

DECAEVOLVE: DECOMPOSE, ADAPT, AND EVOLVE, OR, THREE PILLARS OF EFFECTIVE LLM-BASED SCIENTIFIC EQUATION DISCOVERY

Anonymous authors

Paper under double-blind review

ABSTRACT

Finding mathematical relations underlying natural phenomena and scientific systems has been one of the fundamental tasks in the history of scientific discovery. Recent advancements in evolutionary search with Large Language Models (LLMs), with their embedded scientific knowledge, have shown great promise for this task. However, discovering such mathematical models governing scientific observations still remains significantly challenging, as it requires navigating vast combinatorial hypothesis spaces with an explosion of possible relations. Existing LLM-based approaches overlook the impact of data on the structure of mathematical relations, and treat LLMs as a static hypothesis generator unaware of the observed scientific system. This leads to suboptimal and inefficient exploration of the hypothesis space with over-reliance on LLMs' internal priors. To bridge this gap, we introduce *Decompose, Adapt, and Evolve (DecAEvolve)*, a framework that leverages granular feedback from symbolic term decomposition and LLM refinement through reinforcement learning (RL) fine-tuning to enhance both robustness and efficiency of evolutionary discovery frameworks. DecAEvolve unifies symbolic decomposition with test-time RL adaptation, enabling adaptive rather than static hypothesis generation and reducing error by up to an order of magnitude compared to state-of-the-art baselines. Our experiments across diverse datasets demonstrate that DecAEvolve significantly improves the accuracy of discovered equations and the efficiency of the discovery process compared to the state-of-the-art baselines¹.

1 INTRODUCTION

The emergence of Large Language Models (LLMs) has fundamentally transformed automated problem-solving across diverse domains. Beyond their well-established capabilities in natural language understanding and programming (Achiam et al., 2023; Touvron et al., 2023), LLMs have recently demonstrated remarkable reasoning abilities that enable them to tackle complex optimization and discovery tasks. Their capacity to leverage embedded domain knowledge, interpolate between them, generate structured hypotheses and engage in iterative refinement, positions LLMs as powerful engines for systematic exploration of complex solution spaces towards discovery goals (Romera-Paredes et al., 2024; Novikov et al., 2025; Surina et al., 2025). This potential extends naturally to scientific discovery tasks, where the combination of domain expertise and systematic search/exploration in the hypothesis space can unlock new approaches to longstanding challenges of scientific inquiry (Shojaee et al., 2025a).

Scientific equation discovery—the process of uncovering compact and interpretable mathematical models that govern natural phenomena—represents one of the fundamental tasks in automated scientific discovery, with applications across many fields of science such as physics, biology, and material science (Makke & Chawla, 2024). Traditional approaches in Symbolic Regression (SR) rely on genetic programming and evolutionary strategies (Koza, 1994b; Cava et al., 2021b); however, these approaches often struggle with scalability limitations and inefficient exploration of the vast combinatorial hypothesis space (Virgolin & Pissis, 2022). More recent advances have introduced neural-guided approaches, where deep learning architectures are trained to generate or refine symbolic expressions

¹Code is available at: <https://anonymous.4open.science/r/decaevolve-1215>

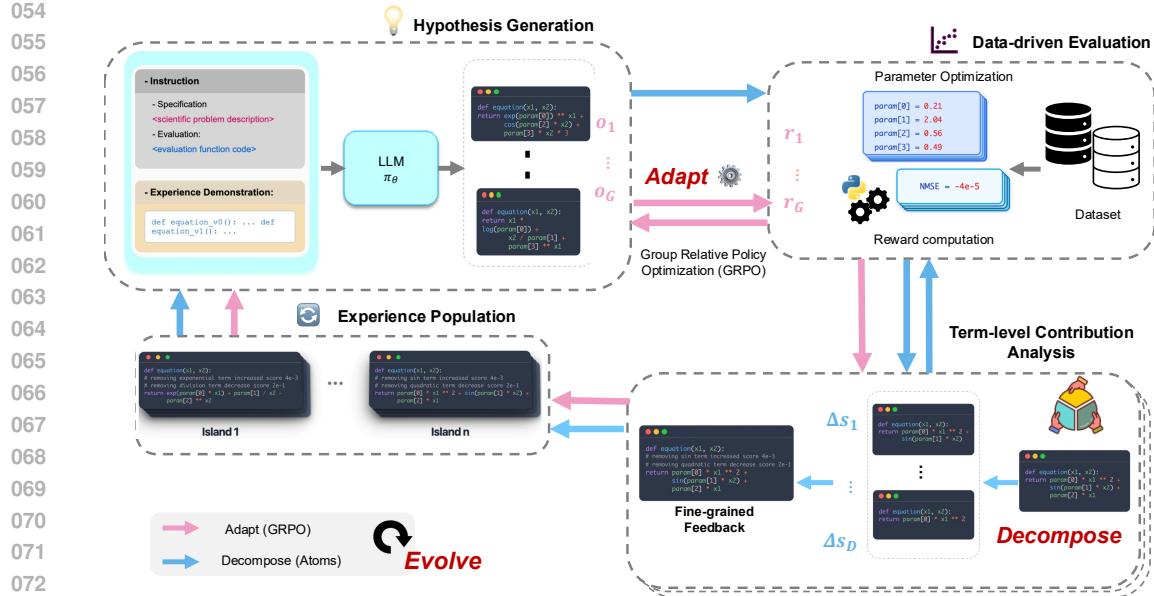


Figure 1: **Overview of the DecAEvolve framework.** The framework integrates *Adaptation* (LLM fine-tuning via reinforcement learning using Group Relative Policy Optimization with data-driven rewards) and *Decomposition* (granular-level feedback through symbolic atomic term analysis) within an *Evolutionary* search process. The adaptation aligns the LLM to the target scientific system beyond its internal priors, while decomposition provides fine-grained guidance for hypothesis refinement. Iterating these three key components enables effective and efficient exploration of the combinatorial hypothesis space in equation discovery.

(Udrescu & Tegmark, 2020b; Bruneton, 2025a), and transformer-based methods that are pre-trained with large-scale synthetic data to directly model symbolic sequences as language generation tasks (Kamienny et al., 2022a; Shojaee et al., 2023a; Meidani et al., 2024b). These developments have demonstrated promising capabilities in data-driven learning, yet are limited in balancing learning and search components and in incorporating scientific prior knowledge into the process of discovery.

Several works have recently introduced promising frameworks to integrate LLMs for scientific equation discovery, leveraging their scientific priors and reasoning capabilities to navigate the complex landscape of mathematical expressions more effectively. Notably, LLM-SR (Shojaee et al., 2025a) combines LLMs' scientific knowledge with multi-island evolutionary search, generating equation hypotheses as Python function skeletons guided by data feedback. LaSR (Grayeli et al., 2024b) introduces a concept learning approach that extracts abstract textual concepts from successful equation hypotheses, using these concepts to guide both evolutionary search (with PySR (Cranmer, 2023)) and LLM-based hypothesis generation. SGA (Ma et al., 2024b) employs a bilevel optimization framework that iteratively combines LLMs for discrete hypothesis generation with physical simulations for continuous parameter optimization. These methods demonstrate this potential by combining LLMs' domain expertise with systematic search strategies, treating equation discovery as a program synthesis problem guided by scientific knowledge (Shojaee et al., 2025b; Reddy & Shojaee, 2025).

Our key insight is that equation discovery benefits from both adaptation (aligning the model with data distributions) and decomposition (understanding which symbolic components matter), neither of which prior LLM frameworks integrate. While classical methods such as SINDy (Brunton et al., 2016) employ sparse regression techniques that implicitly decompose equations through term selection, they lack the adaptive learning and scientific reasoning capabilities that LLMs provide. Current LLM-based discovery methods also exhibit fundamental limitations that constrain their effectiveness. First, they treat LLMs as static hypothesis generators, where the model's parameters remain fixed regardless of the problem domain, nuances of the specific observed system or, insights gained during the search process. This prevents LLMs from adapting their generation strategies based on the specific problem, the data, and the domain-specific requirements. Second, existing approaches mainly provide coarse-grained feedback about solution quality, typically limited to scalar reward signals (e.g., mean squared error) from execution of whole hypothesis that indicate which hypotheses perform well respectively, without revealing why specific mathematical components or patterns drive success. This

108 limited feedback mechanism prevents LLMs from understanding the underlying symbolic structure
 109 of successful solutions and refining their search strategies accordingly.
 110

111 To address these limitations, we introduce **DecAEvolve** (Decompose, Adapt, and Evolve), a novel
 112 framework that enhances the effectiveness and efficiency of LLM-based equation discovery by com-
 113 bining decomposition-based feedback with test-time adaptation within an LLM-guided evolutionary
 114 framework. Our key contributions are as follows:

- 115 • We develop a systematic methodology for providing LLMs with interpretable directional feedback
 116 about which components of their generated hypothesis prove effective. Through structured hypoth-
 117 esis decomposition and evaluations, the contributions of individual terms and their interactions are
 118 quantified and provided as feedback. This enables LLMs to understand not just which hypotheses
 119 succeed, but *why* specific mathematical building blocks are effective, transforming blind generation
 120 into informed iterative refinement.
- 121 • We employ reinforcement learning with Group-Relative Policy Optimization (GRPO) to implicitly
 122 distill the data distribution into the model’s parameters for better hypothesis generation process.
 123 This test-time adaptation/training approach allows the LLM to learn from successful equation
 124 discoveries without directly observing raw data, progressively aligning its hypothesis generation
 125 with the underlying symbolic relationships through reward-weighted gradient updates.
- 126 • We demonstrate that these synergistic contributions dramatically improve search efficiency, requir-
 127 ing significantly fewer iterations to discover accurate symbolic expressions. Our comprehensive
 128 evaluation across multiple benchmarks shows superior performance compared to LLM-SR and
 129 other baselines in both in-domain and out-of-domain settings, validating the effectiveness of our
 130 guided discovery approach.

131 2 PRELIMINARIES

133 **Problem Formulation.** In scientific equation discovery, the goal is to find a compact mathematical
 134 expression (hypothesis) $h(\mathbf{x}; \mathcal{T})$ that approximates an unknown target function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ within
 135 the context of a specific scientific problem \mathcal{T} with dataset of input–output pairs $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$.
 136 The objective is to discover functional relationships such that $f(\mathbf{x}_i) \approx y_i$ for all i , producing
 137 expressions that are both interpretable and capable of generalizing to unseen data. Perform-
 138 ance is typically evaluated using fitness to data with metrics such as mean squared error:
 139 $\text{MSE}(h, \mathcal{D}) = -\frac{1}{n} \sum_{i=1}^n (h(\mathbf{x}_i; \mathcal{T}) - y_i)^2$.

140 **LLM-SR Framework.** A recent advance in this space is LLM-SR (Shojaee et al., 2025a), which
 141 pioneers reframing of symbolic regression as a program synthesis task. Instead of traditional
 142 expression trees, equations are represented as executable Python functions with placeholder
 143 parameters. The framework leverages large language models (LLMs) to iteratively generate equation
 144 program skeletons, drawing on their embedded scientific priors (i.e., built-in knowledge about a
 145 scientific problem without explicit task-specific feedback or fine-tuning) as well as programming, and
 146 reasoning capabilities. In LLM-SR, each proposed equation skeleton first undergoes data-driven pa-
 147 rameter optimization (via BFGS or Adam), and then, high-scoring hypotheses are stored in a dynamic
 148 multi-island experience buffer to guide the program optimization. Subsequent LLM prompts incorpo-
 149 rate these top hypotheses as in-context examples sampled from buffer, enabling iterative refinement of
 150 the search process. More details on LLM-SR, its evolution, multi-island buffer design, and sampling
 151 strategies are provided in Appendix A. Our framework DecAEvolve builds on LLM-SR foundations
 152 and integrate its evolutionary search mechanism with decomposition-based feedback and test-time
 153 training. This guides the LLM beyond its default priors toward regions of hypothesis space that better
 154 align with the actual observed scientific system, enabling more effective and efficient discovery.

155 **Group-Relative Policy Optimization (GRPO).** GRPO (Shao et al., 2024) is an effective reinfor-
 156 cements learning technique that has recently gained attention to fine-tune LLMs with verifiable outcome
 157 rewards towards enhancing models’ reasoning capabilities. The key innovation of GRPO is its use of
 158 group-relative advantages computed from multiple completions per prompt to provide stable reward
 159 signals. For each prompt p with completions $\{h_i\}_{i=1}^G$ from old policy $\pi_{\theta_{\text{old}}}$ and corresponding rewards
 160 $\{r_i\}_{i=1}^G$, the method calculates group-relative advantages $A_i = r_i - b(p)$ where $b(p) = \frac{1}{G} \sum_{j=1}^G r_j$
 161 serves as the per-prompt baseline. This per-prompt normalization across group samples provides
 162 variance reduction and stable credit assignment across candidates. The GRPO objective optimizes:

$$\begin{aligned}
162 \quad \mathcal{L}(\theta) = & -\mathbb{E}_{p, \{h_i\} \sim \pi_{\theta_{\text{old}}}} \\
163 \quad & \frac{1}{G} \sum_{i=1}^G \frac{1}{|h_i|} \sum_{t=1}^{|h_i|} \left\{ \min \left[\frac{\pi_{\theta}(h_{i,t}|p, h_{i,<t})}{\pi_{\theta_{\text{old}}}(h_{i,t}|p, h_{i,<t})} A_{i,t}, \text{clip} \left(\frac{\pi_{\theta}(h_{i,t}|p, h_{i,<t})}{\pi_{\theta_{\text{old}}}(h_{i,t}|p, h_{i,<t})}, 1-\epsilon, 1+\epsilon \right) A_{i,t} \right] \right. \\
164 \quad & \left. + \beta \mathbb{E}_p [KL(\pi_{\theta}(\cdot|p) \parallel \pi_{\text{ref}}(\cdot|p))] \right\}, \tag{1}
\end{aligned}$$

168 where π_{ref} is the frozen reference model, $\beta > 0$ controls the KL regularization strength, and ϵ
169 controls the clipping ratio to avoid excessive single-step updates to the policy. This approach enables
170 efficient adaptation to task-specific rewards while maintaining model stability.
171

172 3 METHOD

174 **DecAEvolve** (Decompose, Adapt, and Evolve) combines adaptation, decomposition, and evolutionary
175 search. Adaptation employs GRPO to align the LLM with data-driven rewards, while decomposition
176 delivers term-level feedback that highlights which symbolic components drive accuracy. Coupled
177 with evolutionary search, these mechanisms form a feedback-driven loop that improves both equa-
178 tion quality and the model’s generation policy. We illustrate the pseudocode for **DecAEvolve** in
179 Algorithm 1 and emphasize its two key contributions: (i) reinforcement-learning-based test-time
180 adaptation and (ii) fine-grained decomposition feedback.
181

182 3.1 TEST-TIME ADAPTATION WITH GRPO

183 In the adaptation stage of DecAEvolve, the generation policy π_{θ}^n is updated to improve its ability to
184 propose valid and accurate symbolic equations which are better aligned with the observed data. After
185 each iteration of generating hypothesis group samples and computing their corresponding rewards, the
186 policy is fine-tuned using GRPO (Shao et al., 2024) objective function (equation 1) on the accumulated
187 dataset $(\mathbf{p}_n, \{h_i^n, r_i^n\}_{i=1}^G)$, where each prompt at online iteration n (\mathbf{p}_n) is paired with multiple
188 corresponding hypothesis completions $\{h_i^n\}$ and their data-driven rewards $\{r_i^n\}$ as: $\{r_i^n\}_{i=1}^G \leftarrow$
189 $\text{Score}_{\mathcal{T}}(\{h_i^n\}_{i=1}^G, \mathcal{D})$. The score function $\text{Score}_{\mathcal{T}}(\cdot)$ for each equation candidate is the negative mean
190 squared error (MSE) with respect to the training data, which is also transformed to a bounded reward
191 between 0 and 1 via exponential transformation: $\text{Score}_{\mathcal{T}}(h, \mathcal{D}) = \exp(-\text{MSE}(h, \mathcal{D}))$. Failed or
192 invalid completions receive a floor reward of 0.01. Fine-tuning is performed using LoRA adapters,
193 enabling efficient parameter updates while maintaining the base model as a reference anchor (π_{ref}).
194 The KL coefficient β in equation 1 ensures the fine-tuned model retains its general capabilities while
195 effectively adapting to the observed scientific system with the help of data-driven reward through
196 GRPO. More implementation details and hyperparameters are provided in Section 4 and Appendix C.
197

198 Notably, the use of GRPO in DecAEvolve differs fundamentally from its use in prior RL-tuning
199 for reasoning literature. Here, GRPO is not used for global model fine-tuning, but for per-system,
200 test-time adaptation that steers the model parameters toward hypotheses better aligned with the
201 observed data of a scientific system. The equation program synthesis optimization is implicitly
202 guided with the reward from how well an equation hypothesis candidate explains the observed
203 data, not from correctness or textual feedback as in typical GRPO setups. Also, the inputs/prompts
204 used in test-time adaptation GRPO here correspond to a fixed scientific system but also include
205 dynamic in-context examples (with decomposition signal) that are sampled from the buffer in an
206 online manner over GRPO iterations. This design helps to align model adaptation naturally with the
207 decomposition-guided evolutionary search happening later during inference.
208

209 3.2 DIRECTIONAL FEEDBACK WITH TERM-LEVEL CONTRIBUTION

210 At the heart of our framework is an iterative discovery process in which the LLM performs a form of
211 self-reflection: it not only generates candidate symbolic equations but also receives feedback about
212 what certain symbolic components succeed or fail. Rather than relying solely on coarse error scores,
213 we introduce a decomposition-based contribution analysis within the self-reflection procedure that
214 quantifies the role of each term and its pairwise interactions, producing interpretable signals that guide
215 subsequent iterations of discovery. A related use of decomposition has also been explored in Liu et al.
216 (2025) for experimental-chemistry hypothesis discovery. Check Appendix D for detailed discussion.
217

218 During the decomposition step, each candidate equation program is parsed into an abstract syntax
219 tree (AST) and decomposed into atomic symbolic terms $\{u_m(x)\}_{m=1}^M$ where M is the total number
220 of terms (see Appendix B.1 for details). To assess the contribution of a given term u_i (or a pair

216 **Algorithm 1.** DecAEvolve

217 **Input:** LLM π_θ ; dataset \mathcal{D} ; task \mathcal{T} ; Number of GRPO steps N ; Group size G ; Evolution
 218 iterations T ; in-context size k ; samples per prompt b

219 **Output:** Best equation h^* with score s^*

220

221 # Stage 1: Test-time adaptation with GRPO and decomposition

222 Initialize buffer: $\mathcal{B} \leftarrow \emptyset$

223 $h^* \leftarrow \text{null}$, $s^* \leftarrow -\infty$

224 **for** $n = 1$ **to** N **do**

225 $\mathcal{H}_n \leftarrow \{h_j\}_{j=1}^b$ where $h_j \sim \pi_\theta(\cdot \mid \mathcal{T})$

226 **for** $h \in \mathcal{H}_n$ **do**

227 $s \leftarrow \text{Score}_{\mathcal{T}}(h, \mathcal{D})$

228 **if** $s > s^*$ **then**

229 $h^* \leftarrow h$, $s^* \leftarrow s$

230 $\{u_m\} \leftarrow \text{Decompose}(h)$

231 **for** each term u_m **do**

232 $c_m \leftarrow \text{TermContribution}(u_m, h, s, \mathcal{D})$

233 $h_{\text{ann}} \leftarrow \text{Annotate}(h, \{u_m, c_m\})$

234 $\mathcal{B} \leftarrow \mathcal{B} \cup \{(h_{\text{ann}}, s)\}$

235 $\mathbf{p}_n \leftarrow \text{MakeFewShotPrompt}(\mathcal{B}, k)$

236 $\{h_i^n\}_{i=1}^G \sim \pi_\theta(\cdot \mid \mathbf{p}_n)$

237 $\{r_i^n\}_{i=1}^G \leftarrow \text{Score}_{\mathcal{T}}(\{h_i^n\}_{i=1}^G, \mathcal{D})$

238 $\pi_\theta \leftarrow \text{GRPO_Update}(\pi_\theta, \{r_i^n\}_{i=1}^G)$

239

240 # Stage 2: Evolutionary search with decomposition

241 $\mathcal{P}_0 \leftarrow \mathcal{B}$

242 **for** $t = 1$ **to** T **do**

243 $\mathbf{p}_t \leftarrow \text{MakeFewShotPrompt}(\mathcal{P}_{t-1}, k)$

244 $\mathcal{H}_t \leftarrow \{h_j\}_{j=1}^b$ where $h_j \sim \pi_\theta(\cdot \mid \mathbf{p}_t)$

245 **for** $h \in \mathcal{H}_t$ **do**

246 $s \leftarrow \text{Score}_{\mathcal{T}}(h, \mathcal{D})$

247 **if** $s > s^*$ **then**

248 $h^* \leftarrow h$, $s^* \leftarrow s$

249 $\{u_m\} \leftarrow \text{Decompose}(h)$

250 **for** each term u_m **do**

251 $c_m \leftarrow \text{TermContribution}(u_m, h, s, \mathcal{D})$

252 $h_{\text{ann}} \leftarrow \text{Annotate}(h, \{u_m, c_m\})$

253 $\mathcal{P}_t \leftarrow \mathcal{P}_{t-1} \cup \{(h_{\text{ann}}, s)\}$

254

255

256

257

258 **Output:** h^* and s^*

258 (u_i, u_j) , we construct ablated hypotheses, denoted $f \setminus u_i(x)$ or $f \setminus u_i, u_j(x)$, by removing the corresponding subtree from the AST and generating a new equation program in which that component no longer participates in the computation. After this symbolic ablation, we re-optimize all remaining parameters of the ablated equation on the training dataset D . This comparison, with each structure re-optimized to its own best-fit parameters, ensures that the performance change is attributable to the structure differences rather than suboptimal tuning and parameterization. We then quantify the contribution of a term by $\Delta_{u_i} = \text{Score}_{\mathcal{T}}(f, D) - \text{Score}_{\mathcal{T}}(f \setminus u_i, D)$ and similarly for pairs $\Delta_{u_i, u_j} = \text{Score}_{\mathcal{T}}(f, D) - \text{Score}_{\mathcal{T}}(f \setminus \{u_i, u_j\}, D)$. These computed contribution signals are then serialized directly into their corresponding programs as *inline comments* (without affecting executability) and stored in the buffer for evolving populations. In the next iteration, the LLM samples from this population, and the decomposition-annotated programs are reused as in-context examples. This design ensures that the model is not only guided by global error metrics but also reflects on explicit evidence of which symbolic building blocks mattered, progressively refining its generation strategy. For more details about the directional feedback mechanism and its implementation, check Appendix B.

270 Algorithm 1 presents the summarized pseudo-code of DecAEvolve. The framework integrates de-
 271 composition throughout the discovery process, beginning with Stage 1 where test-time adaptation
 272 combines GRPO fine-tuning with decomposition feedback: the LLM generates candidate hypotheses,
 273 which are decomposed into symbolic terms, annotated with term-level contributions, and stored
 274 in buffer \mathcal{B} to create decomposition-augmented examples for subsequent GRPO updates. The
 275 Decompose(.) and TermContribution(u_m, \dots) functions refer to term decomposition and contri-
 276 bution score estimation defined in Section 3.2; and the Annotate(.) and MakeFewShotPrompt(.)
 277 functions refer to the prompt updates with decomposition annotations (as in Figure 4), and in-context
 278 examples sampled from buffer (as in (Shojaee et al., 2025a)), respectively. Following adaptation,
 279 Stage 2 performs evolutionary search with decomposition by initializing population \mathcal{P}_0 from the
 280 buffer and iteratively: (i) constructs prompts with in-context few-shot examples from \mathcal{P}_{t-1} , (ii)
 281 generates b candidate hypotheses from the adapted LLM, and (iii) evaluates, decomposes, annotates
 282 with term-level feedback, and updates the population. This unified approach leverages decomposition
 283 for both adaptation and evolutionary refinement, enabling efficient exploration of the equation space
 284 through the synergy of decomposition-guided adaptation and evolutionary search.

285 4 EXPERIMENTS

288 We evaluate DecAEvolve on benchmark datasets for LLM-based scientific equation discovery from
 289 (Shojaee et al., 2025a), covering domains like physics, biology, and materials science:

290 **Nonlinear Oscillator:** Simulates two nonlinear damped oscillators (*Oscillator1*, *Oscillator2*) gov-
 291 erned by second-order differential equations in displacement and velocity. Both systems are designed
 292 with complex but solvable nonlinear structures that differ from standard oscillator models to challenge
 293 LLMs towards discovery through data-driven reasoning.

294 **Bacterial Growth:** Models *E. coli* growth under varying conditions of density, substrate, temperature,
 295 and pH. Novel nonlinear terms designed for temperature and pH introduce complexities that require
 296 exploration and discovery and are hard to recover from LLM recall.

298 **Stress-Strain Behavior:** Captures tensile response of aluminum alloy across temperatures. This
 299 dataset uses experimental measurements, providing a more realistic setting with experimental data
 300 that challenge LLM-based models beyond synthetic formulations.

301 All datasets consist of predefined train, validation, in-domain (ID) test, and out-of-domain (OOD)
 302 test splits. In our experiments, the training split is used for parameter optimization for each equation
 303 skeleton, the validation split is used to compute the feedback score that guides the search in symbolic
 304 space, and the held-out ID and OOD test splits are used solely for evaluation. We compare DecAE-
 305 evolve against state-of-the-art non-LLM symbolic regression (SR) baselines, including approaches
 306 such as GPlearn², PySR³ (Cranmer, 2023), and SINDy (Brunton et al., 2016), deep learning methods
 307 like DSR (Petersen et al., 2021) and uDSR (Landajuela et al., 2022), and pre-trained Transformer
 308 SR models NeSymReS (Biggio et al., 2021) and E2E (Kamienny et al., 2022b) (check Appendix E
 309 for implementation details). In addition, we evaluate against the leading LLM-based SR baseline,
 310 LLM-SR (Shojaee et al., 2025a), under same configurations: 3,000 LLM calls per problem with
 311 sampling temperature $\tau = 0.8$. In both approaches, equation parameters are optimized with the
 312 BFGS solver from SciPy python library and a 30s timeout used for the execution of each hypothesis.
 313 In the GRPO adaptation phase, we use batch size of 16 per device, gradient accumulation 4, learning
 314 rate 10^{-5} , and KL coefficient $\beta = 0.05$. For fine-tuning, we use LoRA adapters with $r = 16$.
 315 Decomposition analysis is also conducted based on the AST extracted from the equation program to
 316 define each term and pairwise term contributions. We conduct experiments on six open-source models
 317 (Qwen2.5-1.5B and Qwen2.5-3B, Qwen2.5-7B, Llama-3.2-1B, Llama-3.1-3B, and Llama-3.1-8B) to
 318 evaluate effectiveness across different model variants as well as the scaling behaviors across different
 319 model capacities within our computational constraints for fine-tuning.

320 For the analysis, we use the normalized mean squared error (NMSE) as in (Shojaee et al., 2025b):

$$321 \text{NMSE} = \frac{\sum_{i=1}^{N_{\text{test}}} (\hat{y}_i - y_i)^2}{\sum_{i=1}^{N_{\text{test}}} (y_i - \bar{y})^2} \text{ on both in-domain (ID) and out-of-domain (OOD) test settings, where } N_{\text{test}}$$

322 ²<https://gplearn.readthedocs.io/en/stable/>

323 ³<https://github.com/MilesCranmer/PySR>

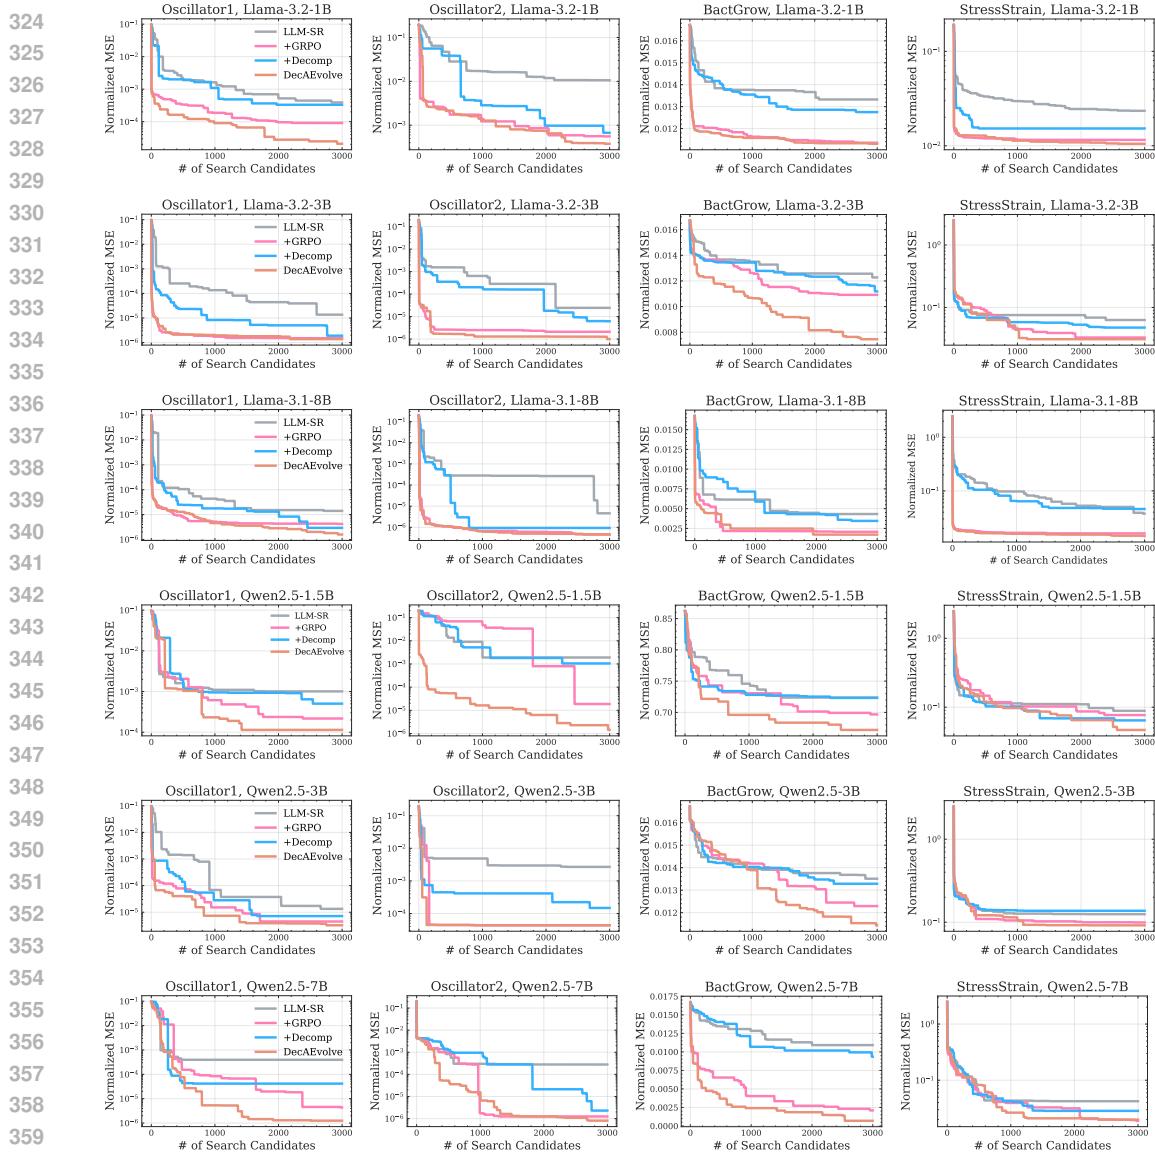


Figure 2: Best-score trajectories of DecAEvolve and its variants against the LLM-SR baseline across benchmark problems. Adaptation (+GRPO) and decomposition (+Decomp) each enhance discovery effectiveness and efficiency, yielding more accurate final equations with fewer search candidates. Their integration in DecAEvolve achieves the best result across all datasets (lower is better). [Each curve shows the average across five runs.](#)

is the test size and \bar{y} the mean target value. NMSE normalizes errors by scale of dataset variance, enabling comparison across datasets.

4.1 RESULTS

To assess the contribution of our proposed framework DecAEvolve and its key components—decomposition and adaptation—on top of the default evolutionary discovery framework, we compared against the LLM-SR baseline under the same LLM backbones across multiple benchmark datasets. Figure 2 reports the discovery trajectories, showing the progression of the best-achieved normalized MSE (NMSE) across the search process. The results highlight three consistent trends. First, both ablated variants improve over the baseline: the **Adaptation** (+GRPO) and **Decomposition** (+Decomp) modules both help to accelerate convergence and lower discovery error. Second, these improvements hold across diverse LLM backbones (Llama-3.2-1B, Llama-3.2-3B, Llama-3.1-8B, Qwen-2.5-1.5B, Qwen-2.5-3B, Qwen-2.5-7B), indicating that the gains are not model-specific but

378 379 380 381 382 383 384 385 386 387 388 389 390 391 392 393 394 395 396	Model	Oscillation 1		Oscillation 2		E. coli growth		Stress-Strain	
		ID↓	OOD↓	ID↓	OOD↓	ID↓	OOD↓	ID↓	OOD↓
GPlearn	0.0155	0.5567	0.7551	3.188	1.081	1.039	0.1063	0.4091	
NeSymReS	0.0047	0.5377	0.2488	0.6472	N/A ($d > 3$)		0.7928	0.6377	
E2E	0.0082	0.3722	0.1401	0.1911	0.6321	1.4467	0.2262	0.5867	
DSR	0.0087	0.2454	0.0580	0.1945	0.9451	2.4291	0.3326	1.108	
uDSR	0.0003	0.0007	0.0032	0.0015	0.3322	5.4584	0.0502	0.1761	
PySR	0.0009	0.3106	0.0002	0.0098	0.0376	1.0141	0.0331	0.1304	
SINDy	0.9888	0.7097	4.6e-16	1.45e-8	1.078	1.039	0.0781	3.5e+15	
LLM-SR (Mixtral)	7.89e-8	0.0002	0.0030	0.0291	0.0026	0.0037	0.0162	0.0946	
LLM-SR (GPT-3.5-turbo)	4.65e-7	0.0005	<u>2.12e-7</u>	3.81e-5	0.0214	0.0264	0.0210	0.0516	
LLM-SR (Llama-3.2-1B)	0.0003	0.1121	0.0105	0.0543	0.0133	0.3544	0.0934	0.3821	
LLM-SR (Llama-3.2-3B)	1.41e-5	0.0014	0.0021	0.0053	0.0122	0.0588	0.0629	0.1672	
LLM-SR (Llama-3.1-8B)	1.36e-5	0.0009	4.61e-6	0.0001	0.0117	0.0240	0.0376	0.0761	
LLM-SR (Qwen2.5-1.5B)	0.0011	0.1233	0.0027	0.0721	0.7237	0.9483	0.1249	0.2435	
LLM-SR (Qwen2.5-3B)	0.0003	0.0168	0.0018	0.0432	0.0135	0.8011	0.0905	0.2085	
LLM-SR (Qwen2.5-7B)	1.33e-5	0.0017	0.0002	0.0011	0.0109	0.1285	0.0423	0.1851	
DecAEvolve (Llama-3.2-1B)	2.09e-5	0.0011	0.0018	0.0136	0.0114	0.0698	0.0704	0.0924	
DecAEvolve (Llama-3.2-3B)	1.57e-6	0.0004	0.0003	0.0005	0.0074	0.0102	0.0311	0.0358	
DecAEvolve (Llama-3.1-8B)	1.37e-6	0.0002	3.64e-7	2.11e-5	0.0019	0.0045	0.0144	0.0322	
DecAEvolve (Qwen2.5-1.5B)	0.0001	0.0784	1.22e-6	0.0012	0.6719	0.9211	0.0916	0.1134	
DecAEvolve (Qwen2.5-3B)	3.23e-6	0.0002	4.36e-5	0.0008	0.0115	0.0454	0.0487	0.1612	
DecAEvolve (Qwen2.5-7B)	1.25e-6	1.51e-5	8.06e-7	<u>1.64e-5</u>	0.0007	0.0012	0.0198	0.0322	

Table 1: Comparison of DecAEvolve with SR baseline models on different scientific benchmark problems, measured by Normalized Mean Squared Error (lower is better) over five runs. **The best performance** for each dataset is in bold, and the second best performance is underlined.

instead stem from the principled design of the framework components that transfer across different backbones. Finally, the full **DecAEvolve** framework, which integrates all three components of evolution, decomposition, and adaptation, consistently delivers the lowest terminal NMSE and the fastest convergence rate in the discovery process.

Table 1 provides a quantitative comparison of DecAEvolve against both non-LLM baselines and the LLM-based baseline LLM-SR across in-domain (ID) and out-of-distribution (OOD) evaluations. We observe that DecAEvolve **mostly** outperforms state-of-the-art non-LLM methods (e.g., PySR, uDSR, **SINDy**) as well as the LLM-SR baseline when evaluated under the same LLM backbones. These improvements are mostly robust across both ID and OOD test sets, demonstrating not only higher accuracy but also stronger generalization to unseen data distributions. Performance gains are particularly pronounced with larger backbones such as Llama-3.1-8B and Qwen-2.5-7B, which is expected given that the success of DecAEvolve relies on two key components: (1) decomposition, which requires sufficient reasoning capacity to interpret granular feedback, and (2) adaptation, which depends on reinforcement learning finetuning to exploit reward signals from observed scientific data. Larger models are better able to leverage both of these mechanisms, resulting in consistently stronger performance. Nevertheless, we also find that DecAEvolve with smaller backbones can achieve results that are competitive with, and in some cases better than, the originally reported LLM-SR performance using much larger models such as Mixtral and GPT-3.5. This underscores the critical role of adaptation in scientific discovery: by tailoring even modestly sized open-source models to the specific scientific system, DecAEvolve can surpass the performance of significantly larger general-purpose models.

Lastly, Figure 3 shows consistent reward improvement during GRPO adaptation across both model scales and all datasets, validating our reinforcement learning fine-tuning approach as test-time adaptation for equation discovery. Notably, we observe some scale-dependent behaviors where smaller models show more noise in their RL and reward improvement process than their larger model counterparts. Interestingly, the smaller model usually matches larger model performance eventually even on complex datasets, suggesting that targeted adaptation through GRPO can help to effectively bridge the capability gap between model scales for scientific discovery tasks.

5 RELATED WORK

Symbolic regression and Scientific Discovery. Early research in symbolic regression and scientific discovery established a foundation for automated equation finding, relying on genetic programming and evolutionary search to explore hypothesis spaces (Koza, 1994a; Cava et al., 2021a). While effective on small problems, these approaches struggled with scalability and tended to rediscover

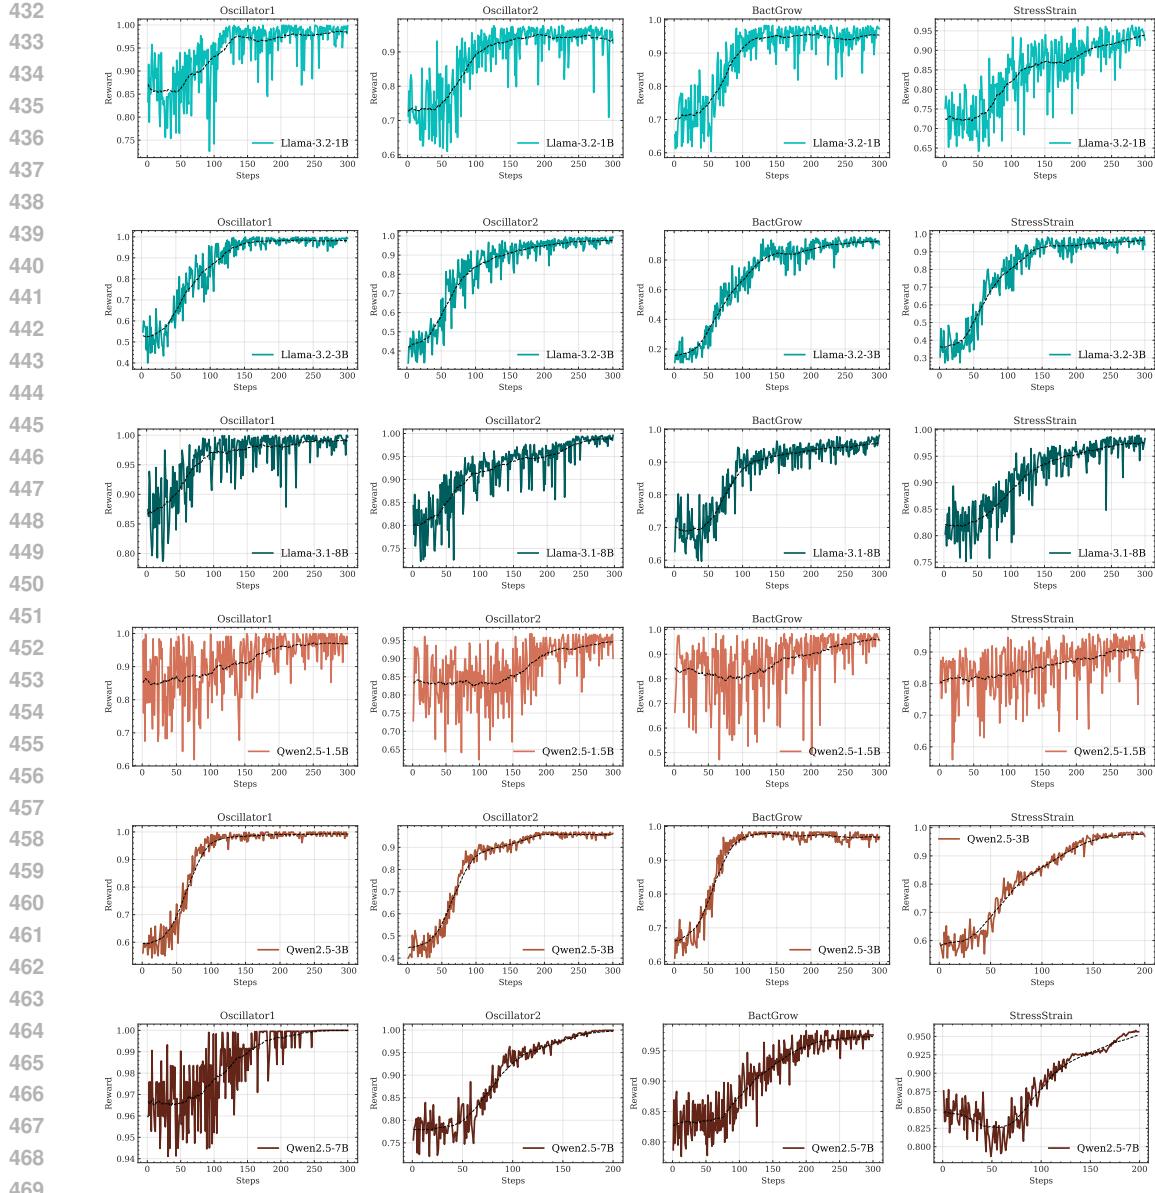


Figure 3: Reward improvements across different models and datasets, showing the success of adaptation with GRPO RL fine-tuning.

shallow functional forms. Later neural-guided methods, such as [sparse regression approaches like SINDy](#) (Brunton et al., 2016), AI Feynman (Udrescu & Tegmark, 2020a), physics-inspired constraints (Bruneton, 2025b), and transformer-based symbolic generation methods (Kamienny & colleagues, 2022; Shojaee et al., 2023b; Meidani et al., 2024a) extended these capabilities but often suffered from poor scalability and limited structural diversity. The rise of large language models shifted this landscape. Works such as LLM-SR (Shojaee et al., 2025a) reframed symbolic regression as program synthesis, allowing models to generate equation skeletons enriched by internal scientific priors. Subsequent frameworks expanded this view: LaSR (Grayeli et al., 2024a) guided search with abstracted concepts extracted from prior successes, while bilevel optimizers (Ma et al., 2024a) combined symbolic hypothesis generation with simulation-driven parameter tuning. Benchmarks such as LLM-SRBench (Shojaee et al., 2025c) highlighted both the promise of these methods and their limitations, showing that LLMs, even when coupled with evolutionary refinement, fails to capture the adaptive strategies that real scientific discovery demands.

486 **Test-time adaptation.** Test-time adaptation has recently emerged as a way to adapt models during
 487 inference, guiding models toward novel distributions without additional offline training. In reasoning
 488 benchmarks such as ARC-AGI (Chollet et al., 2024), gradient-based test-time training (TTT) has
 489 shown great performance in better adapting models to tasks that require more novelty (Akyürek
 490 et al., 2024). The ARC-AGI 2024 report similarly attributes recent state-of-the-art results to pipelines
 491 that incorporate test-time training components into the problem-solving process (Chollet & Team,
 492 2024). Beyond empirical advances, recent theoretical analyses establish conditions under which a
 493 single gradient step at inference provably enhances transformers as in-context learners (Gozeten et al.,
 494 2025). Extending beyond supervised updates, Zuo et al. (2025) introduce test-time reinforcement
 495 learning (TTRL), where models adapt using consensus-based rewards rather than labels, yielding
 496 further improvements across reasoning and math tasks. **Despite the successes and potential benefits of**
 497 **test-time training in tasks that require better adapting to novelty, test-time adaptation remains largely**
 498 **unexplored in scientific discovery frameworks. It is still unclear exactly how inference-time learning**
 499 **can align the priors of a model by leveraging the dynamics of specific scientific system during the**
 500 **evolutionary process of search towards discovery.**

501 **Evolution and Prompt Optimization.** A parallel line of work focuses on evolutionary search and
 502 optimization of prompts rather than model weights, treating instructions and in-context exemplars
 503 as a inference-time search space. Yang et al. (2023) propose OPRO, which frames prompt design
 504 as black-box optimization and iteratively improves instructions through feedback with LLMs as
 505 optimizer. Guo et al. (2025) extend this perspective with EvoPrompt, combining evolutionary
 506 operators such as mutation and crossover with LLMs to explore diverse prompt populations. More
 507 recently, Opsahl-Ong et al. (2024) develop MIPRO, a system that jointly optimizes instructions and
 508 demonstrations in multi-stage LM programs, demonstrating robust improvements without weight
 509 updates. Agrawal et al. (2025) introduce GEPA, which leverages reflective prompt evolution and
 510 self-feedback to surpass reinforcement learning baselines like GRPO, achieving higher efficiency
 511 in both code and reasoning tasks. Surveys on evolution and prompt optimization synthesize these
 512 approaches and position prompt evolution as a label- and compute-efficient alternative to general
 513 RL fine-tuning (Ramnath et al., 2025). **Our framework builds on this motivation of self-evolving**
 514 **optimization via prompting along with the test-time model adaptation to search deeper and more**
 515 **efficient in the large combinatorial hypothesis spaces of scientific discovery.**

516 6 CONCLUSION

518 We introduce DecAEvolve, a framework that enhances LLM-based equation discovery through
 519 granular term-level directional feedbacks, test-time adaptation via GRPO and evolutionary search with
 520 LLMs. Our approach transforms static hypothesis generation into adaptive learning, enabling LLMs
 521 to progressively align with nuances of underlying observed scientific systems through reinforcement
 522 learning model adaptation and interpretable feedback mechanisms. Experimental results across
 523 diverse benchmark datasets demonstrate that DecAEvolve consistently outperforms state-of-the-
 524 art baselines in both discovery accuracy and search efficiency, while maintaining strong out-of-
 525 domain generalization. The success of smaller models through targeted test-time adaptation suggests
 526 promising directions for democratizing scientific discovery tools without requiring large, resource-
 527 intensive models. **Future work could extend our simple decomposition mechanisms to more complex**
 528 **reflection structures and explore better optimization strategies for the evolutionary process. The term-**
 529 **level feedback approach developed here may also prove valuable for systems with highly correlated**
 530 **components and the broader program synthesis tasks requiring iterative refinement in the symbolic**
 531 **space of programs based on component-level understanding.**

540 REFERENCES
541

542 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,
543 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report.
544 *arXiv preprint arXiv:2303.08774*, 2023.

545 Shubham Agrawal et al. Gepa: Reflective prompt evolution can outperform reinforcement learning.
546 *arXiv preprint arXiv:2507.19457*, 2025. URL <https://arxiv.org/abs/2507.19457>.

547 Ekin Akyürek et al. The surprising effectiveness of test-time training for abstract reasoning. In
548 *International Conference on Machine Learning (ICML)*, 2024. URL <https://ekinakyurek.github.io/papers/ttt.pdf>.

549 Luca Biggio, Tommaso Bendinelli, Alexander Neitz, Aurelien Lucchi, and Giambattista Parascandolo.
550 Neural symbolic regression that scales. In Marina Meila and Tong Zhang (eds.), *Proceedings of
551 the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine
552 Learning Research*, pp. 936–945. PMLR, 18–24 Jul 2021.

553 J.-P. Bruneton. Enhancing symbolic regression with quality-diversity and physics-inspired constraints
554 (qdsr). *arXiv preprint*, 2025a. URL <https://arxiv.org/abs/2503.19043>.

555 Jean-Philippe Bruneton. Enhancing symbolic regression with quality-diversity and physics-inspired
556 constraints (qdsr). *arXiv preprint arXiv:2501.01234*, 2025b.

557 Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data
558 by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of
559 sciences*, 113(15):3932–3937, 2016.

560 William La Cava, Patryk Orzechowski, Bogdan Burlacu, Fabrício Olivetti de França, Marco Virgolin,
561 Ying Jin, Michael Kommenda, and Jason H. Moore. Contemporary symbolic regression meth-
562 ods and their relative performance. *Proceedings of the Genetic and Evolutionary Computation
563 Conference (GECCO)*, 2021a.

564 William La Cava, Patryk Orzechowski, Bogdan Burlacu, Fabrício Olivetti de França, Marco Virgolin,
565 Ying Jin, Michael Kommenda, and Jason H. Moore. Contemporary symbolic regression methods
566 and their relative performance, 2021b. URL <https://arxiv.org/abs/2107.14351>.

567 Francois Chollet, Mike Knoop, Gregory Kamradt, and Bryan Landers. Arc prize 2024: Technical
568 report. *arXiv preprint arXiv:2412.04604*, 2024.

569 François Chollet and ARC Prize Team. Arc prize 2024: Technical report. Technical report, ARC
570 Prize, 2024. URL <https://arcprize.org/>.

571 Miles Cranmer. Interpretable machine learning for science with pysr and symbolicregression. jl.
572 *arXiv preprint arXiv:2305.01582*, 2023.

573 Halil Alperen Gozeten, M. Emrullah Ildiz, Xuechen Zhang, Mahdi Soltanolkotabi, Marco Mondelli,
574 and Samet Oymak. Test-time training provably improves transformers as in-context learners, 2025.
575 URL <https://arxiv.org/abs/2503.11842>.

576 Arya Grayeli, Atharva Sehgal, Omar Costilla-Reyes, Miles Cranmer, and Swarat Chaudhuri. Symbolic
577 regression with a learned concept library. In *Advances in Neural Information Processing Systems
578 (NeurIPS)*, 2024a.

579 Arya Grayeli, Atharva Sehgal, Omar Costilla-Reyes, Miles Cranmer, and Swarat Chaudhuri. Symbolic
580 regression with a learned concept library. In *Advances in Neural Information Processing Systems
581 (NeurIPS)*, 2024b. URL <https://arxiv.org/abs/2409.09359>.

582 Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and
583 Yujiu Yang. Evoprompt: Connecting llms with evolutionary algorithms yields powerful prompt
584 optimizers, 2025. URL <https://arxiv.org/abs/2309.08532>.

585 Author Kamienny et al. End-to-end transformer-based equation generation for symbolic regression.
586 In *NeurIPS*, 2022a.

594 Pierre Kamienny and colleagues. End-to-end transformer-based equation generation for symbolic
 595 regression. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
 596

597 Pierre-Alexandre Kamienny, Stéphane d’Ascoli, Guillaume Lample, and Francois Charton. End-
 598 to-end symbolic regression with transformers. In *Advances in Neural Information Processing
 599 Systems*, 2022b.

600 Alan A. Kaptanoglu, Brian M. de Silva, Urban Fasel, Kadierdan Kaheman, Andy J. Goldschmidt,
 601 Jared Callaham, Charles B. Delahunt, Zachary G. Nicolaou, Kathleen Champion, Jean-Christophe
 602 Loiseau, J. Nathan Kutz, and Steven L. Brunton. Pysindy: A comprehensive python package for
 603 robust sparse system identification. *Journal of Open Source Software*, 7(69):3994, 2022. doi:
 604 10.21105/joss.03994. URL <https://doi.org/10.21105/joss.03994>.
 605

606 John R. Koza. *Genetic Programming as a Means for Programming Computers by Natural Selection*,
 607 volume 4. 1994a.

608 John R. Koza. Genetic programming as a means for programming computers by natural selection.
 609 *Statistics and Computing*, 4(2):87–112, 1994b. doi: 10.1007/BF00175355. URL <https://doi.org/10.1007/BF00175355>.
 610

612 Mikel Landajuela, Chak Lee, Jiachen Yang, Ruben Glatt, Claudio P. Santiago, Ignacio Aravena,
 613 Terrell N. Mundhenk, Garrett Mulcahy, and Brenden K. Petersen. A unified framework for deep
 614 symbolic regression. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho
 615 (eds.), *Advances in Neural Information Processing Systems*, 2022.

616 Wanhao Liu, Zonglin Yang, Jue Wang, Lidong Bing, Di Zhang, Dongzhan Zhou, Yuqiang Li,
 617 Houqiang Li, Erik Cambria, and Wanli Ouyang. Moose-chem3: Toward experiment-guided
 618 hypothesis ranking via simulated experimental feedback. *arXiv preprint arXiv:2505.17873*, 2025.

619

620 Pingchuan Ma, Tsun-Hsuan Wang, Minghao Guo, Zhiqing Sun, Joshua B. Tenenbaum, Daniela Rus,
 621 Chuang Gan, and Wojciech Matusik. Llm and simulation as bilevel optimizers: A new paradigm to
 622 advance physical scientific discovery. In *International Conference on Machine Learning (ICML)*,
 623 2024a.

624

625 Pingchuan Ma, Tsun-Hsuan Wang, Minghao Guo, Zhiqing Sun, Joshua B. Tenenbaum, Daniela
 626 Rus, Chuang Gan, and Wojciech Matusik. LLM and simulation as bilevel optimizers: A new
 627 paradigm to advance physical scientific discovery. In *Forty-first International Conference on
 628 Machine Learning*, 2024b. URL <https://openreview.net/forum?id=hz8cFsdz7P>.

629

630 Nour Makke and Sanjay Chawla. Interpretable scientific discovery with symbolic regression: a
 631 review. *Artificial Intelligence Review*, 57(1):2, 2024.

632

633 Kazem Meidani, Parshin Shojaee, Chandan K. Reddy, and Amir Barati Farimani. Snip: Bridging
 634 mathematical symbolic and numeric realms with unified pre-training. In *International Conference
 635 on Learning Representations (ICLR)*, 2024a.

636

637 Kazem Meidani, Parshin Shojaee, Chandan K. Reddy, and Amir Barati Farimani. SNIP: Bridging
 638 mathematical symbolic and numeric realms with unified pre-training. In *The Twelfth International
 639 Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=KZSEgJGPxu>.

640

641 Alexander Novikov, Ngan Vn, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wag-
 642 ner, Sergey Shirobokov, Borislav Kozlovskii, Francisco J. R. Ruiz, Abbas Mehrabian, M. Pawan
 643 Kumar, Abigail See, Swarat Chaudhuri, George Holland, Alex Davies, Sebastian Nowozin, Push-
 644 meet Kohli, and Matej Balog. Alphaevolve: A coding agent for scientific and algorithmic discovery,
 645 2025.

646

647 Krista Opsahl-Ong, Michael J Ryan, Josh Purtell, David Broman, Christopher Potts, Matei Zaharia,
 648 and Omar Khattab. Optimizing instructions and demonstrations for multi-stage language model
 649 programs, 2024. URL <https://arxiv.org/abs/2406.11695>.

648 Brenden K Petersen, Mikel Landajuela Larma, Terrell N. Mundhenk, Claudio Prata Santiago,
 649 Soo Kyung Kim, and Joanne Taery Kim. Deep symbolic regression: Recovering mathemat-
 650 ical expressions from data via risk-seeking policy gradients. In *International Conference on*
 651 *Learning Representations*, 2021.

652 Kiran Ramnath, Kang Zhou, Sheng Guan, Soumya Smruti Mishra, Xuan Qi, Zhengyuan Shen, Shuai
 653 Wang, Sangmin Woo, Sullam Jeoung, Yawei Wang, Haozhu Wang, Han Ding, Yuzhe Lu, Zhichao
 654 Xu, Yun Zhou, Balasubramaniam Srinivasan, Qiaojing Yan, Yueyan Chen, Haibo Ding, Panpan
 655 Xu, and Lin Lee Cheong. A systematic survey of automatic prompt optimization techniques, 2025.
 656 URL <https://arxiv.org/abs/2502.16923>.

657 Chandan K Reddy and Parshin Shojaee. Towards scientific discovery with generative ai: Progress,
 658 opportunities, and challenges. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
 659 volume 39, pp. 28601–28609, 2025.

660 Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog,
 661 M. Pawan Kumar, Emilien Dupont, Francisco J. R. Ruiz, Jordan S. Ellenberg, Pengming
 662 Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein Fawzi. Mathematical discoveries from
 663 program search with large language models. *Nat.*, 625(7995):468–475, January 2024. URL
 664 <https://doi.org/10.1038/s41586-023-06924-6>.

665 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 666 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of
 667 mathematical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.

668 Parshin Shojaee, Kazem Meidani, Amir Barati Farimani, and Chandan Reddy. Transformer-
 669 based planning for symbolic regression. In A. Oh, T. Naumann, A. Globerson,
 670 K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information
 Processing Systems*, volume 36, pp. 45907–45919. Curran Associates, Inc.,
 671 2023a. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/8fffb4e3118280a66b192b6f06e0e2596-Paper-Conference.pdf.

672 Parshin Shojaee, Kazem Meidani, Amir Barati Farimani, and Chandan K. Reddy. Transformer-based
 673 planning for symbolic regression, 2023b. URL <https://arxiv.org/abs/2303.06833>.

674 Parshin Shojaee, Kazem Meidani, Shashank Gupta, Amir Barati Farimani, and Chandan K Reddy.
 675 Llm-sr: Scientific equation discovery via programming with large language models, 2025a. URL
 676 <https://arxiv.org/abs/2404.18400>.

677 Parshin Shojaee, Ngoc-Hieu Nguyen, Kazem Meidani, Amir Barati Farimani, Khoa D Doan, and
 678 Chandan K. Reddy. LLM-SRBench: A new benchmark for scientific equation discovery with large
 679 language models. In *Forty-second International Conference on Machine Learning*, 2025b. URL
 680 <https://openreview.net/forum?id=SyQPiZJVWY>.

681 Parshin Shojaee, Ngoc-Hieu Nguyen, Kazem Meidani, Amir Barati Farimani, Khoa D Doan, and
 682 Chandan K Reddy. Llm-srbench: A new benchmark for scientific equation discovery with large
 683 language models, 2025c. URL <https://arxiv.org/abs/2504.10415>.

684 Anja Surina, Amin Mansouri, Lars Quaedvlieg, Amal Seddas, Maryna Viazovska, Emmanuel Abbe,
 685 and Caglar Gulcehre. Algorithm discovery with llms: Evolutionary search meets reinforcement
 686 learning. *arXiv preprint arXiv:2504.05108*, 2025.

687 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
 688 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand
 689 Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language
 690 models, 2023. URL <https://arxiv.org/abs/2302.13971>.

691 Silviu-Marian Udrescu and Max Tegmark. Ai feynman: A physics-inspired method for symbolic
 692 regression. *Science Advances*, 6(16), 2020a.

693 Silviu-Marian Udrescu and Max Tegmark. Ai feynman: a physics-inspired method for symbolic
 694 regression, 2020b. URL <https://arxiv.org/abs/1905.11481>.

702 Marco Virgolin and Solon P. Pissis. Symbolic regression is np-hard, 2022. URL <https://arxiv.org/abs/2207.01018>.
 703
 704

705 Greg Yang et al. Large language models as optimizers. In *NeurIPS*, 2023. URL <https://arxiv.org/abs/2309.03409>.
 706
 707

708 Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, Xuekai Zhu, Haozhan Li, Yuchen Zhang,
 709 Xinwei Long, Ermo Hua, Biqing Qi, Youbang Sun, Zhiyuan Ma, Lifan Yuan, Ning Ding, and
 710 Bowen Zhou. Ttrl: Test-time reinforcement learning, 2025. URL <https://arxiv.org/abs/2504.16084>.
 711
 712

APPENDIX

A LLM-SR MULTI-ISLAND EVOLUTIONARY BUFFER & SAMPLING

713 LLM-SR (Shojaee et al., 2025a) maintains a multi-island evolutionary experience buffer designed
 714 to preserve structural diversity and prevent premature convergence during the process of symbolic
 715 equation discovery. The buffer is organized as $\mathcal{P}_t = \bigcup_i \mathcal{P}_t^{(i)}$, where each $\mathcal{P}_t^{(i)}$ is an independently
 716 evolving island/population containing pairs of symbolic programs and their corresponding scalar
 717 fitness scores. Each island begins with the initial prompt equation v_0 with score s_0 supplied to
 718 the framework at the beginning: $\mathcal{P}_0^{(i)} = \{(v_0, s_0)\}$. Although all islands start identically, they
 719 diverge over the search process and evolution. At iteration t , the LLM generates a batch of
 720 programs $\mathcal{F}_t = \{f_j\}_{j=1}^b$, each associated with a source island, which is the island from which
 721 its in-context examples were sampled. Each equation program candidates contains a placeholder
 722 vector of parameters which are then optimized with respect to the observed data with the help of
 723 BFGS optimizer in python. After parameter optimization, each program receives a fitness score
 724 $\text{Score}_{\mathcal{T}}(f, \mathcal{D})$ which is computed as negative mean squared error (MSE) with respect to data:
 725
 726

$$\text{Score}_{\mathcal{T}}(f, \mathcal{D}) = -\frac{1}{n} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i)^2$$

727 , where $\mathcal{D} = \{(\mathbf{x}_i, y_i)\} \in \mathbb{R}^d \times \mathbb{R}$ refers to observed datapoints. Notably, a program is inserted *only*
 728 into its source island if its fitness exceeds that island's current best:
 729

$$\mathcal{P}_t^{(i)} \leftarrow \mathcal{P}_{t-1}^{(i)} \cup \{(f, s)\} \text{ if } s > s_{\text{best}}^{(i)}$$

730 This ensures that each island progressively specializes in a distinct region of the hypothesis space
 731 to avoid local minima and encourage diversity in the search process towards discovery. In LLM-SR,
 732 programs are also grouped into clusters within each island using a simple signature corresponding
 733 to score s where all programs with identical fitness scores fall into the same cluster. This prevents
 734 over-representation of structurally similar programs and maintains a lower-level diversity within
 735 each island. After every $T_{\text{reset}} = 4hr$, the algorithm identifies the worst-performing half of the
 736 islands as $\mathcal{W} = \arg \min_i s_{\text{best}}^{(i)}$. For each $i \in \mathcal{W}$, the entire island is replaced by a copy of the
 737 best-performing equation from a randomly selected surviving island. This mechanism is designed
 738 to remove stagnating islands and increases global exploration during the discovery .
 739

740 Each iteration begins by sampling k equations from the multi-island buffer to construct the few-shot
 741 prompt. Sampling follows a hierarchical procedure. First, an island index is sampled uniformly
 742 $i \sim \text{Uniform}\{1, \dots, I\}$. This is mainly to prevent dominant islands from monopolizing the prompt,
 743 encouraging cross-island exploration. Within the selected island $\mathcal{P}_t^{(i)}$, sampling proceeds in two
 744 steps: *Cluster Selection* and *Program Selection*.
 745

746 For each cluster c , we have the mean fitness $s_c = \text{mean}\{s : (f, s) \in c\}$ and clusters are sampled
 747 according to Boltzmann weights as:
 748

$$p(c) = \frac{\exp(s_c/\tau_c)}{\sum_{c'} \exp(s_{c'}/\tau_c)}$$

749 . In this sampling, the temperature τ_c anneals with island size u as
 750

$$\tau_c = T_0 \left(1 - \frac{u \bmod N}{N}\right)$$

756 where $T_0 = 0.1$ and $N = 10,000$. Within a cluster, programs are sampled with preference for better
 757 scores and shorter length. The sampling distribution for each program follows
 758

$$759 \quad p(f_j) \propto \exp(-\tilde{\ell}_j/\tau_p)$$

760 where sampling temperature $\tau_p = 1$. Here, f_j refers to the program, ℓ_j refers to the corresponding
 761 length, and $\tilde{\ell}_j$ is defined as
 762

$$763 \quad \tilde{\ell}_j = \frac{\ell_j - \min_j \ell_j}{\max_j \ell_j + 10^{-6}}$$

765 After sampling in-context examples with above procedure, the k sampled programs are serialized into
 766 a structured few-shot prompt supplied to the LLM. This prompt acts as the guiding context for the
 767 next generation of hypotheses \mathcal{F}_{t+1} , completing the LLM-SR evolutionary search and self-reflection
 768 mechanism.

770 B DETAILED TERM DECOMPOSITION AND CONTRIBUTION ATTRIBUTION

772 We give a complete account of how generated programs are decomposed into symbolic terms and
 773 how single and pairwise contributions are computed in our implementation. This procedure follows
 774 the implementation of evaluator in Shojaee et al. (2025a) with the additional steps: the function
 775 body is parsed into an abstract syntax tree (AST), simple assignment chains are inlined, the returned
 776 expression is decomposed at additive nodes, and ablation is carried out by rewriting only the final
 777 assignment `return` and re-executing in a sandbox. *All annotations are serialized as inline comments*
 778 *in the program without changing executable semantics*. Unless otherwise specified, all ablations
 779 include re-optimizing of the remaining parameters on the dataset to ensure that contributions reflect
 780 structural differences on instead of suboptimal tuning or parameterization

781 Operationally, the evaluator isolates the evolved function body, executes it in a sandbox to obtain a
 782 scalar score, and then reconstructs the returned expression by expanding intermediate assignments
 783 via an assignment map and dependency graph before parsing with Python’s `ast` module. From the
 784 resulting atoms $\{\tau_t\}$, we perform ablation-based attribution: for each term t , we remove τ_t , rebuild
 785 a syntactically valid RHS with correct parentheses, re-execute the modified program, and compute
 786 a marginal contribution $\Delta_t = \text{Score}_{\mathcal{T}}(f, D) - \text{Score}_{\mathcal{T}}(f_{\setminus t}, D)$, where S is the evaluator score
 787 (negative MSE). We analogously compute pairwise signals $\Delta_{t,u}$ by removing (τ_t, τ_u) . The evaluator
 788 writes these results back as inline comments within the function body, so subsequent iterations can
 789 consume structured, term-level feedback rather than a single scalar reward. This AST-centric pipeline
 790 is lightweight, robust to multi-line programs, and provides the granular guidance that underpins our
 791 evolutionary refinement.

792 **Decomposition to atomic terms** Let \hat{y} denote the expression returned by the function (or a final
 793 assigned variable). We parse the equation program skeleton (body of the function) into an AST, build
 794 a line-level assignment map, inline \hat{y} if it is an intermediate variable, and traverse the AST with
 795 the following rules: (i) *addition/subtraction* split terms, (ii) *multiplication/division/power* subtrees
 796 are preserved as atomic units, and (iii) *unary operators and function calls* (e.g., `sin`, `exp`, `np.abs`)
 797 are atomic operators. The resulting model has the form $\hat{y} = f(\mathbf{x}; \text{params})$ with

798 **Single-term ablation and contribution.** After identifying terms from previous step, For each
 799 term u_m , we form the ablated version of the original equation program as $f_{\setminus u_m}$. After ablation,
 800 the remaining equation parameters from placeholder parameter vector `params` ($\text{params}_{\setminus u_m}$) are re-
 801 optimized on dataset \mathcal{D} (using the same BFGS parameter optimization procedure as in (Shojaee et al.,
 802 2025a) with Scipy library in Python). The term contribution is defined as influence function with:

$$804 \quad \Delta_{u_m} \triangleq \text{Score}_{\mathcal{T}}(f, \mathcal{D}) - \text{Score}_{\mathcal{T}}(f_{\setminus u_m}, \mathcal{D}). \quad (2)$$

806 If removal yields invalid outputs, we treat term u_m as essential and assign maximal contribution
 807 under the current score scale.

808 **Pairwise interaction.** Similarl to the single-term contribution score estimation, for a pair (u_m, u_n)
 809 we define ablated function as $f_{\setminus \{u_m, u_n\}}$ and the corresponding contribution score for this pair-wise

```

810
811     def equation(x, v, params):
812         """ Mathematical function for acceleration in damped nonlinear oscillator """
813
814         # Individual Term Contributions:
815         # [Ablation] Removing this term decreases the score by 0.00394026: params[0]*x
816         # [Ablation] Removing this term decreases the score by 0.00000446: params[2]
817         # [Ablation] Removing this term decreases the score by 0.00000001: params[1]*v
818
819         # Term Pair Contributions:
820         # [Ablation] Removing these terms together decreases the score by 0.00414153:
821         #   Term 1: params[0]*x
822         #   Term 2: params[2]
823         # [Ablation] Removing these terms together decreases the score by 0.00394474:
824         #   Term 1: params[0]*x
825         #   Term 2: params[1]*v
826         # [Ablation] Removing these terms together decreases the score by 0.00000450:
827         #   Term 1: params[1]*v
828         #   Term 2: params[2]
829
830         return params[0] * x + params[1] * v + params[2]
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

```

Figure 4: **A simple example of program-level annotations.** Candidate equation for a damped nonlinear oscillator (simple linear here) is annotated in-line with single-term and pairwise ablation contributions (as comments) immediately above the `return` statement. The evaluator computes these contribution scores after re-optimizing remaining parameters via BFGS, decomposing the return expressions, and re-evaluating ablations in a sandbox.

interaction subtree as:

$$\Delta_{u_m, u_n} \triangleq \text{Score}_{\mathcal{T}}(f, \mathcal{D}) - \text{Score}_{\mathcal{T}}(f_{\setminus \{u_m, u_n\}}, \mathcal{D}). \quad (3)$$

As with the single-term case, the remaining parameters after pairwise term ablation are re-optimized after removal to isolate the structural interaction effect without suboptimal parameterization. These values reveal redundancy versus synergy by comparing Δ_{u_m, u_n} against the single-term counterparts Δ_{u_m} and Δ_{u_n} .

Annotation and persistence. After computing $\{\Delta_{u_m}\}$ and $\{\Delta_{u_m, u_n}\}$, we serialize them as human-readable inline comments directly in the equation skeleton programming function body (above the `return` statement), and store the annotated program in the experience buffer. This preserves executable semantics while exposing interpretable, decomposition-based directional feedback that guides subsequent process of discovery towards better sub-terms in the hypothesis space.

B.1 AST-BASED DECOMPOSITION

An Abstract Syntax Tree (AST) is a language-agnostic representation of a program that makes explicit the hierarchical composition of an expression. Internal nodes correspond to operators or function applications (e.g., `+`, `*`, `**`, `np.sin`), and leaves correspond to parameters or input variables. In our setting, we parse each LLM-generated equation skeleton Python program into an AST and split at additive nodes, while preserving multiplicative, divisional, power, and functional subtrees as atomic terms. This yields a linear combination $f(x) = \sum_{m=1}^M u_m(x)$ of symbolic atoms u_m . For example, Figure 5 illustrates this mapping from the generated hypothesis function (left) to its AST (right): a return expression such as $y = t_1 + t_2 + t_3 - t_4$ becomes a top-level sum where each t_i is an intact subtree, enabling principled symbolic decomposition without altering operator precedence.

C DETAILED TEST-TIME ADAPTATION WITH REINFORCEMENT LEARNING

We formulate our test-time training/adaptation procedure as reinforcement learning over a deterministic Markov Decision Process (MDP) $\mathcal{M} = (S, A, R, T)$.

```

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918

```

```

def equation(x1, x2):
    t1 = param[0] * np.sin(x1)
    t2 = param[1] * (x1 ** 2)
    t3 = param[2] * np.cos(x2)
    y = t1 + t2 + t3 - param[3]
    return y

```

AST

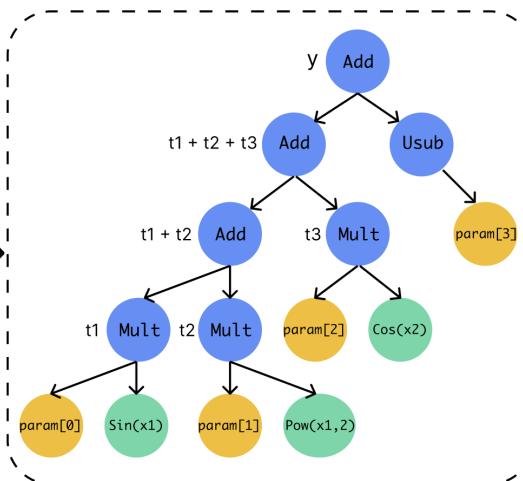


Figure 5: Parsing an equation program into an AST.

States. Similar to the reinforcement learning LLM fine-tuning, the state space S in this setting consists of partial sequences $(p, h_{1:k})$ where p is the prompt and $h_{1:t}$ is a prefix of the generated output until token t .

Actions. At each step of this reinforcement learning test-time adaptation, the action space A corresponds to the next token from vocabulary $h_{t+1} \in V$.

Rewards. In our setting, rewards are assigned based on the execution of the equation program on the observed data and only obtained at terminal states. For a completed program h , we execute the synthesized program on the validation set and compute the reward as noted in Section 3.1 with $r(p, h) = \text{Score}_{\mathcal{T}}(h, \mathcal{D}) = \exp(-\text{MSE}(h, \mathcal{D}))$ with invalid completions included in training and assigned a fixed floor reward of $r = 0.01$.

Implementation details After generating a new batch of candidate equation hypotheses and evaluating their performance, the model updates its policy through GRPO (Shao et al., 2024). At iteration n , we treat the prompt p_n together with its group of completions $h_i^{nG}_{i=1}$ and their associated rewards $\{r_i^n\}_{i=1}^G$ as a single training step. These prompt-completion groups accumulate over iterations and serve as the training data for GRPO updates. We fine-tune using Adam with learning rate 10^{-6} and a warmup-stable-decay schedule (200 warmup steps). Training uses an effective batch size of 64 (16 per device with gradient accumulation 4). Each prompt is sampled with $G = 64$ completions at temperature 0.8 and top- $p = 0.9$. Only LoRA adapter parameters are updated ($r = 16$, $\alpha = 16$, dropout 0.05), while the base model remains frozen as π_{ref} .

D HOW IS DECOMPOSITION DIFFERENT THAN MOOSE-CHEM3?

One of the relevant works from literature with the same high-level motivation of incorporating decomposition in the scientific discovery is MOOSE-Chem3 (Liu et al., 2025). MOOSE-Chem3 targets the discovery and ranking of experimental chemistry hypotheses. Each hypothesis is a natural-language description of a reaction mechanism or material design strategy. Their decomposition procedure is mechanistic: the framework identifies functional chemical components (e.g., polymer matrices, redox pairs, electrode structures, etc.) to build a qualitative similarity measure for experiment-guided hypothesis ranking. This process is domain-specific, mechanistic, and grounded in wet-lab feedback or its simulated analogue. In contrast, DecAEvolve focuses on scientific equation discovery, where hypotheses are explicit symbolic expressions. Our decomposition is quantitative and analytic: an equation is ablated into symbolic terms and operators, and an influence function is used to measure each component’s contribution to prediction error and structural behavior with respect to data. These influence signals directly shape mutation choices, and the structured feedback used during evolution.

918 Although both frameworks use the general idea of breaking hypotheses into components to find
 919 successful components, the decomposition differ fundamentally in (i) the objects being decomposed
 920 (mechanistic chemical descriptions vs. symbolic math relations), (ii) the nature of decomposition
 921 (qualitative functional roles vs. quantitative influence), and (iii) the role of decomposition within the
 922 discovery framework (similarity-based ranking vs. fine-grained feedback for optimization). More-
 923 over, decomposition is only one part of DecAEvolve that consists of a hybrid three-module system
 924 (Decompose, Adapt, Evolve), where term-level analysis interacts directly with RL-based test-time
 925 adaptation, an approach not explored in MOOSE-Chem3 or prior discovery methods.

927 E BASELINE IMPLEMENTATION DETAILS

929 We evaluate DecAEvolve against the baseline methods reported in Shojaee et al. (2025a), including
 930 GPlearn, PySR, DSR, uDSR, NeSymReS, and E2E. These baselines represent diverse SR approaches
 931 spanning genetic programming, reinforcement learning, and pre-trained transformers. For implemen-
 932 tation details and hyperparameters of these methods, we refer readers to Appendix A of Shojaee et al.
 933 (2025a). Below, we provide additional details for the SINDy baseline, which we newly evaluate in
 934 this work.

936 F ADDITIONAL EXPERIMENTS

938 F.1 COMPARISON WITH SINDY BASELINE

940 We evaluate SINDy: Sparse Identification of Nonlinear Dynamics method (Brunton et al., 2016)
 941 as an additional non-LLM baseline using the PySINDy implementation (Kaptanoglu et al., 2022).
 942 SINDy discovers governing equations by performing sparse regression over a predefined library of
 943 candidate functions, assuming dynamics can be expressed as sparse linear combinations of nonlinear
 944 basis functions.

947 **Experimental Setup.** We configured SINDy with a polynomial library (degree 3) augmented
 948 with Fourier terms (frequencies up to 2) to provide a balanced function library without excessive
 949 expansion. The STLSQ optimizer used sparsity threshold 0.1, regularization $\alpha = 0.01$, and 20
 950 maximum iterations.

Model	Oscillation 1		Oscillation 2		E. coli growth		Stress-Strain	
	ID \downarrow	OOD \downarrow						
PySR	0.0009	0.3106	0.0002	0.0098	0.0376	1.0141	<u>0.0331</u>	0.1304
SINDy	0.9888	0.7097	4.62e-16	1.45e-8	1.078	1.039	0.0781	3.52e+15
LLM-SR (Qwen2.5-7B)	<u>1.33e-5</u>	<u>0.0017</u>	0.0002	0.0011	<u>0.0109</u>	<u>0.1285</u>	0.0423	0.1851
DecAEvolve (Qwen2.5-7B)	1.25e-6	1.51e-5	8.06e-7	1.64e-5	0.0007	0.0012	0.0198	0.0322

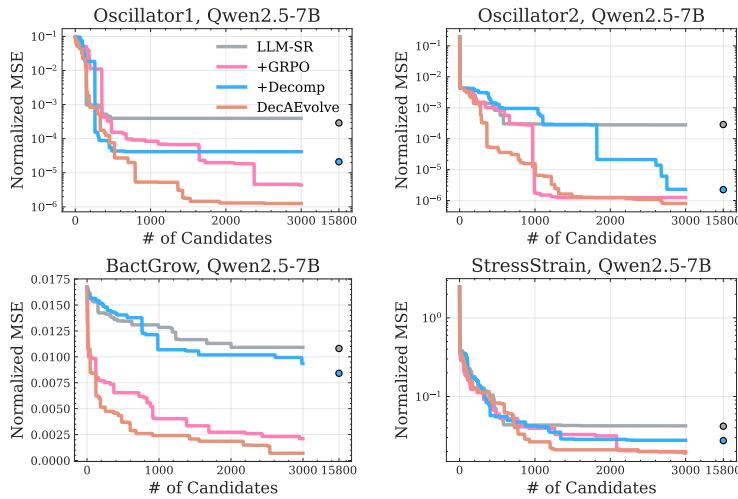
951 Table 2: Comparison of SINDy with representative SR baselines (best non-LLM method PySR, LLM-based
 952 LLM-SR, and our DecAEvolve) using Qwen2.5-7B backbone.

953
 954
 955
 956
 957 **Results and Analysis.** Table 2 shows SINDy achieves the best performance on Oscillation 2 (NMSE
 958 $< 10^{-7}$ on OOD) among all baselines. This is because Oscillation 2 represents a typical dynamical
 959 system, the type SINDy was explicitly designed for, where dynamics follow a linear form with
 960 nonlinear basis functions (polynomials and trigonometrics). Given an appropriate library containing
 961 the true functional forms, SINDy is highly efficient for such problems. However, SINDy exhibits
 962 severe limitations on problems outside its linear-formulation assumption. On Bacterial Growth, it
 963 completely fails (NMSE ≈ 1.08) as the ground-truth involves products of multiple nonlinear terms,
 964 which cannot be represented as linear combinations. On Stress-Strain, while fitting training data
 965 reasonably (NMSE = 0.078), it suffers catastrophic extrapolation failure on OOD data (NMSE
 966 = 3.52×10^{15}) as polynomial approximations diverge outside the training range. On Oscillation 1,
 967 performance degrades significantly (NMSE = 0.99 ID) when dynamics deviate from SINDy’s
 968 assumed form, demonstrating sensitivity to equation structure even within the dynamical systems
 969
 970
 971

972 domain. These results highlight SINDy’s fundamental constraint: performance is entirely determined
 973 by whether the true equation lies within the predefined library’s representational capacity.
 974

978 F.2 ADDITIONAL SAMPLE BUDGET FOR INFERENCE FRAMEWORKS

981 We also conducted additional experiments to run inference search variants (LLM-SR and +Decomp)
 982 for total number of samples used by the GRPO-based variants (DecAEvolve and +GRPO). Specifically,
 983 we run LLM-SR and +Decomp for additional 12800 samples (64×200) with Qwen2.5-7B model
 984 backbone. As it can be observed from Figure 2 and Figure 6, the performance of LLM-SR and
 985 +Decomp is already well-converged by roughly 3000 samples, and allocating an additional 12800
 986 samples offers no meaningful performance improvement. This result shows that DecAEvolve
 987 leverages these samples more effectively for discovery than baselines.



1005 Figure 6: The impact of additional GRPO-matched samples on inference algorithms.

1009 F.3 IMPACT OF PARAMETER RE-OPTIMIZATION IN DECOMPOSITION

1012 We conducted additional ablation study to compare the impact of two strategies of parameter opti-
 1013 mization during the decomposition with structure ablations: (1) *Re-opt Parameters*: After ablating a
 1014 decomposed sub-term, we re-optimize the remaining parameters to obtain the best fit for the modified
 1015 structure; and (2) *Freeze Parameters*: After ablating a term, we keep the remaining parameters fixed
 1016 to the values learned from the original full structure of equation. Symbolic terms of an equation
 1017 are often highly coupled, and modifying part of the structure can shift the optimal values of the
 1018 remaining parameters. This coupling is indeed a general challenge across all equation discovery
 1019 methods. However, our goal is to attribute performance changes to structural differences, not to
 1020 artifacts of suboptimal parameterization. Re-optimizing parameters ensures that each ablated structure
 1021 is evaluated at its own best performance—providing a more faithful estimate of the true contribution
 1022 of each symbolic component. To examine this empirically, we performed a focused ablation study
 1023 using the Qwen2.5-7B backbone across several representative datasets (Oscillator1, Oscillator2,
 1024 BactGrow, and StressStrain). Results are shown in Figure 7. The “Re-opt Params” curves consistently
 1025 achieve lower normalized MSE than the “Freeze Params” curves, particularly in datasets with strong
 1026 nonlinear coupling such as BactGrow and StressStrain. This indicates that re-optimization yields
 1027 more reliable and stable assessments of structural contributions.

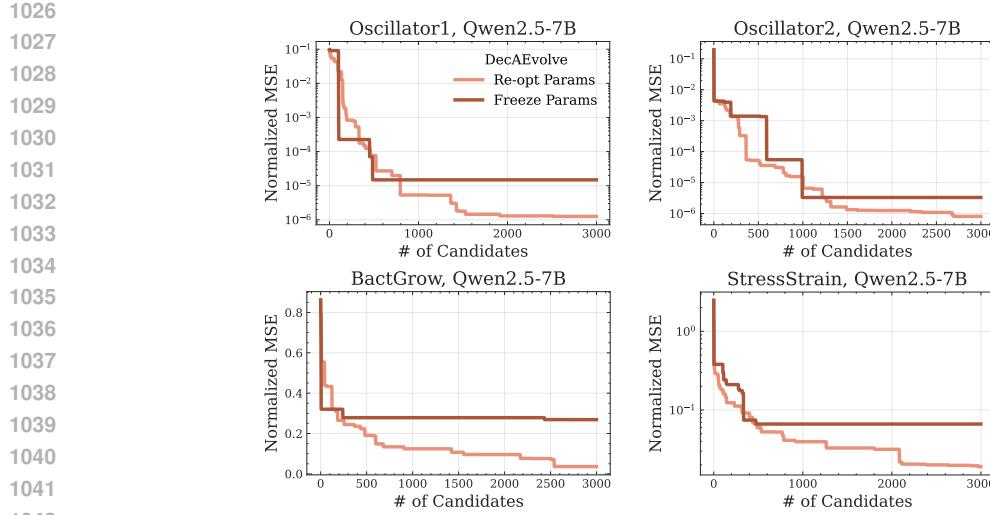


Figure 7: [Ablation results comparing DecAEvolve decomposition with and without re-optimizing the parameters after structure ablation.](#)

G INITIAL INPUT PROMPTS

The prompts in Figs. 8–11 were used for evaluating DecAEvolve on the four regression tasks from the LLM-SR(Shojaee et al., 2025a) benchmark. Each prompt specifies the target problem and a Python function template with placeholder parameters. In all cases, the prompts shown below correspond to the initial call to the LLM. In subsequent iterations, DecAEvolve augments the prompt with examples drawn from top-performing hypotheses in the evolving buffer, enabling in-context learning from previously discovered candidates and their annotated contributions.

```

1080
1081
1082
1083
1084
1085 You are a helpful assistant tasked with discovering mathematical function structures for scientific
1086 systems. You are given an example of the function signature in the first function equation_v0 below.
1087 Your task is to complete the last 'equation' function with your mathematical relationship, considering
1088 the physical meaning and relationships of inputs. Only give me the completion code of the BODY of the
1089 current 'equation' FUNCTION. Do NOT give me a new function. Do NOT give me 'pass' instead of
1090 completion. Just give me the new mathematical relationship with inputs for completion of current
1091 function body. Do NOT use equation_v0 in your completion.
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112 You are a helpful assistant tasked with discovering mathematical function structures for scientific
1113 systems. You are given an example of the function signature in the first function equation_v0 below.
1114 Your task is to complete the last 'equation' function with your mathematical relationship, considering
1115 the physical meaning and relationships of inputs. Only give me the completion code of the BODY of the
1116 current 'equation' FUNCTION. Do NOT give me a new function. Do NOT give me 'pass' instead of
1117 completion. Just give me the new mathematical relationship with inputs for completion of current
1118 function body. Do NOT use equation_v0 in your completion.
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

```

Figure 8: Oscillator I input prompt used for evaluating DecAEvolve.

```

1103
1104
1105
1106
1107
1108
1109
1110
1111
1112 You are a helpful assistant tasked with discovering mathematical function structures for scientific
1113 systems. You are given an example of the function signature in the first function equation_v0 below.
1114 Your task is to complete the last 'equation' function with your mathematical relationship, considering
1115 the physical meaning and relationships of inputs. Only give me the completion code of the BODY of the
1116 current 'equation' FUNCTION. Do NOT give me a new function. Do NOT give me 'pass' instead of
1117 completion. Just give me the new mathematical relationship with inputs for completion of current
1118 function body. Do NOT use equation_v0 in your completion.
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

```

Figure 9: Oscillator II input prompt used for evaluating DecAEvolve.

```

1134
1135
1136
1137
1138
1139 You are a helpful assistant tasked with discovering mathematical function structures for scientific
1140 systems. You are given an example of the function signature in the first function equation_v0 below.
1141 Your task is to complete the last 'equation' function with your mathematical relationship, considering
1142 the physical meaning and relationships of inputs. Only give me the completion code of the BODY of the
1143 current 'equation' FUNCTION. Do NOT give me a new function. Do NOT give me 'pass' instead of
1144 completion. Just give me the new mathematical relationship with inputs for completion of current
1145 function body. Do NOT use equation_v0 in your completion.
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167 You are a helpful assistant tasked with discovering mathematical function structures for scientific
1168 systems. You are given an example of the function signature in the first function equation_v0 below.
1169 Your task is to complete the last 'equation' function with your mathematical relationship, considering
1170 the physical meaning and relationships of inputs. Only give me the completion code of the BODY of the
1171 current 'equation' FUNCTION. Do NOT give me a new function. Do NOT give me 'pass' instead of
1172 completion. Just give me the new mathematical relationship with inputs for completion of current
1173 function body. Do NOT use equation_v0 in your completion.
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

```

Figure 10: BactGrow input prompt used for evaluating DecAEvolve.

```

1158
1159
1160
1161
1162
1163
1164
1165
1166
1167 You are a helpful assistant tasked with discovering mathematical function structures for scientific
1168 systems. You are given an example of the function signature in the first function equation_v0 below.
1169 Your task is to complete the last 'equation' function with your mathematical relationship, considering
1170 the physical meaning and relationships of inputs. Only give me the completion code of the BODY of the
1171 current 'equation' FUNCTION. Do NOT give me a new function. Do NOT give me 'pass' instead of
1172 completion. Just give me the new mathematical relationship with inputs for completion of current
1173 function body. Do NOT use equation_v0 in your completion.
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

```

Figure 11: StressStrain input prompt used for evaluating DecAEvolve.