
Is Programming by Example solved by LLMs?

Wen-Ding Li
Cornell University
wl678@cornell.edu

Kevin Ellis
Cornell University
kellis@cornell.edu

Abstract

Programming-by-Examples (PBE) aims to generate an algorithm from input-output examples. Such systems are practically and theoretically important: from an end-user perspective, they are deployed to millions of people, and from an AI perspective, PBE corresponds to a very general form of few-shot inductive inference. Given the success of Large Language Models (LLMs) in code-generation tasks, we investigate here the extent to which LLMs can be said to have ‘solved’ PBE. We experiment on classic domains such as lists and strings, and an uncommon graphics programming domain not well represented in typical pretraining data. We find that pretrained models are not effective at PBE, but that they can be fine-tuned for much higher performance, provided the test problems are in-distribution. We analyze empirically what causes these models to succeed and fail, and take steps toward understanding how to achieve better out-of-distribution generalization. Collectively these results suggest that LLMs make strong progress toward solving the typical suite of PBE tasks, potentially increasing the flexibility and applicability of PBE systems, while also identifying ways in which LLMs still fall short.

1 Introduction

Programming-by-Example (PBE) systems solve a challenging task: Given input-output examples of a hidden algorithm, they seek to construct the source code of the underlying function [1, 2]. PBE is deployed to millions of users [3, 4, 5, 6], lies near the heart of core AI challenges [7, 8, 9, 10], and is a qualitatively different problem from the bulk of recent work on LLM code generation, because rather than generate source code from natural language [11], PBE is instead fundamentally about few-shot inductive inference: Given a handful of examples, inferring the program that will generalize to new inputs, or which captures the ‘true’ latent regularity, *without* relying on natural-language guidance.

We investigate here the extent to which large language models pretrained on source code can solve PBE. If they can, this unlocks the ability to do PBE in general-purpose Turing complete languages like Python, unlike the restricted domain-specific languages which have so far dominated PBE [4, 12, 13, 14, i.a.], thereby increasing the scope and power of this paradigm. If LLMs cannot perform PBE, then this highlights a deficit of inductive reasoning and problem solving, and suggests LLMs lean too heavily on natural language cues to generate code.

We find that pretrained and instruction-tuned models serve as poor PBE systems, a finding also supported by recent work [15, 16, 12, 17]. But our investigation further finds that LLMs can be fine-tuned for significantly higher performance, provided they are not asked to generalize far beyond the fine-tuning data. To address this failure of generalization we give an algorithm for taking a small unlabeled dataset of problems and adapting the LLM to it, which we find narrows this domain gap.

The resulting recipe allows PBE over Turing-complete languages across three qualitatively different domains (Fig. 1): algorithms on vectors of numbers, string manipulation macros, and graphics programs in LOGO/Turtle. In every case, our final model is at least as effective as custom symbolic search algorithms operating over domain-specific languages, and surpasses powerful closed-source

models such as GPT4 [18]. We also find that the resulting system can cover a broader scope of problems than classic symbolic methods, owing to the use of a Turing-complete language, which, at least theoretically, allows learning any computable function.

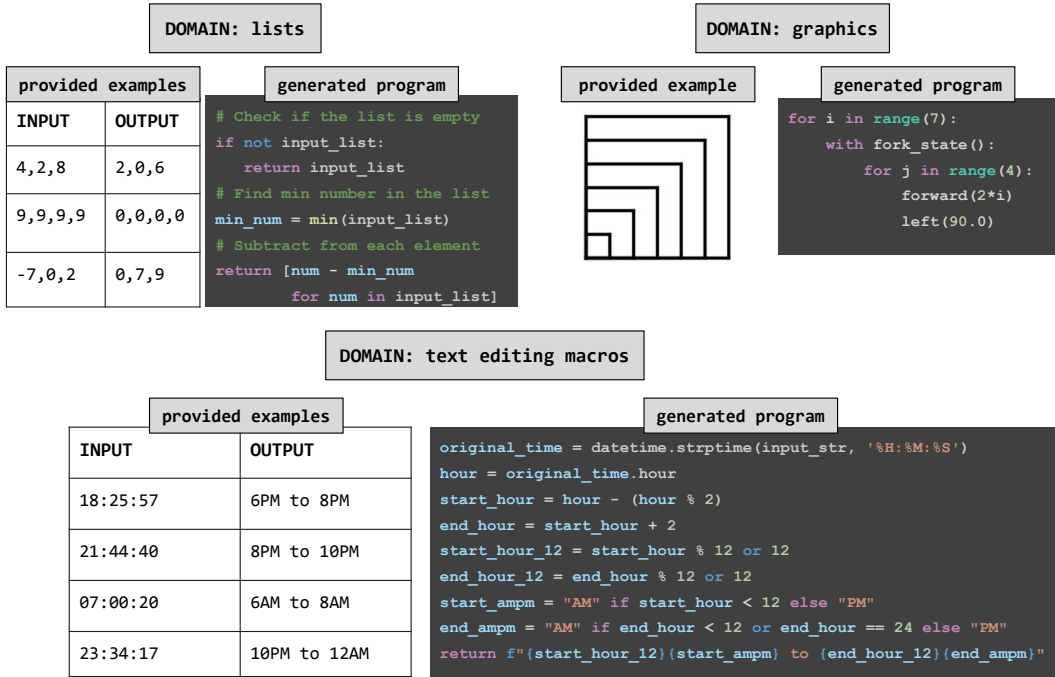


Figure 1: Domains, including standard ones that resemble programs found in pretraining data, as well as a less common graphics domain, which is likely less represented in LLM pretraining data.

2 Background

Programming by Example considers synthesizing a program ρ given a vector of inputs X and corresponding outputs Y . Typically the program is expected to exactly fit the provided examples, $\rho(X_i) = Y_i, \forall i$, where i indexes examples. The program ρ is drawn from a (potentially infinite) language \mathcal{L} . Typically \mathcal{L} is a domain-specific language designed for a specific PBE system, not a general-purpose programming language. For example, the PBE system FlashFill synthesizes string manipulation macros designed to automate common spreadsheet edits [2]. FlashFill’s domain-specific language \mathcal{L} includes commonly occurring regular expressions, together with string slicing and concatenation, and restricted forms of loops. The language \mathcal{L} is also designed to allow polynomial-time construction of programs consistent with input-output examples. FlashFill’s goal, like most PBE systems, is to generalize to hold-out test cases: inputs X' with (hidden) target outputs Y' .

$$\text{Pick a program } \rho \text{ from } \{\rho \in \mathcal{L} : \rho(X_i) = \rho(Y_i), \forall i\}. \quad \text{Succeed if } \rho(X'_j) = \rho(Y'_j), \forall j \quad (1)$$

In its simplest forms, PBE can be accomplished by guess-and-check enumeration until a program is found that is consistent with the examples. Although there exist more sophisticated search algorithms, including those accelerated by neural guidance [19, 20, 21, 22], a key enabler of practical PBE systems is the design of a carefully restricted domain-specific language \mathcal{L} . The domain-specific language effectively hardcodes symbolic knowledge, focusing the system on what programs the human engineer thinks are most promising, but at the expense of the wider set of computable functions expressible in general-purpose languages.

The PBE setup covers other cases as well, such as sequence extrapolation (the inputs are indices into the sequence), as well as data compression (the input is null, and the data is compressed by synthesizing a program that reproduces the output data). Therefore, a truly general solution to PBE—one which could express its solutions in general purpose programming languages, and cover most

practically relevant problems—would be broadly applicable to many inductive inference problems, a point that has been long appreciated [9].

LLMs for solving programming problems have been recently very successful [11, 23, 24, 25, 26]. These systems typically input a prompt describing a problem in natural language, then sample candidate programs, and optionally filter those samples by checking them against input-output test cases, with the goal of passing holdout tests:

Draw $\rho_k \sim p_{LM}(\cdot|\text{prompt})$. Pick a $\rho \in \{\rho_k : \rho_k(X_i) = \rho_k(Y_i), \forall i\}$. Success: $\rho(X'_j) = \rho(Y'_j), \forall j$

Unlike PBE, the primary driver of program generation is a natural language prompt, although input-outputs may also be in the prompt [27, 28]. Recent work using LLMs to synthesize programs solely from examples has either obtained negative results [16, 12], or focused on simple and/or nonstandard problems [29, 30, 31], leaving the extent to which PBE is ‘solved’ by LLMs an open question.

3 Methods

Basic prompting is the most straightforward way of performing PBE with a pre-trained model: Given input-output examples (X, Y) a prompt is constructed and K programs are generated. Programs are filtered by the I/O examples, and a random satisfying program is returned:

$$\text{Sample } \rho_k \sim p_{LM}(\cdot|\text{prompt}(X, Y)), \text{ for } k \text{ from } 1..K \quad (2)$$

$$\text{Pick a } \rho \text{ from } \{\rho_k : \rho_k(X_i) = \rho_k(Y_i), \forall i\} \quad (3)$$

Fine-tuning improves the above approach in a conceptually straightforward way. Given a dataset comprising tuples of programs and I/O examples, $\{(\rho, X, Y)\}$, we fine-tune the LM to predict a program from its input-outputs. But this dataset is hard to come by: Although there are web-scale corpora of naturally occurring source code, there is no analogous dataset of runnable code snippets paired with representative input-output examples, and this data deficit is especially true for new or unusual applications of PBE, such as the graphics programs we consider.

To assemble a large dataset of (ρ, X, Y) triples we start with a small manually-constructed seed dataset, $\mathcal{D}_{\text{seed}}$, and then randomly generate new programs ρ and inputs X by prompting an LLM with members of $\mathcal{D}_{\text{seed}}$. The output Y comes from running ρ on X . The seed dataset effectively defines a prior over (ρ, X) , notated \mathcal{G} in Fig. 2. We sample from \mathcal{G} to collect many program-input pairs, but use program execution to predict Y , not an LLM. The resulting dataset, which we call $\mathcal{D}_{\text{tune}}$, is used to train an LLM to generate programs when prompted with input-outputs. As this fine-tuned LLM effectively learns to do probabilistic inference in the graphical model shown in Fig. 2 (right), we write this fine-tuned LLM as $q_\theta(\rho|X, Y)$. This inference network is trained to maximize

$$\max_{\theta} \log q_\theta(\rho|X, Y), \text{ where } (\rho, X) \sim \mathcal{G}(\mathcal{D}_{\text{seed}}) \text{ and } Y = \rho(X) \quad (4)$$

This method is closely related to self-instruct [32] and wake-sleep [33]. Like self-instruct, we use prompting to bootstrap a large dataset from a small manually-constructed one. Our method differs by using the LLM to generate a hidden latent variable (the program) while a different generative process produces an observed variable (the program outputs). Like wake-sleep, we use samples from a generative model to train an inference network, but we do not further train the generative model itself. Next, we will see that bringing the method much closer to wake-sleep by updating the generative model plays an important role when deploying the system on out-of-distribution problems.

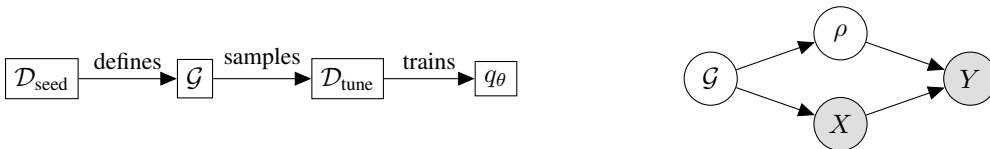


Figure 2: Left: Data generation pipeline. Right: The fine-tuned network q_θ learns to do inference in a graphical model where the prior over programs, \mathcal{G} , is defined by prompting an LLM with example code in $\mathcal{D}_{\text{seed}}$, while the likelihood $p(Y|\rho, X)$ is defined by program execution.

Adaptation. One of the most powerful features of source code as a representation is its ability to efficiently express a wide range of computations. Therefore it is of interest to study the ability of fine-tuned LLMs to extrapolate to PBE problems outside the distribution of the fine-tuning data.

We consider a basic approach to adapting to a different distribution of problems, assuming access to problems drawn from the testing distribution, but without labeled program solutions. This mimics the deployment of PBE systems to end-users who may have their own idiosyncratic distribution of problems they care about, and who do not provide ground-truth programs, but who can provide feedback on if a generated program has correct behavior. This means we have an unlabeled dataset $\mathcal{D}_{\text{adapt}}$ comprising input-outputs (X, Y) , as well as a labeled seed dataset $\mathcal{D}_{\text{seed}}$ comprising triples (ρ, X, Y) . Adaptation proceeds by iterating between pretraining with $\mathcal{G}(\mathcal{D}_{\text{seed}})$, testing on $\mathcal{D}_{\text{adapt}}$, and adding back into $\mathcal{D}_{\text{seed}}$ any program solutions found on the adaptation problems, which then become seeds for the next iteration. This produces a sequence of fine-tuned models, indexed below by i :

$$\begin{aligned}
 \text{train model:} \quad & \theta^i = \arg \max_{\theta} \log q_{\theta}(\rho|X, Y), \text{ where } (\rho, X) \sim \mathcal{G}(\mathcal{D}_{\text{seed}}^i) \text{ and } Y = \rho(X) \\
 \text{run inference:} \quad & \rho_k^{X,Y} \sim q_{\theta^i}(\rho|X, Y) \text{ for } (X, Y) \in \mathcal{D}_{\text{adapt}} \text{ and } k \text{ from } 1..K \\
 \text{update seed:} \quad & \mathcal{D}^{i+1} = \mathcal{D}^i \cup \left\{ (\rho_k^{X,Y}, X, Y) : (X, Y) \in \mathcal{D}_{\text{adapt}}, k \in [K] \text{ if } \rho_k^{X,Y}(X) = Y \right\}
 \end{aligned} \tag{5}$$

The equations can be seen as a wake-sleep algorithm where “dreaming” corresponds to training q on fantasy data (first equation) while “waking” corresponds to running inference and updating the generative model \mathcal{G} (by updating the seed, second pair of equations). Ideally, each cycle of this wake-sleep adaptation solves more out-of-distribution problems, which tugs the generative model \mathcal{G} toward the target distribution, unlocking solutions to more out-of-distribution problems, etc. This hinges on each iteration actually solving new problems from the unlabeled dataset. Theoretically this is guaranteed given enough inference-time compute (large K above). We explore in Sec. 4.3 the extent to which this holds in practice.

4 Experiments

We study different LLM-approaches to programming-by-examples across three domains (Fig. 1):

1. **List functions** is a PBE domain meant to model a “programmer’s assistant”. It concerns discovering algorithms that transform lists of numbers, given input-output examples. This problem statement has a long history within program synthesis [13, 34], and was popularized within machine learning by DeepCoder [35]. We consider two modern list function datasets created by Rule et al. 2024 [17] and Shi et al. 2023 [12], which both involve higher-order functions and nontrivial procedures such as map, filter, and sort. Rule et al. was recently added to BigBench [36].
2. **Text editing** is a domain where a program synthesizer assists an end-user edit their spreadsheets or other documents. From string-to-string examples, the system generates edit macros for tasks such as reformatting dates, extracting fields from semistructured text, etc. [2, 37, 38, 4]. Text editing is the most prominent commercial success of PBE: The FlashFill PBE system ships in Microsoft Excel and is used by many millions of people [6]. We consider two text editing datasets: SyGuS problems [22]—which are easier—and PROSE [39] problems, which constitute the most challenging dataset of its kind [38].
3. **LOGO/Turtle graphics** is a domain whose goal is to synthesize a program that generates a target image.¹ Systems of this kind can be used both for high-level visual reasoning and for helping artists make structured edits to images [40, 41]. We use a dataset of geometric designs expressed as LOGO/Turtle [42] programs—where the programs move a simulated pen over a canvas—taken from Wong et al. [43]. To allow the LLM to visually perceive the input image, we convert the image to ASCII-art style strings; see Fig. 5 and Appendix. A.1.3.

¹This is PBE with a single example and null input, effectively compressing the image into a program.

4.1 How well does the fine-tuned model perform?

We prepare seed datasets for each domain, synthetically generate a large training set, and then fine-tune a DeepSeekCoder LLM [44] that was pretrained on source code.² For list functions we seed with 50 problems from Rule et al. 2024; For text editing, we consider seeding with either SyGuS or a 40-problem subset of PROSE; for LOGO we seed with 200 training-set problems in Wong et al. [43].

The resulting fine-tuned models are surprisingly effective within their respective PBE domains. On list functions our finetuned model surpasses the best symbolic search baselines reported in Rule et al. (Fig. 3a), surpasses the best neurosymbolic search method from Shi et al. (Appendix Fig. 10), and surpasses GPT4. It also solves 100% of the list to list benchmark problems from λ^2 (a well-known symbolic synthesizer), shown in Appendix Tbl. 4: although plausibly, many λ^2 problems are in the pretraining data. On text editing, it surpasses the performance of FlashFill and approaches the level of FlashFill++ (Tbl. 1, Fig. 3b). On LOGO, it solves 90% of the test set (Fig. 3c), surpassing systems such as DreamCoder [45], which introduced the first version of these LOGO problems. It also solves more problems than LILO and Regal [43, 46], which are LOGO program synthesizers that input natural language describing how the image should be drawn. In contrast, our model does not use any language clues, generating purely from the image. In addition to quantitatively solving more problems, we note that there are qualitative improvements to the breadth of problems that can be solved in the first place because the LLM can generate Turing-complete code spanning a much broader space of computations (Fig. 4).

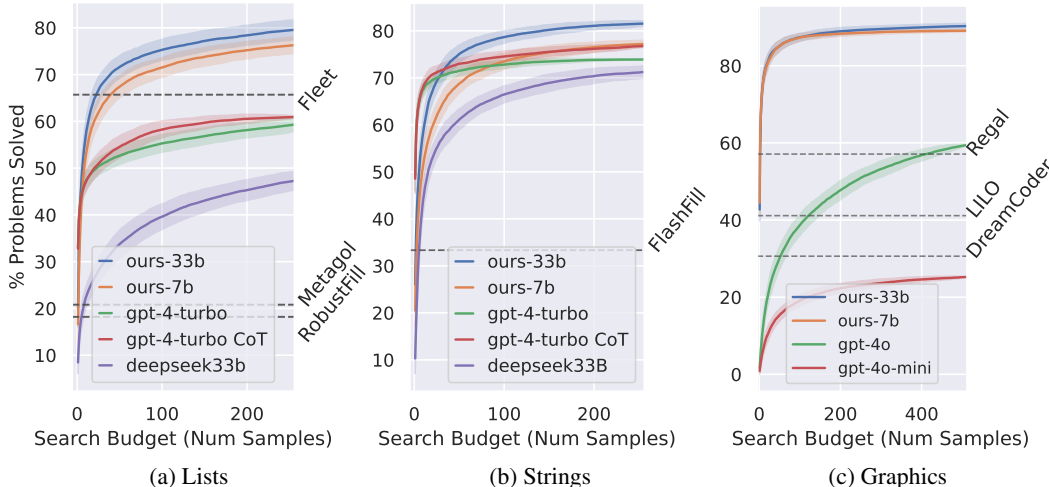


Figure 3: Test set performance. A problem is solved if the predicted program generates correct outputs on the holdout inputs. Metagol [47], RobustFill [20], and Fleet [48] results taken from [17]

There are caveats to the above results. First, the fine-tuned model essentially never produces a correct program on the first try: It requires tens or hundreds of samples, each of which is compared against the ground-truth input-outputs, and discarded if it contradicts the examples. On a GPU like the one we use (an Nvidia A6000) this rejection sampling takes on the order of a few minutes to solve a given problem. However, compared to classic enumerative program synthesizers [13, 49], or even compared to those with neural guidance [35, 22], proposing a few thousand programs is relatively little, and could not plausibly cover a significant fraction of the exponentially large search space.

	gen. accuracy	oracle accuracy
FlashFill	33%	—
FlashFill++	—	≈100%
ours, 33B	82%	88%

Table 1: Generalization accuracy: % problems where the program makes correct predictions on every holdout test. Oracle accuracy: % problems where a correct program was generated (even if incorrect programs were also generated that also passed the training input-outputs). FlashFill++ [38] only reports oracle accuracy.³

²We prefer DeepSeek because it is roughly LLaMA-level, but has fully open training details.

³The FlashFill results were obtained using Microsoft Excel for Mac.

INPUT	OUTPUT	generated program
Mary had a little lamb Its fleece was white...	1:Mary had a little lamb 2:Its fleece was white...	<pre>lines = text.splitlines() result = "" for i, line in enumerate(lines): result += f"{i + 1}:{line}\n" return result[:-1]</pre>
Twinkle, twinkle, ... How I wonder what you... Up above the world so... Like a diamond in the...	1:Twinkle, twinkle, ... 2:How I wonder what you... 3:Up above the world so... 4:Like a diamond in the...	

INPUT	OUTPUT	generated program
NY	New York	<pre>states = { "AL": "Alabama", "AK": "Alaska", }</pre>
CA	California	
AK	Alaska	

Figure 4: PBE with LLMs allows using general-purpose programming languages which can mix string and numerical operations in ways not allowed by domain-specific languages [38] (top), and allows world knowledge to inform code generation (bottom). I/Os and code partly elided for space.

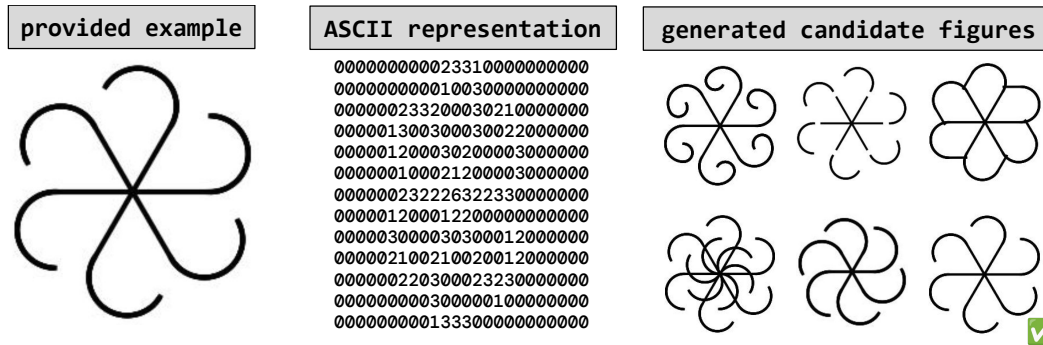


Figure 5: ASCII representation of LOGO graphics. Average pixel intensity indicated by numbers 0-9

The second caveat is that the model degrades when tested out-of-distribution. An example of this degradation is illustrated in Fig. 6, which tests the LOGO graphics model on hand drawings (after training on clean computer graphics). On the out-of-distribution hand drawing the model mostly samples programs that do not fit the data, but its accuracy does not fall to zero, meaning that with enough compute budget, it does actually generate reasonable programs. This foreshadows the results in Sec. 4.3, which more systematically studies out-of-distribution behavior.

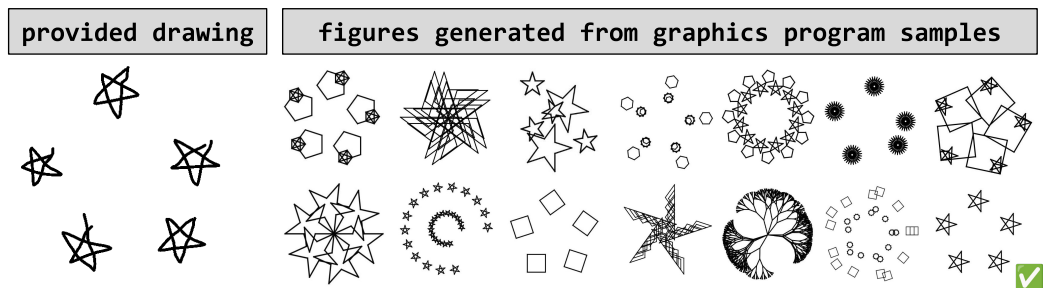


Figure 6: Example out-of-distribution LOGO test: inferring a graphics program from a hand drawing. See also Appendix Fig. 12

4.2 What causes the fine-tuned model to succeed or fail?

Classic symbolic approaches to PBE, when they are based on enumeration, tend to succeed whenever the target program is syntactically small. Approaches based on clever dynamic programming, such as the FlashFill family [4], succeed when the program is representable in the domain-specific language. What predicts success for these LLM approaches?

To answer this question we investigate several hypotheses. First, potentially the success is determined by program size, and degrades as programs grow longer. Second, as a more refined notion of size, we instead measure the description length *under the prior*, which for a program ρ , is $-\log p_{\text{LM}}(\rho|\mathcal{G}(\mathcal{D}_{\text{seed}}))$. Description length under the prior would be a good predictor of success if the fine-tuned model engages in blind guess-and-check: simply learning the distribution $\mathcal{G}(\mathcal{D}_{\text{seed}})$, and sampling from this prior while ignoring the input-outputs. Third, one possibility is that success is predicted by description length *under the approximate posterior* ($-\log q_{\theta}(\rho|X, Y)$), which would be the case if the fine-tuned model attends closely to the input-outputs and reshapes its distribution accordingly, instead of defaulting to the prior. To test these hypotheses we calculate the average compute budget needed to solve each problem, and compare it with these different variables. Fig. 7 shows that posterior description length is more predictive than program size and prior description length: unlike classical methods, metrics of program length correlate poorly with problem difficulty, and there is no evidence that the fine-tuned model’s behavior can be characterized as blind guess-and-check. (See also Fig. 5).

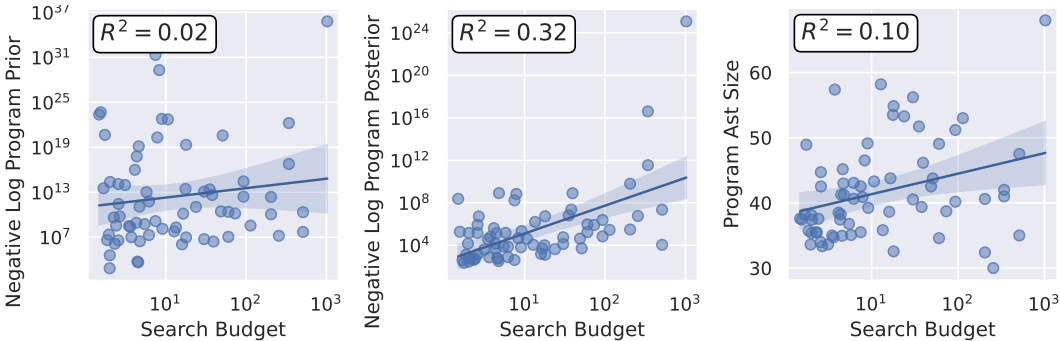


Figure 7: Compute budget needed to solve a problem is best predicted by description length under the approximate posterior, *not* program size or prior description length, suggesting that the fine-tuned model is not engaging in blind guess-and-check.

4.3 Out-of-distribution generalization

One advantage of classic symbolic PBE methods is that they do not make statistical assumptions about their test problems. Indeed, some classic methods can, within their domains, synthesize programs perfectly (i.e. always find a program that fits the training input-outputs). In contrast, neural networks can struggle to generalize beyond the training distribution.

We therefore consider train/test splits that force the model to generalize beyond the distribution of its training data (beyond $\mathcal{D}_{\text{seed}}$). On text editing, we seed with SyGuS problems, and perform out-of-distribution testing on PROSE problems (PROSE is much harder than SyGuS). On list functions, we seed with problems from Rule et al. 2024 and test on Shi et al. 2023 (the Shi dataset contains unusual combinators, such as Scan). On LOGO, we seed with short programs (≤ 12 lines of code), and test on long programs (> 12 lines of code). Using these splits we also measure the ability of the adaptation method in Sec. 3 to improve out-of-distribution generalization.⁴

Fig. 8 shows that there is nontrivial degradation when testing out of distribution. For example, a 7B model seeded with PROSE problems and tested on a different subset of PROSE has an accuracy of 76% (Fig. 3b), but this degrades to 59% when seeded with SyGuS problems, which follow a different distribution and are generally simpler and easier than PROSE (Fig. 8b).

⁴We work here with 7B models because Sec. 4.1 found that fine-tuned 33B models are only slightly better than 7B, and 7B is cheaper to run.

We further perform the adaptation method described in Sec. 3 in order to measure the extent to which it can narrow these domain gaps. In every case it allows solving more out-of-distribution problems, increasing absolute performance by around 10% or more in all domains, which is a relative increase of about 16% for text/list and a relative increase of about 190% for LOGO (approximately tripling the number of solved LOGO problems).

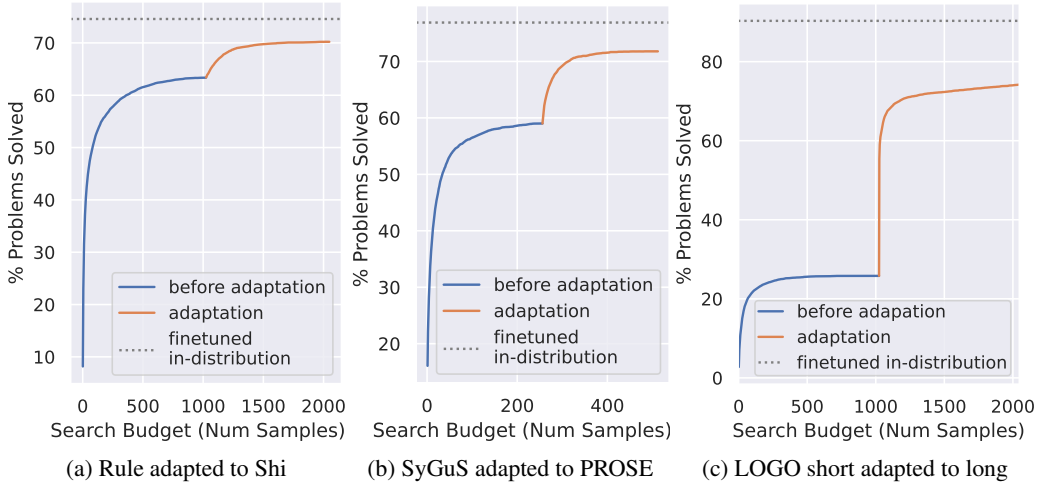


Figure 8: Out-of-distribution generalization and adaptation to new test distribution.

To better understand the dynamics of adaptation, we visualize the specific problems solved before and after adaptation on LOGO graphics (Fig. 9). Before adaptation, only a handful of out-of-distribution problems are solvable, and only with a significant search budget. Adaptation allows the system to quickly solve similar out-of-distribution problems in the future, but does not allow the system to generalize to problems very unlike those originally solvable by the fine-tuned model. In principle, expanding the inference-time compute budget should allow successful adaptation (large K in Eq. 5). Another more compute-efficient approach would be to increase the amount of adaptation data by introducing ‘steppingstone’ problems in the adaptation set that give a gentler transition from the original training distribution.

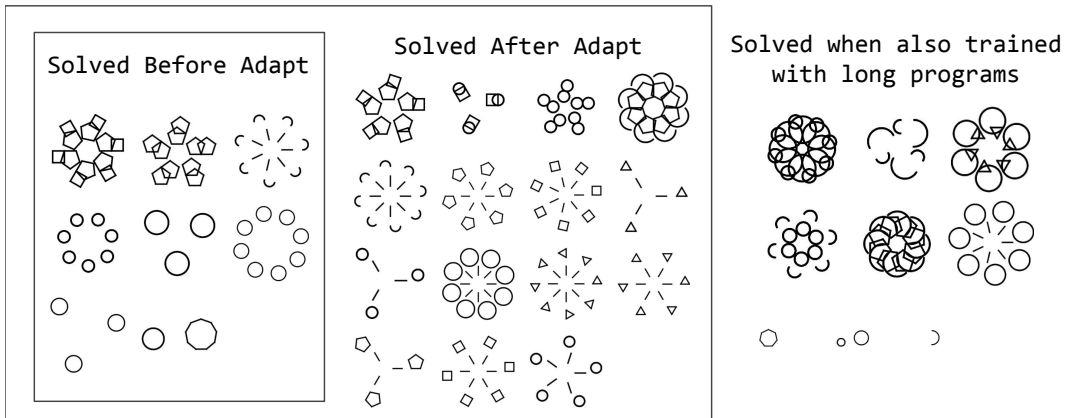


Figure 9: Out-of-distribution LOGO problems (requiring long programs with > 12 lines of code). We show example problems that are solved by the original model fine-tuned on short programs, which then become training data for the next round of adaptation. Adaptation allows consistent solving of problems similar to those that the original fine-tuned model could sometimes solve, but is not a panacea: Problems dissimilar to those solved by the initial model are not ever correctly generated, despite the fact that they are solvable by a model fine-tuned in-distribution.

5 Related Work

Automatic data generation with LLMs, such as self-instruct [32], WizardCoder [50], and many others [51, 52, 53, i.a.], works by prompting an LLM to produce outputs which are then used for later learning stages such as fine-tuning. These approaches are applied recursively to their own output: Previously generated data is incorporated into future prompts. We similarly generate a dataset $\mathcal{D}_{\text{tune}}$ by prompting an LLM with $\mathcal{D}_{\text{seed}}$, but (1) do not recursively prompt the LLM with its own outputs and (2) combine the LLM generations with program execution to make program outputs. This gives a different mathematical interpretation to our data generator. First, the programs are samples from a prior, \mathcal{G} , defined by $\mathcal{D}_{\text{seed}}$, which would not be a valid interpretation if the LLM was repeatedly fed its own outputs. Second, there is an observation model or likelihood function, $p(Y|\rho, X)$, which is defined not by the LLM, but by a Python interpreter. In this way, our data generator constructs training examples for fine-tuning that teach the network how to ‘invert’ the execution process of the Python interpreter.

Machine learning applied to PBE has often sought to accelerate search: to find *any program at all* consistent with the I/O examples [19, 20, 45, 21, 35, 22, 12, 54, 55, 56], which is nontrivial due to the combinatorial nature of the search, even after confining to a domain-specific programming language. A complementary line of research explores inductive biases that favor programs likely to generalize to new inputs, such as learning a prior or ranking function [57, 41, 58, 59]. Our work should be seen within the tradition of learning to search for programs. We show that finetuned models serve as an effective yet simple foundation for accelerating search in PBE, allowing search to be tractable over much richer and more expressive languages such as Python.

Classic PBE. Traditional approaches to programming-by-examples operate by symbolically searching or solving for programs consistent with the input-output examples [13, 49, 2, 1, 37, 6, i.a.]. They use domain-specific programming languages that are designed to either enable efficient search and/or bias the system toward functions that are likely to generalize new inputs. Search for programs can even be polynomial time when this domain-specific language has a special structure (roughly, when every function can be ‘inverted’), a key enabler of FlashFill, the first commercial success of PBE [4, 2].

LLMs as inductive reasoners. Using an LLM to perform inductive reasoning—to generate abstract hypotheses from concrete specific examples—has been explored by several recent works [29, 30, 60, 61], all of which has found significant value in translating these hypotheses into programs, and all of which have worked by prompting pretrained GPT-style models. Our work can be seen as helping answer a natural question posed by these previous works: Given that LLMs can generate hypotheses from examples, can they produce programs of the nature and complexity demanded by PBE? We find this is largely the case after fine-tuning, both for classic PBE domains and unusual ones.

Self-Debugging, Refinement, and Self-repair. One way of improving the code generation abilities of an LLM is to have it attempt to debug its own code whenever the initially generated code does not pass the provided test cases [62, 63, 64, 65, 66, 24, 67]. We did not explore this strategy, however, because a more basic approach that simply regenerated a new program from scratch already surpassed the prior state of the art (both symbolic and neural baselines), provided we finetune. However, further pushing the boundary of PBE may benefit from self-debugging strategies.

Ranking LLM-generated code. Past work considers a variety of ways to select an output from a collection of LLM-sampled programs [23, 59, 68, 11], many of which are more sophisticated than simply filtering by the examples, which is what we do here. Like with self-debugging, integrating these techniques should be synergistic with our approach.

6 Limitations

Our work has important limitations. From an engineering perspective, using a 7B-33B neural network to perform PBE is not practical for most end-users, who may be doing PBE on their laptop or desktops in order to accomplish small one-off tasks. For this reason, true deployment to end-users may require

investigating the effectiveness of much smaller neural networks (not an LLM), and it may also be valuable to study the effect of network compression and distillation upon our finetuned models.

From the perspective of understanding where and why our system succeeds and fails, we have shown that neither program size nor likelihood under the prior suffice to predict success, finding the posterior likelihood is a better predictor, albeit an imperfect one. Although this allows us to discard the hypothesis that the system is merely sampling from the prior, it also just pushes the question back one stage further: What exactly about specific problems causes the neural network’s approximate posterior to put more or less probability mass on correct solutions? While in classic PBE one can obtain sharp answers as to why a certain problem was solved or not, this is a much harder question with neural networks, whose workings are more opaque.

7 Discussion

PBE with fine-tuned LLMs is surprisingly effective, surpassing many of the best neural and symbolic baselines we know of, even for uncommon domains such as LOGO graphics. Why is that? Fundamentally, the neural network only needs to act as a heuristic proposer of solutions, because we can check against the input-outputs. Therefore, one possible explanation is that the tendency of language models to over-generate, hallucinate, and cover the long tail of possibilities is actually an asset, instead of a liability. And although there is a degree of degradation on out-of-sample problems, the degradation is not so severe that out-of-distribution problems become utterly unsolvable: Instead, they merely become harder to solve, a phenomenon that allows adaptation to work in the first place.

Simultaneously one should be hesitant about claiming that PBE is ‘solved.’ Optimistically, current PBE benchmarks exist to test the frontier of what is possible, and so doing well on those benchmarks might just mean that the frontier has moved. More realistically, determining if an AI system truly works in the wild requires more than just pushing benchmark numbers, which can be misleading when those benchmarks do not capture the long tail of naturally-occurring tests. Furthermore, all AI systems present tradeoffs, and a neural system’s unpredictability, high computational cost, and out-of-distribution fragility should be weighed against whatever high benchmark numbers they may achieve. Despite these caveats, we are optimistic about the promise of tuning LLMs for PBE, and believe that it has the potential to dramatically expand the scope of solvable problems and even solvable domains.

Acknowledgements. We are grateful for assistance from Joshua Rule in the processing of the list functions data, and for feedback from Yewen Pu on the manuscript. This work was supported by an NSF CAREER grant as well as gifts from Google and Cisco.

References

- [1] Henry Lieberman. *Your wish is my command: Programming by example*. Morgan Kaufmann, 2001.
- [2] Sumit Gulwani. Automating string processing in spreadsheets using input-output examples. In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL ’11, page 317–330, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450304900. doi: 10.1145/1926385.1926423. URL <https://doi.org/10.1145/1926385.1926423>.
- [3] Sumit Gulwani. Automating string processing in spreadsheets using input-output examples. *SIGPLAN Not.*, 46(1):317–330, jan 2011. ISSN 0362-1340. doi: 10.1145/1925844.1926423. URL <https://doi.org/10.1145/1925844.1926423>.
- [4] Oleksandr Polozov and Sumit Gulwani. Flashmeta: A framework for inductive program synthesis. *SIGPLAN Not.*, 50(10):107–126, oct 2015. ISSN 0362-1340. doi: 10.1145/2858965.2814310. URL <https://doi.org/10.1145/2858965.2814310>.
- [5] Xinyun Chen, Petros Maniatis, Rishabh Singh, Charles Sutton, Hanjun Dai, Max Lin, and Denny Zhou. Spreadsheetcoder: Formula prediction from semi-structured context. In *International Conference on Machine Learning*, pages 1661–1672. PMLR, 2021.

- [6] Sumit Gulwani, José Hernández-Orallo, Emanuel Kitzelmann, Stephen H Muggleton, Ute Schmid, and Benjamin Zorn. Inductive programming meets the real world. *Communications of the ACM*, 58(11):90–99, 2015.
- [7] François Chollet. On the measure of intelligence, 2019.
- [8] Stephen H Muggleton, Ute Schmid, Christina Zeller, Alireza Tamaddoni-Nezhad, and Tarek Besold. Ultra-strong machine learning: comprehensibility of programs learned with ilp. *Machine Learning*, 107(7):1119–1140, 2018.
- [9] J Solomono Raymond. A formal theory of inductive inference i. *Information and Control*, 7: 1–22, 1964.
- [10] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [11] Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. Competition-level code generation with alphacode. *Nature*, 2022.
- [12] Kensen Shi, Hanjun Dai, Wen-Ding Li, Kevin Ellis, and Charles Sutton. Lambdabeam: Neural program search with higher-order functions and lambdas. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=qVMPXrX4FR>.
- [13] John K. Feser, Swarat Chaudhuri, and Isil Dillig. Synthesizing data structure transformations from input-output examples. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI ’15, page 229–239, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450334686. doi: 10.1145/2737924.2737977. URL <https://doi.org/10.1145/2737924.2737977>.
- [14] Nathanaël Fijalkow, Guillaume Lagarde, Théo Matricon, Kevin Ellis, Pierre Ohlmann, and Akarsh Nayan Potta. Scaling neural program synthesis with distribution-based search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6623–6630, 2022.
- [15] Kensen Shi, Joey Hong, Yinlin Deng, Pengcheng Yin, Manzil Zaheer, and Charles Sutton. Exedec: Execution decomposition for compositional generalization in neural program synthesis. In *International Conference on Learning Representations*, 2024.
- [16] Ansong Ni, Miltiadis Allamanis, Arman Cohan, Yinlin Deng, Kensen Shi, Charles Sutton, and Pengcheng Yin. Next: Teaching large language models to reason about code execution, 2024.
- [17] Joshua S. Rule, Steven T. Piantadosi, Andrew Cropper, Kevin Ellis, Maxwell Nye, and Joshua B. Tenenbaum. Symbolic metaprogram search improves learning efficiency and explains rule learning in humans. *Nature Communications*, 2024.
- [18] OpenAI. Gpt-4 technical report, 2023.
- [19] Ashwin Kalyan, Abhishek Mohta, Oleksandr Polozov, Dhruv Batra, Prateek Jain, and Sumit Gulwani. Neural-guided deductive search for real-time program synthesis from examples. *arXiv preprint arXiv:1804.01186*, 2018.
- [20] Jacob Devlin, Jonathan Uesato, Surya Bhupatiraju, Rishabh Singh, Abdel-rahman Mohamed, and Pushmeet Kohli. Robustfill: Neural program learning under noisy i/o. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML’17, page 990–998. JMLR.org, 2017.
- [21] Xinyun Chen, Chang Liu, and Dawn Song. Execution-guided neural program synthesis. In *International Conference on Learning Representations*, 2018.
- [22] Kensen Shi, Hanjun Dai, Kevin Ellis, and Charles Sutton. Crossbeam: Learning to search in bottom-up program synthesis. In *International Conference on Learning Representations*, 2021.

- [23] Bei Chen, Fengji Zhang, Anh Nguyen, Daoguang Zan, Zeqi Lin, Jian-Guang Lou, and Weizhu Chen. Codet: Code generation with generated tests. *ICLR*, 2023.
- [24] Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven Hoi. CodeRL: Mastering code generation through pretrained models and deep reinforcement learning. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=WaGvb70zySA>.
- [25] Eric Zelikman, Qian Huang, Gabriel Poesia, Noah Goodman, and Nick Haber. Parsel: Algorithmic reasoning with language models by composing decompositions. *Advances in Neural Information Processing Systems*, 36:31466–31523, 2023.
- [26] Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, and Jacob Steinhardt. Measuring coding challenge competence with apps. *NeurIPS*, 2021.
- [27] 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.
- [28] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and LINGMING ZHANG. Is your code generated by chatGPT really correct? rigorous evaluation of large language models for code generation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=1qv610Cu7>.
- [29] Linlu Qiu, Liwei Jiang, Ximing Lu, Melanie Sclar, Valentina Pyatkin, Chandra Bhagavatula, Bailin Wang, Yoon Kim, Yejin Choi, Nouha Dziri, and Xiang Ren. Phenomenal yet puzzling: Testing inductive reasoning capabilities of language models with hypothesis refinement. *ICLR*, 2024.
- [30] Kevin Ellis. Human-like few-shot learning via bayesian reasoning over natural language. *NeurIPS*, 2023.
- [31] Ruocheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah D. Goodman. Hypothesis search: Inductive reasoning with language models. *ICLR*, 2024.
- [32] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions, 2023.
- [33] Geoffrey E Hinton, Peter Dayan, Brendan J Frey, and Radford M Neal. The "wake-sleep" algorithm for unsupervised neural networks. *Science*, 268(5214):1158–1161, 1995.
- [34] Peter-Michael Osera and Steve Zdancewic. Type-and-example-directed program synthesis. *ACM SIGPLAN Notices*, 50(6):619–630, 2015.
- [35] Matej Balog, Alexander L. Gaunt, Marc Brockschmidt, Sebastian Nowozin, and Daniel Tarlow. Deepcoder: Learning to write programs. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=ByldLrqlx>.
- [36] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*, 2022.
- [37] Tessa Lau, Steven A Wolfman, Pedro Domingos, and Daniel S Weld. Programming by demonstration using version space algebra. *Machine Learning*, 53:111–156, 2003.
- [38] José Cambronero, Sumit Gulwani, Vu Le, Daniel Perelman, Arjun Radhakrishna, Clint Simon, and Ashish Tiwari. Flashfill++: Scaling programming by example by cutting to the chase. *Proceedings of the ACM on Programming Languages*, 7(POPL):952–981, 2023.

- [39] Microsoft prose public benchmark suite, 2022. Available at <https://github.com/microsoft/prose-benchmarks>.
- [40] Jiayuan Mao, Xiuming Zhang, Yikai Li, William T. Freeman, Joshua B. Tenenbaum, and Jiajun Wu. Program-Guided Image Manipulators. In *International Conference on Computer Vision*, 2019.
- [41] Kevin Ellis and Sumit Gulwani. Learning to learn programs from examples: Going beyond program structure. In *IJCAI*, pages 1638–1645, 2017.
- [42] David D. Thornburg. Friends of the turtle. *Compute!*, March 1983.
- [43] Catherine Wong, Kevin Ellis, Joshua B. Tenenbaum, and Jacob Andreas. Leveraging language to learn program abstractions and search heuristics. In *ICML*, 2021.
- [44] Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y. K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. Deepseek-coder: When the large language model meets programming – the rise of code intelligence, 2024.
- [45] Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sablé-Meyer, Lucas Morales, Luke Hewitt, Luc Cary, Armando Solar-Lezama, and Joshua B. Tenenbaum. *DreamCoder: Bootstrapping Inductive Program Synthesis with Wake-Sleep Library Learning*, page 835–850. Association for Computing Machinery, New York, NY, USA, 2021. ISBN 9781450383912. URL <https://doi.org/10.1145/3453483.3454080>.
- [46] Elias Stengel-Eskin, Archiki Prasad, and Mohit Bansal. Regal: Refactoring programs to discover generalizable abstractions. *arXiv preprint arXiv:2401.16467*, 2024.
- [47] Stephen H. Muggleton, Dianhuan Lin, and Alireza Tamaddoni-Nezhad. Meta-interpretive learning of higher-order dyadic datalog: Predicate invention revisited. *Mach. Learn.*, 100(1):49–73, jul 2015. ISSN 0885-6125. doi: 10.1007/s10994-014-5471-y. URL <https://doi.org/10.1007/s10994-014-5471-y>.
- [48] Steven Piantadosi. Fleet. <https://github.com/piantado/Fleet/>, 2023. [Online GitHub repository].
- [49] Rajeev Alur, Rastislav Bodik, Garvit Juniwal, Milo MK Martin, Mukund Raghothaman, Sanjit A Seshia, Rishabh Singh, Armando Solar-Lezama, Emina Torlak, and Abhishek Udupa. *Syntax-guided synthesis*. IEEE, 2013.
- [50] Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. Wizardcoder: Empowering code large language models with evol-instruct, 2023.
- [51] Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. Zerogen: Efficient zero-shot learning via dataset generation. *arXiv preprint arXiv:2202.07922*, 2022.
- [52] Ajay Patel, Colin Raffel, and Chris Callison-Burch. Datadreamer: A tool for synthetic data generation and reproducible llm workflows, 2024.
- [53] Yu Meng, Jiaxin Huang, Yu Zhang, and Jiawei Han. Generating training data with language models: Towards zero-shot language understanding. *Advances in Neural Information Processing Systems*, 35:462–477, 2022.
- [54] Kavi Gupta, Peter Ebert Christensen, Xinyun Chen, and Dawn Song. Synthesize, execute and debug: learning to repair for neural program synthesis. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS ’20, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- [55] Shraddha Barke, Hila Peleg, and Nadia Polikarpova. Just-in-time learning for bottom-up enumerative synthesis. *Proc. ACM Program. Lang.*, 4(OOPSLA), nov 2020. doi: 10.1145/3428295. URL <https://doi.org/10.1145/3428295>.

- [56] Kevin Ellis, Maxwell Nye, Yewen Pu, Felix Sosa, Joshua B. Tenenbaum, and Armando Solar-Lezama. *Write, Execute, Assess: Program Synthesis with a REPL*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- [57] Rishabh Singh and Sumit Gulwani. Predicting a correct program in programming by example. In *27th International Conference on Computer Aided Verification (CAV 2015)*, July 2015. URL <https://www.microsoft.com/en-us/research/publication/predicting-a-correct-program-in-programming-by-example/>.
- [58] Percy Liang, Michael I. Jordan, and Dan Klein. Learning programs: a hierarchical bayesian approach. In *Proceedings of the 27th International Conference on International Conference on Machine Learning, ICML'10*, page 639–646, Madison, WI, USA, 2010. Omnipress. ISBN 9781605589077.
- [59] Jeevana Priya Inala, Chenglong Wang, Mei Yang, Andres Codas, Mark Encarnación, Shuvendu Lahiri, Madanlal Musuvathi, and Jianfeng Gao. Fault-aware neural code rankers. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 13419–13432. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/5762c579d09811b7639be2389b3d07be-Paper-Conference.pdf.
- [60] Wasu Top Piriyaakulkij and Kevin Ellis. Doing experiments and revising rules with natural language and probabilistic reasoning. *CogSci*, 2024.
- [61] Ruocheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah Goodman. Hypothesis search: Inductive reasoning with language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=G7UtIGQmjm>.
- [62] Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128*, 2023.
- [63] Sean Welleck, Ximing Lu, Peter West, Faeze Brahman, Tianxiao Shen, Daniel Khashabi, and Yejin Choi. Generating sequences by learning to self-correct. *ICLR*, 2023.
- [64] Theo X. Olausson, Jeevana Priya Inala, Chenglong Wang, Jianfeng Gao, and Armando Solar-Lezama. Is self-repair a silver bullet for code generation?, 2023.
- [65] Noah Shinn, Beck Labash, and Ashwin Gopinath. Reflexion: an autonomous agent with dynamic memory and self-reflection. *arXiv preprint arXiv:2303.11366*, 2023.
- [66] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36, 2024.
- [67] Hao Tang, Keya Hu, Jin Peng Zhou, Sicheng Zhong, Wei-Long Zheng, Xujie Si, and Kevin Ellis. Code repair with llms gives an exploration-exploitation tradeoff. *arXiv*, 2024.
- [68] Ansong Ni, Sridhar Iyer, Dragomir Radev, Ves Stoyanov, Wen-tau Yih, Sida I Wang, and Xi Victoria Lin. Lever: Learning to verify language-to-code generation with execution. In *Proceedings of the 40th International Conference on Machine Learning (ICML'23)*, 2023.
- [69] jinaai/jina-embeddings-v2-base-code. <https://huggingface.co/jinaai/jina-embeddings-v2-base-code>. Accessed: 2024-05-22.

A Appendix / supplemental material

A.1 Experiment Details

We used a temperature of 1.0 for sampling in our experiments unless otherwise stated. All experiments were performed on single-node machines (8xA6000 or 8xA100, etc.) without a multi-node distributed computing setup.

A.1.1 List Tasks

We selected 50 problems from Rule et al. as our seed set, reserving the remaining problems for testing. To ensure comparability with Rule et al., we tested on the first 100 problems (excluding those in the seed set), resulting in 77 test problems. We consistently used 10 input-output examples, with the remaining 54 examples serving as a held-out test set. When filtering duplicate synthetic data, we employed an open code embedding model[69] available on Hugging Face. As a sanity check, we also use 10 list to list problems from λ^2 benchmark and shown the model can effectively solve them in Table.4

A.1.2 String Tasks

We utilized 100 string-to-string/null transformation problems from the prose-benchmark. When available, we used 10 input-output examples, always reserving at least one example as a hold-out test set. This ensures the generalization of our synthesized programs, as the benchmark did not provide held-out data.

For the FlashFill baseline, we used Microsoft Excel for Mac version 16.8. We opened each individual xlsx file containing the test problem examples and manually triggered the FlashFill function by pressing Ctrl+E.

A.1.3 Logo Tasks

To facilitate graphics input inference with code language models, we converted logo graphics into ASCII-represented strings, as shown in Figure 5. For each input image, we cropped a 512x512 section from the center and then divided it into 32x32 blocks, each with a size of 16x16. We counted the number of black pixels in each block, calculated their density ($\frac{\text{num black pixels}}{16 \cdot 16}$), and quantized this density value into 10 levels, represented by the ASCII numbers 0-9. By representing each block with an ASCII number, an input image is represented with a string of 32 lines, and each line has 32 numbers.

For the turtle graphics program, we adopted the Python turtle graphics program from Regal[46] with a minor modification of changing the ‘embed’ function to use a ‘with’ context manager instead, calling it ‘fork_state’. This allows for equivalent but more readable code.

For the GPT-4o and GPT-4o mini multimodal baselines, we directly use images as inputs. (See Section A.5.8 for the prompt template.)

A.2 Synthetic Dataset Generation and Training Parameters

We present the dataset generation and training parameters in Table. 2 and Table. 3.

A.3 Adaptation Implementation Details

For adaptation experiments, we generally followed the settings described above, with a few specific differences detailed below.

A.3.1 String Tasks

To induce a domain gap (easier problems in Sygus compared to the harder, noisier problems in the Prose Benchmark), we first fine-tuned a model using Sygus problems and then tested it on the Prose Benchmark. Due to the noisy nature of the Prose Benchmark (some problems have very few examples), we adopted a setting where we utilized all the test cases to select which problems were

	List	String
Seed Dataset Source	Rule et al.	Prose (FlashFill++) Problems
Seed Dataset Size	50	40
Synthetic Data Generator	deepseekcoder-33b-instruct	deepseekcoder-33b-instruct
Synthetic Dataset Size	10k	10k
Sampling Temperature	0.8	1.0
Similarity Filter	code embedding model	-
Filter Ratio	around 1/3 (threshold=0.93)	-
Synthetic Data Prompt	4-shot examples	10-shot examples
LoRA Finetuning		
Model Used	deepseekcoder-1.5-7b-instruct	deepseekcoder-1.5-7b-instruct
LoRA Rank	1024	256
LoRA α	1024	256
Learning Rate	2.00E-04	2.00E-04
LR Schedule	cosine	cosine
Warmup Steps	10	10
Epoch	1	1
Batchsize	32	16
LoRA		
33b Model Used for FT	deepseekcoder-33b-instruct	deepseekcoder-33b-instruct
LoRA	256	128
LoRA α	256	128
Learning Rate	2.00E-04	2.00E-04
LR Schedule	cosine	cosine
Warmup Steps	10	10
Epoch	1	1
Batchsize	32	32

Table 2: List task and String task synthetic dataset generation and finetuning parameters.

solved and then used them as the seed programs for adaptation. This resulted in 64 solved problems out of the 100 problems in the benchmark.

A.3.2 List Tasks

To obtain the finetuned in-distribution result in Fig. 8a, we fine-tuned on a synthetic dataset generated by seeding with 20 out of 100 problems from LambdaBeam, and tested on the remaining 80 problems.

A.3.3 LOGO Tasks

For LOGO adaptation experiments, we induce domain gap by using the shorter programs ($\text{LoC} \leq 12$) of the training set, and tested on the longer programs ($\text{LoC} > 12$). The shorter programs training seed consists of around 80% problems (156 out of 200) from the original training set. The test set consists of 31 problems out of 111 problems from the original test set.

A.4 Model Performance on LambdaBeam Benchmark

We present the results of both our 7B and 33B models on the LambdaBeam benchmark in Figure 10. We observed that even without fine-tuning for this specific benchmark, and instead fine-tuned for the list-to-list problems from Rule et al., our models performed exceptionally well, surpassing the state-of-the-art results specifically designed for the LambdaBeam problems [12].

A.5 Prompts Used in the Experiments

A.5.1 Synthetic Data Generation Prompt

List

Logo	
Seed Dataset Source	Regal Python Logo Programs
Seed Dataset Size	200
Synthetic Data Generator	deepseekcoder-33b-instruct
Synthetic Dataset Size	32k
Similarity Filter	-
Filter Threshold	-
Synthetic Data Prompt	6-shot examples
LoRA Finetuning	
Model Used	deepseekcoder-1.5-7b-instruct
LoRA Rank	512
LoRA α	512
Learning Rate	2.00E-04
LR Schedule	cosine
Warmup Steps	20
Epoch	3
Batchsize	64
LoRA Finetuning	
Model Used	deepseekcoder-33b-instruct
LoRA Rank	512
LoRA α	512
Learning Rate	2.00E-04
LR Schedule	cosine
Warmup Steps	50
Epoch	3
Batchsize	64

Table 3: Logo task synthetic datasets generation and finetuning parameters

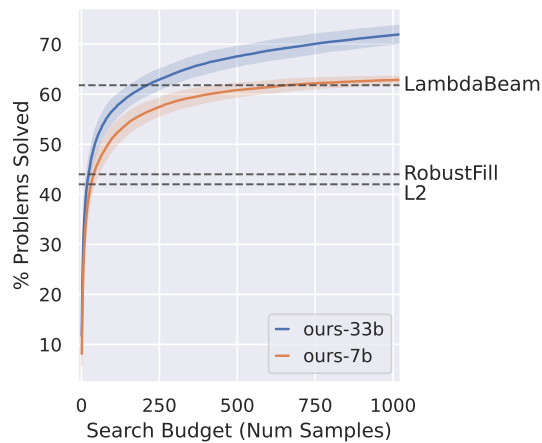


Figure 10: Performance on LambdaBeam problems from Shi et al. 2023

You are a CS professor. You are providing a set of challenging and diverse integer list to integer list function puzzle for your student to solve.

Puzzle 1:

Python function: input a list of integers and return a list of integers

```
'''python
{PROGRAM EXAMPLE 1}
'''
```

Test cases:

...

Puzzle 2:

Python function: input a list of integers and return a list of integers

```
'''python
{PROGRAM EXAMPLE 2}
'''
```

Test cases:

...

Following the above format, please provide 3 functions each follow by 10 random test cases to check the function's correctness full coverage

String

Excel just introduce a new feature that allows user to use Python to perform data transformation.

Please generate a csv file with two columns. The first column contains the input data and the second column contains the output data.

Following a accompanying python function which showcasing the transformation of the input data to the output data.

Here are 10 challenging examples showcaing the features

Example 1

```
'''csv
{INPUT OUTPUT EXAMPLE 1}
'''
```

Here is the Python function that help transform the first column to the second column.

```
'''python
{PYTHON EXAMPLE 1}
'''
```

Example 2

```
'''csv
{INPUT OUTPUT EXAMPLE 2}
'''
```

Here is the Python function that help transform the first column to the second column.

```
'''python
{PYTHON EXAMPLE 2}
'''
```

...

Following the above format, please provide a CSV file with two columns, containing between 5 to 10 rows of data showing a transformation from the first column to the second column. This csv data should illustrate a challenging and complex example, similar to the above examples. Following that, create a Python function designed to process this data. Be aware that this function will be tailored to not only accommodate the current data but also any future data that follows the same format or structure.

Logo

Your task is to draw simple black and white graphics with the custom library. DO NOT USE THE BUILT-IN TURTLE LIBRARY. You will use a custom turtle library, similar to the built-in turtle library, which is sufficient for all tasks.

Here are all the available functions in the custom turtle library:

- forward(x): move forward x pixels
- left(theta): rotate left by theta degrees
- right(theta): rotate right by theta degrees
- penup(): stop drawing
- pendown(): start drawing
- teleport(x, y, theta): move to position (x, y) with angle theta
- heading(): get the current angle of the turtle
- isdown(): check if the pen is down
- forward(x): Move forward x pixels.
- left(theta): Rotate left by theta degrees.
- right(theta): Rotate right by theta degrees.
- penup(): Stop drawing.
- pendown(): Start drawing.
- teleport(x, y, theta): Move to position (x, y) with angle theta.
- heading(): Get the current angle of the turtle.
- isdown(): Check if the pen is down.
- with fork_state(): A context manager that runs the code in the block using the current context and restores the original state afterwards. Allows you to nest programs. Internally, fork_state saves the turtle state (is_down, x, y, heading), executes the block, then restores the original state.

Graphic 1

Python program: draw an interesting graphic using our own custom turtle library

the following program draws ...

```
{PROGRAM EXAMPLE 1}
```

Graphic 2

Python program: draw an interesting graphic using our own custom turtle library

the following program draws ...

```
{PROGRAM EXAMPLE 2}
```

...

Following the above format, please provide 5 more programs using our custom drawing library.

A.5.2 Prompt Template for Finetuning and Zero-Shot Experiments

A.5.3 List

```
Implement the function solve_puzzle that takes a list of
integers and returns a list of integers. The function
should satisfy the following assertions
assert solve_puzzle(...) == ...
assert solve_puzzle(...) == ...
assert solve_puzzle(...) == ...
...
```

A.5.4 List (Chain-of-Thought)

```
Implement the function solve_puzzle that takes a list of
integers and returns a list of integers. The function
should satisfy the following assertions:
assert solve_puzzle(...) == ...
assert solve_puzzle(...) == ...
assert solve_puzzle(...) == ...
...
Please observe the relation between input and output and think
step by step. Output the function in a markdown format in
the end.
Solution:
```

A.5.5 String

```
Implement the function edit_text that takes a string and
returns a string. The function transforms the input string
to the output string. The function should satisfy the
following assertions:
assert edit_text(...) == ...
assert edit_text(...) == ...
assert edit_text(...) == ...
...
```

A.5.6 String (Chain-of-Thought)

```
Please implement the function edit_text that takes a string
input and returns a modified string.
Note that you can import re, datetime, or any built-in Python
library to solve the problem.
The function should satisfy the following test cases:
assert edit_text(...) == ...
assert edit_text(...) == ...
assert edit_text(...) == ...
...
Please reason through the problem and think step by step, and
finally implement the function and output the full function
implementation in a markdown code block in the end.
```

A.5.7 Logo

```
Here is a gray scale images representing with integer values
0-9.
{CONVERTED IMAGE STRING}...
```

```
...
Please write a Python program that generates the image using
our own custom turtle module
```

A.5.8 Logo (Multimodal Few-shot)

```
Given the following custom Turtle graphics-like library, Please
use it to write a program that draws the given image.
'''python
from myturtle import Turtle
from myturtle import HALF_INF, INF, EPS_DIST, EPS_ANGLE
turtle = Turtle()
def forward(dist):
    turtle.forward(dist)
def left(angle):
    turtle.left(angle)
def right(angle):
    turtle.right(angle)
def teleport(x, y, theta):
    turtle.teleport(x, y, theta)
def penup():
    turtle.penup()
def pendown():
    turtle.pendown()
def position():
    return turtle.x, turtle.y
def heading():
    return turtle.heading
def isdown():
    return turtle.is_down
def fork_state():
    """
    Fork the current state of the turtle.
    Usage:
    with fork_state():
        forward(100)
        left(90)
        forward(100)
    """
    return turtle._TurtleState(turtle)
'''

Below are some example programs that draw the given images.
<IMAGE>
Here is a program that draws the above image.
The figure is like a medium 8 gon.
'''python
for i in range(8):
    forward(4)
    left(45.0)
'''

<IMAGE>
Here is a program that draws the above image:
The figure is like 5 sided snowflake with a medium circle and a
medium semicircle as arms.
'''python
for j in range(5):
    with fork_state():
        penup()
        forward(2)
```

```

left(0.0)
pendown()
for i in range(HALF_INF):
    forward(EPS_DIST*2)
    left(EPS_ANGLE)
for i in range(HALF_INF):
    forward(EPS_DIST*2)
    left(EPS_ANGLE)
penup()
forward(2)
left(0.0)
pendown()
for i in range(HALF_INF):
    forward(EPS_DIST*2)
    left(EPS_ANGLE)
forward(0)
left(72.0) "
'''
<IMAGE>
Here is a program that draws the above image.
The figure is like 7 concentric circles.
'''python
for j in range(8):
    for i in range(HALF_INF):
        forward(EPS_DIST*j)
        left(EPS_ANGLE)
    for i in range(HALF_INF):
        forward(EPS_DIST*j)
        left(EPS_ANGLE)
'''
<TEST_IMAGE>
Output the program that draws the following image. Reason about
the given image and write a program that draws it in a
markdown code block.

```

Name	Description	ours-7B	ours-33B
Dedup	Remove duplicate elements from a list.	✓	✓
Reverse	Reverse a list.	✓	✓
Droplast	Drop the last element in a list.	✓	✓
Dropmax	Drop the largest number(s) in a list.	✓	✓
Dupli	Duplicate each element of a list.	✓	✓
Evens	Remove the odd numbers from a list.	✓	✓
Multfirst	Replace every item in a list with the first item.	✓	✓
Multlast	Replace every item in a list with the last item.	✓	✓
Shiftl	Shift all elements in a list to the left.	✓	✓
Shiftr	Shift all elements in a list to the right.	✓	✓

Table 4: 10 list \mapsto list functions from λ^2 [13]

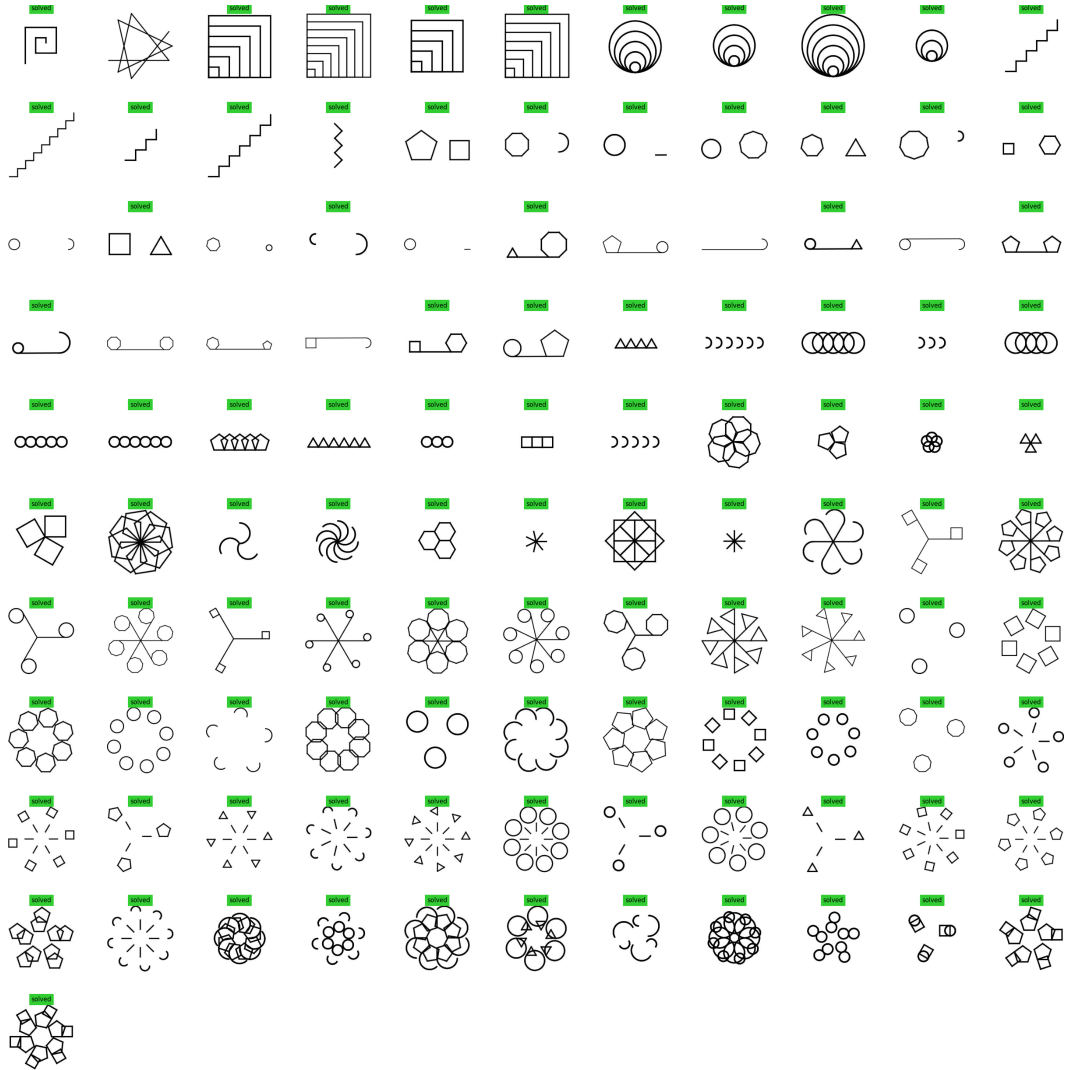


Figure 11: All LOGO test problems: problems solved by our finetuned 33b model are marked as green

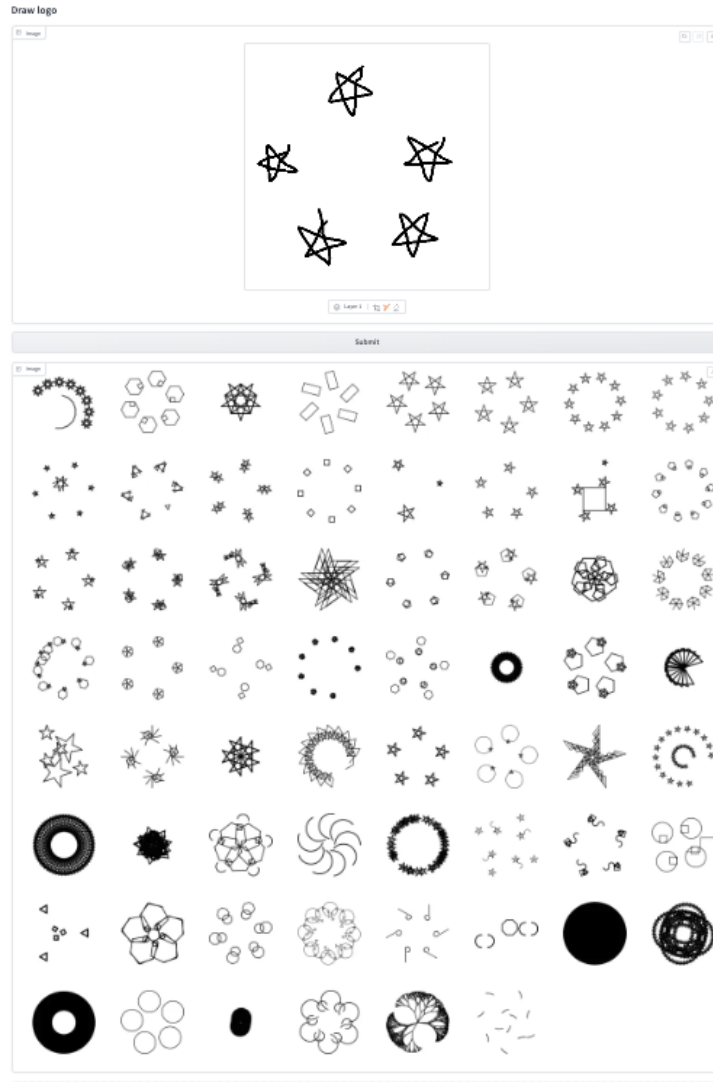


Figure 12: Hand-drawn LOGO test showing every generated sample. We built a graphical interface to allow users to draw images as input. The sample budget for this demo is 64.

NeurIPS Paper Checklist

1. Claims

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

Answer: [Yes]

Justification: We provide experiment results to support our claims.

Guidelines:

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

2. Limitations

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

Answer: [Yes]

Justification: We discuss several limitations in the limitation section.

Guidelines:

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

3. Theory Assumptions and Proofs

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

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

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

4. Experimental Result Reproducibility

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

Answer: [Yes]

Justification: We provide training settings and prompts used in the appendix.

Guidelines:

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

5. Open access to data and code

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

Answer: [No]

Justification: We provided the detailed training settings and prompts in the appendix.

Guidelines:

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

6. Experimental Setting/Details

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

Answer: [Yes]

Justification: We provide experimental settings and details in the appendix.

Guidelines:

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

7. Experiment Statistical Significance

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

Answer: [Yes]

Justification: We provide error bars or confidence intervals when applicable.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

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

Answer: [Yes]

Justification: We provide the information in the appendix.

Guidelines:

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

9. Code Of Ethics

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

Answer: [Yes]

Justification: The paper conform with the NeurIPS Code of Ethics.

Guidelines:

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

10. Broader Impacts

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

Answer:[NA]

Justification: The paper focuses on narrow domains of PBE problems and has limited impacts on society.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

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

Answer: [NA]

Justification: The paper poses a low risk.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

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

Answer: [Yes]

Justification: We credit and follow the license terms properly.

Guidelines:

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

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

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

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

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

14. **Crowdsourcing and Research with Human Subjects**

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

Answer: [NA]

Justification: The paper does not involve crowdsourcing or human subjects

Guidelines:

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

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

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

Answer: [NA]

Justification: The paper does not involve crowdsourcing or research with human subjects.

Guidelines:

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