LLMs. The program is subjected to rigorous code execution and automatic evaluation, a process previously demonstrated to be highly effective in mitigating incorrect suggestions often observed from LLMs (Tian et al., 2024). The synergy between an iterative real-world peer review environment and program-based verification empowers AlphaResearch to continuously explore novel research ideas and verify them via program execution. Once the generated optimal program surpasses current human-best achievements, these validated novel ideas could form feasible algorithms, thereby pushing the boundaries of human research forward.

To compare AlphaResearch with human researchers on novel algorithm discovery, we construct **AlphaResearchComp**, a simulated discovery **comp**etition between research agents and human researchers, by collecting 8 open-ended research problems and their best-of-human records (shown in Appendix D). Our results demonstrate that AlphaResearch surpasses human researchers on two problems but fails on the other six. The novel algorithms discovered by AlphaResearch not only surpass best-of-human performance but also significantly outperform the state-of-the-art results achieved by AlphaEvolve. Specifically, AlphaResearch optimizes the result of "Packing Circles (n=32)" problem to 2.939, where the goal is to pack n disjoint circles inside a unit square so as to maximize the sum of their radii, surpassing the results of best-of-human and previous SoTA results achieved by AlphaEvolve (as shown in Appendix B). These entirely novel ideas and algorithms constitute the most advanced solutions currently present in the human knowledge base, demonstrating the feasibility of employing LLMs to advance the frontiers of human knowledge. The six failure modes in AlphaResearchComp demonstrate the challenges for the autonomous algorithm discovery

Algorithm 1 AlphaResearch

Require: initial idea i_0 , initial program p_0 , initial result r_0 , model \mathcal{A} , evaluation program $\mathcal{E}(\cdot)$, maximum iteration rounds n,

```
111
                                                                                                                                      ▶ Initialization
             1: \tau_0 \leftarrow (i_0, p_0, r_0), r_{best} = 0
112
             2: for k = 1 to n do do

    States Sampling

             3:
                      (i_t, p_t, r_t) \sim \mathbb{P}(\cdot | \tau_{k-1})
113
             4:
                                                                                                                 ⊳ New Idea Generation (Eq. 1)
                      i_k \sim \mathbb{P}_{\mathcal{A}}(\cdot|i_t \oplus p_t \oplus r_t)
114
                      if \mathcal{RM}(i_k) < threshold then
             5:
115
                           continue
                                                                                                                  116
             7:
                      end if
117
             8:
                      p_k \sim \mathbb{P}_{\mathcal{A}}(\cdot|p_t \oplus i_k)
                                                                                                                  ⊳ Program Generation (Eq. 2)
                      r_k \leftarrow \mathcal{E}(p_k)
                                                                                                                     ▶ Program-based Execution
118
            10:
                      if r_k > r_{best} then
119
            11:
                           (i_{best}, p_{best}, r_{best}) = (i_k, p_k, r_k)
120
            12:
                      end if
121
            13:
                                                                                                                     \tau_k \leftarrow \tau_{k-1} \oplus i_k \oplus p_k \oplus r_k
122
            14: end for
            15: return (i_{best}, p_{best}, r_{best})
123
```

124 125 126

108

109

110

with research agents. We analyze the benefits and remaining challenges of autonomous research agents for knowledge discovery, providing valuable insights for future work.

127 128 129

130 131

132

133

134

135

136

137

138

139

140

ALPHARESEARCH

2.1 Overview

AlphaResearch discovers out-of-boundary novel algorithms by continuously optimizing the research outcome from the dual reward that synergizes rigorous program verification and a simulated real-world peer review environment. As shown in Figure 1, given initial idea i_0 and program p_0 , AlphaResearch runs the program p_0 with execution, producing r_0 , which represents the initial overall rating. The triplet (i_0, p_0, r_0) will be fed to AlphaResearch for subsequent processing, including newer idea generation, code implementation, and program-based execution. When reaching a point where execution output r_n surpasses the previous rating, AlphaResearch will save the triplet $(i_{best}, p_{best}, r_{best})$ as the best record. We repeat the process until r_{best} surpasses the best-of-human score, or the maximum round is reached. The resulting trajectory is denoted as $\tau = i_0 p_0 r_0 \dots i_{n-1} p_{n-1} r_{n-1} i_n p_n r_n$, where n is the total rounds.

141 142 143

144 145

146

147

148

149 150

2.2 ACTIONS

New Idea Generation. For each step k, AlphaResearch start with generating a new idea i_k based on a sampled previous step (i_t, p_t, r_t) from previous trajectory $\tau_{k-1} = i_0 p_0 r_0 \dots i_{k-1} p_{k-1} r_{k-1}$. This process can be denoted as:

$$i_k \sim \mathbb{P}_{\mathcal{A}}(\cdot|i_t \oplus p_t \oplus r_t)$$
 (1)

where \oplus means concatenation, t is the sampled step from trajectory τ_{i-1} . We use a reward model to filter out high-quality ideas overall. If $\mathcal{RM}(i_n)$ outputs a negative score, we cease the subsequent actions in this round.

151 152 153

154

Program-based Verification. After obtain the fresh idea, AlphaResearch generates new program p_k based on the previous implementation p_t and new idea i_k next:

155 $p_k \sim \mathbb{P}_{\mathcal{A}}(\cdot|p_t \oplus i_k)$ 156

(2)

157 158

and yield the evaluation result r_k by verifying p_k with code executor $r_k \leftarrow \mathcal{E}(p_k)$. Then, we update the trajectory τ_k with the newly generated idea i_k , program p_k and result r_k :

159 160 161

$$\tau_k \leftarrow \tau_{k-1} \oplus i_k \oplus p_k \oplus r_k \tag{3}$$

We repeat the above interaction process until k reaches the maximum rounds n and get the best result $(i_{best}, p_{best}, r_{best})$ as final output.

Table 1: Dataset for reward model training. We use the end of author-reviewer rebuttal period as the latest knowledge date.

Split	Train	Test
Records	ICLR	ICLR
Range	2017~2024	2025
Num	24,445	100
Start Date	2016-11	2024-10
End Date	2023-12	2024-12

Table 2: Evaluation results of RM. We use the more recent date between the model release date and the dataset cutoff as the latest date.

Reward Model	Cutoff	Acc
Random (theoretical)	-	50.0%
Human Annotator	-	65.0%
Deepseek-V3-0324	2025-03	39.0%
Qwen2.5-7B-Instruct	2024-09	37.0%
AlphaResearch-RM-7B	2024-09	72.0%

2.3 ENVIRONMENT

2.3.1 REWARD FROM REAL-WORLD RESEARCH RECORDS

Existing autonomous idea generation process suffers from a trade-off where highly novel research ideas may lack feasibility (Guo et al., 2025; Si et al., 2025). To address this gap and ensure the feasibility of idea candidates, we train a reward model with ideas from real-world peer-review information to simulate the real-world peer-review environment.

Dataset for reward model. To train our reward model (RM) to identify good ideas, we collect all ICLR peer review records from 2017 to 2024 as our training set. We sample a subset of ICLR 2025 records as a test set, where the dates of train and test are disjoint, which prevents knowledge contamination between the train and test split. We also select Qwen2.5-7B-Instruct as our base model, whose release date 2024-09 is earlier than the ICLR 2025 author-reviewer rebuttal period 2024-10. For each record in the training dataset, we extract the abstract part as RM input and wrap the average peer-review overall ratings with \boxed{} at RM output. We fine-tune Qwen2.5-7B-Instruct with the RM pairs, yielding the AlphaResearch-RM-7B model.

Can LLMs identify good ideas? To simplify the RM evaluation, we binarize the RM output score according to the ICLR Reviewer Guide, where overall rating > 5.5 records are regarded as a positive score and ≤ 5.5 records are negative. We compute the binary classification accuracy and evaluate three models (Deepseek-V3-0324, Qwen2.5-Coder-Instruct, and AlphaResearch-RM-7B) on the AlphaResearch-RM test set. Table 2 presents the evaluation results that eliminate the knowledge contamination, highlighting the following observations: (1) Both Deepseek-V3-0324 and Qwen2.5-7B-Instruct have lower than 50% accuracy when identifying the good ideas from ICLR 2025 records. (2) After fine-tuned with ideas from previous ICLR peer-review information, AlphaResearch-RM-7B demonstrates 72% binary classification accuracy on unseen ICLR 2025 ideas, significantly outperforming baseline models and human annotators. Based on these observations, we use the fine-tuned AlphaResearch-RM-7B as the final RM to simulate a real-world peer-review environment and filter out good ideas generated by AlphaResearch.

2.3.2 REWARD FROM PROGRAM-BASED EXECUTION

Inspired by AlphaEvolve (Novikov et al., 2025), we construct an automatic evaluation process with a code executor where each new program p_k generated by AlphaResearch will be captured and evaluated. The evaluation program $\mathcal{E}(\cdot)$ includes two modules: (i) **Verification** module that validates whether p_k conforms to the problem constraints. (ii) **Measurement** module that output the score r_k of program performance. The program output r_k will be injected into the idea generation prompt (if sampled), thereby participating in the optimization process for fresh ideas. These programs and results are stored in a candidate pool, where the primary goal is to optimally resurface previously explored ideas in future generations. The verifiable reward by code executor significantly simplifies the action spaces of AlphaResearch, thereby enhancing the efficiency of the discovery process.

3 ALPHARESEARCHCOMP

Table 3: Problem overview in AlphaResearchComp. More information are shown at Appendix D.

Problem	Human Best	Human Researcher
packing circles (n=26)	2.634	David Cantrell (2011)
packing circles (n=32)	2.936	Eckard Specht (2012)
minimizing max-min distance raio (d=2, n=16)	12.89	David Cantrell (2009)
third autocorrelation inequality	1.4581	Carlos Vinuesa (2009)
spherical code (n=30)	0.67365	Hardin & Sloane (1996, 2002)
autoconvolution peak minimization (upper bound)	0.755	Matolcsi–Vinuesa (2010)
littlewood polynomials (n=512)	32	Rudin-Shapiro (1959/1952)
MSTD (n=30)	1.04	Hegarty (2006/2007)

Problems collection. To evaluate AlphaResearch, we curate a set of 8 frontier program-based research tasks spanning geometry, number theory, harmonic analysis, and combinatorial optimization. These problems were selected based on the following principles:

- Well-defined objectives. Each task has a precise mathematical formulation with an objective function that admits rigorous automatic evaluation.
- **Known human-best baselines.** For every problem, we provide the best-known human result from the literature. These represent conjectured best-known values rather than proven optima, ensuring ample room for further improvement.

The curated problems are either inherited from prior work (e.g., AlphaEvolve) or collected from online repositories and domain experts. Each problem's baseline is supported by verifiable resources in the corresponding field. This design enables AlphaResearch to demonstrate both the *reproducibility* of established mathematical results and the *potential for discovery* beyond current human-best achievements. Detailed definitions, baseline values, and references for each problem are provided in the Appendix D.

Initialization strategy. After obtaining the research problems of AlphaResearchComp, we construct diverse initial states for each problem with the following strategies: (1) For the "Packing Circles" (n=26) and "Packing Circles" (n=32) problems, we initialize them with null programs (r_0 = 0) to simulate researches starting from scratch. (2) For the "Littlewood Polynomials" and "MSTD (n=30)" problems, we directly adopt the best-known solutions (r_0 = r_{human}) from human researchers to emulate improvements upon established methods. (3) For the remaining problems, we employ a moderate initialization strategy ($0 < r_0 < r_{human}$) to ensure sufficient room for the research agent to explore. This initialization strategy simulates a variety of real-world scenarios for the research agent, thereby facilitating a thorough evaluation process.

Metrics. For benchmarks like code generation with good verification techniques (e.g., unit tests), pass@k (Chen et al., 2021) is a metric denoting that at least one out of k i.i.d. task trials is successful, which captures the ability of LLMs to solve easy-to-verified problems. For open-ended real-world algorithm discovery tasks, we propose a new metric - excel@best (excel at best), defined as the percentage excess on baseline (best of human level) results:

$$excel@best = \mathbb{E}_{Problems} \left[\frac{|r_{best} - r_{human}| \cdot \mathbb{I}_d}{r_{human}} \right]$$
(4)

where r_{human} indicates the results of human's best level. \mathbb{I}_d indicates the optimization direction where $\mathbb{I}_d = 1$ represents that higher score is better and $\mathbb{I}_d = -1$ represents lower.

4 EXPERIMENTS

4.1 SETUP

We select o4-mini, a strong but cost-efficient LLM as our research agent and run AlphaResearch on each problem to get the best algorithm. We perform supervised finetuning

Table 4: Results on AlphaResearchComp. ↑ inidicates that higher score is better and ↓ for lower.

Problem	Human	AlphaR init	esearch best	Excel@best
packing circles (n=26) ↑	2.634	0	2.636	0.32%
packing circles (n=32) ↑	2.936	0	2.939	0.10%
minimizing max-min distance ratio ↓	12.89	15.55	12.92	-0.23%
third autocorrelation inequality ↓	1.458	35.746	1.546	-6.03%
spherical code (d=3, n=30) ↑	0.6736	0.5130	0.6735	-0.01%
autoconvolution peak minimization ↓	0.755	1.512	0.756	-0.13%
littlewood polynomials (n=512) ↓	32	32	32	0
MSTD (n=30) ↑	1.04	1.04	1.04	0

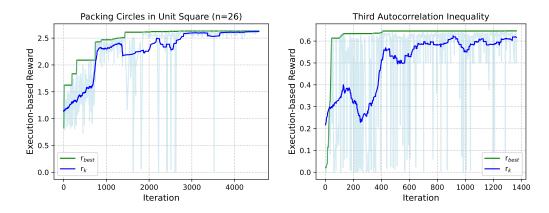


Figure 2: Execution-based reward of AlphaResearch on packing circles (n=26) problem (left) and third autocorrelation inequality problem (right).

on <code>Qwen-2.5-7B-Instruct</code> (Yang et al., 2025) with the collected ICLR records, yielding AlphaResearch-RM-7B. We do not compute loss on paper information, only on the average rating scores within <code>\boxed{}</code>. For fine-tuning hyperparameters, we train our model with a learning rate of 1e-5 warmed up linearly for 100 steps. We train all the models in bfloat16 precision with Pytorch Fully Shard Data Parallel (FSDP) and set a global batch size to 128 for 2 epochs. All other settings not mentioned in this paper follow the default values of Huggingface Trainer ¹.

4.2 RESULTS

LLMs can discover new algorithms themselves. Table 4 presents the results of AlphaResearch-Comp on 8 algorithms discovery problems. AlphaResearch achieved a 2/8 win rate (excel@best > 0) against human researchers, with one notable success: the algorithm discovered by AlphaResearch for "*Packing Circles*" problem reaches the best-of-known performance (2.636 for n=26, 2.939 for n=32), outperforming human researchers (2.634 for n=26, 2.936 for n=32) and AlphaEvolve (2.635 for n=26, 2.937 for n=32), where case (n = 32) is shown in Figure 7.

LLMs can refine their research ideas autonomously. AlphaResearch discovers advanced algorithms by iteratively proposing and verifying new research ideas. As shown in Table 2, 6/8 problems demonstrate consistent improvement throughout the discovery process. Figure 2 presents two examples of the reward trend in AlphaResearch, where the execution-based reward initially grows rapidly, then slowly plateaus for optimal performance seeking. This improvement trend emphasizes the autonomous discovery ability of research agents.

The discovery of superhuman algorithms remains challenging for LLMs. As illustrated in Table 2, despite exhibiting continuous reward growth, AlphaResearch's performance still under-

Ihttps://huggingface.co/docs/transformers/main_classes/trainer

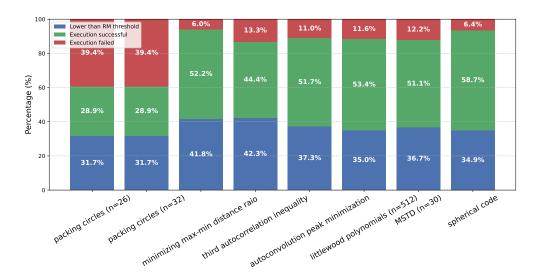


Figure 4: Reward overview during the discovery process. Each action in AlphaResearch will obtain 3 kinds of reward: (1) idea scrapping due to a lower RM score than the threshold (2) idea execution successes (3) idea execution fails.

performs human researchers in 6 out of 8 problems. We initialize AlphaResearch with the best-known solution from human researchers on "Littlewood polynomials" and "MSTD(n=30)" problems, where AlphaResearch didn't show an increase in execution-based rewards. This indicates that current LLMs still struggle to consistently find better algorithms than human researchers.

4.3 ABLATIONS AND ANALYSIS

Execution-only agent against AlphaResearch.

To compare AlphaResearch with execution-only agents, we utilize AlphaResearch-RM-7B to evaluate the novelty of ideas generated by the execution-only agent and ideas produced by AlphaResearch. As illustrated in Figure 3, the ideas generated by AlphaResearch generally achieve higher scores than execution-only research agents. This illustrates that AlphaResearch tends to generate better ideas to get higher external rewards, thus facilitating a more effective research optimization process.

Analysis of the discovery process. We analyze the reward distribution in AlphaResearch discovery process. As shown in Figure 4, approximately 30%~40% of newly proposed ideas fall below the RM threshold and are thus discarded. The remaining ideas are executed, with the success rate of execution largely depending on the inherent characteristics of the problems. For example, the execu-

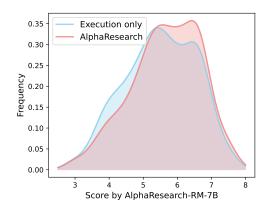


Figure 3: The idea comparison between the execution-only research agent and AlphaResearch, where AlphaResearch-RM-7B is used.

tion success rate on "Packing Circles" problem is 28.9%, whereas it reaches 51.7% on the "Third Autocorrelation Inequality" problem. Figure 2 illustrates the execution-based rewards for these two examples in AlphaResearch. Despite the substantial variations in execution success rates, the execution-based rewards in both cases exhibit a consistent increasing trend. These findings demonstrate the interactions between LLM-based autonomous research agents and real-world environments.



Figure 5: The impact of real-world peer review environment on execution results. AlphaResearch-RM-7B filters 151 bad ideas, where 108 ideas fail to execute and 43 are successful.

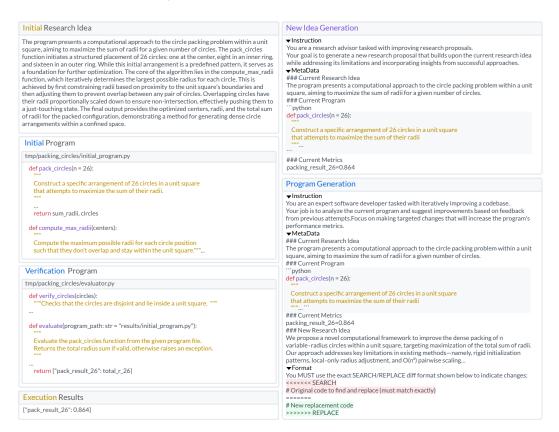


Figure 6: We show an example of a formatted task of AlphaResearch.

The impact of real-world peer-review environment. To assess the effectiveness of reward from a simulated real-world peer-view environment, we ablate AlphaResearch-RM-7B at the first 400 iterations on "Packing Circles" problem. Figure 5 presents the execution results of w/ and w/o AlphaResearch-RM-7B during the discovery process. Compared to the baseline without RM, AlphaResearch-RM-7B successfully filtered 151 ideas below the threshold. This process yielded 108 correct rejections of execution failures while making 43 erroneous rejections of viable ideas. AlphaResearch attained an accuracy of 71.5% (108/151), a result that aligns closely with its performance on the AlphaResearch-RM test set, as shown in Table 2 This outcome effectively demonstrates the model's generalization capabilities and the efficacy of incorporating feedback from a simulated real-world peer-review environment.

4.4 CASE STUDY

We select the successful example from AlphaResearch to better understand the discovery process. We'll consider the problem "Packing Circles" where the goal is to pack n disjoint circles inside a unit square so as to maximize the sum of their radii, shown in Figure 6. We first initialize AlphaResearch with an original research proposal and a related program that returns a list of circles (x, y, r) as output, as shown in Appendix D.4. The verification program first employs $verify_circles$

function to check if the outputs of the initial program meet the problem constraints (e.g., all circles are inside a unit square) and evaluate function to output the sum of their radii. The metadata, including: (1) research ideas, (2) programs, (3) execution results, are subsequently preserved as candidates which represent the end of one step. At the next step, AlphaResearch will sample from the candidate pool and generate a new idea to improve the research proposals from the sampled metadata. After generating the new research ideas, AlphaResearch will further generate a patch to modify the existing program if the idea obtains a positive score from AlphaResearch-RM. The new program is then evaluated by the same verification program, thereby generating new metadata. We select the best program and idea as the final solution of AlphaResearch in this iterative process.

5 RELATED WORK

LLMs for New Ideas. Several recent works explored methods to improve research idea generation, such as iterative novelty refinement (Wang et al., 2024a; Baek et al., 2024). These works focus on improving the research idea over vanilla prompting but critically miss an effective verification method. To promote more reliable AI-generated research ideas, many studies have proposed solutions from different perspectives, such as comparisons with any human expert (Si et al., 2024), using LLMs for executing experiments by generating code with human-curated research problems (Huang et al., 2024; Tian et al., 2024), and executing LLM-generated research ideas with LLM-generated programs (Li et al., 2024; Lu et al., 2024; Aygün et al., 2025). These works either use automatic program evaluation or a misaligned LLM evaluator method, which presents a challenge for their scalability to real-world advanced algorithm discovery. Our AlphaResearch presents a more feasible direction by combining program execution with RM training from real-world peer-reviewed research records.

LLMs for Code Generation. In autonomous research agents, code generation serves as a fundamental step. Previous models (Guo et al., 2024; Yu et al., 2023; Hui et al., 2024) and benchmarks (Chen et al., 2021; Yu et al., 2025) for code generation are in a longstanding pursuit of synthesizing code from natural language descriptions. SWE-Bench (Jimenez et al., 2024) introduces the problems in real-world software development. Many studies on SWE-Bench have greatly contributed to the emergence of coding agents like SWE-Agent (Yang et al., 2024) and OpenHands (Wang et al., 2025). These agent frameworks greatly facilitate the training of agentic LLMs like Kimi-K2 (Team et al., 2025) and GLM-4.5 (Zeng et al., 2025). The surge of these models on SWE-Bench underscores a critical need to reassess the future directions of coding agent research. Our AlphaResearchComp benchmark shows that testing LLMs on open-ended research for algorithm discovery is a promising direction to adapt language models to real-world tasks.

6 Discussion

Limitations and future directions. Although AlphaResearch successfully discovers novel algorithms, we hope to expand its coverage to more realistic applications like accelerating tensor computations. Second, our experiments aim to establish the simplest and most straightforward approaches for algorithm discovery. Future research should pay more attention to augmenting the research agents with useful external tools and the application to more complex problems. Lastly, the training of RM in AlphaResearch is based on small models (e.g., Qwen-2.5-7B-Instruct) and 24,445 ICLR peer review records. Enhancing the reward model parameter and dataset size are two important directions which is left for future research.

Conclusion. We present AlphaResearch, an autonomous research agent that synergistically combines new idea generation with program-based verification for novel algorithm discovery. Our approach demonstrates the potential of employing LLM to discover unexplored research areas, enabling language models to effectively tackle complex open-ended tasks. We construct AlphaResearchComp, including 8 open-ended algorithmic problems, where AlphaResearch outperforms human researchers in 2/8 algorithmic problems but lags behind in the remaining 6 problems. Our systematic analysis of the benefits and remaining challenges of autonomous algorithm discovery provides valuable insights for future research, contributing to the development of more advanced and capable research agents.

REFERENCES

- Eser Aygün, Anastasiya Belyaeva, Gheorghe Comanici, Marc Coram, Hao Cui, Jake Garrison, Renee Johnston Anton Kast, Cory Y McLean, Peter Norgaard, Zahra Shamsi, et al. An ai system to help scientists write expert-level empirical software. *arXiv preprint arXiv:2509.06503*, 2025.
- Jinheon Baek, Sujay Kumar Jauhar, Silviu Cucerzan, and Sung Ju Hwang. ResearchAgent: Iterative Research Idea Generation over Scientific Literature with Large Language Models. *ArXiv*, abs/2404.07738, 2024.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Yu Wu, YK Li, et al. Deepseek-coder: When the large language model meets programming—the rise of code intelligence. *arXiv* preprint arXiv:2401.14196, 2024.
- Sikun Guo, Amir Hassan Shariatmadari, Guangzhi Xiong, Albert Huang, Myles Kim, Corey M Williams, Stefan Bekiranov, and Aidong Zhang. Ideabench: Benchmarking large language models for research idea generation. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 5888–5899, 2025.
- Qian Huang, Jian Vora, Percy Liang, and Jure Leskovec. MLAgentBench: Evaluating Language Agents on Machine Learning Experimentation. In *ICML*, 2024.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*, 2024.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=chfJJYC3iL.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R Narasimhan. SWE-bench: Can language models resolve real-world github issues? In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=VTF8yNQM66.
- Ruochen Li, Teerth Patel, Qingyun Wang, and Xinya Du. MLR-Copilot: Autonomous Machine Learning Research based on Large Language Models Agents. *ArXiv*, abs/2408.14033, 2024.
- Yong Lin, Shange Tang, Bohan Lyu, Jiayun Wu, Hongzhou Lin, Kaiyu Yang, Jia LI, Mengzhou Xia, Danqi Chen, Sanjeev Arora, and Chi Jin. Goedel-prover: A frontier model for open-source automated theorem proving. In *Second Conference on Language Modeling*, 2025. URL https://openreview.net/forum?id=x2y9i2HDjD.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.
- Alexander Novikov, Ngân Vũ, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco JR Ruiz, Abbas Mehrabian, et al. Alphaevolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint arXiv:2506.13131*, 2025.
 - OpenAI. Gpt-5. 2025. URL https://openai.com/index/introducing-gpt-5/.

- Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Chen Bo Calvin Zhang, Mohamed Shaaban, John Ling, Sean Shi, et al. Humanity's last exam. *arXiv* preprint *arXiv*:2501.14249, 2025. URL https://arxiv.org/abs/2501.14249.
- Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. Can Ilms generate novel research ideas. 2024.
 - Chenglei Si, Tatsunori Hashimoto, and Diyi Yang. The ideation-execution gap: Execution outcomes of llm-generated versus human research ideas. *arXiv preprint arXiv:2506.20803*, 2025.
 - Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, et al. Kimi k2: Open agentic intelligence. *arXiv* preprint arXiv:2507.20534, 2025.
 - Minyang Tian, Luyu Gao, Shizhuo Dylan Zhang, Xinan Chen, Cunwei Fan, Xuefei Guo, Roland Haas, Pan Ji, Kittithat Krongchon, Yao Li, Shengyan Liu, Di Luo, Yutao Ma, Hao Tong, Kha Trinh, Chenyu Tian, Zihan Wang, Bohao Wu, Yanyu Xiong, Shengzhu Yin, Min Zhu, Kilian Lieret, Yanxin Lu, Genglin Liu, Yufeng Du, Tianhua Tao, Ofir Press, Jamie Callan, E. A. Huerta, and Hao Peng. SciCode: A Research Coding Benchmark Curated by Scientists. *ArXiv*, abs/2407.13168, 2024.
 - Trieu H Trinh, Yuhuai Wu, Quoc V Le, He He, and Thang Luong. Solving olympiad geometry without human demonstrations. *Nature*, 625(7995):476–482, 2024.
 - Qingyun Wang, Doug Downey, Heng Ji, and Tom Hope. Scimon: Scientific inspiration machines optimized for novelty. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 279–299, 2024a.
 - Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert Brennan, Hao Peng, Heng Ji, and Graham Neubig. Openhands: An open platform for AI software developers as generalist agents. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=OJd3ayDDoF.
 - Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhu Chen. MMLU-pro: A more robust and challenging multitask language understanding benchmark. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024b. URL https://openreview.net/forum?id=y10DM6R2r3.
 - An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
 - John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems*, 37:50528–50652, 2024.
 - Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, et al. Justice or prejudice? quantifying biases in llm-as-a-judge. arXiv preprint arXiv:2410.02736, 2024.
 - Zhaojian Yu, Xin Zhang, Ning Shang, Yangyu Huang, Can Xu, Yishujie Zhao, Wenxiang Hu, and Qiufeng Yin. Wavecoder: Widespread and versatile enhancement for code large language models by instruction tuning. *arXiv preprint arXiv:2312.14187*, 2023.
 - Zhaojian Yu, Yilun Zhao, Arman Cohan, and Xiao-Ping Zhang. HumanEval pro and MBPP pro: Evaluating large language models on self-invoking code generation task. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 13253–13279, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025. findings-acl.686. URL https://aclanthology.org/2025.findings-acl.686/.

Aohan Zeng, Xin Lv, Qinkai Zheng, Zhenyu Hou, Bin Chen, Chengxing Xie, Cunxiang Wang, Da Yin, Hao Zeng, Jiajie Zhang, et al. Glm-4.5: Agentic, reasoning, and coding (arc) foundation models. arXiv preprint arXiv:2508.06471, 2025.

APPENDIX CONTENTS

1	Introduction		
2	Alpl	naResearch	3
	2.1	Overview	3
	2.2	Actions	3
	2.3	Environment	4
		2.3.1 Reward from Real-world Research Records	4
		2.3.2 Reward from Program-based Execution	4
3	Alpl	haResearchComp	4
4	Exp	eriments	5
	4.1	Setup	5
	4.2	Results	6
	4.3	Ablations and Analysis	7
	4.4	Case Study	8
5	Rela	ated Work	9
6	Disc	eussion	9
A	The	Use of Large Language Models	14
В	Exa	mples	14
C	Pro	mpts	15
D	Cur	ated Problems and Human-Best Values	17
	D.1	Spherical Code (S^2 , $n=30$)	17
	D.2	Littlewood Polynomials.	18
	D.3	Sum vs. Difference Sets (MSTD).	20
	D.4	Packing Circle in a Square (variable radii)	21
	D.5	Minimizing Max/Min Distance Ratio ($d=2, n=16$)	26
	D.6	Autoconvolution Peak Minimization (L^{∞})	28
	D.7		31
	D.8	Third-Order Autocorrelation Inequality (C_2 Upper Bound)	31

A THE USE OF LARGE LANGUAGE MODELS

During the preparation of this manuscript, we utilized large language models (LLMs) for grammar checking and writing suggestions to enhance the readability and clarity of the content.

B EXAMPLES

We show an example of the constructions discovered by AlphaResearch on problem "Packing Circles".

AlphaEvolve

702

703 704

705

706

708

710 711 712

713

714

715

716

717

718

719

720

721

722

723

724

725

726 727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

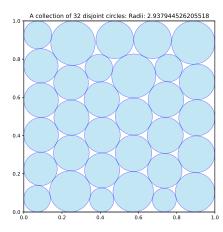
743 744

754

```
packing_circles_alphaevolve = np.array([[0.09076163, 0.40381803, 0.090761620923837],
                                          [0.07310993, 0.92689178, 0.07310821268917801], [0.08745017, 0.22570576,
                                       0.087381421261857],\; [0.24855246,\; 0.30880277,\; 0.093428060657193],\; [0.4079865,\; 0.06300614,\; 0.06300614],\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614],\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.06300614,\; 0.063006144,\; 0.063006144,\; 0.063006144,\; 0.063006144,\; 0.063006144,\; 0.063006144,\; 0.063006144,\; 0.063006144,\; 0.063006144,\; 0
                                     0.063006133699386], [0.47646318, 0.90136179, 0.09863820013617901], [0.89604966,
                                       0.10309934,\ 0.10309932969006601],\ [0.9066386,\ 0.68096117,\ 0.09336139066386],\ [0.08962002,\ 0.1030934],\ 0.10309932969006601],\ [0.9066386,\ 0.68096117,\ 0.09336139066386],\ [0.9066386,\ 0.68096117,\ 0.09336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.68096117,\ 0.909336139066386],\ [0.9066386,\ 0.9066386],\ [0.9066386,\ 0.9066386],\ [0.9066386,\ 0.9066386],\ [0.9066386,\ 0.9066386],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.90638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.90638],\ [0.9066386,\ 0.906638],\ [0.9066386,\ 0.906638],\ [0.906638
                                              0.76509474,\ 0.0895289910471],\ [0.06973669,\ 0.06965159,\ 0.06965158303484101],
                                          [0.40979823, 0.21756451, 0.09156283084371601], [0.25742466, 0.88393887,
                                       0.11606111839388701], \; [0.09064689, \; 0.58506214, \; 0.090482500951749], \; [0.90294698, \; 0.090482500951749], \; [0.90294698, \; 0.090482500951749], \; [0.90294698, \; 0.90294698], \; [0.90294698, \; 0.90294698], \; [0.90294698, \; 0.90294698], \; [0.90294698, \; 0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.90294698], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948], \; [0.902948
                                       0.30231577, 0.09623644037635501], [0.57265603, 0.10585396, 0.105853949414604],
                                        [0.74007588, 0.40129314, 0.09435083056491601], [0.57539962, 0.71183255,
                                        0.115160168483982], [0.7367635, 0.21592191, 0.09104997089500201], [0.41096972,
                                       0.40263617,\ 0.093512520648747],\ [0.88664452,\ 0.88667032,\ 0.113317128668286],\ [0.57582722,\ 0.88667032,\ 0.113317128668286],\ [0.88681],\ [0.88681],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.88881],\ [0.888
                                                0.49961748, 0.09705531029446801], [0.24962585, 0.49417195, 0.09194421080557799],
                                        [0.90546338, 0.49309632, 0.094507120549287], [0.67381348, 0.90149423,
                                        0.09850576014942301], [0.24310147, 0.1077195, 0.10771948922805], [0.40815297, 0.5886157,
                                     0.09248833075116601], [0.24737889, 0.6771266, 0.090994980900501], [0.75801377, 0.7532924, 0.07192969280703], [0.73526642, 0.06243992, 0.062439303756069], [0.57415412, 0.30715219,
                                              0.095403150459684], \; [0.39239379, \; 0.75259664, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.7439361, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.07223814277618501], \; [0.743961, \; 0.072238141], \; [0.745961, \; 0.072238142701], \; [0.755961, \; 0.07223814
                                     0.58879735, 0.09316663068333611)
```

AlphaResearch

```
packing_circles_alpharesearch = np.array([[(0.1115677319034151, 0.11156773191787371,
                                  0.11156438489140026), (0.09380224787136374, 0.3161654253705352, 0.09379943380606216),
                                  0.6962443020287629,\ 0.09657032835808858),\ (0.10365512530384222,\ 0.8963448746980195,
                                 0.10365201565567386), (0.3334956594919712, 0.09664441783072292, 0.0966415184920332),
                                  (0.26448615440016093,\ 0.9376113341122044,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.06238679422590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.5287192731314015,\ 0.0623867942590162),\ (0.52871927313140150162),\ (0.52871927313140150162),\ (0.52871927313140150162),\ (0.528719273140150162),\ (0.528719273140150162),\ (0.528719273140150162),\ (0.528719273140150162),\ (0.528719273140150162),\ (0.528719273140150162),\ (0.528719273140150162),\ (0.52871927313140150162),\ (0.52871927313140150162),\ (0.52871927313140150162),\ (0.52871927313140150162),\ (0.52871927313140150162),\ (0.528719273140150162),\ (0.528719273140150162),\ (0.528719273140150162),\ (0.528719273140150162),\ (0.528719273
                                 0.09859146596680078,\ 0.09858850822808951),\ (0.591325020569507,\ 0.9366833118077788,
                                 0.0633147886877468), \\ (0.7427106948954978, \\ 0.11611889563206494, \\ 0.11611541209023483), \\ (0.7427106948954978, \\ 0.11611889563206494, \\ 0.11611541209023483), \\ (0.7427106948954978, \\ 0.11611889563206494, \\ 0.11611541209023483), \\ (0.7427106948954978, \\ 0.11611889563206494, \\ 0.11611541209023483), \\ (0.7427106948954978, \\ 0.7427106948954978, \\ 0.7427106948954978, \\ 0.7427106948954978, \\ 0.7427106948954978, \\ 0.7427106948954978, \\ 0.7427106948954978, \\ 0.7427106948954978, \\ 0.74271069489549, \\ 0.74271069489549, \\ 0.74271069489, \\ 0.74271069489, \\ 0.7427106948, \\ 0.7427106948, \\ 0.7427106948, \\ 0.7427106948, \\ 0.7427106948, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.742710694, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\ 0.74271064, \\
                                   0.08942314561430993), \quad (0.9094700615258342, \quad 0.41468336419923396, \quad 0.09052722258939731), \quad (0.9052722258939731), \quad (0.90527222589731), \quad (0.905272225897222589721), \quad (0.90527222589721), \quad (0.9052722258721), \quad (0.9052722258721), \quad
                                  0.09216300786888813), \quad (0.5896628759126524, \quad 0.5965222415947758, \quad 0.09365298106148348), \quad (0.5896628759126524, \quad 0.5966222415947758, \quad 0.09365298106148348), \quad (0.5896628759126524, \quad 0.5966222415947758, \quad 0.09365298106148348), \quad (0.5896628759126524, \quad 0.5966222415947758, \quad 0.09365298106148348), \quad (0.5896628759126524, \quad 0.59662224415947758, \quad 0.09365298106148348, \quad 0.096628759126524, \quad 0.096628759126524, \quad 0.096628759126524, \quad 0.0966287591266524, \quad 0.0966287591264, \quad
                                  (0.26303074890883915,\ 0.783747668079202,\ 0.09148238826692158),\ (0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.42710033854875884,\ 0.4271003385487584,\ 0.42710033854854,\ 0.42710033854854,\ 0.4271003385484,\ 0.4271003385484,\ 0.4271003385484,\ 0.4271003385484,\ 0.4271003385484,\ 0.4271003385484,\ 0.4271003385484,\ 0.4271003385484,\ 0.4271003385484,\ 0.4271003385484,\ 0.4271003385484,\ 0.4271003385484,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.427100338544,\ 0.4271003444,\ 0.42
                                  0.28662965969327264, 0.1151473780101257), (0.7511102582575875, 0.5051558281448295,
                                  0.09185177348783963), (0.4273023330525072, 0.8937703360976411, 0.10622647700018645)
                                  0.6918664604322906, 0.09567746779211372), (0.2572363869779693, 0.4085253312744954,
                                 (0.42560864125756626,\ 0.49898110459434486,\ 0.09720528992590773),\ (0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.7533817110763772,\ 0.753381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.75381711076372,\ 0.753817111076372,\ 0.753817111076372,\ 0.753817111076372,\ 0.753817111076372,\ 0.75381711076
                                 0.32263902019589896, 0.09067643144615074), (0.5903729314333418, 0.7817733747765757, 0.09159665425215473), (0.7515568081174837, 0.6905957415401818, 0.09358581053778628),
                                  0.24805951545091487, 0.07133567304015336)]])
```



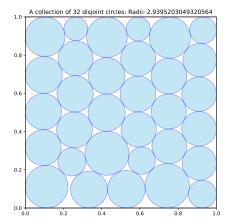


Figure 7: New construction of AlphaResearch (right) improving the best known AlphaEvolve (right) bounds on packing circles to maximize their sum of radii. Left: 32 circles in a unit square with sum of radii ≥ 2.9379 . Right: 32 circles in a unit square with sum of radii ≥ 2.9395

C PROMPTS

Prompt for New Program Generation

You are an expert software developer tasked with iteratively improving a codebase. Your job is to analyze the current program and suggest improvements based on the current proposal and feedback from previous round. Focus on making targeted changes that will increase the program's performance metrics.

Previous Proposal:

{previous proposal}

Previous Program:

{previous program}

Previous Performance Metrics:

{previous result}

Current Proposal

{proposal}

Task

Suggest improvements to the program that will lead to better performance on the specified metrics.

You MUST use the exact SEARCH/REPLACE diff format shown below to indicate changes:

Example of valid diff format:

```
1  <<<<<< SEARCH
2  for i in range(m):
3   for j in range(p):</pre>
```

```
810
811
        4
                    for k in range(n):
        5
                        C[i, j] \leftarrow A[i, k] * B[k, j]
812
        6
813
        7
814
        8
815
           # Reorder loops for better memory access pattern
       9
816
       10
           for i in range(m):
817
       12
               for k in range(n):
818
                    for j in range(p):
       13
819
       14
                        C[i, j] += A[i, k] * B[k, j]
820
       15
821
           >>>>> REPLACE
822
```

You can suggest multiple changes. Each SEARCH section must exactly match code in the current program.

Be thoughtful about your changes and explain your reasoning thoroughly.

IMPORTANT: Do not rewrite the entire program - focus on targeted improvements.

Prompt for New Idea Generation

You are a research advisor tasked with evolving and improving research proposals. Your goal is to generate a new research proposal that builds upon the current proposal while addressing its limitations and incorporating insights from successful approaches.

Based on the following information, generate an improved research proposal:

Focus on:

- 1. Identifying weaknesses in the current approach based on performance metrics
- 2. Proposing novel improvements that could enhance performance
- 3. Learning from successful inspirations while maintaining originality
- 4. Ensuring the new proposal is implementable
- Current Proposal:

{proposal}

- Current Program:

{program}

- Current Metrics:

{results}

Please generate a new research proposal that:

- 1. Addresses the limitations shown in the current metrics
- 2. Incorporates insights from successful approaches
- 3. Proposes specific technical improvements
- 4. Maintains clarity and technical rigor

Return the proposal as a clear, concise research abstract.

Prompt for AlphaResearch-RM-7B

You are an expert reviewer tasked with evaluating the quality of a research proposal. Your goal is to assign a score between 1 and 10 based on the proposal's clarity, novelty, technical rigor, and potential impact. Here are the criteria:

- 1. Read the following proposal carefully and provide a score from 1 to 10.
- 2. Score 6 means slightly higher than the borderline, 5 is slightly lower than the borderline. Write the score in the \boxed{}. {proposal}

D CURATED PROBLEMS AND HUMAN-BEST VALUES

We summarize the ten problems used in the ALPHARESEARCH benchmark. For each item we state the objective, the current human-best value at the benchmark's default parameters, and whether this value is proved optimal or only best-known.

D.1 SPHERICAL CODE $(S^2, n = 30)$.

Problem Description: Place n=30 points on the unit sphere in \mathbb{R}^3 to *maximize* the minimal pairwise angle θ_{\min} .

Human Best: $\theta_{\rm min} \approx 0.673651 \text{ radians } (\approx 38.5971^{\circ}).$

Initial Proposal

Problem definition. Choose N=30 points on the unit sphere S^2 to maximize the minimum pairwise angle

$$\theta_{\min} = \min_{i < j} \arccos(\langle p_i, p_j \rangle).$$

Constraints.

- Points are unit vectors (rows normalized).
- Metric is θ_{\min} in radians.

Optimization goal. Maximize θ_{\min} . The evaluator returns $\{\text{score}, \theta_{\min}, N, \text{dimension}\}$, with $\text{score} = \theta_{\min}$.

Best-known reference (for N = 30 on S^2):

```
\cos(\theta^*) \approx 0.7815518750949873 \quad \Rightarrow \quad \theta^* \approx 0.6736467551690225 \text{ rad.}
```

Reference table: Henry Cohn's spherical codes data (https://cohn.mit.edu/spherical-codes).

Best-known results (human).

- On S^2 (3D), small N optima coincide with symmetric polyhedra (e.g., tetrahedron, octahedron, icosahedron).
- ullet For larger N, best codes come from numerical optimization; exact optimality is only known in limited cases.

Algorithmic goal. Construct codes with larger θ_{\min} . The baseline seeds with symmetric configurations and uses farthest-point max-min. Stronger methods include:

- Energy minimization,
- Projected gradient / coordinate descent,
- Stochastic max-min refinement.

```
918
           if n == 6: # octahedron
919
              return np.array([[1,0,0],[-1,0,0],[0,1,0],[0,-1,0],[0,0,1],[0,0,-1]], dtype=float)
           if n == 8: # cube vertices
               V = np.array([[sx, sy, sz] for sx in (-1,1) for sy in (-1,1) for sz in (-1,1)],
921
                   dtype=float)
              return _normalize_rows(V)
922
           if n == 12: # icosahedron (one realization)
923
              phi = (1+np.sqrt(5))/2
               V = []
924
              for s in (-1,1):
                 V += [[0, s, phi],[0, s, -phi],[s, phi,0],[s, -phi,0],[phi,0, s],[-phi,0, s]]
925
              V = np.array(V, dtype=float)
926
              return _normalize_rows(V)
927
           return None
928
        def farthest_point_greedy(n, seed=None, rng=np.random.default_rng(0)):
929
           Greedy max min on S^2: start from seed, then add points that maximize min angle.
930
931
           def random_unit(k):
              X = rng.normal(size=(k,3)); return _normalize_rows(X)
932
933
           if seed is None:
              P = random_unit(1) # start with one random point
934
           else:
935
              P = _normalize_rows(seed)
           while len(P) < n:
936
              # generate candidates and pick the one with largest min angle to current set
937
              C = random\_unit(2000) \# candidates per iteration (tune as needed)
              # cosines to existing points
938
              cos = C @ P.T
939
               # min angle to set -> maximize this
              min\_ang = np.arccos(np.clip(np.max(cos, axis=1), -1.0, 1.0))
940
              idx = np.argmax(min_ang)
941
              P = np.vstack([P, C[idx:idx+1]])
           return P
942
943
        def main():
           n = 30
944
           seed = seed_platonic(n)
945
           pts = farthest_point_greedy(n, seed=seed, rng=np.random.default_rng(42))
           print(f"n={n}, points={len(pts)}")
946
           return pts
947
        if __name__ == "__main__":
948
           points = main()
949
           np.save("points.npy", points)
951
        # Ensure compatibility with evaluators that expect a global variable
952
           points # type: ignore[name-defined]
953
         except NameError:
           points = main()
954
```

D.2 LITTLEWOOD POLYNOMIALS.

955956957958959960961962963

964 965 966

967 968

969

970

971

Problem Description For coefficients $c_k \in \{\pm 1\}$ and $P_n(t) = \sum_{k=0}^{n-1} c_k e^{ikt}$, minimize $\|P_n\|_{\infty} = \sup_{t \in \mathbb{R}} |P_n(t)|$.

Human Best: the Rudin–Shapiro construction gives $||P_n||_{\infty} \le \sqrt{2n}$. At the benchmark setting n = 512, this yields $||P_{512}||_{\infty} \le 32$ (so the "larger-is-better" score $1/||P_n||_{\infty}$ is $\ge 1/32 = 0.03125$). Sharper constants are known for special families, but $\sqrt{2n}$ remains a clean baseline.

Initial Proposal

972

973 974

975976

977978979

980

981 982

983

984 985

986

987

988

993

994 995

996 997

998

999

1000 1001

1002

Choose coefficients $c_k \in \{\pm 1\}$ for

$$P(z) = \sum_{k=0}^{n-1} c_k z^k, \qquad |z| = 1,$$

so as to minimize the supremum norm

$$||P||_{\infty} = \max_{|z|=1} |P(z)|.$$

Constraints.

- Coefficients c_k are restricted to ± 1 .
- The metric $\|P\|_{\infty}$ is estimated by FFT sampling on an equally spaced grid (denser grid \to tighter upper bound).

Optimization Goal. The evaluator returns:

$$\text{score} = \begin{cases} \frac{1}{\|P\|_{\infty}}, & \text{if valid}, \\ -1.0, & \text{otherwise}. \end{cases}$$

Notes on Bounds. For the Rudin–Shapiro construction of length n, a classical identity gives

$$||P||_{\infty} \le \sqrt{2n}$$
.

For the benchmark default n = 512, this yields

$$||P||_{\infty} \le \sqrt{1024} = 32,$$

so

score =
$$\frac{1}{32}$$
 = 0.03125.

```
1003
         def rudin_shapiro(n: int):
1004
1005
            First n signs of the Rudin-Shapiro sequence.
1006
            a = np.ones(n, dtype=int)
1007
            for k in range(n):
               x, cnt, prev = k, 0, 0
1008
               while x:
1009
                  b = x \& 1
                  if b & prev: # saw '11'
cnt ^= 1
1010
1011
                  prev = b
                  x >>= 1
1012
               a[k] = 1 if cnt == 0 else -1
1013
            return a
1014
         def random littlewood(n: int, seed=0):
1015
            rng = np.random.default_rng(seed)
            return rng.choice([-1, 1], size=n).astype(int)
1016
1017
         def main():
1018
           n = 512
            c = rudin_shapiro(n)
1019
           print(f"n={n}, coeffs={len(c)}")
1020
            return c
1021
        if __name__ == "__main__":
    coeffs = main()
1022
1023
         # Ensure compatibility with evaluators that expect a global variable
1024
            coeffs # type: ignore[name-defined]
1025
         except NameError:
            coeffs = main()
```

1028

1029 1030

1031

1032

1033 1034

1035

1038 1039

1036 1037

1040 1041

> 1042 1043 1044

1045 1046

1047 1048 1049

1050 1051 1052

1053 1054 1055

1056 1057 1058

1059

1062 1063

1064

1065 1066

1067 1068 1069

1070

1071

1072

1077 1078 1079

D.3 SUM VS. DIFFERENCE SETS (MSTD).

Problem Description For a finite set $A \subset \mathbb{Z}$, maximize |A+A|/|A-A|.

Human Best: MSTD sets exist; the smallest possible size is |A| = 8 (classification up to affine equivalence is known). For larger |A|, extremal ratios remain open; our benchmark instance reports a representative value (≈ 1.04 for |A| = 30).

Initial Proposal

Objective. Classical MSTD (enforced): Given $A \subset \{0, 1, \dots, N-1\}$ represented by a 0/1 indicator array of length N, maximize the ratio

$$R = \frac{|A+A|}{|A-A|}.$$

- Score: score = R (higher is better).
- Comparisons should be made under the same N.

Default setup.

- N = 30.
- Evaluator enforces A = B (classical setting). If a pair (A, B) is provided, B is ignored and A is used.

Known best for N=30 (baseline). Conway's MSTD set

$$A = \{0, 2, 3, 4, 7, 11, 12, 14\}$$

yields $R \approx 1.04$. This is the baseline included in initial_program.py. Better ratios may exist for N = 30; pushing R upwards is the optimization goal.

Notes.

- R > 1 is rare and indicates sum-dominance.
- The ratio depends strongly on N; do not compare ratios across different N without a normalization scheme.
- If cross-N comparison is necessary, consider reporting both R and N, or use $\log R$ as an auxiliary measure.

```
def main():
  N = 30
   \mbox{\#} Conway MSTD set example; we take A=B for classical MSTD
   A = [0, 2, 3, 4, 7, 11, 12, 14]
   B = A[:]
   A_ind = np.zeros(N, dtype=int); A_ind[A] = 1
   B_ind = np.zeros(N, dtype=int); B_ind[B] = 1
   return A_ind, B_ind
# Ensure globals for evaluator
   A_indicators; B_indicators # type: ignore[name-defined]
except NameError:
   A_indicators, B_indicators = main()
```

D.4 PACKING CIRCLE IN A SQUARE (VARIABLE RADII).

Problem Description In the unit square, place n disjoint circles (radii free) to maximize the sum of radii $\sum r_i$.

Best-known: for n = 26, $\sum r_i = 2.634$ (Cantrell, 2011); for n = 32, $\sum r_i = 2.936$ (Specht, 2012).

Initial Proposal

1080

1081 1082

1083

1085

1086

1088 1089

1090

1091

1092

1093

1095

1096

1097

1099

1100

1101

11021103

1104

Problem definition. Given an integer n, place n disjoint circles in the unit square $[0,1]^2$ to maximize the total sum of radii.

Objective and metric.

- Score: score = $\sum_{i=1}^{n} r_i$ (larger is better).
- Validity: circles must be pairwise disjoint and fully contained in the unit square.

Notes on records.

- This variable-radius "sum of radii" objective is not the classical equal-radius packing; authoritative SOTA tables are not standardized.
- Values reported in code or experiments should be treated as benchmarks rather than literature SOTA.

Goal. Create algorithms that increase the total sum of radii for $n \in \{26, 32\}$ under the above validity constraints.

```
1105
1106
        from concurrent.futures import ThreadPoolExecutor
1107
        def pack_circles(n, square_size=1.0):
1108
1109
           Pack n disjoint circles in a unit square using uniform tiling approach.
           Returns the sum of radii and list of circles (x, y, r).
1110
1111
           def max_circle_radius(x, y, circles, square_size=1.0, skip_idx=None):
1112
1113
              Compute the maximum radius for a circle centered at (x, y) that:
              - Stays within the unit square [0, square_size] \times [0, square_size].
1114
              - Does not overlap with existing circles.
1115
              skip_idx: if provided, index in circles[] to ignore (self).
1116
              # Distance to nearest boundary of the unit square
1117
              r_max = min(x, y, square_size - x, square_size - y)
1118
              # Check distance to existing circles, exit early if r_max \rightarrow 0
1119
              \# early exit if r_max is tiny, and avoid needless sqrt
              for idx, (cx, cy, cr) in enumerate(circles):
    if skip_idx == idx:
1120
1121
                 if r_max <= 1e-8:
1122
                    break
1123
                 dx = x - cx
                 dy = y - cy
1124
                 sep = r_max + cr
1125
                  if dx*dx + dy*dy < sep*sep:
1126
                     # only compute sqrt when we know we can shrink
                    dist = math.sqrt(dx*dx + dy*dy)
1127
                    r_max = min(r_max, dist - cr)
              return max(r_max, 0.0)
1128
1129
           def uniform_tiling_circles(n, square_size=1.0):
1130
              Uniformly tile the square with circles using optimal grid placement.
1131
              if n <= 0:
1132
                  return []
1133
              circles = []
```

```
1134
1135
               # Calculate optimal grid dimensions
               \# For n circles, find the best grid layout (rows x cols)
1136
               best_layout = None
1137
              best_total_radius = 0
1138
               # Try different grid configurations
1139
               for rows in range (1, \min(n + 1, 20)):
                  cols = math.ceil(n / rows)
1140
                  if cols > 20: # Limit grid size
1141
                     continue
1142
                  # Calculate spacing
1143
                  spacing_x = square_size / (cols + 1)
                  spacing_y = square_size / (rows + 1)
1144
1145
                  # Use the smaller spacing to ensure circles fit
                 min_spacing = min(spacing_x, spacing_y)
1146
1147
                  # Calculate maximum radius for this layout
1148
                 max_radius = min_spacing / 2
1149
                  # Ensure radius doesn't exceed boundaries
                 max_radius = min(max_radius,
1150
                             spacing_x / 2 - 1e-6,
spacing_y / 2 - 1e-6)
1151
1152
                  if max_radius <= 0:</pre>
1153
                     continue
1154
                  # Place circles in uniform grid
1155
                  temp_circles = []
                  count = 0
1156
1157
                  for row in range(rows):
                     for col in range(cols):
1158
                        if count >= n:
1159
                          break
1160
                        x = spacing_x * (col + 1)
1161
                        y = spacing_y * (row + 1)
1162
                        # Ensure circle stays within bounds
1163
                        if (x - max_radius >= 0 and x + max_radius <= square_size and
                           y - max_radius >= 0 and y + max_radius <= square_size):
1164
1165
                           temp_circles.append((x, y, max_radius))
                           count += 1
1166
1167
                     if count >= n:
                        break
1168
1169
                  # Calculate total radius for this layout
                  total_radius = len(temp_circles) * max_radius
1170
1171
                  if total_radius > best_total_radius and len(temp_circles) == n:
                     best_total_radius = total_radius
1172
                     best_layout = temp_circles
1173
               # If we found a valid layout, return it
1174
               if best_layout:
1175
                  return best_layout
1176
               # Fallback: use hexagonal packing for better density
1177
               return hexagonal_packing(n, square_size)
1178
           def hexagonal_packing(n, square_size=1.0):
1179
               Use hexagonal close packing for better space utilization.
1180
1181
              circles = []
1182
               # Estimate number of rows and columns for hexagonal packing
1183
               # Hexagonal packing has rows offset by sqrt(3)/2 * diameter
1184
               rows = int(math.sqrt(n * 2 / math.sqrt(3))) + 2
1185
1186
               count = 0
               row = 0
1187
               while count < n and row < rows:
```

```
1188
                  # Calculate y position for this row
1189
                  y = (row + 0.5) * (square_size / (rows + 1))
1190
                  # Number of circles in this row
1191
                  if row % 2 == 0:
                     cols = int(math.sqrt(n)) + 1
1192
                  else:
1193
                     cols = int(math.sqrt(n))
1194
                  spacing_x = square_size / (cols + 1)
1195
                  for col in range(cols):
1196
                     if count >= n:
1197
                        break
1198
                     if row % 2 == 0:
1199
                        x = spacing_x * (col + 1)
                     else:
1200
                        x = \text{spacing}_x * (\text{col} + 1) + \text{spacing}_x / 2
1201
                     \ensuremath{\text{\#}} Calculate maximum radius for this position
1202
                     r = max_circle_radius(x, y, circles, square_size)
1203
                     if r > 0
1204
                        circles.append((x, y, r))
1205
                        count += 1
1206
                  row += 1
1207
1208
               return circles
1209
            def optimize_placement(n, square_size=1.0):
1210
               Optimize circle placement using uniform tiling with radius maximization.
1211
               circles = []
1212
1213
               # First, try hexagonal packing for high initial density
               hex_circles = hexagonal_packing(n, square_size)
1214
               if len(hex_circles) == n:
1215
                  # Ensure maximum radii for hex layout with stronger refinement
                  hex_refined = refine_circles(hex_circles, square_size, iterations=20)
1216
                  return hex_refined
1217
               # Fallback to uniform grid placement
1218
               grid_circles = uniform_tiling_circles(n, square_size)
1219
               if len(grid_circles) == n:
                  return grid_circles
1220
1221
               # If uniform tiling didn't work perfectly, use adaptive approach
               # Calculate optimal radius based on density
1222
               area_per_circle = (square_size * square_size) / n
estimated_radius = math.sqrt(area_per_circle / math.pi) * 0.9 # Conservative estimate
1223
1224
               # Create grid with optimal spacing
1225
               spacing = estimated_radius * 2.1 # Include gap
1226
               cols = int(square_size / spacing)
1227
               rows = int(square_size / spacing)
1228
               actual_spacing_x = square_size / (cols + 1)
1229
               actual_spacing_y = square_size / (rows + 1)
1230
               count = 0
1231
               for row in range (rows):
                  for col in range(cols):
1232
                     if count >= n:
1233
                        break
1234
                     x = actual\_spacing\_x * (col + 1)
1235
                     y = actual_spacing_y * (row + 1)
1236
                     # Calculate maximum possible radius
1237
                     r = max_circle_radius(x, y, circles, square_size)
1238
                     if r > 0:
1239
                        circles.append((x, y, r))
1240
                        count += 1
1241
                  if count >= n:
                     break
```

```
1242
1243
              # If we still need more circles, use remaining space
              remaining = n - len(circles)
1244
              if remaining > 0:
1245
                  # Place remaining circles in remaining spaces
1246
                  for i in range(remaining):
                     # Try different positions systematically
1247
                    best_r = 0
                    best_pos = (0.5, 0.5)
1248
1249
                     # Fine grid search (increased resolution)
                    grid points = 100
1250
                     for gx in range(1, grid points):
1251
                        for gy in range(1, grid_points):
                          x = gx / grid_points
1252
                          y = gy / grid_points
1253
                           r = max_circle_radius(x, y, circles, square_size)
1254
                           if r > best r:
1255
                              best_r = r
                             best_pos = (x, y)
1256
1257
                     if best r > 0:
                       circles.append((best_pos[0], best_pos[1], best_r))
1258
1259
              return circles
1260
           def refine_circles(circles, square_size, iterations=80, perturb_interval=3):
1261
              Iteratively grow each circle to its maximum radius under non-overlap constraints.
1262
              Includes randomized update order, periodic micro-perturbation to escape
1263
              local minima, and a final local-center-perturbation pass for densification.
1264
              for it in range(iterations):
1265
                 # randomize update order to avoid sweep-order bias
                 indices = list(range(len(circles)))
1266
                 random.shuffle(indices)
1267
                  for i in indices:
                    x, y, _ = circles[i]
1268
                     # Compute maximal feasible radius here, skipping self
1269
                     r = max_circle_radius(x, y, circles, square_size, skip_idx=i)
                     circles[i] = (x, y, r)
1270
                  # Periodic micro-perturbation: jiggle a few circles
1271
                  if it % perturb_interval == 0 and len(circles) > 0:
                     subset = random.sample(indices, min(5, len(circles)))
1272
                     for j in subset:
1273
                        x0, y0, r0 = circles[j]
                        dx = random.uniform(-0.03, 0.03)
1274
                        dy = random.uniform(-0.03, 0.03)
1275
                        nx = min(max(x0 + dx, 0), square_size)
                       ny = min(max(y0 + dy, 0), square\_size)
1276
                        # Compute maximal radius skipping self
1277
                        nr = max_circle_radius(nx, ny, circles, square_size, skip_idx=j)
                        if nr > r0:
1278
                          circles[j] = (nx, ny, nr)
1279
              # Full local center-perturbation phase for final densification
              for i in range(len(circles)):
1280
                 x, y, r = circles[i]
1281
                 best_x, best_y, best_r = x, y, r
                 delta = 0.1
1282
                  for _ in range(20):
1283
                    dx = random.uniform(-delta, delta)
                    dy = random.uniform(-delta, delta)
1284
                    nx = min(max(x + dx, 0), square\_size)
1285
                    ny = min(max(y + dy, 0), square\_size)
                     # Compute maximal radius skipping self
1286
                    nr = max_circle_radius(nx, ny, circles, square_size, skip_idx=i)
1287
                     if nr > best r:
                       best_x, best_y, best_r = nx, ny, nr
1288
                    else:
1289
                       delta *= 0.9
                 circles[i] = (best_x, best_y, best_r)
1290
1291
              # Physics-inspired soft relaxation to escape persistent overlaps
              for i in range(len(circles)):
1292
                 x, y, r = circles[i]
1293
                 fx, fy = 0.0, 0.0
1294
                  for j, (xj, yj, rj) in enumerate(circles):
                    if i == j:
1295
                       continue
                    dx = x - xj
```

```
1296
                     dy = y - yj
1297
                     d = (dx*dx + dy*dy) ** 0.5
                     overlap = (r + rj) - d
1298
                     if overlap > 0 and d > 1e-8:
                        fx += dx / d * overlap
fy += dy / d * overlap
1299
1300
                  # Nudge the center by 10\% of the computed net "repulsive" force
1301
                  nx = min(max(x + 0.1 * fx, 0), square\_size)
                  ny = min(max(y + 0.1 * fy, 0), square\_size)
1302
                  nr = max_circle_radius(nx, ny, circles, square_size, skip_idx=i)
1303
                  circles[i] = (nx, ny, nr)
               return circles
1304
1305
            def multi_start_optimize(n, square_size, starts=None):
1306
               Parallel multi-start global \rightarrow local optimization using ThreadPoolExecutor.
1307
               Number of starts adapts to problem size: max(100, 10*n).
1308
               if starts is None:
1309
                  if n <= 50:
                     starts = max(200, n * 20)
1310
                  else:
1311
                     starts = max(100, n * 10)
               # precompute hexagonal packing baseline
1312
               hex_circ = hexagonal_packing(n, square_size)
hex_sum = sum(r for _, _, r in hex_circ)
1313
               best_conf = None
1314
               best_sum = 0.0
1315
               # single trial: seed \rightarrow refine \rightarrow score
1316
               def single_run(_):
1317
                  conf0 = optimize_placement(n, square_size)
                  conf1 = refine_circles(conf0, square_size, iterations=40)
1318
                  s1 = sum(r for _, _, r in conf1)
return s1, conf1
1319
1320
               # dispatch trials in parallel
1321
               with ThreadPoolExecutor() as executor:
                  for score, conf in executor.map(single_run, range(starts)):
1322
                     if score > best_sum:
1323
                        best_sum, best_conf = score, conf.copy()
                     # early exit if near the hex-baseline
1324
                     if best_sum >= hex_sum * 0.995:
1325
1326
               return best_conf
1327
            # Use multi-start global \rightarrow local optimization (adaptive number of starts)
1328
            circles = multi_start_optimize(n, square_size)
1329
            # Quick 2-cluster remove-and-reinsert densification (extended iterations)
1330
            for _ in range(8):
1331
               # remove the two smallest circles to create a larger gap
               smallest = sorted(range(len(circles)), key=lambda i: circles[i][2])[:2]
1332
               removed = [circles[i] for i in smallest]
1333
               # pop in reverse order to keep indices valid
               for i in sorted(smallest, reverse=True):
1334
                  circles.pop(i)
1335
               # refine the remaining configuration briefly
               circles = refine_circles(circles, square_size, iterations=8)
1336
               # reinsert each removed circle with more sampling
1337
               for x_old, y_old, _ in removed:
                  best_r, best_pos = 0.0, (x_old, y_old)
1338
                  for _ in range(500):

x = random.uniform(0, square_size)
1339
                     y = random.uniform(0, square_size)
1340
                     r = max_circle_radius(x, y, circles, square_size)
1341
                     if r > best_r:
                        best_r, best_pos = r, (x, y)
1342
                  circles.append((best_pos[0], best_pos[1], best_r))
1343
               # final local polish after reinsertion
               circles = refine_circles(circles, square_size, iterations=5)
1344
            # end 2-cluster remove-and-reinsert densification
1345
            # Calculate total radius
1346
            total_radius = sum(circle[2] for circle in circles)
1347
            return total_radius, circles
1348
```

D.5 MINIMIZING MAX/MIN DISTANCE RATIO (d = 2, n = 16).

Problem Description For n points in $[0,1]^2$, minimize $R = \frac{\max_{i \neq j} \|x_i - x_j\|}{\min_{i \neq j} \|x_i - x_j\|}$

Best-known: $R^2 \approx 12.890$ (Cantrell, 2009), i.e., $R \approx 3.590$.

Initial Proposal

1350

1351 1352

1353 1354

1355 1356

1357 1358

1359

1360

1361 1362 1363

1364

1365

1366

1367

1369

1370

13711372

137313741375

Problem. Arrange n points in $[0,1]^d$ to optimize the dispersion / packing–covering trade-off. The benchmark metric is

```
ratio = \frac{\min \ pairwise \ distance}{\max \ pairwise \ distance},
```

so that larger ratio is better (values in (0, 1]).

Evaluator. Given a program exposing max_min_dis_ratio (n,d), we obtain configurations for (n,d)=(16,2) and (14,3), then report ratio for each case.

Baseline algorithm. The initial program employs:

- Enhanced simulated annealing with adaptive cooling,
- Neighbor-repulsion moves,
- Periodic smoothing via k-NN weighted averages,
- A local refinement stage.

KD-tree acceleration is used for nearest-neighbor queries; hyperparameters adapt to dimension.

```
1376
        from scipy.spatial.distance import pdist
1377
        from scipy.spatial import cKDTree
1378
        # (Removed) smooth_points smoothing logic is now inlined to reduce indirection
1379
1380
        def calculate_distances(points):
1381
           """Calculates min, max, and ratio of pairwise Euclidean distances using scipy pdist."""
           if points.shape[0] < 2:</pre>
1382
              return 0.0, 0.0, 0.0
1383
           distances = pdist(points, metric='euclidean')
           eps = 1e-8
1384
           min_dist = max(np.min(distances), eps)
1385
           max_dist = np.max(distances)
           ratio = max_dist / min_dist
1386
           return min dist, max dist, ratio
1387
        # (Removed) perturb_point now inlined directly where used
1388
1389
        def update_temperature(temperature, cooling_rate, accept_history, iteration, total_iters,
1390
             initial_temperature, window_size=100):
1391
           Adaptive cooling with acceptancerate feedback and periodic reheating.
1392
           window = accept_history[-min(len(accept_history), window_size):]
1393
           rate = sum(window) / len(window)
           # gentler correction: slow/fast cooling factors reduced
1394
           if rate < 0.2:
1395
              adj = 1.02
1396
           elif rate > 0.8:
              adj = 0.98
1397
           else:
              adj = 1.0
1398
           temperature *= cooling_rate * adj
1399
           \# removed periodic reheating to maintain smoother cooling schedule
           \# if (iteration + 1) % (total_iters // 4) == 0:
1400
           # temperature = initial_temperature
1401
           return temperature
1402
        def max_min_dis_ratio(n: int, d: int, seed=None):
1403
           Finds n points in d-dimensional space to minimize the max/min distance ratio
```

```
1404
           using simulated annealing.
1405
1406
           Args:
              n (int): Number of points.
1407
              d (int): Dimensionality of the space.
1408
           Returns:
           tuple: (best_points, best_ratio)
"""
1409
1410
1411
           # Adaptive hyperparameters based on dimensionality
           iterations = 3000 if d <= 2 else 6000 # increased sweeps for improved convergence
1412
           initial_temperature = 10.0
1413
           cooling_rate = 0.998 if d <= 2 else 0.996 # slower cooling for extended exploration
           perturbation_factor = 0.15 if d <= 2 else 0.12 # tuned smaller steps in 3D for better
1414
                local refinement
1415
            # relaxation factor for post-acceptance repulsive adjustment
            # relaxation factor removed; using inline 0.1 * perturbation factor below
1416
1417
           # 1. Initial State: reproducible random generator
1418
           rng = np.random.default_rng(seed)
           # uniform random initialization in [0,1]^d for simplicity
1419
           current_points = rng.random((n, d))
1420
            _, _, current_ratio = calculate_distances(current_points)
1421
1422
           best_points = np.copy(current_points)
           best_ratio = current_ratio
1423
           temperature = initial_temperature
1424
           accept_history = []
1425
           window_size = 50 # window for stagnation detection and adaptive injection
            # smoothing_interval remains, but smoothing_strength is fixed inlined above
1426
           smoothing\_interval = max(10, iterations // (20 if d <= 2 else 30)) # more frequent
1427
                smoothing in 3D for improved uniformity
1428
           for i in range(iterations):
1429
              # Build KD-tree once per iteration for neighbor queries
              tree = cKDTree(current_points)
1430
              \# optional smoothing step using distance-weighted neighbor smoothing
1431
              if (i + 1) % smoothing_interval == 0:
                  # choose neighbor count based on dimension
1432
                 k_{smooth} = 6 if d > 2 else 4
1433
                  _, idxs = tree.query(current_points, k=k_smooth+1)
                  neighbors = current_points[idxs[:,1:]] # exclude self
1434
                  # compute inverse-distance weights
1435
                 diffs = neighbors - current_points[:, None, :]
                 dists = np.linalg.norm(diffs, axis=2) + 1e-6
1436
                  weights = 1.0 / dists
1437
                  weights /= weights.sum(axis=1, keepdims=True)
                 neighbor_means = (neighbors * weights[..., None]).sum(axis=1)
1438
                 blend = 0.6 \text{ if } d > 2 \text{ else } 0.7
1439
                 current_points = np.clip(current_points * blend + neighbor_means * (1 - blend), 0.0,
                      1.0)
1440
                     _, current_ratio = calculate_distances(current_points)
1441
                  if current_ratio < best_ratio:</pre>
                    best_points = current_points.copy()
1442
                    best_ratio = current_ratio
              # 2. Generate Neighboring State: Perturb a random point
1444
              # Simplify scaling: rely on temperature to adjust step-size instead of best_ratio
1445
              # dynamic perturbation decays sublinearly with temperature for finer local moves
              perturbation_strength = perturbation_factor * ((temperature / initial_temperature) **0.6
1446
                   + 0.15)
1447
              # Choose a random point to perturb
1448
              point_to_perturb_index = rng.integers(0, n)
1449
1450
              old_point = current_points[point_to_perturb_index].copy()
              # Increase repulsivemove frequency in low dimensions
1451
              # dynamic repulsion probability: stronger at high temperature, tapering off as we cool
1452
              if d > 2:
                 # reduce repulsion frequency in 3D for finer refinement
1453
                 repulsion_prob = float(np.clip(temperature / initial_temperature, 0.2, 0.8))
              else:
1454
                 repulsion_prob = float(np.clip(temperature / initial_temperature + 0.1, 0.5, 0.95))
1455
              # start with a random jitter
1456
              # random jitter inlined for readability
              candidate = old_point + rng.uniform(-perturbation_strength, perturbation_strength,
1457
                   size=old_point.shape)
              if n > 1 and rng.random() < repulsion_prob:</pre>
```

```
1458
                  # compute nearest neighbor via KD-tree for efficiency (reusing prebuilt tree)
1459
                  _, nn_idxs = tree.query(old_point, k=2)
                  nn_idx = nn_idxs[1]
1460
                  vec = old_point - current_points[nn_idx]
1461
                  norm = np.linalg.norm(vec)
                  if norm > 1e-8:
1462
                     dir_vec = vec / norm
1463
                    candidate = old_point + perturbation_strength * dir_vec
               # keep the point in [0,1]^d
1464
              current_points[point_to_perturb_index] = np.clip(candidate, 0.0, 1.0)
1465
               _, _, candidate_ratio = calculate_distances(current_points)
1466
               # Acceptance criterion
1467
               delta = candidate_ratio - current_ratio
               accept = (delta < 0) or (rng.random() < np.exp(-delta / temperature))
1468
1469
               if accept:
                  current_ratio = candidate_ratio
1470
                  # Post-acceptance repulsive relaxation to improve local spacing
1471
                  # reuse prebuilt KD-tree for repulsive relaxation
                  dists, idxs_nn = tree.query(current_points[point_to_perturb_index], k=2)
1472
                  dir_vec = current_points[point_to_perturb_index] - current_points[idxs_nn[1]]
                  norm = np.linalg.norm(dir_vec)
                  if norm > 1e-8.
1474
                     \# push away from nearest neighbor adjustment = 0.1 * perturbation_factor * dir_vec / norm
                     current_points[point_to_perturb_index] = np.clip(
1476
                        current_points[point_to_perturb_index] + adjustment, 0.0, 1.0
1477
                     # update ratio and best points after relaxation
1478
                     _, _, relaxed_ratio = calculate_distances(current_points)
1479
                     current_ratio = relaxed_ratio
                     if relaxed_ratio < best_ratio:</pre>
1480
                        best_points = current_points.copy()
1481
                        best_ratio = relaxed_ratio
                  # also keep the standard bestcheck for the candidate move
1482
                  if current_ratio < best_ratio:</pre>
1483
                     best_points = current_points.copy()
                     best_ratio = current_ratio
1484
1485
                  current_points[point_to_perturb_index] = old_point
1486
               # Update temperature with adaptive schedule
1487
               accept_history.append(accept)
               temperature = update_temperature(temperature, cooling_rate, accept_history, i,
1488
                   iterations, initial_temperature)
1489
               # periodic mild reheating for 3D to escape deep minima
               if d > 2 and (i + 1) % (iterations <math>// 3)
                  temperature = max(temperature, initial_temperature * 0.3)
1491
               # random injection to escape plateaus: reinitialize one point every 20% of iterations
1492
               # random injection only if weve stagnated (low acceptance in recent window)
1493
               if (i + 1) % max(1, iterations // 5) == 0 and len(accept_history) >= window_size \
                 and sum(accept_history[-window_size:]) / window_size < 0.1:</pre>
                  j = rng.integers(0, n)
1495
                  current_points[j] = rng.random(d)
                  _, _, current_ratio = calculate_distances(current_points)
1496
1497
           # Local refinement stage: fine-tune best solution with small Gaussian perturbations
           refine_iters = max(100, iterations // 20)
1498
            for _ in range(refine_iters):
1499
               idx = rng.integers(0, n)
              old_point = best_points[idx].copy()
1500
               perturb = rng.normal(0, perturbation_factor * 0.05, size=d)
1501
               best_points[idx] = np.clip(old_point + perturb, 0.0, 1.0)
               _, _, refined_ratio = calculate_distances(best_points)
1502
               if refined ratio < best ratio:
1503
                 best_ratio = refined_ratio
               else:
1504
                 best_points[idx] = old_point
1505
           return best points, best ratio
1506
```

D.6 Autoconvolution Peak Minimization (L^{∞}) .

1508 1509

1510 1511 **Problem Description** For nonnegative densities f supported on $[-\frac{1}{2}, \frac{1}{2}]$ with $\int f = 1$, define

$$\mu_{\infty} = \sup_{t} (f * f)(t).$$

The exact optimum is unknown.

Human Best:

$$0.64 \le \mu_{\infty} \le 0.75496.$$

The lower bound is due to Cloninger–Steinerberger, and the upper bound comes from explicit step-function constructions of Matolcsi–Vinuesa (rescaled to unit support).

Initial Proposal

Problem definition. Let

$$\mathcal{F} = \left\{ f \in L^1\left(\left[-\frac{1}{2}, \frac{1}{2}\right]\right) : f \ge 0, \int_{-1/2}^{1/2} f(x) \, dx = 1 \right\},$$

and define

$$(f * f)(t) = \int_{\mathbb{R}} f(x) f(t - x) dx.$$

We seek to minimize the peak value of the autoconvolution:

$$\mu_{\infty} = \inf_{f \in \mathcal{F}} \|f * f\|_{\infty}.$$

Constraints.

- Nonnegative density.
- Unit mass $(L^1 = 1)$.
- Support length 1 (here taken as $[-\frac{1}{2}, \frac{1}{2}]$).

In the implementation, f is represented by nonnegative step heights on a uniform grid and normalized to unit integral.

Optimization goal. Minimize

$$\mu_{\infty} = \max_{t} (f * f)(t).$$

Smaller values are better.

Best-known human results. In this standard setup, the best currently published bounds are

$$0.64 \le \mu_{\infty} \le 0.75496 \, .$$

The upper bound traces to work of Matolcsi–Vinuesa (after normalizing support length to 1), and the lower bound to Cloninger–Steinerberger.

Algorithmic goal. Create an algorithm that constructs feasible densities with progressively smaller μ_{∞} . The baseline program generates simple analytical candidates (box, triangle, cosine-squared, Gaussian) on a uniform grid, normalizes to unit mass, and computes autoconvolution via FFT to measure μ_{∞} . It serves as a starting point for more advanced search/optimization methods.

References.

- E. P. White, An optimal L^2 autoconvolution inequality, Canadian Mathematical Bulletin (2024).
- M. Matolcsi and C. Vinuesa, *Improved bounds on the supremum of autoconvolutions*, J. Math. Anal. Appl. 372 (2010), 439–447.
- A. Cloninger and S. Steinerberger, On suprema of autoconvolutions with an application to Sidon sets, Proc. Amer. Math. Soc. 145 (2017), 3191–3200.

```
1566
1567
         This program generates step heights for a probability density function
        that minimizes the maximum value of its autoconvolution.
1568
1569
         import numpy as np
1570
         from typing import Dict
1571
        def evaluate_C1_upper_std(step_heights: np.ndarray) -> Dict[str, float]:
1572
1573
            Standard-normalized C1 (support [-1/2, 1/2], dx=1/K).
            - Project to feasible set: h >= 0 and f = 1 (L1 normalization).
1574
            - Objective: mu_inf = max_t (f*f)(t) (smaller is better). Returns: {"valid", "mu_inf", "ratio"(=mu_inf), "integral"(=1.0), "K"}
1575
1576
            h = np.asarray(step_heights, dtype=float)
1577
            if h.size == 0 or np.any(h < 0):
    return {"valid": 0.0, "mu_inf": float("inf"), "ratio": float("inf")}</pre>
1578
            K = int(len(h))
1579
            dx = 1.0 / K
1580
            integral = float(np.sum(h) * dx)
1581
            if integral <= 0:</pre>
               return {"valid": 0.0, "mu_inf": float("inf"), "ratio": float("inf")}
1582
            h = h / integral # f = 1
1583
            F = np.fft.fft(h, 2*K - 1) # linear autoconvolution via padding
1584
            conv = np.fft.ifft(F * F).real
1585
            conv = np.maximum(conv, 0.0) # clamp tiny negatives
1586
            mu\_inf = float(np.max(conv) * dx)
1587
            return {"valid": 1.0, "mu_inf": mu_inf, "ratio": mu_inf, "integral": 1.0, "K": float(K)}
1588
         def make_candidate(K: int, kind: str = "cos2") -> np.ndarray:
1589
            Simple candidate builder on [-1/2, 1/2] (NOT normalized here).
1590
1591
               K: Number of discretization points
1592
               kind: Type of candidate function ("box", "triangle", "cos2", "gauss")
1593
1594
            Step heights array
1595
            x = np.linspace(-1.0, 1.0, K)
1596
            if kind == "box":
1597
              h = np.ones(K)
            elif kind == "triangle":
1598
              h = 1.0 - np.abs(x)
1599
               h[h < 0] = 0.0
            elif kind == "cos2":
1600
               h = np.cos(np.pi * x / 2.0) ** 2
1601
            elif kind == "gauss":
              h = np.exp(-4.0 * x**2)
1602
            else:
1603
               raise ValueError(f"unknown kind={kind}")
            return h
1604
1605
        def main():
1606
            Main function that generates step heights for autoconvolution minimization.
1607
            Returns:
1608
            numpy.ndarray: Step heights array
1609
            K = 128
1610
            kind = "cos2" # Change this to try different candidates (box/triangle/cos2/gauss)
1611
            step_heights = make_candidate(K, kind)
1612
            # Evaluate the result to verify it's valid
1613
            result = evaluate_C1_upper_std(step_heights)
            print(f"Generated {kind} candidate with K={K}, mu_inf={result['mu_inf']:.6f}")
1614
1615
            return step heights
1616
```

D.7 THIRD AUTOCORRELATION INEQUALITY.

Problem Description Let C_3 be the largest constant such that $\max_{|t| \le 1/2} |(f * f)(t)| \ge C_3 (\int_{-1/4}^{1/4} f)^2$ for all (signed) f.

Best-known: classical 1.4581 upper bound.

D.8 Third-Order Autocorrelation Inequality (C_3 Upper Bound)

Initial Proposal

1620

1621 1622

1623

1624 1625

1626 1627

1628 1629

1630 1631

1632

1633 1634

1635 1636 1637

1638

1639

1640

1641

1642

1643

1644 1645

1646 1647 **Problem.** For piecewise-constant nonnegative functions on a fixed support with unit mass, we evaluate an upper bound $C_{\text{upper_bound}}$ derived from the maximum of the autoconvolution (normalized by squared L^1 mass). The benchmark score is

$$\mathrm{score} \ = \ \frac{1}{C_{\mathrm{upper_bound}}},$$

so that larger score indicates a smaller upper bound and hence a better result.

Evaluator. The evaluator calls find_better_c3_upper_bound() from the target program to obtain step heights, computes the normalized autoconvolution maximum, and returns $1/C_{\rm upper_bound}$.

Baseline algorithm. A simple genetic algorithm over height sequences serves as the baseline search method. The algorithm includes:

- Tournament selection,
- One-point crossover,
- Gaussian mutation.

```
1648
         import scipy.integrate
1649
1650
         def calculate_c3_upper_bound(height_sequence):
1651
            N = len(height_sequence)
            delta_x = 1 / (2 * N)
1652
1653
            def f(x):
               if -0.25 <= x <= 0.25:
1654
                  index = int((x - (-0.25)) / delta_x)
1655
                   if index == N:
                      index -= 1
1656
                   return height_sequence[index]
1657
               else:
                  return 0.0
1658
1659
            integral_f = np.sum(height_sequence) * delta_x
1660
            integral\_sq = integral\_f**2
1661
            if integral_sq < 1e-18:</pre>
1662
               return 0.0
1663
            t_{points} = np.linspace(-0.5, 0.5, 2 * N + 1)
1664
            max\_conv\_val = 0.0
1665
            for t_val in t_points:
1666
               lower_bound = \max(-0.25, t_val - 0.25)
upper_bound = \min(0.25, t_val + 0.25)
1667
               if upper_bound <= lower_bound:</pre>
1669
                   convolution\_val = 0.0
               else:
1670
                  def integrand(x):
1671
                      return f(x) * f(t_val - x)
1672
                   convolution_val, _ = scipy.integrate.quad(integrand, lower_bound, upper_bound,
1673
                        limit=100)
```

```
1674
              if abs(convolution_val) > max_conv_val:
1675
                 max_conv_val = abs(convolution_val)
1676
           return max_conv_val / integral_sq
1677
        def genetic_algorithm(population_size, num_intervals, generations, mutation_rate,
1678
             crossover_rate):
1679
           population = np.random.rand(population_size, num_intervals) * 2 - 1
1680
1681
           best_solution = None
           best_fitness = 0.0
1682
1683
           for gen in range (generations):
1684
              fitness_scores = np.array([calculate_c3_upper_bound(individual) for individual in
1685
                   population])
1686
              current best idx = np.argmax(fitness scores)
1687
              if fitness_scores[current_best_idx] > best_fitness:
                 best_fitness = fitness_scores[current_best_idx]
1688
                 best_solution = population[current_best_idx].copy()
1689
                 # print(f"Generation {gen}: New best fitness = {best_fitness}")
1690
1691
              new_population = np.zeros_like(population)
              for i in range(population_size):
1692
1693
                 competitors_indices = np.random.choice(population_size, 2, replace=False)
                 winner_idx = competitors_indices[np.argmax(fitness_scores[competitors_indices])]
1694
                 new_population[i] = population[winner_idx].copy()
1695
              for i in range(0, population_size, 2):
1696
                 if np.random.rand() < crossover_rate:</pre>
1697
                    parent1 = new_population[i]
                    parent2 = new_population[i+1]
1698
                     crossover_point = np.random.randint(1, num_intervals - 1)
1699
                    new_population[i] = np.concatenate((parent1[:crossover_point),
                         parent2[crossover_point:]))
1700
                    new_population[i+1] = np.concatenate((parent2[:crossover_point],
1701
                         parent1[crossover_point:]))
1702
              for i in range(population_size):
1703
                 if np.random.rand() < mutation_rate:</pre>
                    mutation_point = np.random.randint(num_intervals)
1704
                    new_population[i, mutation_point] += np.random.normal(0, 0.1)
1705
                    new_population[i, mutation_point] = np.clip(new_population[i, mutation_point],
1706
                         -2, 2)
1707
              population = new_population
1708
1709
           return best_solution
1710
        def find_better_c3_upper_bound():
1711
           NUM_INTERVALS = 4
1712
           POPULATION_SIZE = 2
1713
           GENERATIONS = 10
           MUTATION_RATE = 0.1
1714
           CROSSOVER_RATE = 0.8
1715
           height_sequence_3 = genetic_algorithm(POPULATION_SIZE, NUM_INTERVALS, GENERATIONS,
1716
                MUTATION_RATE, CROSSOVER_RATE)
1717
           return height_sequence_3
1718
1719
```