

000 001 002 003 004 005 Neuro-Symbolic AI for Analytical Solutions of Differential 006 Equations 007 008

009
010 Anonymous authors
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027

Abstract

012 Analytical solutions to differential equations offer exact insight but are
013 rarely available because discovering them requires expert intuition or ex-
014 haustive search in large combinatorial spaces. We introduce SIGS, a neuro-
015 symbolic framework that automates this process. SIGS uses a formal gram-
016 mar to generate only syntactically and physically valid building blocks,
017 embeds these expressions into a continuous latent space, and then searches
018 this space to assemble, score, and refine candidate closed-form solutions
019 by minimizing a physics-based residual. This design unifies symbolic rea-
020 soning with numerical optimization; the grammar constrains candidate solu-
021 tion blocks to be proper by construction, while the latent search makes
022 exploration tractable and data-free. Across a range of differential equations,
023 SIGS recovers exact solutions when they exist and finds highly accurate ap-
024 proximations otherwise, outperforming tree-based symbolic methods, tra-
025 ditional solvers, and neural PDE baselines in accuracy and wall-clock ef-
026 ficiency. These results are a step forward, integrating symbolic structure
027 with modern ML to discover interpretable, closed-form solutions at scale.
028
029

1 Introduction

030 The understanding of physical processes has been a long-standing effort for scientists and
031 engineers. A key step in this endeavor is to translate physical insights (laws) into precise
032 mathematical relationships that capture the underlying phenomena. These relationships
033 are then tested through experiments that either validate the proposed hypothesis or sug-
034 gest refinements. Among such mathematical formulations, differential equations (DEs) are
035 especially ubiquitous across disciplines, as they describe how physical quantities evolve over
036 time and space. Analytical solutions, closed-form expressions satisfying governing equations
037 and boundary/initial conditions, not only validate theory against experiment but also re-
038 veal intrinsic properties such as stability, periodicity, and symmetries. Classical analytical
039 methods are inherently compositional: they assemble solutions from elementary building
040 blocks such as eigenfunctions, basis expansions, or Green’s functions.
041

042 Unlike the inverse problem of discovery the governing equations given measurements of the
043 solution, which has been widely considered by adapting symbolic regression Petersen et al.
044 (2019b); Landajuela et al. (2022); Petersen et al. (2021); Yu et al. (2025); Kamienny et al.
045 (2022); Biggio et al. (2021); Vastl et al. (2022) to this setting, the forward problem of discov-
046 ering analytical solutions to DEs, considered here, is less explored. In this context, proposed
047 approaches include genetic programming and its variants (Tsoulos & Lagaris, 2006; Seaton
048 et al., 2010; Kamali et al., 2015; Boudouaoui et al., 2020). Lately, symbolic approaches have
049 been enriched with machine learning components to overcome this combinatorial complex-
050 ity. Lample & Charton (2019) train neural networks on sequence representations of trees
051 in order to solve simple explicit ODEs. Wei et al. (2024) propose SSDE, a methodology
052 that employs a recurrent neural network to generate symbolic candidates, guided by a re-
053inforcement learning policy constrained by the governing equations and conditions. As a
baseline, they considered the accuracy of fitting symbolic solutions to functions obtained
by physics-informed neural networks, relying on deep symbolic regression (Petersen et al.,
2019a). Cao et al. (2024) use transfer learning to lift genetic programming results from

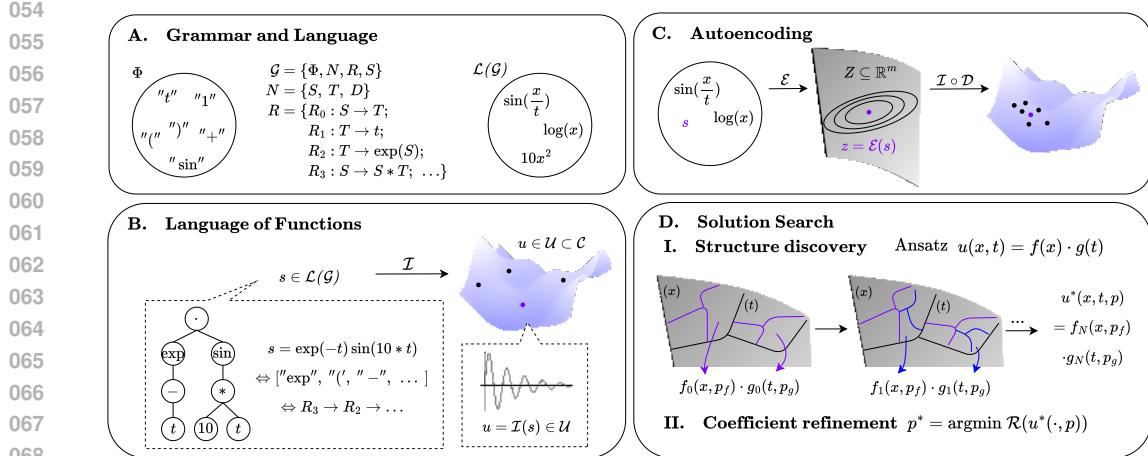


Figure 1: Overview over the proposed Symbolic Iterative Grammar Solver (SIGS). A. Terminal symbols Φ and rules R , together with non-terminals N and starting symbol S , form the grammar \mathcal{G} which generates the mathematical expressions in the library $\mathcal{L}(\mathcal{G})$. B. Each expression $w \in \mathcal{L}(\mathcal{G})$ is identified with a function u in the finite set of candidate functions \mathcal{U} . C. The encoder \mathcal{E} and decoder \mathcal{D} of the Grammar Variational Autoencoder (GVAE, (Kusner et al., 2017)) embed the finite $\mathcal{L}(\mathcal{G})$ into the continuous latent space Z . D. Given a differential equation and system conditions, a structure search is performed over $z \in Z$ using iterative clustering, followed by a separate optimizations of the constants in the final structure, optimizing for lowest residual \mathcal{R} of the corresponding candidate function $u = \mathcal{I} \circ \mathcal{D}(z) \in \mathcal{U}$.

one-dimensional problems to higher dimensions (HD-TLGP). Unlike symbolic regression, where the primitives of differential operators are chosen from a dictionary of fundamental operations such as curl or divergence, there exists no principled way to systematically choose components in solution discovery to combine and get mathematically proper and physically plausible solutions. As a result, solution discovery methods have tended towards two extremes: (i) unconstrained search, which faces combinatorial explosion, sensitivity to initialization, and lack of principled incorporation of domain knowledge; or (ii) narrow pre-training, which biases discovery toward limited problem classes and hinders generalization. A principled middle ground is missing.

This raises the key question: Can we design a framework that generalizes across PDEs while systematically constraining the search to mathematically admissible, physically meaningful solutions?

We answer affirmatively with the Symbolic Iterative Grammar Solver (SIGS). At its core, SIGS casts solution discovery as a hierarchical, grammar-guided composition of analytic atoms (eigenfunctions and related sub-expressions). The hierarchy operates at two levels. At the top level, an Ansatz specifies the structural form of candidate solutions: e.g. the Ansatz $f(x) \times g(t)$ restricts the search to a product of spatial, temporal, or combined terms. At the lower level, each placeholder function in the Ansatz is instantiated with concrete atoms drawn from a grammar; e.g. $f(x)$ is replaced by $\sin(x)$. Using formal grammars (Hopcroft & Ullman, 1979), elementary functions act as terminals while operations such as addition or exponentiation act as production rules. This formalism generalizes classical construction techniques, providing a principled way to generate only admissible expressions and to systematically explore the solution space defined by the PDE.

To overcome the combinatorial complexity of assembling such expressions, SIGS embeds grammar-generated candidates into a continuous latent manifold using a Grammar Variational Autoencoder Kusner et al. (2017). We further impose a novel topological regularization within the GVAE, ensuring that latent neighborhoods map smoothly to valid expressions and that clusters of candidates form convex regions. This embedding transforms discrete tree search into quasi-continuous optimization: instead of enumerating oper-

108 ators and integers, we navigate the latent manifold, progressively refining the search around
 109 promising regions. The final constants are then optimized with gradient descent, yielding
 110 exact or approximate analytical solutions.

111 Our key contributions are as follows.
 112

- 113 • A grammar-based framework (SIGS) that efficiently balances computational
 114 complexity with generality by composing solution units through hierarchical
 115 Ansatz+atom combinations, modeled through formal grammars.
- 116 • The Topological Grammar VAE (TGVAE), which encodes admissible solutions onto
 117 a smooth latent manifold for efficient search.
- 118 • An efficient and task-agnostic approach that employs compositionality of solutions
 119 to solve a broad selection of PDEs, without the need for numerical data.
- 120 • State-of-the-art performance on recent benchmarks, including recovery of exact so-
 121 lutions and symbolic approximations of PDEs lacking closed-form solutions.

124 2 Method

126 Problem setup We consider the generic form of a time-dependent partial differential equa-
 127 tion (PDE) as (Molinaro et al., 2024),

$$\begin{aligned} \partial_t u + \mathbb{D}(u) &= \mathbf{f}, & \forall (\mathbf{x}, t) \in \Omega \times [0, T], \\ u(\mathbf{x}, 0) &= u_0(\mathbf{x}), & \forall \mathbf{x} \in \Omega, \\ \mathbb{B}[u](\mathbf{x}, t) &= g, & \forall (\mathbf{x}, t) \in \partial\Omega \times [0, T], \end{aligned} \quad (1)$$

132 where $\Omega \subset \mathbb{R}^d$ is the spatial domain, $u \in \mathcal{U} \subseteq \mathcal{C}(\Omega \times (0, T))$ is the space-time continuous
 133 solution, $\mathbf{f} \in \mathcal{U}$ is a forcing term, $u_0 \in H^s(\Omega)$ an initial condition, $\mathbb{B}[u](\mathbf{x}, t)$ denotes the
 134 boundary conditions, and $\partial\Omega$ is the boundary of the domain. The differential operator
 135 can include higher-order derivatives, $\mathbb{D}(u) = \mathbb{D}(\xi, u, \partial_{tt}u, \nabla_{\mathbf{x}}u, \nabla_{\mathbf{x}}^2u, \dots)$, where $\xi \in \mathbb{R}^{d_{\xi}}$ are
 136 PDE parameters. We remark that Equation 1 represents a very general form of differential
 137 equations as the solution $u = u(x, t)$ is a function of both space and time. By setting
 138 $u = u(t)$, we recover general ODEs, while setting $u = u(x)$ enables us to recover time-
 139 independent PDEs from the same overall formulation. Henceforth, we use PDEs of the
 140 form of Equation 1, as the objects for which we discover analytical solutions. We call the
 141 collection of $\mathbf{f}, \mathbb{B}[u]$, and u_0 the system conditions that need to be specified in order to solve
 142 a given PDE. We define the symbolic form of a PDE as:

$$\mathcal{S}(u) = \partial_t u + \mathbb{D}(u) - \mathbf{f}, \quad \forall u \in \mathcal{U}.$$

144 We formulate solving PDEs as an iterative computational process, where given a domain
 145 discretization, a set of boundary and initial conditions, and the symbolic form of the PDE
 146

$$(\Omega, \mathbb{B}[u], u_0, S) \xrightarrow{\mathcal{D}(z)} u^i.$$

148 The method searches for a parameterization z of $u_z \in \mathcal{U}$ that minimizes the loss,
 149

$$\mathcal{R}(u) = \|\mathcal{S}(u)\|^2 + \|u(0, x) - u_0^i\|^2 + \|\mathbb{B}[u] - g\|^2, \quad (2)$$

150 where we generally use equal weighting between the residual terms. We restate our goal as
 151 finding an analytical expression u^* that minimizes the residual $\mathcal{R}(u)$, yielding an analytical
 152 solution in case $\mathcal{R}(u^*) = 0$ and an analytical approximate solution if $0 < \mathcal{R}(u^*) \ll 1$.
 153

154 **Grammar Construction.** Analytic expressions are commonly represented as trees, with
 155 internal node labels denoting unary or binary expressions (e.g. "sin", "+") and leaves
 156 denoting constants or variables. However, care must be taken when generating such trees to
 157 avoid exponential complexity and the generation of syntactically wrong expressions Virgolin
 158 & Pissis (2022); Kissas et al. (2024). To alleviate this issue, we consider a Context-Free
 159 Grammar (CFG) (Chomsky, 1956; Hopcroft & Ullman, 1979) as a principled way to generate
 160 exactly the classes of atoms included in an Ansatz. A CFG is defined as $\mathcal{G} = \{\Phi, N, R, S\}$,
 161 where Φ is the set of terminal symbols, N is the set of non-terminal symbols and $\Phi \cap N = \emptyset$,

R is a finite set of production rules and $S \in N$ is the starting symbol. Each rule $r \in R$ is a map $\alpha \rightarrow \beta$, where $\alpha \in N$, and $\beta \in (\Phi \cup N)^*$ (see Fig. 1A). A language $\mathcal{L}(\mathcal{G})$ is defined as the set of all possible terminal strings that can be derived by applying the production rules of the grammar starting from S , or all possible ways that the nodes of a derivation tree can be connected starting from S as $\mathcal{L}(\mathcal{G}) = \{w \in \Phi^* \mid S \rightarrow^* w\}$, where \rightarrow^* implies $T \geq 0$ applications of rules in R . Each expression is equivalently represented by the string w (as a sequence of symbols), by the list of rules applied to generate w from S , and by a derivation tree that represents the syntactic structure of string $w \in \mathcal{L}(\mathcal{G})$ according to grammar \mathcal{G} . We define an interpretation map $\mathcal{I} : \mathcal{L}(\mathcal{G}) \rightarrow \mathcal{U}$, which assigns to each syntactic expression $w \in \mathcal{L}(\mathcal{G})$ semantic meaning in terms of a function $u_w : D \rightarrow \mathbb{R}$. The set of all functions represented by the grammar is $\mathcal{U}(\mathcal{G}) = \{u_w : D \rightarrow \mathbb{R} \mid u_w = \mathcal{I}(w), w \in \mathcal{L}(\mathcal{G})\}$. We refer to u_w as u in the future to simplify the notation.

Compositional Ansatz. When using SIGS on a specific problem, the user may specify a structural Ansatz F that outlines the compositional nature of the proposed solution. For example, one could specify spatiotemporal separability as $u(x, t) = \sum_{j=1}^K a_j T_j(t) \phi_j(x)$, leaving the spatial eigenfunctions ϕ_j and temporal factors T_j , to be chosen by SIGS. In addition to ϕ_j and T_j , the user may include atoms that encode physical mechanisms at the expression level; such as transport phases, $kx - \omega t$; viscous shock profiles, $\tanh((x_0 + x - ct)/\nu)$; or other motifs known to describe the dynamics of interest exactly or approximately. Localized atoms such as Gaussians can also be included to capture spatially confined phenomena. The Ansatz may include hybrid factors that mix space and time, allowing $u(x, t) = \sum_{j=1}^K a_j T_j(t) \phi_j(x) \psi_j(x, t)$ which relaxes separability while retaining controlled, interpretable compositions.

Searching the resulting high-dimensional combinatorial spaces requires a trade-off between generality and complexity. We embed atoms (sub-trees) instead of primitives (unary, binary operators, reals, and variables) to decrease the combinatorial complexity of solutions. In the full Ansatz generality, the solution construction could be performed by considering a number of arbitrary combinations between atoms. This approach would result in a combinatorial explosion, partially losing the benefit of considering atoms. For this reason, we assume that the solutions can be described exactly (or sufficiently well) by the chosen Ansatz. To include the Ansatz into the grammar, we denote by $A : \{\mathcal{L}(\mathcal{G})\}^L \rightarrow \mathcal{L}(\mathcal{G})$ the assembly map that composes the individual components into the final solution following the Ansatz. This restricted function class is obtained by activating only those nonterminals and production rules that implement the user’s Ansatz and its permitted atom categories, and by enforcing the assembly production dictated by A . The Ansatz thus induces a restriction on the language $\mathcal{L}_A(\mathcal{G}) = \{A(w^1, \dots, w^L) : w^c \in \mathcal{L}_c(\mathcal{G})\}$ for the component classes c required by the Ansatz. In all cases, A realizes the user’s choice by assembling requested categories into a single symbolic candidate that is then scored by the PDE residual. In summary, the Ansatz specifies which families of atoms and couplings are admissible, the CFG generates those atoms and couplings, and the interpretation map turns each derivation into a candidate function over which SIGS optimizes the PDE residual.

Grammar Variational Autoencoders. To make the search more efficient, we embed $w \in \mathcal{L}_A(\mathcal{G})$ into a low-dimensional continuous manifold using a Grammar Variational Autoencoder (Kusner et al., 2017). The encoder is defined as $q_\phi(z|w)$ and the decoder $p_\theta(w|z)$, for $z \in Z$ and $w \in \mathcal{L}_A(\mathcal{G})$. The GVAE is trained by minimizing the objective:

$$\mathcal{L} = \mathcal{L}_{\text{recon}} + \gamma \text{KL}(q(z|w) \parallel p(z)),$$

where $\mathcal{L}_{\text{recon}}$ the cross-entropy loss between the predicted and the baseline grammar rules, and $\text{KL}(q(z|w) \parallel p(z))$ the KL divergence between the encoder and the prior distributions.

Training the GVAE does not require numerical data, only expressions $w \in \mathcal{L}$. In practice, we handle numerical matrices with entries $\{0, 1\}$, encoding which rules are employed in which order to generate w , and impose grammar relations through masking parts of these matrices to only allow related elements to interact. These grammar masks are also required for training the model.

216 Geometry Regularization. When we sample the latent manifold, we often evaluate latent
 217 vectors in regions with little or no support from the training distribution, and can also
 218 get trapped in topological artifacts of the latent space. In both cases, the decoder produces
 219 degenerate outputs. For this reason, we impose a geometry-aware regularizer that constrains
 220 the search inside a data-supported enclosure, removes small topological artifacts at the
 221 working resolution, and smooths the decoder so that small latent moves produce predictable
 222 output changes.

223 We augment the GVAE objective with three regularizers (details in App. A.3). A convex-
 224 enclosure loss $\mathcal{L}_{\text{Hull}}$ that discourages latents from leaving the data-supported region esti-
 225 mated from training codes Gonzalez (1985); Rockafellar (2015). A persistent-homology
 226 loss \mathcal{L}_{ph} that suppresses small spurious loops/gaps in the latent cloud at a fixed work-
 227 ing scale Edelsbrunner & Harer (2010). A decoder-smoothness loss $\mathcal{L}_{\text{smooth}}$ that penalizes
 228 large second-order changes in the decoder, so nearby latents decode to predictably similar
 229 functions (Hutchinson, 1989). We combine these losses with the reconstruction and the
 230 KL loss to define the regularized loss of the TGVAE (Topological Grammar Variational
 231 Autoencoder):

$$\mathcal{L} = \mathcal{L}_{\text{recon}} + \gamma \text{KL}(q(z|w) \| p(z)) + \mathcal{L}_{\text{topo}}, \quad \mathcal{L}_{\text{topo}} = \mathcal{L}_{\text{Hull}} + \mathcal{L}_{\text{ph}} + \mathcal{L}_{\text{smooth}}.$$

232 Solution Discovery. The solution discovery is split in two stages (see Fig. 1D, and details
 233 in App. B): In the structure search, we iteratively explore the latent space for a candidate
 234 function included in the structural Ansatz while minimizing the PDE residual, and then op-
 235 timize its numerical constants in a separate stage. For searching, we consider a deterministic
 236 encoding $\mathcal{E}(w) = \mu_{\phi}(w) \in Z$ and decoding $\mathcal{D} : Z \rightarrow \mathcal{L}_A(\mathcal{G})$ obtained by the argmax decoding
 237 under the grammar mask. Composing with \mathcal{I} , we have $\mathcal{I} \circ \mathcal{D} : Z \rightarrow \mathcal{U}_A(\mathcal{G})$, so each $z \in Z$
 238 corresponds to a function $u = \mathcal{I}(\mathcal{D}(z)) \in \mathcal{U}_A(\mathcal{G})$. Let $\tau : \mathcal{L}(\mathcal{G}) \rightarrow \mathcal{T}$ be a semantic map that
 239 assigns tags, e.g. variables, deterministically computed from the parse tree w , computed
 240 once after training and used for any downstream solution problem. For a given differential
 241 equation, we choose the admissible tag set, e.g. any function with x, y arguments, and re-
 242 strict the search to the type-constrained latent subspace $Z' = \{z \in Z : \tau(\mathcal{D}(z)) \in \mathcal{T}' \subseteq \mathcal{T}\}$.
 243

244 Let $\kappa : Z' \rightarrow \{1, \dots, m\}$ be a clustering map in the latent space and denote the clusters
 245 $C_j = \kappa^{-1}(j)$. We cluster a given subspace based on $z \in Z'$, and then solve a discrete
 246 selection problem to choose the cluster that contains the most promising solution forms for
 247 each T and ϕ , $j^* = \arg \min_{1 \leq j \leq m} [\inf_{z \in Z_j \subseteq C_j} \mathcal{R}(\mathcal{D}(z))]$, where Z_j can be constructed by either
 248 only the expressions from the training set that fall in Z' or the expressions together with
 249 samples from the generative model. Within the best cluster C_{j^*} , a global latent search is
 250 performed:

$$z^* = \arg \min_{z \in C_{j^*}} \mathcal{R}(\mathcal{D}(z)),$$

251 either by a global optimizer or iterative clustering, performing discrete selection, and sam-
 252 pling from the most promising cluster until $\mathcal{R}(\mathcal{D}(z))$ drops below a threshold. The solution
 253 takes a parametric form $u^*(\cdot, p)$, including constants p that are only represented in the
 254 grammar with limited precision. Thus, we perform a parameter refining step. We consider
 255 a gradient based method (Adam, Kingma & Ba, 2014), and minimize the loss until a ter-
 256 mination criterion is triggered, $\mathcal{R}(u) \leq 10^{-8}$. The loss $\mathcal{R}(u)$ is augmented here by the hull
 257 loss $\mathcal{R}'(u) = \mathcal{R}(u) + \mathcal{L}_{\text{hull}}$ to penalize whichever latent falls out of the hull defined during
 258 training.

259

260 3 Experiments and Results

261

262 We conduct comprehensive experiments to evaluate SIGS against state-of-the-art symbolic
 263 methods for solving PDEs. Our evaluation comprises three components: cross-validation
 264 on benchmarks sourced from the literature (Table 12), assessment on more complex PDE
 265 problems with and without known analytical solutions (Table 14), and an ablation study
 266 that examines how topology-aware regularization improves sampling efficiency. Details on
 267 our implementation of the grammar and GVAE can be found in Appendix A.

270 Experimental Setup. Our benchmark suite comprises seven PDEs of hyperbolic, parabolic,
 271 and elliptic families. Four problems admit known analytical solutions: viscous Burgers', 1D
 272 Diffusion, 1D Wave, and 2D Damped Wave equations. For the case with no known analytic
 273 solution, we consider three Poisson problems with superposition of different numbers of
 274 Gaussian source terms to test the approximation capabilities of the method.

275 We compare against two recent symbolic discovery methods: HD-TLGP Cao et al. (2024),
 276 and SSDE Wei et al. (2024). Both of these methods sample discrete trees by combinatorially
 277 combining elements of a user defined dictionary. Moreover, HD-TLGP, considers an
 278 Ansatz where the solution is separable in dimension, e.g. $g(x, y) = f(x)g(y)$, as well as the
 279 solution in one dimension as prior knowledge. SSDE considers a recursive single-variable
 280 decomposition Ansatz, e.g. $u(x, y) = g(x, f(y, c))$ and couples reinforcement learning with a
 281 hierarchical approach that resolves each recursion depth sequentially. Both methods search
 282 for expressions satisfying differential equations directly through physics-aware losses similar
 283 to $\mathcal{R}(u)$. The efficiency of these discovery methods lies both in the way they sample trees and
 284 the way they compose solutions. For this reason, we consider two evaluation protocols for
 285 HD-TLGP which considers an Ansatz that is similar to SIGS (details are given in Appendix
 286 C.3.2). Protocol 1 is the same as in the original work: The algorithm is fed a dictionary of
 287 primitives, sin, cos, log, etc., to compose solutions. In Protocol 2, the algorithm is fed atoms
 288 from SIGS instead of primitive functions from a dictionary, see Appendix C.3 for more
 289 details. The objective is to show that naive search in the space of models does not work even
 290 when the algorithm combines atoms, which are mathematically proper subtrees. For SSDE,
 291 we tailor the dictionary of terms for each problem to contain only the primitives, meaning
 292 functions, and variables contained in the solution. For example, if $u(x) = \sin(\pi x) + \cos(\pi y)$
 293 the dictionary contains only sin, x, y , cos and integers. In this way, we show that for sophis-
 294 ticated search methods, if the dictionary considers primitives instead of atoms, the method
 295 cannot find an admissible solution when we consider complex problems. The complete prim-
 296 itive specifications appear in Appendix C.3.1. Neural baselines (PINNs (Raissi et al., 2019),
 297 FBPINNs (Moseley et al., 2023)) and numerical solvers (FEniCS; Alnæs et al., 2015, ; see
 298 details in Appendix C.4) are included for reference. For the Poisson-Gauss problems, no
 299 analytical solutions is available. Therefore, we assume the FEniCS with P4 elements on
 300 a 128×128 mesh as the ground truth. We perform a mesh convergence study to confirm
 301 the convergence of the solution at the chosen resolution. Complete problem specifications,
 302 analytical solutions, and discovered symbolic forms relevant to all the problems in the suite
 303 appear in Appendix C, accompanied by additional figures in Appendix C.7.

303 3.1 Experiments

304 Cross-validation on benchmarks from literature. First, we test SIGS on a subset of prob-
 305 lems considered by Cao et al. and Wei et al. to show how combining the grammar-atoms
 306 approach together with adaptive search is more accurate than alternative approaches. For
 307 this purpose, we chose one-dimensional Poisson and Advection PDEs (HD-TLGP), and a
 308 two-dimensional Wave PDE (SSDE). We consider exactly the same problem specification,
 309 that is, the domain, boundary, and initial conditions, for the comparison. To make the meth-
 310 ods comparable, we impose an Ansatz within SIGS that considers a different function per
 311 variable, e.g. $u(x, t) = g(x) f(t)$. The results are presented in Table 1. While the baseline
 312 approaches achieve high accuracy (HD-TLGP: 4.36×10^{-4} for the Poisson, and 1.01×10^{-2}
 313 for the Wave, SSDE: 1.04×10^{-16}), SIGS achieves exact solutions on all problems, as it
 314 contains π as a symbol and does not approximate it numerically.

315
 316 Table 1: We compare the accuracy, in terms of relative L_2 error against the exact solution,
 317 of SIGS and baselines on a collection of PDEs presented in the HD-TLGP and SSDE papers.
 318

319 Problem (method)	320 Original Method	321 SIGS (ours)
Poisson (HD-TLGP)	4.36×10^{-4}	exact solution
Advection (HD-TLGP)	1.01×10^{-2}	exact solution
Wave (SSDE)	1.04×10^{-16}	exact solution

324 Comparison for Complex PDEs with known solutions. What makes the following collection
 325 of experiments complicated is not only that the solution contains many terms, but also
 326 that the method needs to find solutions that are very precise. For example, even if an
 327 algorithm discovers a solution that describes a viscous shock for the Burgers equation,
 328 slight imprecision in the location of the shock will result in a very large relative L_2 error
 329 against the exact solution. This phenomenon also holds true for the damped wave, as the
 330 problem is sensitive to the coefficients governing the diffusion time. For SIGS, we consider
 331 general solution Ansätze of the forms: We present the results in Table 2. We observe that

332 333 334 335 336 337 Burgers: $u(x, t) = a\psi(x, t)$ 338 Wave: $u(x, y, t) = a\phi^1(x)\phi^2(y)T(t)$ 339 PG-2/3/4: $u(x, y) = \sum_{i=1}^K a_i\psi_i(x, y)\varphi_i(x, y)$ 340 Advection: $u(x, t) = a\psi(x, t)$ 341 Shallow Waters: $\rho(x, y, t) = \psi^1(x, y, t)\psi^2(x, y, t)T(t)$ 342 $u(x, y, t) = \psi^3(x, y, t)\rho(x, y, t)$ 343 $v(x, y, t) = \psi^4(x, y, t)\rho(x, y, t)$	344 345 346 347 348 SIGS recovers exact analytical solutions, achieving machine precision on all problems with 349 relative errors ranging from 6.64×10^{-14} to 1.22×10^{-13} . The discovered expressions match 350 analytical forms up to numerical precision, see Appendix C.5. 351 Both HD-TLGP and the SSDE methods fail to find a solution within the time budget that 352 is accurate or close to the exact, see Appendix C.5. HD-TLGP in the case that we consider 353 atoms in the dictionary, Protocol 1, returns relative L_2 errors in the range 2.04 – 423.40%, 354 demonstrating the importance of the optimization method in discovering an accurate so- 355 lution. Protocol 2 performs worse, with errors in the ranges of 35.68 – 178.77% which 356 shows how the results deteriorate without atoms. SSDE produces errors in the range of 357 $45.62 - 5.87 \times 10^3\%$ even though the primitives are tailored for each problem. Requiring 358 complex and precise solutions, translates to most of the loss landscape being flat with a 359 very high value except for a small area where the loss is small. The failure of SSDE can 360 almost certainly be attributed to the reinforcement learning algorithm failing to find this 361 small region, as in the classic sparse-rewards problem. This result indicates how sophisti- 362 cated optimizers fail completely when the dictionary does not contain elements that support 363 aggressive exploration of the space of candidate models. Neural methods achieve moderate 364 accuracy (2.56-6.09), while numerical solvers (FEniCS) present very accurate results. A 365 visual comparison of the predictions of different methods are provided in Figure 3. 366 367 368 369 370 371 Table 2: Comparison of methods on PDEs with known analytical solutions. Reported are 372 relative L^2 errors. 373
---	--

PDE Problem	SIGS	HD-TLGP P1	HD-TLGP P2	SSDE	PINNs	FBPINNs	FEniCS
Burgers	6.64×10^{-14}	2.04	35.68	45.62	6.09	28.26	8.69×10^{-3}
Diffusion	7.16×10^{-13}	33.34	79.73	5.87×10^3	2.56	55.54	2.26×10^{-3}
Damping Wave	1.22×10^{-13}	423.30	178.77	1.19×10^3	5.56	71.36	2.28×10^{-2}

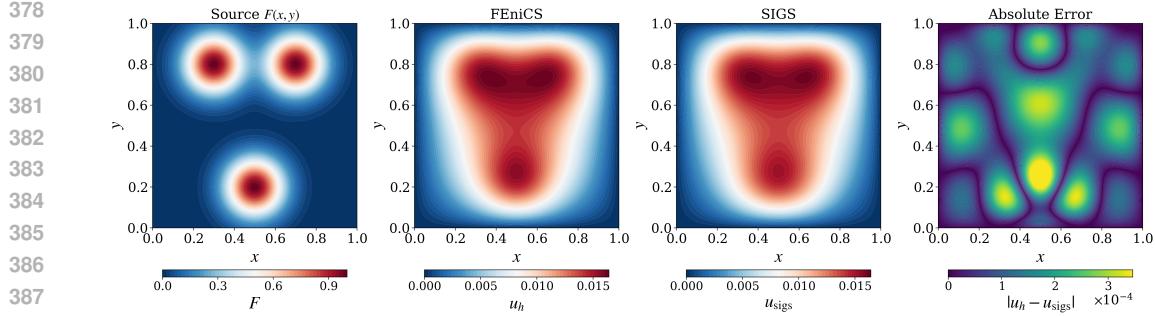


Figure 2: From left to right: source term $F(x, y)$ for the Poisson–Gauss problem; finite-element solution u_h (FEniCS); symbolic approximation u_{sig} (SIGS); absolute error $|u_h - u_{\text{sig}}|$

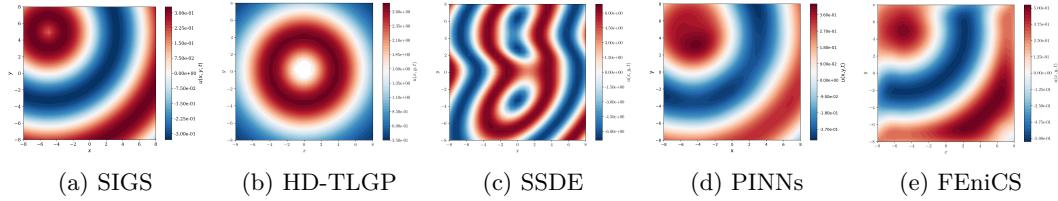


Figure 3: Comparison of different methods for solving the damped wave equation at $t = 2.5$. All methods show the same physical domain $x, y \in [-8, 8]$ with wave center at $(-5, 5)$. Parameters: $k = 0.5$, $\omega = 0.4$, $\alpha = 0.45$.

Symbolic approximation without known solutions. For the PDEs we considered so far, we manufactured, and therefore had access to, the exact solution. This allowed us to make educated guesses about the form of the Ansatz. In this example, we test how well SIGS and the baselines approximate the solution when an exact solution and a strong prior on the Ansatz does not exist. We investigate the Poisson equations with a Gaussian forcing term, which do not admit an exact analytical solution. For SIGS we choose the Ansatz as $u(x, y) = \sum_{j=1}^N \phi_j(x, y) \psi_j(x, y)$, with $j \in [3, 4, 8]$ for PG2, PG3, PG4, respectively. Here $\psi_j(x, y)$ are eigenfunctions of the elliptic operator $\psi(x, y) = \sin(\pi x) \sin(\pi y)$ that impose the homogeneous Dirichlet boundary conditions, and ϕ_j is randomly sampled from the available atoms. We present the results for all the methods in Table 3. SIGS achieves 1 – 3% relative L_2 errors with improving accuracy as complexity increases and number of modes increases (from 2.66% for 2 Gaussians to 1.05% for 4 Gaussians), which suggests that SIGS correctly leverages the superposition of the Gaussian atoms. HD-TLGP failed to find a solution within the time budget, and produced NaNs in our tests, probably due to numerical instabilities for Protocol 1 while for Protocol 2 generates errors exceeding 10^7 . SSDE achieves errors in the range 58-70%, which translates to missing the precise superposition of Gaussians.

The results support that successfully discovering the solution of complex PDEs requires a combination of structured atoms and a global(to explore)-local(to discover precise arrangements) optimization algorithm. Moreover, Table 4 shows how the approach of SIGS is practically viable as the solutions are found in seconds to minutes.

3.2 Ablation Studies

Atoms vs. Primitives. We previously stated that considering atoms or combining primitives can have a decisive effect in finding PDE solutions. We test this hypothesis by considering the damped wave PDE with the Ansatz $u(x, y, t) = \sum_{j=1}^N \psi_j(x, y, t) \phi_j(x) T_j(t)$ and instead of considering atoms, we sample ϕ, T, ψ with uniform probability over rules of the same grammar as before. If we sample a function with the correct arity, e.g. ψ containing

432
433
434
435
436
437
438
439
440

Table 3: Approximation on Poisson-Gauss problems without analytical solutions. Relative L2 errors against FEniCS references.

Problem	SIGS	HD-TLGP P1	HD-TLGP P2	SSDE
PG-2	2.66	200.9	98.94	69.29
PG-3	1.54	NaN	5.61×10^7	69.64
PG-4	1.05	NaN	5.45×10^7	58.70

441
442
443
444Table 4: Wall-clock time (CPU). SIGS reports time-to- ε ; others report time-at-termination. Notation: \checkmark reached ε ; \dagger hit budget / failed to reach ε . HD-TLGP budget: 20 generations, SSDE budget: 25 generations.

Problem	SIGS \checkmark	HD-TLGP P1 \dagger	HD-TLGP P2 \dagger	SSDE \dagger	PINNs \dagger	FEniCS \checkmark
Burgers	13.5 sec	> 239 m36 sec	> 200 m57 sec	> 6 m34 sec	8.8 sec	2.2 sec
Diffusion	39.2 sec	> 192 m41 sec	> 181 m49 sec	> 8 m6 sec	2 m2 sec	1.4 sec
Damping Wave	30.2 sec	> 88 m40 sec	> 37 m8 sec	> 6 m19 sec	29.5 sec	3.4 sec
PG-2	1 m30 sec	> 182 m25 sec	> 90 m43 sec	> 11 m45 sec	n/a	19.3 sec
PG-3	1 m51 sec	> 120 m7 sec	> 97 m16 sec	> 11 m4 sec	n/a	6.9 sec
PG-4	1 m23 sec	> 145 m32 sec	> 80 m50 sec	> 12 m32 sec	n/a	3.4 sec

453
454
455
456
457
458
459
460

x, y, t , we consider the function admissible. We sampled 50,000 functions, out of which only 133 were admissible, which means that it would be impossible to start and adaptive optimization procedure due to the admissible sampling rate being so low. Moreover, the admissible functions with the lowest loss provides $\mathcal{R}(u) \approx 366\%$ relative L^2 error to the exact solution. This clearly demonstrates the necessity of atoms, and thus the embedding, to the whole process.

461
462
463
464
465
466
467
468
469
470
471
472
473
474

TGVAE vs. vanilla GVAE. We measure sampling efficiency using a race-to- k -valid benchmark, which counts the total attempts required to generate $k = 1000$ syntactically valid expressions by sampling random latent vectors $z \in Z$. To assess the quality of the latent space, the latent vectors are decoded to analytical expressions w , which are rejected if they fail to meet the grammar-based and mathematical consistency checks in Section A.1.3. We expect both VAEs to be more stable in regions surrounding $z_i = \mathcal{E}(w_i)$. Hence, we only consider latent vectors z with a minimal distance of $\tau = 0.8$ away from any training sample z_i in terms of the Mahalanobis norm (App. C.6). We sample 15,000 admissible latent vectors and split them into ten disjoint sets. We provide each set to the GVAE and the TGVAE and count the total decode attempts required to obtain 1,000 valid expressions. The GVAE required 1486.2 ± 19.5 attempts, while our topology-regularized TGVAE needed 1433.2 ± 27.3 attempts, a $3.56\% \pm 1.81$ relative reduction. This indicates that geometric regularization (hull loss, persistent homology, smoothness penalties) yields a more navigable latent space with fewer degenerate decodes.

475
476

4 Discussion and Conclusion

477
478
479
480
481
482
483
484
485

Discussion. This work advances solution discovery for PDEs by demonstrating that grammar-guided neuro-symbolic methods can reliably and efficiently recover analytical solutions. SIGS consistently improves the state-of-the-art, both in accuracy and speed, often by several orders of magnitude. Its success stems from two complementary design choices: (i) constructing a latent manifold of solution components, which enables smooth and efficient exploration of admissible expressions; and (ii) employing a hierarchical Ansatz+atom approach that reduces search complexity by structuring the solution space into manageable placeholders, later refined into concrete symbolic elements. This is in contrast to the baselines explored in this work, which do not address the combinatorial explosion inherent in symbolic solution discovery. HD-TLGP (Cao et al., 2024) transfers structures from one-

486 dimensional solutions to higher dimensions, but still relies on stochastic recombination of
 487 primitives, which quickly becomes intractable as complexity grows. SSDE (Wei et al., 2024)
 488 instead uses reinforcement learning to guide the construction of candidate solutions, but
 489 its flat search space remains prohibitively large without strong priors. As our experiments
 490 show, both methods degrade sharply when such priors are absent. In contrast, the hierar-
 491 chical Ansatz+atom design of SIGS separates global structure from local symbolic details,
 492 making tractable what would otherwise be an unmanageable search. In this way, SIGS not
 493 only advances but fundamentally redefines the state-of-the-art for solution discovery. Be-
 494 yond these empirical gains, we view SIGS as part of a broader shift toward neuro-symbolic
 495 foundation models for PDEs. Current foundation approaches (Herde et al., 2024; Hao et al.,
 496 2024; Sun et al., 2024; Alkin et al., 2024; Shen et al., 2024) rely on extensive pretraining and
 497 often serve as black-box predictors for downstream tasks. In contrast, SIGS requires only a
 498 one-time pretraining step to construct its manifold, after which it transfers directly to new
 499 problems without retraining. Moreover, it produces analytical expressions that incorporate
 500 physical priors (e.g., eigenfunctions), yielding interpretable solutions rather than opaque
 501 approximations. This suggests that grammar-based neuro-symbolic models could comple-
 502 ment or even provide an alternative to purely data-driven foundation models in scientific
 503 computing.

504 Limitations. Despite these contributions, SIGS faces two main limitations. First, scalabil-
 505 ity to complex engineering problems remains challenging. PDEs involving discontinuities,
 506 multiscale structure, or turbulence may require grammars enriched with special functions
 507 that cannot be easily decomposed into smaller atoms, or long expressions that increase
 508 search complexity. Hybrid approaches that combine symbolic structures with numerical
 509 bases (e.g., POD-derived eigenfunctions, or Neural Operators) may provide a path forward,
 510 particularly for multiscale phenomena, as well as for problems with irregular geometries
 511 or boundary conditions. Second, the framework depends on the joint design of grammar,
 512 Ansatz, and latent space. A richer Ansatz can offset a simpler grammar, while a more expres-
 513 sive grammar requires larger latent spaces and more sophisticated optimization. Currently,
 514 the Ansatz still reflects human expert choices. This can be advantageous in domains with
 515 strong theoretical foundations (e.g., Burgers or Poisson equations), but limits applicability
 516 in less understood settings. A promising direction is to leverage large language models (e.g.,
 517 Romera-Paredes et al., 2024) to automate Ansatz construction, learning general solution
 518 structures directly from governing equations.

519 Conclusion. In this work, we introduced the Symbolic Iterative Grammar Solver (SIGS),
 520 a grammar-guided neuro-symbolic framework for discovering analytical solutions to differ-
 521 ential equations. By unifying classical compositional methods with modern latent-space op-
 522 timization through the Topological Grammar VAE, SIGS systematically explores the space
 523 of admissible solutions, enabling efficient search and refinement of closed-form expressions.
 524 Our approach achieves state-of-the-art performance on recent benchmarks, recovering exact
 525 solutions when available, and producing interpretable symbolic approximations for PDEs
 526 without known closed form solutions. These results highlight the potential of grammar-based
 527 neuro-symbolic methods as a scalable and interpretable alternative to purely data-driven
 528 approaches, opening new directions for automated solution discovery in scientific computing.

529
 530
 531
 532
 533
 534
 535
 536
 537
 538
 539

540 References

541

542 Benedikt Alkin, Andreas Fürst, Simon Schmid, Lukas Gruber, Markus Holzleitner, and
543 Johannes Brandstetter. Universal physics transformers. arXiv preprint arXiv:2402.12365,
544 2024.

545 Martin Alnæs, Jan Blechta, Johan Hake, August Johansson, Benjamin Kehlet, Anders Logg,
546 Chris Richardson, Johannes Ring, Marie E. Rognes, and Garth N. Wells. The FEniCS
547 Project Version 1.5. Archive of Numerical Software, 3(100), December 2015. ISSN 2197-
548 8263. doi: 10.11588/ans.2015.100.20553.

549

550 Luca Biggio, Tommaso Bendinelli, Alexander Neitz, Aurelien Lucchi, and Giambattista
551 Parascandolo. Neural symbolic regression that scales. In International Conference on
552 Machine Learning, pp. 936–945. PMLR, 2021.

553 Yassine Boudouaoui, Hacene Habbi, Celal Ozturk, and Dervis Karaboga. Solving differential
554 equations with artificial bee colony programming. Soft Computing, 24:17991–18007, 2020.

555

556 Lulu Cao, Yufei Liu, Zhenzhong Wang, Dejun Xu, Kai Ye, Kay Chen Tan, and Min Jiang.
557 An interpretable approach to the solutions of high-dimensional partial differential equa-
558 tions. Proceedings of the AAAI Conference on Artificial Intelligence, 38(18):20640–20648,
559 Mar. 2024. doi: 10.1609/aaai.v38i18.30050. URL <https://ojs.aaai.org/index.php/AAAI/article/view/30050>.

560

561 Noam Chomsky. Three models for the description of language. IRE Transactions on Infor-
562 mation Theory, 2(3):113–124, 1956. doi: 10.1109/TIT.1956.1056813.

563

564 Herbert Edelsbrunner and John Harer. Computational topology: an introduction. American
565 Mathematical Soc., 2010.

566

567 Teofilo F Gonzalez. Clustering to minimize the maximum intercluster distance. Theoretical
568 computer science, 38:293–306, 1985.

569

570 Zhongkai Hao, Chang Su, Songming Liu, Julius Berner, Chengyang Ying, Hang Su, An-
571 anima Anandkumar, Jian Song, and Jun Zhu. Dpot: Auto-regressive denoising operator
572 transformer for large-scale pde pre-training. arXiv preprint arXiv:2403.03542, 2024.

573

574 Maximilian Herde, Bogdan Raonić, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Em-
575 manuel de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for
576 pdes. arXiv preprint arXiv:2405.19101, 2024.

577

578 John E. Hopcroft and Jeffrey D. Ullman. Introduction to Automata Theory, Languages,
579 and Computation. Addison-Wesley, Reading, MA, 1979. ISBN 978-0201029888.

580

581 Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for
582 laplacian smoothing splines. Communications in Statistics-Simulation and Computation,
583 18(3):1059–1076, 1989.

584

585 MZM Kamali, N Kumaresan, and Kuru Ratnavelu. Solving differential equations with ant
586 colony programming. Applied Mathematical Modelling, 39(10-11):3150–3163, 2015.

587

588 Pierre-Alexandre Kamienny, Stéphane d’Ascoli, Guillaume Lample, and Francois Char-
589 ton. End-to-end symbolic regression with transformers. In Alice H. Oh, Alekh Agarwal,
590 Danielle Belgrave, and Kyunghyun Cho (eds.), Advances in Neural Information Processing
591 Systems, 2022. URL https://openreview.net/forum?id=GoOuIrDHG_Y.

592

593 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv
594 preprint arXiv:1412.6980, 2014.

595

596 Georgios Kissas, Siddhartha Mishra, Eleni Chatzi, and Laura De Lorenzis. The language of
597 hyperelastic materials. Computer Methods in Applied Mechanics and Engineering, 428:
598 117053, 2024.

594 Matt J Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar variational
 595 autoencoder. In International conference on machine learning, pp. 1945–1954. PMLR,
 596 2017.

597 Guillaume Lample and François Charton. Deep learning for symbolic mathematics. arXiv
 598 preprint arXiv:1912.01412, 2019.

600 Mikel Landajuela, Chak Shing Lee, Jiachen Yang, Ruben Glatt, Claudio P Santiago, Ignacio
 601 Aravena, Terrell Mundhenk, Garrett Mulcahy, and Brenden K Petersen. A unified frame-
 602 work for deep symbolic regression. Advances in Neural Information Processing Systems,
 603 35:33985–33998, 2022.

604 Roberto Molinaro, Samuel Lanthaler, Bogdan Raonić, Tobias Rohner, Victor Armegioiu,
 605 Zhong Yi Wan, Fei Sha, Siddhartha Mishra, and Leonardo Zepeda-Núñez. Generative ai
 606 for fast and accurate statistical computation of fluids. arXiv preprint arXiv:2409.18359,
 607 2024.

608 Ben Moseley, Andrew Markham, and Tarje Nissen-Meyer. Finite basis physics-informed neu-
 609 ral networks (fbpinns): a scalable domain decomposition approach for solving differential
 610 equations. Advances in Computational Mathematics, 49(4):62, 2023.

612 Brenden K Petersen, Mikel Landajuela, T Nathan Mundhenk, Claudio P Santiago, Soo K
 613 Kim, and Joanne T Kim. Deep symbolic regression: Recovering mathematical expressions
 614 from data via risk-seeking policy gradients. arXiv preprint arXiv:1912.04871, 2019a.

615 Brenden K Petersen, Mikel Landajuela, T Nathan Mundhenk, Claudio P Santiago, Soo K
 616 Kim, and Joanne T Kim. Deep symbolic regression: Recovering mathematical expressions
 617 from data via risk-seeking policy gradients. arXiv preprint arXiv:1912.04871, 2019b.

619 Brenden K Petersen, Claudio P Santiago, and Mikel Landajuela. Incorporating domain
 620 knowledge into neural-guided search. arXiv preprint arXiv:2107.09182, 2021.

621 Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural net-
 622 works: A deep learning framework for solving forward and inverse problems involving
 623 nonlinear partial differential equations. Journal of Computational physics, 378:686–707,
 624 2019.

625 Ralph Tyrell Rockafellar. Convex analysis:(pms-28). Princeton university press, 2015.

627 Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Ba-
 628 log, M. Pawan Kumar, Emilien Dupont, Francisco J. R. Ruiz, Jordan S. Ellenberg, Peng-
 629 ming Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein Fawzi. Mathematical discoveries
 630 from program search with large language models. Nature, 625(7995):468–475, January
 631 2024. ISSN 1476-4687. doi: 10.1038/s41586-023-06924-6.

632 Tom Seaton, Gavin Brown, and Julian F Miller. Analytic solutions to differential equations
 633 under graph-based genetic programming. In European Conference on Genetic Program-
 634 ming, pp. 232–243. Springer, 2010.

636 Junhong Shen, Tanya Marwah, and Ameet Talwalkar. Ups: Efficiently building founda-
 637 tion models for pde solving via cross-modal adaptation. In ICML 2024 AI for Science
 638 Workshop, 2024.

639 Jingmin Sun, Yuxuan Liu, Zecheng Zhang, and Hayden Schaeffer. Towards a foundation
 640 model for partial differential equation: Multi-operator learning and extrapolation. arXiv
 641 preprint arXiv:2404.12355, 2024.

642 Ioannis G. Tsoulos and Isaac E. Lagaris. Solving differential equations with genetic
 643 programming. Genetic Programming and Evolvable Machines, 7:33–54, 2006. URL
 644 <https://api.semanticscholar.org/CorpusID:1719377>.

646 Martin Vastl, Jonáš Kulhánek, Jiří Kubalík, Erik Derner, and Robert Babuška. Symformer:
 647 End-to-end symbolic regression using transformer-based architecture. arXiv preprint
 arXiv:2205.15764, 2022.

648 Marco Virgolin and Solon P Pissis. Symbolic regression is np-hard. arXiv preprint
649 arXiv:2207.01018, 2022.
650

651 Shu Wei, Yanjie Li, Lina Yu, Min Wu, Weijun Li, Meilan Hao, Wenqiang Li, Jingyi Liu,
652 and Yusong Deng. Closed-form symbolic solutions: A new perspective on solving partial
653 differential equations. arXiv preprint arXiv:2405.14620, 2024.
654 Wolfram Research, Inc. Mathematica. Wolfram Research, Inc., Champaign, IL, 2024. URL
655 <https://www.wolfram.com/mathematica/>. Version 14.1.
656

657 Karin L Yu, Eleni Chatzi, and Georgios Kissas. Grammar-based ordinary differential equa-
658 tion discovery. arXiv preprint arXiv:2504.02630, 2025.
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 703 704 705 706 707 Appendix

708	A Grammar and GVAE	14
709	A.1 Library Generation	14
710	A.2 GVAE Model and Training Details	17
711	A.3 Geometry Regularization	19
712	B Iterative Search and Refinement	19
713	B.1 Notation and Setup	19
714	B.2 Target loss: Discretized PDE residual	20
715	B.3 Stage I: Structure discovery by iterative clustering	20
716	B.4 Stage II: Coefficient Refinement	21
717	C Experiments	21
718	C.1 Problem Definitions	21
719	C.2 Solution Ansatz specific to our experiments	23
720	C.3 Configuration of baseline methods	25
721	C.4 FEniCS Validation for Reference Solutions	26
722	C.5 Discovered Symbolic Expressions	27
723	C.6 Details on the ablation Study	27
724	C.7 Solution Visualizations	30
725	C.8 Additional results during the review process	33
726	C.9 Additions to the computational performance assessment	40
727		
728		
729		
730		
731	A Grammar and GVAE	
732	A.1 Library Generation	
733		
734	We construct symbolic component libraries that serve as input vocabulary for the TGVAE	
735	architecture in the discovery of DE solutions. The fundamental challenge in symbolic regression	
736	for DEs is that naive search over arbitrary mathematical expressions is computationally	
737	intractable and often produces physically meaningless results. Our library generation	
738	approach addresses this by creating curated collections of symbolic components that	
739	correspond directly to different atoms that compose solutions of a wide range of differential	
740	operators. The library consists of individual analytical building blocks rather than complete	
741	DE solutions. Each component represents a fundamental mathematical pattern such as spatial	
742	eigenfunctions $\sin(k\pi x)$, temporal factor $e^{-\lambda t}$, or their separable product $\sin(k\pi x)e^{-\lambda t}$	
743	that naturally arises in the decomposition of certain operator. This modular design enables	
744	the neural architecture to learn complex solution structures through principled combinations	
745	of mathematically meaningful primitives, rather than searching over the vast space of arbitrary	
746	symbolic expressions. The key insight driving our approach is that different DE	
747	operators admit characteristic families of atoms that reflect their underlying mathematical	
748	structure. This principled approach transforms the solution discovery task from an open-ended	
749	search problem into a structured exploration of mathematically principled solution	
750	components.	
751	A.1.1 Atom Generation	
752	The atoms of the library represent temporal factors, eigenfunctions of operators, expressions	
753	that describe dynamics of interest, and random compositions.	
754		
755	Assume a bounded Lipschitz domain $\Omega \subset \mathbb{R}^d$ and a linear second order operator with	
	homogeneous boundary conditions, that is self-adjoint and non-negative, e.g. $\mathcal{S} = -\Delta$ with	

756 Table 5: Closed-form Laplacian eigenfamilies $A = -\Delta$ on common domains. Here Ω is the
 757 spatial domain, ϕ the eigenfunction, and μ the eigenvalue in $A\phi = \mu\phi$. For Neumann on
 758 boxes the constant mode has $\mu_0 = 0$.

760 Domain / BC	761 Eigenfunction ϕ	762 Eigenvalue μ	763 Index set
764 Rectangles / periodic boxes $\Omega = \prod_{d=1}^D [0, L_d]$			
765 Periodic (torus)	$\phi_{\mathbf{k}}(x) = \exp\left(i 2\pi \sum_{d=1}^D \frac{k_d}{L_d} x_d\right)$	$\mu_{\mathbf{k}} = 4\pi^2 \sum_{d=1}^D \frac{k_d^2}{L_d^2}$	$\mathbf{k} \in \mathbb{Z}^D$
766 Dirichlet	$\phi_{\mathbf{k}}(x) = \prod_{d=1}^D \sin\left(\frac{k_d \pi}{L_d} x_d\right)$	$\mu_{\mathbf{k}} = \pi^2 \sum_{d=1}^D \frac{k_d^2}{L_d^2}$	$k_d \in \mathbb{N}$
767 Neumann	$\phi_{\mathbf{k}}(x) = \prod_{d=1}^D \cos\left(\frac{k_d \pi}{L_d} x_d\right)$	$\mu_{\mathbf{k}} = \pi^2 \sum_{d=1}^D \frac{k_d^2}{L_d^2}$	$k_d \in \mathbb{N}_0$
768	(includes constant mode $\mathbf{k} = \mathbf{0}$ with $\mu_0 = 0$)		
769 Disks and balls (Dirichlet conditions)			
770 2-D disk, radius R	$\phi_{mn}(r, \theta) = J_m\left(\frac{j_{mn}}{R} r\right) \times \begin{cases} \cos(m\theta), \\ \sin(m\theta), \end{cases}$	$\mu_{mn} = \frac{j_{mn}^2}{R^2}$	$m \in \mathbb{Z}_{\geq 0}, n \in \mathbb{N}$
771 3-D ball, radius R	$\phi_{\ell mn}(r, \theta, \varphi) = j_{\ell}\left(\frac{\alpha_{\ell n}}{R} r\right) Y_{\ell m}(\theta, \varphi)$	$\mu_{\ell n} = \frac{\alpha_{\ell n}^2}{R^2}$	$\ell \in \mathbb{Z}_{\geq 0}, m \leq \ell, n \in \mathbb{N}$

772 Dirichlet, Neumann, or periodic boundary conditions. Then there exists an L^2 -orthonormal
 773 eigen basis $\{\phi\}_{j \geq 1} \subset L^2(\Omega)$ and eigenvalues $\{\mu\}_{j \geq 1} \subset [0, \infty)$ with $\mathcal{S}\phi_j = \mu_j \phi_j$. For example,
 774 the Diffusion equation $u_t - \kappa \Delta u = 0$ has the scalar ODE $T'(t) + \kappa \mu_j T_j(t) = 0$ as the temporal
 775 rule, with solution $T_j(t) = e^{-\kappa \mu_j t}$. For the spatial rule, we can consider the rectangle box
 776 $\Omega = \prod_{d=1}^D [0, L_d]$ with Dirichlet boundary conditions. Given indices $k \in \mathbb{N}^D$, the grammar
 777 produces:

$$\phi_k(x) = \prod_{d=1}^D \sin(k_d \pi x_d / L_d), \quad \mu_k = \pi^2 \sum_{d=1}^D \frac{k_d^2}{L_d^2}.$$

778 Composing together with the temporal model, we get $u(x, t) = \sum_k a_k e^{\kappa \mu_j t} \phi_k(x)$, which solves
 779 $u_t + \kappa \Delta u = 0$ exactly. The amplitudes a_j are drawn from prior $a_j \sim \mathcal{N}(0, \sigma^2 \rho)$ where ρ
 780 decays exponentially to control the regularity or the spectrum. This construction generalizes
 781 for multiple classes of known operators, see Table 6. For constant coefficient operators on
 782 separable geometries we have explicit $\{\phi_k, \mu_k\}$ eigenfunctions as shown in Table 5.

783 As we discussed, the grammar can also produce expressions that describe dynamics of
 784 interest such as viscous shocks $\tanh(\frac{(u_l - u_r)(x - st - x_0)}{4\nu})$, transport $g(kx - \omega t)$, heat kernels
 785 $\frac{1}{(4\pi kt)^{d/2}} \exp(\frac{\|x\|^2}{4kt})$, Gaussian bumps $\exp(\frac{\|x - x_0\|^2}{2k})$ and others. Moreover, atoms are polyno-
 786 mials, and combinations of the above.

787 A.1.2 Formal Grammar Specifics

788 The grammar $\mathcal{G} = (V, \Sigma, R, S)$ contains 51 production rules that provide the complete
 789 symbolic vocabulary for DE eigenfunction families used in the experiments herein. The
 790 grammar systematically generates expressions through the application of production rules
 791 R , including

- 803 • compositional rules $S \rightarrow S + T \mid S \times T \mid S/T \mid S - T \mid T \mid -T$ that build complex
 804 mathematical structures,
- 805 • function application rules $T \rightarrow (S) \mid (S)^2 \mid \sin(S) \mid \exp(S) \mid \log(S) \mid \cos(S) \mid$
 806 $\sqrt{S} \mid \tanh(S)$ that provide the transcendental functions essential for eigenfunction
 807 representation,
- 808 • variable and monomial specifications $T \rightarrow T^D \mid \pi \mid x \mid y \mid t \mid x^2 \mid x^3 \mid y^2 \mid y^3$ that
 809 capture spatial and temporal dependencies,

810
 811 Table 6: Modal time factors for common PDE families. Here A is a nonnegative self-adjoint
 812 spatial operator with eigenpairs $A\phi_{\mathbf{k}} = \mu_{\mathbf{k}}\phi_{\mathbf{k}}$, $\mu_{\mathbf{k}} \geq 0$. Projecting the PDE onto $\phi_{\mathbf{k}}$ yields
 813 the scalar ODE for $T_{\mathbf{k}}(t)$ shown in the middle column and its solution in the right column.

PDE family	Modal ODE (after projection)	Temporal factor $T_{\mathbf{k}}(t)$
Heat / diffusion	$T'_{\mathbf{k}} + \kappa \mu_{\mathbf{k}} T_{\mathbf{k}} = 0$	$e^{-\kappa \mu_{\mathbf{k}} t}$
Stokes (divergence-free)	$T'_{\mathbf{k}} + \nu \mu_{\mathbf{k}} T_{\mathbf{k}} = 0$	$e^{-\nu \mu_{\mathbf{k}} t}$
Undamped wave	$T''_{\mathbf{k}} + c^2 \mu_{\mathbf{k}} T_{\mathbf{k}} = 0$	$\cos(c\sqrt{\mu_{\mathbf{k}}} t)$ or $\sin(c\sqrt{\mu_{\mathbf{k}}} t)$
Damped wave (telegraph)	$T''_{\mathbf{k}} + 2\gamma T'_{\mathbf{k}} + c^2 \mu_{\mathbf{k}} T_{\mathbf{k}} = 0$	Underdamped $c^2 \mu_{\mathbf{k}} > \gamma^2$: $e^{-\gamma t} (C_1 \cos(\omega_{\mathbf{k}} t) + C_2 \sin(\omega_{\mathbf{k}} t))$, $\omega_{\mathbf{k}} = \sqrt{c^2 \mu_{\mathbf{k}} - \gamma^2}$. Critical $c^2 \mu_{\mathbf{k}} = \gamma^2$: $e^{-\gamma t} (C_1 + C_2 t)$. Overdamped $c^2 \mu_{\mathbf{k}} < \gamma^2$: $C_1 e^{-(\gamma - \sqrt{\gamma^2 - c^2 \mu_{\mathbf{k}}}) t} + C_2 e^{-(\gamma + \sqrt{\gamma^2 - c^2 \mu_{\mathbf{k}}}) t}$.
Biharmonic diffusion	$T''_{\mathbf{k}} + \kappa \mu_{\mathbf{k}}^2 T_{\mathbf{k}} = 0$	$e^{-\kappa \mu_{\mathbf{k}}^2 t}$
Damped plate/beam	$T''_{\mathbf{k}} + 2\gamma T'_{\mathbf{k}} + c^2 \mu_{\mathbf{k}}^2 T_{\mathbf{k}} = 0$	As for damped wave, with $c^2 \mu_{\mathbf{k}}$ replaced by $c^2 \mu_{\mathbf{k}}^2$
Klein-Gordon (damped)	$T''_{\mathbf{k}} + 2\gamma T'_{\mathbf{k}} + (c^2 \mu_{\mathbf{k}} + m^2) T_{\mathbf{k}} = 0$	As for damped wave, with $c^2 \mu_{\mathbf{k}}$ replaced by $c^2 \mu_{\mathbf{k}} + m^2$
Fractional diffusion	$T'_{\mathbf{k}} + \kappa \mu_{\mathbf{k}}^s T_{\mathbf{k}} = 0$, $s \in (0, 1]$	$e^{-\kappa \mu_{\mathbf{k}}^s t}$
Reaction-diffusion (linear part)	$T'_{\mathbf{k}} + (\kappa \mu_{\mathbf{k}} - \rho) T_{\mathbf{k}} = 0$	$e^{-(\kappa \mu_{\mathbf{k}} - \rho) t}$
Allen-Cahn (linearized)	$T'_{\mathbf{k}} + (\kappa \mu_{\mathbf{k}} - \alpha) T_{\mathbf{k}} = 0$	$e^{-(\kappa \mu_{\mathbf{k}} - \alpha) t}$
Cahn-Hilliard (linearized)	$T'_{\mathbf{k}} + M \mu_{\mathbf{k}} (\mu_{\mathbf{k}} + \sigma) T_{\mathbf{k}} = 0$	$e^{-M \mu_{\mathbf{k}} (\mu_{\mathbf{k}} + \sigma) t}$
Kuramoto-Sivashinsky (linearized)	$T'_{\mathbf{k}} + (\nu \mu_{\mathbf{k}}^2 - \kappa \mu_{\mathbf{k}}) T_{\mathbf{k}} = 0$	$e^{-(\nu \mu_{\mathbf{k}}^2 - \kappa \mu_{\mathbf{k}}) t}$
Maxwell in a PEC cavity	$\varepsilon T''_{\mathbf{k}} + \sigma T'_{\mathbf{k}} + c^2 \mu_{\mathbf{k}} T_{\mathbf{k}} = 0$	Vector modes; as damped wave (if $\sigma = 0$: \cos / \sin with $\omega_{\mathbf{k}} = c\sqrt{\mu_{\mathbf{k}}}$)
Isotropic linear elasticity	$T''_{\mathbf{k}} + \omega_{B,\mathbf{k}}^2 T_{\mathbf{k}} = 0$, $B \in \{T, L\}$	Two branches: $\omega_{T,\mathbf{k}} = c_T \sqrt{\mu_{\mathbf{k}}}$, $\omega_{L,\mathbf{k}} = c_L \sqrt{\mu_{\mathbf{k}}}$; $T = \cos / \sin$

833
 834 • numeric construction $T \rightarrow D \mid D.D \mid -D \mid -D.D \mid TD$ with digit generation
 835 $D \rightarrow D0 \mid D1 \mid \dots \mid D9 \mid 0 \mid 1 \mid \dots \mid 9$,
 836 • and scientific notation $D \rightarrow e-1 \mid e-2 \mid e-3 \mid e-4$ for numerical stability across multiple
 837 scales.

838 The terminal alphabet hence encompasses

$$\Sigma = \{x, y, t, \pi\} \cup \{\sin, \cos, \exp, \log, \tanh, \sqrt{\cdot}, (\cdot, \cdot)\} \cup \{+, -, \times, /, \wedge\} \cup \{0, 1, \dots, 9\} \cup \{e-1, \dots, e-4\}.$$

841 A.1.3 Mathematical checks on generated functions

843 Each generated component undergoes rigorous symbolic validation to guarantee syntactic
 844 and mathematical sense of the generated expressions. In case a generated expression does
 845 not satisfy the checks, it is rejected, and a new one is generated.

847 Syntactic requirements. We enforce strict variable presence requirements where ODE problems
 848 must contain $\{x\} \subseteq \text{Vars}(u)$, spatial DE problems require $\{x, y\} \subseteq \text{Vars}(u)$, and spatiotemporal
 849 problems need $\{x, t\} \subseteq \text{Vars}(u)$ or $\{x, y, t\} \subseteq \text{Vars}(u)$. Function domain restrictions
 850 prevent undefined operations through logarithmic function constraints $\log(f) \Rightarrow f > 0$ on $\bar{\Omega}$, square root function
 851 requirements $\sqrt{f} \Rightarrow f \geq 0$ for spatial components, and division safety ensuring denominators remain bounded away from zero. To ensure symbolic rather
 852 than constant generation, we forbid purely numeric arguments to transcendental functions
 853 so that $\sin(\alpha), \cos(\alpha), \exp(\alpha), \log(\alpha) \notin \text{Lang}(\mathcal{G}_{\mathbb{D}})$ for $\alpha \in \mathbb{R}$. Integer powers are restricted to
 854 degree ≤ 3 to preserve $H^1(\Omega)$ membership on bounded domains, ensuring that for polynomials
 855 $u = \sum_{|\beta| \leq 3} c_{\beta} x^{\beta}$ we have $\|u\|_{H^1(\Omega)} < \infty$ when Ω is bounded. Function compositions are
 856 validated for smoothness preservation where admissible functions $f \in \{\sin, \cos, \exp, \tanh\}$
 857 applied to arguments g with controlled growth maintain C^{∞} regularity on bounded domains.

859 Boundary conditions. For boundary condition compatibility, homogeneous Dirichlet conditions
 860 $u|_{\partial\Omega} = 0$ are enforced by multiplying spatial components with boundary-vanishing
 861 envelopes such as $\psi_{\text{env}}(x, y) = \sin\left(\frac{\pi x}{L_x}\right) \sin\left(\frac{\pi y}{L_y}\right)$ for rectangular domains, ensuring $u_{\text{modified}} \in$
 862 $H_0^1(\Omega)$. Neumann compatibility for problems requiring $\frac{\partial u}{\partial n}|_{\partial\Omega} = 0$ uses cosine spatial modes
 863 that naturally satisfy zero normal derivative conditions.

864 Table 7: GVAE architecture summary. LN: LayerNorm over $[C, L]$; all linear/conv layers
 865 use bias=False unless noted.

866

867 Block	868 Layer	869 Dims / Kernel / Len	870 Act/Norm
871 Input	872 Tensor	873 $C=53, L=72$	874 –
875 Encoder	876 Conv1D	877 $53 \rightarrow 64, k=2, L: 72 \rightarrow 71$	878 ELU
	879 Conv1D	880 $64 \rightarrow 128, k=3, L: 71 \rightarrow 69$	881 ELU
	882 Conv1D	883 $128 \rightarrow 256, k=4, L: 69 \rightarrow 66$	884 ELU
885 Decoder (positional)	886 Linear	887 $256 \times 66 = 16, 896 \rightarrow 256$	888 ELU
	889 Heads	890 $256 \rightarrow 32 (\mu), 256 \rightarrow 32 (\log \sigma^2)$	891 –
	892 Linear	893 $32 \rightarrow 512$	894 ELU
895	896 GRU	897 input=512, hidden=512, layers=1	898 –
	899 TimeDense	900 $512 \rightarrow 53$ (per position, $L=72$)	901 –
902 Latent dim / samples		z -dim = 32; decoder samples per input = 1	

878

879 Table 8: Lightning module summary (train mode).

880

881 Name	882 Type	883 Params	884 Mode
885 model	886 GrammarVAE	887 6.1 M	888 train
889 Total trainable params			890 6.1 M (24.495 MB)

891

892 Constants. Numerical stability is maintained through exponential scaling control using
 893 scientific notation coefficients with mantissa $m \in [0.001, 999]$ and exponent $e \in [-4, 4]$ to
 894 prevent overflow and underflow. Floating point precision involves rounding numeric literals
 895 to 3 decimal places for most components and 6 decimal places for wave modes, converting
 896 to rational representations when possible to avoid precision degradation.

897

898 Uniqueness of expressions. Expression canonization includes converting fractional powers
 899 to $\sqrt{\cdot}$ notation when $p = 1/2$, transforming reciprocal notation $x^{-1} \mapsto 1/x$, and simplifying
 900 coefficients such as $(2 \times 3)x \mapsto 6x$. Uniqueness is enforced through syntactic equivalence
 901 classes where we define $s \sim s'$ if their canonized forms coincide after symbolic simplification,
 902 maintaining exactly one representative per equivalence class $[s] \in \mathcal{L}(\mathcal{G}_{\mathbb{D}})/\sim$ using a global
 903 hash table that tracks all generated canonical forms.

904

905 A.2 GVAE Model and Training Details

906

907 We employ a Grammar Variational Autoencoder (GVAE) that maps one-hot sequences of
 908 CFG production rules to a continuous latent space and decodes back to valid rule sequences.
 909 Inputs are $x \in \mathbb{R}^{B \times C \times L}$ with $C=53$ rules and $L=72$ time steps (dataset shape $N \times C \times L =$
 910 **23, 682** $\times 53 \times 72$); targets are $y = \arg \max_c x \in \{0, \dots, 52\}^{B \times L}$.

911

912 Architecture. The encoder stacks three valid (no-pad) 1D convolutions with ELU activations,
 913 followed by a linear layer and two bias-free heads producing $\mu, \log \sigma^2 \in \mathbb{R}^{32}$. The
 914 decoder is non-autoregressive (“positional”): it lifts $z \in \mathbb{R}^{32}$ to a hidden state, runs a GRU
 915 across positions, then applies a time-distributed linear projection to rule logits. Shapes and
 916 hyperparameters are summarized in Table 7. Lightning reports 6.1 M trainable parameters
 917 (model size 24.495 MB; see Table 8).

918

919 Losses and regularization. The objective is

920

$$921 \mathcal{L} = \mathcal{L}_{\text{recon}} + \beta(t) \underbrace{\text{KL}(q(z|x) \parallel p(z))}_{\text{latent}} + \gamma(e) \left(0.8 \mathcal{L}_{\text{Hull}} + 0.8 \mathcal{L}_{\text{ph}} + 10^{-4} \mathcal{L}_{\text{smooth}} \right),$$

922

923 with $\mathcal{L}_{\text{recon}}$ being equal to cross entropy loss of $\mathcal{L}_{\text{recon}}(\overline{\text{logits}}, y)$, where $\overline{\text{logits}}$ being the
 924 mean over decoder samples (here = 1). The KL weight uses a linear warmup $\beta(t) =$

918
919
920
Table 9: Training hyperparameters and $\mathcal{L}_{\text{topo}}$ term weights. ReduceLROnPlateau monitors
the balanced ELBO.

921 Item	922 Value	923 Details
924 Optimizer	925 AdamW	926 $lr = 3 \times 10^{-4}$, weight decay = 10^{-5}
927 Batch size (train/val)	928 64 / 64	929 4 dataloader workers
930 Precision	931 16-mixed (AMP)	932 global grad clip = 1.0
933 Scheduler	934 ReduceLROnPlateau	935 factor = 0.2, patience = 5
936 Epochs / Early stop	937 200 / 10	938 monitor (validation's set ELBO)
939 KL warmup $\beta(t)$	940 to 1.0 by 7000 updates	941 $\beta_0=0.01$
942 Topo loss activation	943 at val-acc $\geq 20\%$	944 ramp γ over 5 epochs
945 Topo loss schedule	946 train/val every 50 / 12	947 sparse to limit cost
948 Topo loss weights	949 $w_{\text{Hull}}=0.8$, $w_{\text{ph}}=0.8$	950 $w_{\text{smooth}}=10^{-4}$
951 Ph settings	952 max points = 24, max dim = 1	953 Rips on CPU, scales {0.10, 0.50}
954 Hull directions	955 $K=256$	956 fixed $U_K \subset \mathbb{S}^{d-1}$

$\beta_0 + (1 - \beta_0) \min\left(\frac{t}{7000}, 1\right)$ with $\beta_0=0.01$. The geometric topological block activates once validation sequence-exact accuracy reaches 20%, then ramps $\gamma(e)$ from 0 to 1 over 5 epochs. Topological loss' terms (Hull, ph@scale on CPU, smooth) are computed in fp32 and scheduled sparsely (train every 50 steps and validate every 12 batches). Upon $\mathcal{L}_{\text{topo}}$ activation, the LR scheduler's best-score baseline is reset to the new balanced ELBO.

939
Data and splits. We train on an HDF5 corpus of one-hot sequences under a typed CFG.
940 Random split with seed 42 into train/val/test of 70%/20%/10% yields the counts in Table 10.
941

942
Table 10: Dataset and splits for GVAE training ($C=53$, $L=72$).
943

944 Split	945 # Sequences
946 Train	947 16,578
948 Val	949 4,736
950 Test	951 2,368

950
Environment and software. Experiments ran on an NVIDIA RTX 5080 Laptop GPU
951 (16 GB VRAM). Key versions are summarized in Table 11.

952
Table 11: Compute environment.
953

954 Component	955 Spec
956 CPU	957 Intel Core Ultra 9 275HX, 24C/24T @ 2.7 GHz
958 RAM	959 32 GB
960 GPU	961 NVIDIA GeForce RTX 5080 Laptop GPU 16 GB VRAM)
962 Python	963 3.10.18
964 PyTorch / Lightning	965 2.7.1+cu128 / 2.5.2
966 CUDA / cuDNN	967 12.8 / 90800

962
Training procedure and metrics. We train on a single GPU with AMP and gradient clipping.
963 The primary validation metric is the val ELBO, combining CE, KL (with warmup),
964 and $\mathcal{L}_{\text{topo}}$ (when enabled). We also log CE, KL, ELBO variants, Topo loss components, and
965 sequence-exact accuracy. Early stopping halts after 10 epochs without improvement in val
966 elbo full.

967
Runtime observations. Before the activation of $\mathcal{L}_{\text{topo}}$, epochs take a few seconds. Around
968 the activation point (48th epoch), training duration is ~ 3 s, however at the (50th epoch)
969 that $\mathcal{L}_{\text{topo}}$ starts getting calculated training duration raises to ~ 262 s. This spike is ex-
970 pected. The $\mathcal{L}_{\text{topo}}$ builds a Vietoris–Rips complex and computes persistent homology on
971 the CPU. Constructing distance matrices and boundary operators. Differentiating through

972 them (ph/smooth-Hessian), dominates wall-clock time and introduces CPU↔GPU synchronization overhead.
 973
 974

975 Decoding & evaluation settings. Non-autoregressive (positional) decoding with one latent
 976 sample, max length 72, vocabulary size 53. When applicable, a CFG mask enforces per-step
 977 validity. Report sequence-exact accuracy, validity rate, CE, KL, and ELBO on validation/
 978 test sets.
 979

980 A.3 Geometry Regularization

981 Here, we provide details on the additional loss terms added to the GVAE loss to form the
 982 Topological GVAE (TGVAE).
 983

984 Convex hull loss. Let $z \in \mathbb{R}^d$ be a latent vector and $Z = \{z_i\}_{i=1}^B$ the current batch. We
 985 maintain a reservoir of latent vectors $R_t \subset \mathcal{R}^d$ using a farthest point insertion with distance
 986 $\delta > 0$ (Gonzalez, 1985). At each iteration t , we freeze R_t^{prev} , compute losses, and update
 987 the reservoir only if a latent is not δ close to any z seen thus far. Considering fixed unit
 988 directions $\{n_k\}_{k=1}^K \in \mathbb{S}^{d-1}$ and a support function $h_Z(n) = \sup_{z \in Z} \langle n, z \rangle$ (Rockafellar, 2015)
 989 we define:
 990

$$991 \mathcal{L}_{\text{Hull}}(R_t^{\text{prev}}, Z_t) = \frac{1}{BK} \sum_{i=1}^B \sum_{k=1}^K [\langle n_k, z_i \rangle - h_{R_t^{\text{prev}}}(n_k)]^2 +, \quad \text{where } [\cdot]_+ = \max\{\cdot, 0\}.$$

992 If $\mathcal{L}_{\text{Hull}} = 0$ then every z_i lies in the an explicit convex enclosure of the frozen reservoir
 993 $\cap_{d=1}^D \{z : \langle n_m, z \rangle \leq h_{R_t^{\text{prev}}} \}$. Inside those bounds, we remove small spurious loops and holes
 994 at the working resolution set by δ .
 995

996 Persistent homology loss. Let $P_t = R_t^{\text{prev}} \cup Z_t$, and $V_k(P)$ the persistence diagram of
 997 Vietoris-Rips homology across the scale ϵ (Edelsbrunner & Harer, 2010). We set the working
 998 radius $r = \sqrt{2}\delta$ using the clamped lifetime $\ell_r(b, d) = \max\{0, \min(d, r) - \min(b, r)\}$, and
 999 define the persistent homology loss:
 1000

$$1002 \mathcal{L}_{\text{ph}}(P_t) = \sum_{(b, d) \in V_1(P_t)} \ell_r(b, d)^2 + a_0 \sum_{(b, d) \in V_0(P_t)} \ell_r(b_0, d_0)^2,$$

1003 which suppresses small loops H_1 and micro-clusters H_0 at resolution r .
 1004

1005 Smoothing loss. We also penalize large Hessian energies to prevent sharp decode curva-
 1006 tures, to ensure that small moves in the latent space produce stable changes. For decoder
 1007 $\mathcal{D}_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^M$, we set $f(z) = \mathbf{1}^\top \mathcal{D}_\theta$ and $H_f(z) = \nabla^2 f(z)$, with $v \sim \mathcal{N}(0, I_d)$ we define:
 1008

$$1009 \mathcal{L}_{\text{smooth}} = \mathbb{E}_{z \in Z_t} \|H_f(z)v\|_2^2,$$

1010 which is estimated using the probes as in Hutchinson (1989).
 1011

1012 B Iterative Search and Refinement

1013 This appendix details the two-stage search procedure outlined in Section 2. We provide
 1014 the mathematical formulation and implementation details for both structure discovery and
 1015 coefficient refinement.
 1016

1017 B.1 Notation and Setup

1018 When using SIGS on a specific problem, the user may specify a structural Ansatz F con-
 1019 sisting of the compositional nature the proposed solutions should follow. For example, one
 1020 could specify spatio-temporal separability as $u(x, t) = \sum_{j=1}^K \phi_j(x)\psi_j(t)$, leaving K spatial
 1021 and K temporal functions, overall $L = 2K$ components, to be chosen by SIGS. We denote
 1022 by $A : \{\mathcal{L}(\mathcal{G})\}^L \rightarrow \mathcal{L}(\mathcal{G})$ the assembly map that composes the single components into the
 1023 1024 1025

final solution following the Ansatz, and refer to the component indices as $\mathbb{N}_L = \{1, \dots, L\}$. Recall the TGVAE encoder $\mathcal{E} : \mathcal{L}(\mathcal{G}) \mapsto Z$ and decoder $\mathcal{D} : Z \rightarrow \mathcal{L}(\mathcal{G})$ that assign strings (functions) from the language to latent vectors in $Z \subset \mathbb{R}^d$. For a given component $c \in \mathbb{N}_L$, the Ansatz specifies the variables c should contain. This prior knowledge is incorporated by filtering the library to contain only valid component candidates $\mathcal{L}^{(c)}$ and restricting Z to the component-specific latent set $Z^{(c)} = \{z_i^{(c)}\}_{i=1}^{N_c}$ by applying the encoder \mathcal{E} to $\mathcal{L}^{(c)}$.

1033 B.2 Target loss: Discretized PDE residual

1035 Based on the continuous augmented PDE residual $\mathcal{R}(u)$ from Equation equation 2, we
1036 formulate its discretized form

$$1037 \quad R(u) = \frac{1}{|\mathcal{M}|} \sum_{x \in \mathcal{M}} \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (\mathcal{S}[u](x, t))^2 \quad (3)$$

$$1040 \quad + \beta_1 \frac{1}{|\mathcal{M}_{IC}|} \sum_{x \in \mathcal{M}} (u(x, 0) - u_0(x))^2 \quad (4)$$

$$1042 \quad + \beta_2 \frac{1}{|\mathcal{M}_{BC}|} \sum_{x \in \mathcal{M}_{\partial\Omega}} \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (\mathbb{B}[u](x, t) - g(x, t))^2, \quad (5)$$

1045 where \mathcal{M} is the discretization grid inside the domain, $\mathcal{M}_{BC}/\mathcal{M}_{IC}$ are the discretization
1046 points on the domain boundary and initial conditions, respectively, and \mathcal{T} is the time dis-
1047cretization to evaluate the PDE and boundary operators on. For any candidate decoded
1048 function $u_w = \mathcal{I}(\mathcal{D}(z))$, we use R as the target metric throughout all steps of the solution
1049 search pipeline. Within our experiments, we choose the spatial discretization to be 128 and
1050 128 as a time discretization for all the problems except Damping Wave where we use 64 and
1051 64 respectively.

1052 B.3 Stage I: Structure discovery by iterative clustering

1054 Component-wise libraries. The Ansatz function A specifies the number of components
1055 as well as which variables should be present per component (and possibly other syntactic
1056 requirements). We therefore filter the initial library \mathcal{L} for each component to retain only
1057 viable candidate expressions to obtain $\mathcal{L}^{(c)}$, and the corresponding encoded latent vectors
1058 $Z^{(c)} = \{\mathcal{E}(w) : w \in \mathcal{L}^{(c)}\}$.

1060 Initial clustering. We then iteratively partition the latent subspaces $Z^{(c)}$ for each of the
1061 components separately into K^c clusters by k-means clustering, sample from each the clus-
1062 ters, and assemble solution candidates. Let $\mathbb{N}_{K^{(c)}} = \{1, \dots, K^c\}$ denote the cluster indices
1063 for component c , and by $\mathcal{K} = \mathbb{N}_{K^{(1)}} \times \dots \times \mathbb{N}_{K^{(L)}}$ the cluster index set of all possible
1064 cross-component cluster combinations. For example, in case of spatio-temporal separability
1065 $u(x, t) = f(x) \cdot g(t)$, spatial and temporal components are separated into component-wise
1066 libraries, clustered, and sampled independently and solutions are assembled from pairs of
1067 clusters (k_x, k_t) .

1068 Cluster selection. We sample M tuples of cluster indices $k_i = (k_i^{(1)}, \dots, k_i^{(L)}) \in \mathcal{K}$, where
1069 $k_i^{(c)}$ denotes the index of the cluster used for the c -th component in the i -th sample. For
1070 each component c , we choose a latent vector $z_i^{(c)}$ from the current encoded library vectors
1071 in cluster $k_{i,c}$, decode $w_i^{(c)} = \mathcal{D}(z_i^{(c)})$, and assemble $w_i = A(w_i^{(1)}, \dots, w_i^{(L)})$. Then, we
1072 evaluate the discretized residual for each candidate, $r_i = R(\mathcal{I}(w_i))$. Finally, we select the
1073 candidate with minimal residual and record the cluster indices $k^* = k_i = (k_i^{(1)}, \dots, k_i^{(L)})$ of
1074 the candidate with minimal residual $w^* = w_i$ as the current best clusters.

1077 Iterative subclustering. Each of the component-wise subclusters selected as current best
1078 clusters in the previous iteration are partitioned into K_c sub-clusters by k-means clustering
1079 on the latent vectors. The cluster selection is repeated (sample combinations of clusters,

1080 decode and assemble expressions, evaluate the residual, choose best cluster combination)
 1081 and the best cluster combination k^* is updated from the new, refined clusters. Iteratively,
 1082 the size of the resulting clusters shrinks, focusing in on the final best cluster combination.
 1083 This procedure is repeated until a target residual is reached, $r^* \leq \varepsilon_{\text{struct}}$, or an evaluation
 1084 budget on the number of iterations is exhausted.
 1085

1086 Generation of additional latent vectors. As the size of the latent clusters decreases, there
 1087 are fewer latent vectors of the initial training library z_i . New latent vectors can be generated
 1088 for these clusters, further exploring the latent space beyond what the GVAE has seen during
 1089 training. We generate these samples by convex interpolation of its members with small
 1090 isotropic jitter (decodable latent interpolation).
 1091

1092 B.4 Stage II: Coefficient Refinement

1093 Given the best symbolic structure w^* from Stage I, we freeze the form and expose only
 1094 its numeric literals as trainable parameters $p \in \mathbb{R}^P$, where we protect constants such as
 1095 π , e , and integer exponents. We minimize the PDE residual from $R(u)$ on the resulting
 1096 parametric function family $u^*(\cdot; p) = \mathcal{I}(w^*(p))$ to obtain the best constants
 1097

$$1098 \quad p^* = \arg \min_p R(u^*(\cdot; p))$$

1100 and the corresponding final (exact or approximate) solution $u^*(\cdot) = u^*(\cdot; p^*)$.
 1101

1102 Implementation. We compile $u(\cdot; p)$ in float64 JAX, obtain the required derivatives by
 1103 automatic differentiation to evaluate $\mathcal{S}[u]$ on the grids $(\mathcal{M}, \mathcal{T})$, and compute $R(u)$ batched
 1104 over all points.
 1105

1106 We parse the numeric literals in w^* to form \bar{p} . For single-start, set $p^{(0)} = \bar{p}$. For multi-start,
 1107 draw

$$1108 \quad p^{(0,s)} \sim \mathcal{N}(\bar{p}, \text{diag}((\eta|\bar{p}|)^2)), \quad s = 1, \dots, S,$$

1109 and optimize all starts in parallel (JAX vmap). We use Adam (Optax) with exponential
 1110 learning-rate decay and JIT. Early stopping triggers when $\sqrt{R(u)} < \varepsilon_{\text{tol}}$ or a budget is
 1111 reached.
 1112

1113 Algorithm 1 SIGS: Symbolic Iterative Grammar Solver (overview)

1114 Require: Grammar \mathcal{G} , assembly map A , trained TGVAE $(\mathcal{E}, \mathcal{D})$, discretized residual R ,
 1115 budgets (M, T_{\max}) , thresholds $(\varepsilon_{\text{struct}}, \varepsilon_{\text{tol}})$
 1116 Ensure: Refined symbolic solution u^* with coefficients p^*
 1117 1: Stage 0 (amortized): Initial_Sampling $\rightarrow (w^*, z^*, k^*, r^*)$
 1118 2: Stage I (structure): Subcluster_Refine(w^*, z^*, k^*, r^*) $\rightarrow (w^*, z^*, k^*, r^*)$
 1119 3: Stage II (coeffs): Coefficient_Refinement($w^*, \varepsilon_{\text{tol}}$) $\rightarrow p^*$
 1120 4: return $w^*(p^*)$
 1121

1122

1123

1124 C Experiments

1125

1126 C.1 Problem Definitions

1127

1128 We evaluate our method on five representative PDEs spanning steady-state and time-
 1129 dependent settings. Following the general formulation in equation 1, we specify the dif-
 1130 ferential operator \mathbb{D} in Table 14, together with the computational domain and mesh used
 1131 to evaluate the residual $\mathcal{R}(u)$ and to compute discretized solutions with FEM and PINN
 1132 methods. The forcing term \mathbf{f} and initial/boundary conditions for each test problem are spec-
 1133 ified in the following, where we distinguish cases with known (manufactured) and unknown
 analytic solutions.

1134 Algorithm 2 Stage 0 (amortized): Library clustering and initial assembly

1135

1136 Require: Assembly map A , encoder/decoder $(\mathcal{E}, \mathcal{D})$, discretized residual R , draw budget M

1137 Ensure: Best candidate (w^*, z^*, k^*, r^*)

1138 1: For each component $c \in \mathbb{N}_L$: enforce variable constraints \mathcal{C}_c from the Ansatz to filter
the library $\mathcal{L}^{(c)}$

1139 2: Encode and cluster: $Z^{(c)} = \{\mathcal{E}(w) : w \in \mathcal{L}^{(c)}\}$, partition into K^c clusters; let $\mathbb{N}_{K^{(c)}} =$
 $\{1, \dots, K^c\}$

1140 3: Initialize incumbent $r^* \leftarrow +\infty$

1141 4: for $i = 1$ to M do

1142 5: Sample cluster tuple $k_i = (k_{i,1}, \dots, k_{i,L}) \in \mathbb{N}_{K^{(1)}} \times \dots \times \mathbb{N}_{K^{(L)}}$

1143 6: for each c do

1144 7: Draw $z_i^{(c)}$ from cluster $k_{i,c}$; decode $w_i^{(c)} = \mathcal{D}(z_i^{(c)})$

1145 8: end for

1146 9: Assemble $w_i = A(w_i^{(1)}, \dots, w_i^{(L)})$

1147 10: Score $r_i = R(\mathcal{I}(w_i))$

1148 11: if $r_i < r^*$ then

1149 12: $(w^*, z^*, k^*, r^*) \leftarrow (w_i, [z_i^{(1)}; \dots; z_i^{(L)}], k_i, r_i)$

1150 13: end if

1151 14: end for

1152 15: return (w^*, z^*, k^*, r^*)

1154

1155 Algorithm 3 Stage I: Focused subclustering and structure refinement

1156

1157 Require: Incumbent (w^*, z^*, k^*, r^*) from Stage 0, encoder/decoder $(\mathcal{E}, \mathcal{D})$, residual R , as-
sembly A , budgets (T_{\max}) , threshold $\varepsilon_{\text{struct}}$

1158 Ensure: Updated (w^*, z^*, k^*, r^*)

1159 1: For each $c \in \mathbb{N}_L$: restrict to latents in cluster $K^{(c),*}$ and partition into H_c subclusters

1160 2: $t \leftarrow 0$

1161 3: while $r^* > \varepsilon_{\text{struct}}$ and $t < T_{\max}$ do

1162 4: Sample subcluster tuple $h = (h_1, \dots, h_L) \in \mathbb{N}_{H^{(1)}} \times \dots \times \mathbb{N}_{H^{(L)}}$

1163 5: for each c do

1164 6: if subcluster h_c is too small then

1165 7: generate samples in h_c by convex interpolation plus small isotropic jitter

1166 8: end if

1167 9: Draw $z^{(c)}$ from subcluster h_c

1168 10: Decode $w^{(c)} = \mathcal{D}(z^{(c)})$

1169 11: end for

1170 12: Assemble $w = A(w^{(1)}, \dots, w^{(L)})$

1171 13: Score $r = R(\mathcal{I}(w))$

1172 14: if $r < r^*$ then

1173 15: $(w^*, z^*, k^*, r^*) \leftarrow (w, [z^{(1)}; \dots; z^{(L)}], h, r)$

1174 16: end if

1175 17: $t \leftarrow t + 1$

1176 18: end while

1177 19: return (w^*, z^*, k^*, r^*)

1178

C.1.1 Construction of known analytical solutions

1181 For the four problems with known analytical solutions, we employ the method of man-
1182 ufactured solutions to construct the test problems and ensure exact error quantification.
1183 Given a chosen analytical solution u_{true} , we construct the forcing term via $\mathbf{f} = -\mathbb{D}[u_{\text{true}}]$
1184 to guarantee that u_{true} satisfies the PDE exactly. Initial and boundary conditions are then
1185 prescribed from u_{true} to complete the well-posed problem formulation.

1186 The specific analytical solutions are detailed in Table 15. These solutions are chosen to ex-
1187 hibit diverse mathematical behaviors: the Burgers’ equation features a smooth shock profile
with nonlinear advection, the diffusion equation uses a multi-mode separated solution with

1188 Algorithm 4 Stage II: Multi-start coefficient refinement (JAX)
1189
1190 Require: Best structure w^* , residual R , tolerance ε_{tol} , starts S , noise scale η
1191 Ensure: Optimized coefficients p^* and refined $w^*(p^*)$
1192 1: Parse numeric literals in w^* to get \bar{p}
1193 2: for $s = 1$ to S do
1194 3: Initialize $p^{(0,s)} \sim \mathcal{N}(\bar{p}, \text{diag}((\eta|\bar{p}|)^2))$
1195 4: end for
1196 5: Optimize all starts with Adam (JAX, float64, JIT) and exponential LR decay; at each
1197 step evaluate $R(\mathcal{I}(w^*(p^{(t,s)})))$
1198 6: Early-stop when $\sqrt{R} < \varepsilon_{\text{tol}}$ or budget reached; keep the best $p^* =$
1199 $\arg \min_s R(\mathcal{I}(w^*(p^{(s)})))$
1200 7: return p^* and $w^*(p^*)$

1201 Table 12: Canonical problems reproduced from prior work. Dimension notation: $n+mD$
1202 denotes n spatial and m temporal variables.
1203

Problem (paper)	Operator \mathbb{D}	Dim	Domain	Mesh	Ground truth u^*
Poisson1 (HDTLGP)	$u_{xx} + u_{yy}$	2D	$[0, 1]^2$	64^2	$\sin(\pi x) \sin(\pi y)$
Advection3 (HDTLGP)	$u_t + u_x + u_y$	2+1D	$[0, 1]^2 \times [0, 2]$	$64^2 \times 64$	$\sin(x-t) + \sin(y-t)$
Wave2D (SSDE)	$u_{tt} - (u_{xx} + u_{yy})$	2+1D	$[-1, 1]^2 \times [0, 1]$	$8^2 \times 8$	$e^{x^2} \sin(y) e^{-0.5t}$

1209
1210 exponential decay, the wave equation employs a truncated Fourier series, and the damping
1211 wave incorporates both temporal decay and spatial wave propagation in two dimensions.
1212

1213 C.1.2 Problem without known analytic solution

1214
1215 The Poisson–Gauss problem represents a realistic scenario where no analytical solution is
1216 available, making it particularly valuable for assessing method performance in practical
1217 applications. The problem consists of the steady-state Poisson equation $\nabla^2 u = \mathbf{f}$ on the
1218 unit square $[0, 1]^2$ with homogeneous Dirichlet boundary conditions $u = 0$ on $\partial[0, 1]^2$.

1219 The forcing term \mathbf{f} is constructed as a superposition of n isotropic Gaussian sources:

$$1220 \quad \mathbf{f}(x, y) = \sum_{i=1}^n \exp \left(-\frac{(x - \mu_{x,i})^2 + (y - \mu_{y,i})^2}{2\sigma^2} \right) \quad (6)$$

1223 with fixed width $\sigma = 0.1$ and deterministically chosen centers:

- 1224 • PG-2: $(0.3, 0.8), (0.7, 0.2)$
- 1225 • PG-3: $(0.3, 0.8), (0.7, 0.2), (0.5, 0.2)$
- 1226 • PG-4: $(0.3, 0.8), (0.7, 0.2), (0.5, 0.2), (0.4, 0.6)$

1228 This configuration creates localized source regions with smooth spatial variation, testing
1229 the method’s ability to capture multi-scale features and handle problems without ground
1230 truth solutions. For evaluation on this problem, we rely on mesh convergence studies and
1231 physics-based consistency checks rather than direct error computation against an analytical
1232 reference.

1234 C.2 Solution Ansatz specific to our experiments

1235
1236 Our framework generates eigenfunction components for five distinct operator classes, each
1237 producing characteristic mathematical patterns with specific parameter ranges that ensure
1238 physical relevance and numerical stability.

- 1239 • Wave operators $\mathbb{D}_{\text{wave}} = \nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2}$ generate oscillatory eigenmodes
1240 $a_k \sin(k\pi x) \cos(ck\pi t)$ where mode indices $k \in \{1, 2, \dots, K\}$ determine spatial har-
1241 monic frequencies, wave speeds $c \in [0.1, 0.8]$ control temporal oscillation rates,

1242 Table 13: Closed-form parity on canonical problems from prior work. We list the operator,
 1243 domain, evaluation mesh, the ground-truth solution u^* , the closed form printed in the
 1244 original paper, and the expression found by SIGS, together with relative L^2 errors (discrete,
 1245 uniform grid).

1246

1247 Problem (source)	1248 Baseline expression	1249 SIGS expression
1248 Poisson1 (HDTLGP)	$\sin(3.141 x) \sin(3.142 y)$	$\sin(\pi x) \sin(\pi y)$
1249 Advection3 (HDTLGP)	$-\sin(0.9838t - x) - \sin(0.9979t - y)$	$\sin(x - t) + \sin(y - t)$
1250 Wave2D (SSDE)	$e^{x^2 - 0.5t} \sin(y)$	$e^{x^2} \sin(y) e^{-0.5t}$

1252

1253 Table 14: Summary of benchmark problems. Dimension notation: $n+mD$ denotes n spatial
 1254 and m temporal dimensions.

1255

1256 Problem	1257 Operator \mathbb{D}	1258 Dim	1259 Domain	1260 Mesh	1261 Key Parameters
Burgers'	$u_t + uu_x - \nu u_{xx}$	1+1D	$[-5, 5] \times [0, 2]$	128 \times 128	$\nu = 0.01$
Diffusion	$u_t - \kappa u_{xx}$	1+1D	$[0, 1.397] \times [0, 1]$	128 \times 128	$\kappa = 0.697$
Damping wave	$u_{tt} + u_t - c^2(u_{xx} + u_{yy})$	2+1D	$[-8, 8]^2 \times [0, 4]$	32 \times 32 \times 32	$c = 0.8$
Poisson-Gauss	$-(u_{xx} + u_{yy})$	2D	$[0, 1]^2$	100 \times 100	

1262

1263 and amplitude coefficients $a_k = \frac{m \times 10^e}{k} \times \frac{\pi}{K}$ use scientific notation with mantissa
 1264 $m \in [5, 9]$ and exponential damping to ensure numerical stability across multiple
 1265 scales.

1266

- 1267 Diffusion operators $\mathbb{D}_{\text{diff}} = \nabla^2 - \frac{\partial}{\partial t}$ produce separable heat modes
 $\frac{2M_0}{L} \sin\left(\frac{(2n+1)\pi x}{L}\right) e^{-\frac{(2n+1)^2 \pi^2 D t}{L^2}}$ where amplitude coefficients $M_0 \in [1, 3]$ set initial
 1268 magnitudes, domain lengths $L \in [0.1, 1.5]$ determine spatial scales, diffusivities
 1269 $D \in [0.01, 1]$ control temporal decay rates, and odd harmonic indexing $(2n+1)$
 1270 corresponds to homogeneous Dirichlet boundary conditions with mode numbers
 1271 $n \in \{0, 1, 2\}$ generating the first three eigenmodes.
- 1272 Viscous Burgers operators $\mathbb{D}_{\text{Burgers}} = u \frac{\partial u}{\partial x} - \nu \frac{\partial^2 u}{\partial x^2}$ create shock transition
 1273 profiles consisting of average components $\frac{u_L + u_R}{2}$ and shock components
 $\frac{u_L - u_R}{2} \tanh\left(\frac{(x - x_0 - st)(u_L - u_R)}{4\nu}\right)$ where left asymptotic states $u_L \in [1, 3]$ and right
 1274 asymptotic states $u_R \in [-1, 1]$ define the shock amplitude, propagation speeds
 1275 $s \in [0.1, 2]$ control shock movement, initial positions $x_0 \in [-1, 1]$ set shock locations,
 1276 and kinematic viscosities $\nu \in [0.01, 1]$ determine shock width.
- 1277 Poisson-Gauss operators $\mathbb{D}_{\text{Poisson}} = \nabla^2$ with source terms generate localized Gaussian
 1278 profiles $e^{-\alpha((x-x_0)^2 + (y-y_0)^2)}$ for superposition of Gaussian source terms and
 1279 polynomial harmonic functions for steady-state equilibrium configurations, where
 1280 decay parameters α control Gaussian widths, center coordinates (x_0, y_0) determine
 1281 localization, and multiple Gaussians can be superposed as source terms.
- 1282 Outgoing damped wave operators $\mathbb{D}_{\text{out-wave}} = \nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} + \gamma \frac{\partial}{\partial t}$ combine
 1283 envelope functions $h\left(e^{((x-x_0)^2 + (y-y_0)^2)/(w(1+t))} + 1\right)^{-1}$, oscillatory kernels
 $\cos(k\sqrt{(x-x_0)^2 + (y-y_0)^2} - ct)$, and decay factors e^{-at} where amplitudes $h \in$
 1284 $[0.01, 0.5]$, envelope width parameters $w \in [0.3, 1.0]$, radial wave numbers $k \in$
 1285 $[0.5, 4.0]$, phase velocities $c \in [0.1, 1.0]$, temporal decay rates $a \in [0.02, 0.8]$, and
 1286 center coordinates $(x_0, y_0) \in [-6, 6]^2$ control the composite spatiotemporal structure.

1287

1288 Our grammar-based approach produces eigenfunction components at different structural
 1289 levels including elementary eigenmodes corresponding to individual spatial harmonics
 $\phi_k(x)$ and temporal factors $\psi_k(t)$, separable products representing complete eigenfunctions
 1290 $\phi_k(x)\psi_k(t)$ generated when the grammar produces expressions containing both spatial and

1296 Table 15: Analytical (manufactured) solutions for benchmark problems.
1297

1298 Problem	1299 Analytical Solution u_{true}	1300 Constants
Burgers'	$0.86 + 0.6 \tanh(25.8t - 30.0x + 9.9)$	–
Diffusion	$A[\sin\left(\frac{\pi x}{L}\right) e^{-\frac{\pi^2 \kappa}{L^2} t} - \sin\left(\frac{3\pi x}{L}\right) e^{-\frac{9\pi^2 \kappa}{L^2} t} + \sin\left(\frac{5\pi x}{L}\right) e^{-\frac{25\pi^2 \kappa}{L^2} t}]$	$A = 3.974, L = 1.397$
Damping wave	$e^{-\alpha t} \cos(\omega t - KR(x, y)),$ where $R(x, y) = \sqrt{(hx + 1)^2 + (hy - 1)^2}$	$h = 0.2, K = 2.5, \omega = 0.4,$ $\alpha = 0.45$

1305
1306 temporal variables, and composite structures like non-separable patterns $\cos(\sqrt{x^2 + y^2}/t)$
1307 that the grammar can generate through its compositional rules but cannot be factorized.
1308 For ODE problems and linear spatial DEs such as Poisson and Laplace equations where
1309 component structure is simpler, we supplement operator-informed generation with proba-
1310 bilistic grammar expansion using a context-free grammar that recursively builds expression
1311 trees by selecting binary operations with probability 0.6, unary functions with probability
1312 0.3, and terminal symbols with probability 0.1.
1313

1314 C.3 Configuration of baseline methods
13151316 C.3.1 SSDE Primitive Sets
1317

1318 To ensure fair comparison, SSDE receives primitive sets derived from the same structural
1319 Ansatz used by SIGS. For problems expecting separated variable forms (e.g., $f(x) \cdot g(t)$ for
1320 spatiotemporal PDEs), we provide SSDE with functions that appear in the corresponding
1321 variable-specific clusters within SIGS’s grammar.
1322

1323 Table 16: SSDE primitive sets derived from SIGS’s structural Ansatz

1324 Problem	1325 Expected Form	1326 Variables	1327 Function Set
Burgers	$f(x, t)$	(x, t)	$\{+, -, \times, \div, \exp, \tanh, \sin, \cos\}$
Diffusion	$\sum_i f_i(x) \cdot g_i(t)$	(x, t)	$\{+, -, \times, \div, \exp, \tanh, \sin, \cos, \log\}$
Damping Wave	$f(x, y) \cdot g(t)$	(x, y, t)	$\{+, -, \times, \div, \exp, \sin, \cos, \sqrt{\cdot}\}$
PG-2/3/4	$f(x, y)$	(x, y)	$\{+, -, \times, \div, \exp, \log, x^n, \sin, \cos\}$

1330
1331 Rationale for Primitive Selection. The primitive sets are determined by analyzing which
1332 functions appear in SIGS’s variable-specific clusters:
1333

- 1334 • For separated forms $f(x) \cdot g(t)$: We include functions from both the spatial cluster
1335 (containing x) and temporal cluster (containing t)
- 1336 • For spatiotemporal problems: $\{\sin, \cos\}$ from spatial modes, $\{\exp\}$ from temporal
1337 decay, $\{\tanh\}$ for shock profiles (Burgers-specific)
- 1338 • For wave equations: Exclude \exp since the temporal cluster for waves contains only
1339 oscillatory functions
- 1340 • For spatial-only problems (PG): Include functions from the (x, y) spatial cluster
1341

1342 This ensures both methods access identical function spaces, SIGS through its structured
1343 grammar clusters and SSDE through explicit primitive specification. The key difference lies
1344 in search strategy: SIGS restricts combinations to physically motivated forms, while SSDE
1345 explores all possible tree compositions.
1346

1347 SSDE Hyperparameters. All problems use consistent RL hyperparameters: learning rate
1348 0.0005, entropy weight 0.07, batch size 1000, 200,000 training samples, and expression length
1349 constraints between 4 and 30 tokens (extended to 60 for Diffusion due to its multimodal
structure requiring more complex expressions).

1350 C.3.2 HD-TLGP Protocols
13511352 We evaluate HD-TLGP under two protocols that parallel the conditions for SSDE and SIGS.
13531354 Protocol 1: Knowledge-Based Initialization. HD-TLGP receives problem-specific solution
1355 components in its knowledge base:
1356

- 1357 • Diffusion: First mode with exact amplitude $A \sin(\pi x/L) \exp(-\pi^2 D t/L^2)$, templates
1358 for modes 3 and 5, and 2-3 mode combinations
- 1359 • Burgers: Core shock $\tanh(\alpha(x - x_0 - st))$ and scaled variant
- 1360 • Damping Wave: Radial motif $\cos(k\sqrt{(x - x_0)^2 + (y - y_0)^2} - \omega t)$ and separable tem-
1361 plate $\sin(\pi x) \sin(\pi y) \cos(\omega t) \exp(-\gamma t)$
- 1362 • PG-2/3/4: Boundary mask $\sin(\pi x) \sin(\pi y)$, individual Gaussians for each center,
1363 and sum-of-Gaussians template

1364 These components test whether genetic programming can extend partial solutions (Dif-
1365 fusion/Wave), refine parametric forms (Burgers), or combine spatial structures (Damping
1366 Wave, PG).
13671368 Protocol 2: Primitive-Only Discovery. HD-TLGP starts from random expressions using
1369 exactly the same primitive set as SSDE for each problem:
1370

- 1371 • 1D problems (Diffusion, Wave, Burgers): $\{+, -, \times, \div, \sin, \cos, \exp, \tanh\}$
- 1372 • Damping Wave: $\{+, -, \times, \div, \sin, \cos, \exp, \tanh, \sqrt{}\}$
- 1373 • PG-2/3/4: $\{+, -, \times, \div, \sin, \cos, \exp, \log, \sqrt{}\}$

1374 No knowledge base components are provided, requiring complete discovery from elementary
1375 functions. This ensures all symbolic methods explore identical function spaces.
13761377 Implementation Details. Population size 200 (1D) or 50 (2D), crossover 0.6, mutation 0.6,
1378 KB transfer 0.6 (Protocol 1 only), maximum 25 generations or 120-1200 seconds, local
1379 optimization enabled for constant tuning, peephole simplification for expression reduction.
13801381 C.4 FEniCS Validation for Reference Solutions
13821383 For the Poisson-Gauss problems lacking analytical solutions, we establish numerical ground
1384 truth through rigorous finite element analysis. We solve the Poisson equation $-\nabla^2 u = f$
1385 with homogeneous Dirichlet boundary conditions on the unit square, where f consists of
1386 superposed Gaussian sources:
1387

$$1388 f(x, y) = \sum_{i=1}^n \exp\left(-\frac{(x - \mu_{x,i})^2 + (y - \mu_{y,i})^2}{2\sigma^2}\right)$$

1389 Verification Methodology. To validate our FEniCS reference solutions, we employ three
1390 convergence criteria:
1391

- 1392 1. Mesh convergence: Solutions computed on progressively refined meshes (16×16
1393 through 128×128) with P4 elements
- 1394 2. Energy balance: The weak form identity $a(u_h, u_h) = L(u_h)$ must hold to machine
1395 precision, where $a(u, v) = \int_{\Omega} \nabla u \cdot \nabla v \, dx$ and $L(v) = \int_{\Omega} f v \, dx$
- 1396 3. Residual minimization: The strong-form PDE residual $\| -\nabla^2 u_h - f \|_{L^2}$ decreases
1397 at the expected rate $O(h^{p+1})$

1404 Table 17: FEniCS Convergence Study for Poisson-Gauss (PG-2) Problem using P2 Elements
1405

1406	Problem	Mesh	DOF	Runtime (s)	PDE Residual ($\mathbf{R}(\mathbf{u})$)
1407	PG-2	32×32	4225	0.0672	2.0525e−02
		64×64	16641	0.2103	1.0238e−02
		100×100	40401	0.5271	6.5487e−03
		128×128	66049	0.9463	5.1154e−03
		256×256	263169	3.5499	2.5573e−03
		512×512	1050625	14.9201	1.2787e−03
		1024×1024	4198401	57.0198	6.3970e−04

1414
1415 Validation Results. Table 17 shows the convergence study for the Poisson-Gauss (PG-2)
1416 problem using $P2$ elements. This study focuses on validating the stability and efficiency of
1417 the Finite Element solution. The *Runtime* shows the computational cost increases proportionally
1418 with the DOF. Most importantly, the *PDE Residual* ($R(u)$) demonstrates stable, clear convergence,
1419 decreasing by a factor of approximately 2 with each mesh refinement. This confirms a consistent
1420 $O(h^1)$ decay rate for the residual, proving that the solution is systematically converging to satisfy the strong form of the Poisson equation and the Dirichlet
1421 boundary conditions.

1422 Evaluation of Discovered Expressions. Symbolic expressions discovered by SIGS and base-
1423 line methods are evaluated against these FEniCS references through Galerkin projection.
1424 Given a discovered expression $u_{\text{sym}}(x, y)$, we compute its projection onto the finite element
1425 space and measure the relative L^2 error:

$$1426 \text{Error} = \frac{\|u_h^{\text{FEM}} - \Pi_h u_{\text{sym}}\|_{L^2}}{\|u_h^{\text{FEM}}\|_{L^2}}$$

1427 where Π_h denotes the L^2 projection operator onto the $P4$ finite element space. This provides
1428 a rigorous, mesh-independent measure of solution quality for problems without analytical
1429 ground truth.

1430 C.5 Discovered Symbolic Expressions

1431 Tables 18–20 present the symbolic expressions discovered by each method. The structural
1432 differences are immediately apparent: SIGS produces compact, physically interpretable ex-
1433 pressions that directly reflect PDE solution structures—separated variables for diffusion,
1434 traveling waves for Burgers, and properly masked Gaussians for Poisson problems. In
1435 contrast, both HD-TLGP and SSDE generate deeply nested compositions of elementary func-
1436 tions. HD-TLGP Protocol 1, despite receiving solution components, wraps them in other
1437 operations (e.g., $\sin(\sin(\cos(\exp(\cdot))))$ around the Burgers shock), while Protocol 2 often col-
1438 lapses to trivial constants for complex problems. SSDE consistently produces expressions
1439 with extreme nesting depth—up to 500+ operations for Damping Wave, that represent
1440 brute-force fitting rather than discovery of underlying mathematical structure. These ex-
1441 pressions, while potentially achieving low training error fail to generalize and provide no
1442 insight into the PDE dynamics.

1443 C.6 Details on the ablation Study

1444 Mahalanobis distance. Given $x \in \mathbb{R}^d$, mean μ , covariance $\Sigma \succ 0$, the Mahalanobis distance
1445 is $d_M(x, \mu) = \sqrt{(x - \mu)^\top \Sigma^{-1} (x - \mu)}$. For each model, we compute training encoder means
1446 $\{\mu_i\}_{i=1}^N$, estimate Σ from these means, and define $d_{\min}(z) = \min_i d_M(z, \mu_i)$. Our filter
1447 accepts a candidate z iff $d_{\min}^{(\text{with})}(z) \geq \tau$ and $d_{\min}^{(\text{w/o})}(z) \geq \tau$ (we use $\tau=0.8$).

1458

1459

1460

Table 18: Symbolic expressions discovered by SIGS

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

Table 18: Symbolic expressions discovered by SIGS

Problem	Discovered Expression
Burgers	$0.86 + 0.6 \tanh(30(x - 0.33 - 0.86t))$
Diffusion	$3.974(\sin(2.15\pi x)e^{-3.21\pi^2 t} + \sin(0.71\pi x)e^{-0.36\pi^2 t} + \sin(3.58\pi x)e^{-8.93\pi^2 t})$
Damping Wave	$\cos(0.5\sqrt{(x + 5.0)^2 + (y - 5)^2} - 0.4t)e^{-0.45t}$
PG-2	$\sin(\pi x) \sin(\pi y) \cdot \left[0.0080 \exp\left(\frac{0.424((x-0.923)^2 + (y-0.760)^2)}{2.136-0.573^2}\right) + 0.0251 \exp\left(\frac{-1.071((x-0.794)^2 + (y-0.054)^2)}{2.245-0.201^2}\right) + 0.0105 \exp\left(\frac{-1.110((x-0.248)^2 + (y-0.496)^2)}{1.862-0.185^2}\right) \right]$
PG-3	$\sin(\pi x) \sin(\pi y) \cdot \left[0.0079 \exp\left(\frac{0.461((x-0.500)^2 + (y+0.217)^2)}{2.152-0.508^2}\right) + 0.0137 \exp\left(\frac{-0.816((x-0.750)^2 + (y-0.873)^2)}{1.898-0.138^2}\right) + 0.0137 \exp\left(\frac{-0.851((x-0.250)^2 + (y-0.873)^2)}{2.505-0.123^2}\right) + 0.0206 \exp\left(\frac{-1.092((x-0.500)^2 + (y-0.043)^2)}{1.738-0.221^2}\right) \right]$
PG-4	$\sin(\pi x) \sin(\pi y) \cdot \left[0.0068 \exp\left(\frac{-1.489((x-0.731)^2 + (y-0.502)^2)}{1.553-0.195^2}\right) + 0.0112 \exp\left(\frac{-1.123((x-0.500)^2 + (y-0.124)^2)}{1.804-0.159^2}\right) + 0.0294 \exp\left(\frac{-0.031((x-0.665)^2 + (y-0.887)^2)}{2.025-0.584^2}\right) + 0.0069 \exp\left(\frac{-0.992((x-0.267)^2 + (y-0.502)^2)}{1.664-0.155^2}\right) + 0.0286 \exp\left(\frac{-1.024((x-0.501)^2 + (y+0.276)^2)}{1.569-0.190^2}\right) \right]$

Table 19: Symbolic expressions discovered by SSDE

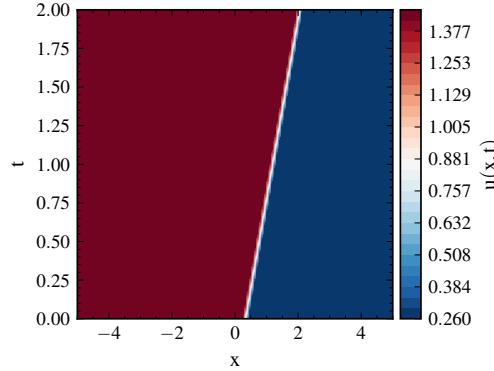
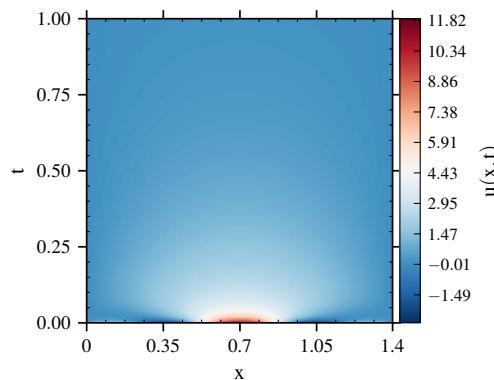
Problem	Expression Found
Burgers	$\exp(\tanh(-1743.845x - 76821.176)) / \sin(\exp(\tanh(\exp(x)))) \cdot \exp(-\tanh(-t \exp(t) - 3t + \tanh(-1743.845x - 76821.176)) / \sin(\exp(\tanh(\exp(x)))))$
Diffusion	$\cos(t + x + \cos(112.185x^3 \tanh(x^2 + x) - 118.201x^3 + 8.824x) - \tanh(t)) / (2t + x) / \cos(-t + x + \cos(112.185x^3 \tanh(x^2 + x) - 118.201x^3 + 8.824x)) + \cos(112.185x^3 \tanh(x^2 + x) - 118.201x^3 + 8.824x) / \cos(-t + x + \cos(112.185x^3 \tanh(x^2 + x) - 118.201x^3 + 8.824x)))$
Damping Wave	Expression with 500+ operations including nested functions, (full expression exceeds reasonable display length)
PG-2	$\sin(x(-0.02582816 + \frac{0.01654789}{\cos(x^2 + \cos(y(y + \exp(x \cdot \cos(y + 0.42096704)))))}))$
PG-3	$x(-0.802y^2 + y)(1.092\cos(\sin(\sin(\sin(x)))) - 0.82463681261637)$
PG-4	$x(\log(\sin(\sin(\cos(\sin(\sin(\sin(\sin(\sin(x)))))))) + 0.383) - 0.007$

1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524

Table 20: Symbolic expressions discovered by HD-TLGP

Problem	Protocol	Expression Found
Burgers	P1	$\sin(\sin(\cos(\exp(-\tanh(0.996 \tanh(25.8t - 30.0x + 9.9)))))) + 0.569$
	P2	$\tanh(\exp(-\cos(\tanh(\tanh(x) + 0.5) + \tanh(1.649 \exp(-x)))) + \tanh(\sin(x + 48.558) + \tanh(x)))$
Diffusion	P1	$3.974 \cdot (\exp(88.121t + 3.974 \exp(-3.525t) \sin(63.495/x) \cdot \sin(2.249x)) \sin(2.249x) + \sin(11.244x)) \exp(-91.646t)$
	P2	$\sin(\tanh(\tanh(x))) + 1.840 \tanh(t + \sin(0.540 \cdot \exp(\frac{0.5}{\sin(\cos(\sin(x+\pi)) / ((1.623t - 5.100)(-t+2x+1.0))}) \tanh(1.019 \cos(1.649t + 0.824))))))$
Damping Wave	P1	$\exp(\cos(0.542t - 0.357((x - 0.135)(x + 0.269) + (y - 0.940)(y + 0.495))^{1/2}))$
	P2	Complex nested expression with 150+ operations including imaginary unit
PG-2	P1	$0.0181 \sin(\exp(1.638(26.282 + \exp(-0.992/(x^2 - 40x + y^2 - 26.598y + 531.969))^{1/2}) - 1.638 \exp(-0.191/(7.560x \cdot \exp(1.670x) + 117.703x - 370.656y + 61.605 \exp(1.670x) - 1833.212)) \sin(x - y)))$
	P2	0.000105
PG-3	P1	$(y - 0.468)(y - 0.440)(\sin((\tanh(y) + 9.870) \exp(0.0120 / ((x - 0.538)(x - 0.032) + (y - 1.662)(y - 1.548))))^{1/2}$
	P2	$9944.705 \sin(x) + 1.218 \times 10^{-12}$
PG-4	P1	$12945.616 / \sin(\exp(\exp(\exp(-1.095 \exp(3.535 / ((x + 0.0183)^2 + (y - 0.389)(y - 0.368)))) / (-1.527y + 1.527(x - 0.764)^2 \exp(3.535 / ((x + 0.0183)^2 + (y - 0.389)(y - 0.368))) - 5.261))))$
	P2	9505.982

1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565

1566 C.7 Solution Visualizations
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586Figure 4: Contour plot of the learned solution $u(x, t)$ for the Burgers equation. The horizontal axis represents the spatial domain $x \in [-5, 5]$, the vertical axis represents the temporal domain $t \in [0, 2]$, and the colormap indicates the solution magnitude ranging from 0.26 to 1.46. The solution is computed on a 128×128 discretization grid1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612Figure 5: Contour plot of the learned solution $u(x, t)$ for the Diffusion equation. The horizontal axis represents the spatial domain $x \in [0, 1.4]$, the vertical axis represents the temporal domain $t \in [0, 1]$, and the colormap indicates the solution magnitude ranging from -1.5 to 11.9. The solution is computed on a 128×128 discretization grid.1613
1614
1615
1616
1617
1618
1619

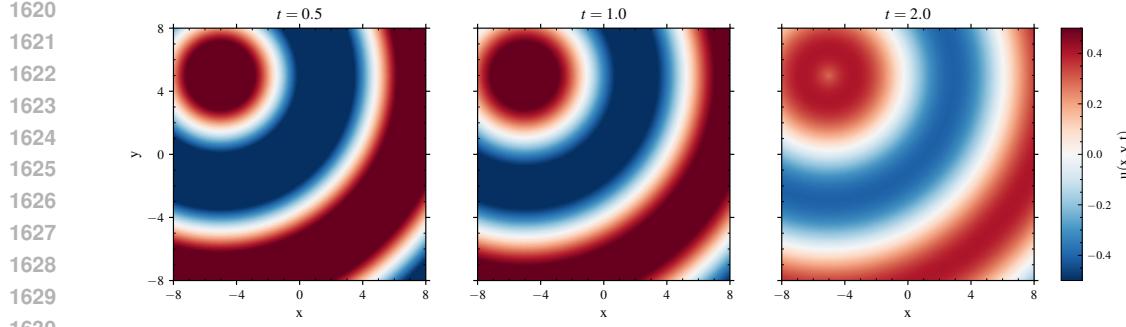


Figure 6: Contour plots of the learned solution $u(x, y, t)$ for the Damped Wave equation at time instances $t \in \{0.5, 1.0, 2.0\}$. The spatial domain is $(x, y) \in [-8, 8]^2$, and the colormap indicates the solution magnitude ranging from -0.5 to 0.5 . The solution is computed on a 128×128 discretization grid.

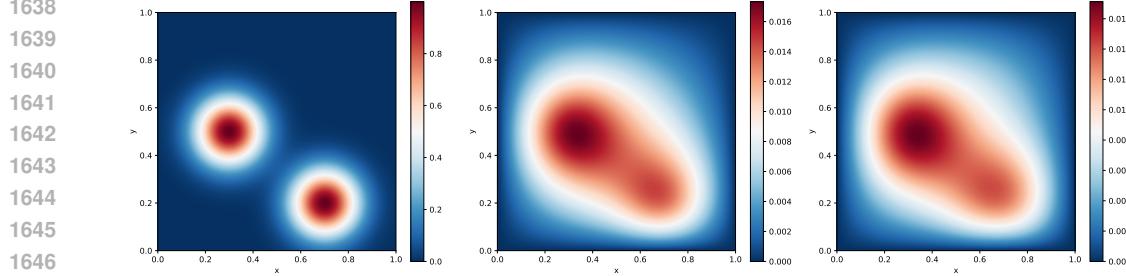
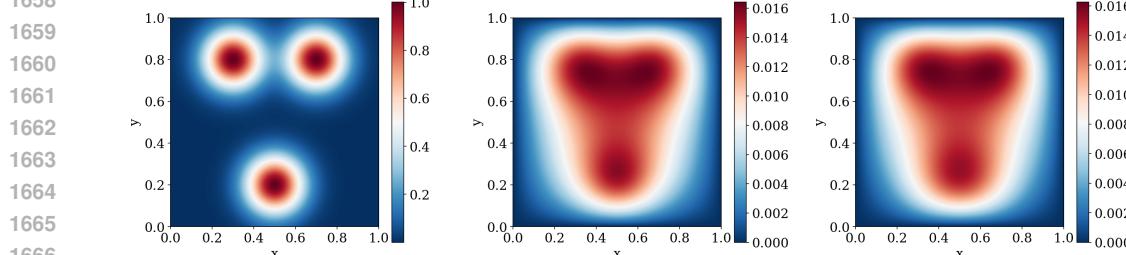


Figure 7: Comparison of numerical approximation and symbolic solution for the Poisson equation with 2 Gaussian source centers. (Left) Source term $F(x, y)$ consisting of 2 Gaussian functions centered at $(0.3, 0.5)$ and $(0.7, 0.2)$ with $\sigma = 0.1$. (Right) Solution obtained using the SIGS method. (Middle) Reference solution obtained by Finite Element Method (FEM) solution computed using FEniCS on a 100×100 mesh with P2 elements. The spatial domain is $(x, y) \in [0, 1]^2$, visualized on a 400×400 grid. The colormap indicates solution magnitude with maximum values of approximately 0.035 .



1668
1669
1670
1671
1672
1673

Figure 8: Comparison of numerical approximation and symbolic solution for the Poisson equation with 3 Gaussian source centers. (Left) Source term $F(x, y)$ consisting of 3 Gaussian functions centered at $(0.3, 0.8)$, $(0.7, 0.8)$, and $(0.5, 0.2)$ with $\sigma = 0.1$. (Middle) Reference solution obtained by Finite Element Method (FEM) solution computed using FEniCS on a 100×100 mesh with P2 elements. (Right) Solution obtained using the SIGS method. The spatial domain is $(x, y) \in [0, 1]^2$, visualized on a 400×400 grid. The relative L^2 error between FEM and SIGS solutions is approximately 1% .

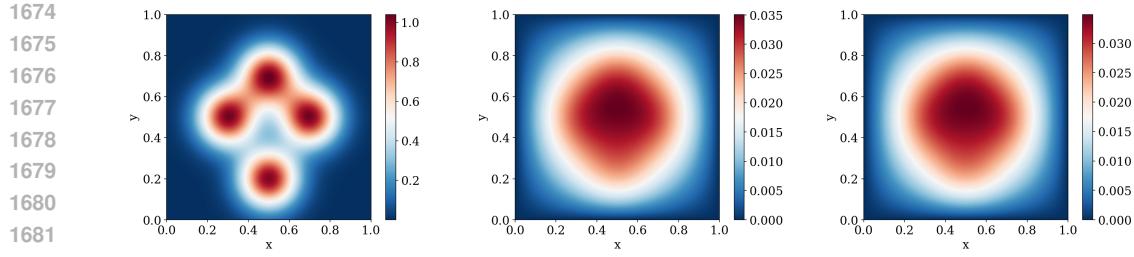


Figure 9: Comparison of numerical approximation and symbolic solution for the Poisson equation with 4 Gaussian source centers. (Left) Source term $F(x, y)$ consisting of 4 Gaussian functions centered at $(0.3, 0.5)$, $(0.7, 0.5)$, $(0.5, 0.2)$, and $(0.5, 0.7)$ with $\sigma = 0.1s$. (Middle) Reference solution obtained by Finite Element Method (FEM) solution using FEniCS on a 100×100 mesh with P2 element. (Right) Solution obtained using the SIGS method. The spatial domain is $(x, y) \in [0, 1]^2$, visualized on a 400×400 grid. The symmetric arrangement of sources produces a cross-like pattern in the solution field.

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

1728 C.8 Additional results during the review process
17291730 C.8.1 Korteweg-de Vries (KdV) Equation
17311732 We are going to study how grammar misspecification, e.g. missing primitive functions,
1733 and the choice of the Ansatz affect the solution capabilities of SIGS. For this purpose, we
1734 consider the Korteweg-de Vries (KdV) equation. The KdV equation is a PDE which models
1735 one-dimensional nonlinear dispersive nondissipative waves, or solitons, defined as:
1736

1737
$$u_t + 6uu_x + u_{xxx} = 0, \quad (7)$$

1738 The one-soliton setup of this problem considers a single wave as it moves across the domain,
1739 and has an analytic solution of the form
1740

1741
$$u(x, t) = \frac{2}{\cosh^2(x - 4t)}, \quad (8)$$

1742 where \cosh is not included in the grammar. We define the domain $\Omega \times T \in [-10, 10] \times [0, 1]$,
1743 and the Ansatz as
1744

1745
$$u(x, t) = \sum_{k=0}^2 \phi(x, t)^k. \quad (9)$$

1746 SIGS found the following solution:
1747

1748
$$0.0003734 \cdot \tanh(0.001355 \cdot t - 0.0006419 \cdot x + 0.0002936)$$

1749
$$-2.0 \cdot \tanh^2(4.0 \cdot t - 1.0 \cdot x + 2.237e - 7) + 2.0.$$

1750 This is a very interesting result. If we apply trigonometric identities, we see that:
1751

1752
$$\frac{2}{\cosh^2(x - 4t)} = 2 \operatorname{sech}^2(x - 4t) = 2(1 - \tanh^2(x - 4t)) = 2 - 2 \tanh^2(x - 4t), \quad (10)$$

1753

1754 which is very close to what SIGS returns, with relative L_2 error $\approx 6.6 \times 10^{-6}$. Despite
1755 the fact that \cosh is missing from the grammar, we can find an equivalent form very fast,
1756 wall-clock time is 36 sec.1757 C.8.2 Shallow Water Equations
17581759 The SWE equations are a hyperbolic system of PDEs which describe the flow below a
1760 pressure surface in a fluid. We use the method of manufactured solutions to construct
1761 analytical equations for the density and velocities. The solutions are coupled because of the
1762 dependence of the velocities on the density. The density $\rho(x, y, t)$ is modeled as:
1763

1764
$$\rho(x, y, t) = 1 + h \exp\left(-\frac{r}{w(1+t)}\right) \cos(k\sqrt{r} - ct) e^{-\alpha t},$$

1765

1766 where $r = (x - x_0)^2 + (y - y_0)^2$ the center of the droplet, h represents the amplitude, w is
1767 the Gaussian width, k is the wave number, c is the wave speed, and α is the decay rate. The
1768 terms x_0 and y_0 define the initial center of the wave. The velocity components $u_x(x, y, t)$
1769 and $u_y(x, y, t)$ are derived using the linear shallow waters theory as:
1770

1771
$$u_x(x, y, t) = \rho(x, y, t) \cdot \frac{x \cdot c}{H \cdot \sqrt{r}},$$

1772
$$u_y(x, y, t) = \rho(x, y, t) \cdot \frac{y \cdot c}{H \cdot \sqrt{r}},$$

1773

1774 where H is a velocity scaling factor. The governing equations for the shallow water system
1775 are:
1776

1777 Mass conservation:
$$\frac{\partial \rho}{\partial t} + \frac{\partial(\rho u_x)}{\partial x} + \frac{\partial(\rho u_y)}{\partial y} = f_\rho(x, y, t),$$

1778 x-momentum:
$$\frac{\partial(\rho u_x)}{\partial t} + \frac{\partial(\rho u_x^2 + \frac{1}{2}g\rho^2)}{\partial x} + \frac{\partial(\rho u_x u_y)}{\partial y} = f_x(x, y, t),$$

1780 y-momentum:
$$\frac{\partial(\rho u_y)}{\partial t} + \frac{\partial(\rho u_x u_y)}{\partial x} + \frac{\partial(\rho u_y^2 + \frac{1}{2}g\rho^2)}{\partial y} = f_y(x, y, t).$$

1781

The forcing terms $f_x(x, y, t)$ and $f_y(x, y, t)$ ensure the manufactured solutions remain valid over the specified domain and time interval by compensating for natural wave decay and dissipation. The parameters of these expressions ranges provide sufficient flexibility to capture diverse wave behaviors $h \sim U(0.5, 2.0)$, $w \sim U(0.3, 3.0)$, $k \sim U(0.5, 8.0)$, $c \sim U(0.3, 3.0)$, $\alpha \sim U(0.02, 0.8)$, $x_0, y_0 \sim U(-6.0, 6.0)$, $H \sim U(1.0, 5.0)$. To discover a new solution of the shallow water equations, we consider $h = 0.97$, $w = 0.88$, $k = 1.78$, $c = 1.28$, $\alpha = 0.72$, $(x_0, y_0) = (3.77, 2.34)$, $H = 4.46$, $(x, y) \in [-10, 10]$, $t \in [0, 5]$, and $r_0 = (x - 3.77)^2 + (y - 2.34)^2$. We consider periodic boundary conditions and initial conditions the function values at time $t = 0$.

We consider Ansatze for each equation:

$$\begin{aligned} \rho(x, y, t) &= f(x, y, t)g(x, y, t)h(t), \\ u(x, y, t) &= \rho(x, y, t)s_x(x, y), \\ v(x, y, t) &= \rho(x, y, t)s_y(x, y), \end{aligned}$$

where f, g, s_x, s_y atoms from the grammar. We consider ρ as part of the ansatz of u, v due to the dependency between the velocity and the density we discussed earlier. We report both the true manufactured solution and the solution found by SIGS, which has the correct structural form, in terms of the values of θ in Table 21. The errors reported are:

$$\begin{aligned} \text{Rel}L^2(\rho) &\approx 1.8731 \times 10^{-4} (0.0187\%), \\ \text{Rel}L^2(u) &\approx 2.2310 \times 10^{-4} (0.0223\%), \\ \text{Rel}L^2(v) &\approx 4.1783 \times 10^{-4} (0.0418\%). \end{aligned} \tag{11}$$

By the problem definition, we know that the initial condition has a local support around a center r_0 which means that it is not physically meaningful for different atoms to have different centers. For this reason, the optimization centers are made identical for the local optimization by choosing the center of the traveling wave which is the dominant physical feature of the problem. The overall optimization time to get the solution is 3 minutes and 23 seconds.

	A_ρ	A_{sx}	A_{sy}	c_{px}	$c_{\rho y}$	c_{sx}	c_{sy}	c_{syx}	c_{syy}	σ	decay	freq	phase
True	1.142	1.142	1.142	0.4	-0.4	0.4	-0.4	0.4	-0.4	1.1	0.6	2.6	0.7
Init	1.64	1.91	0.93	1.02	-0.49	1.02	-0.49	1.02	-0.49	1.96	0.75	1.92	0.11
SIGS	1.142	1.14	1.14	0.39	-0.4	0.40	-0.39	0.39	-0.39	1.09	0.59	2.60	0.70

Table 21: Values of the parameters θ in the solutions to the SWE system: The true manufactured solution (True), the values returned after the SIGS structural search with values of c adjusted (Init), and the final values found after the parameter optimization (SIGS).

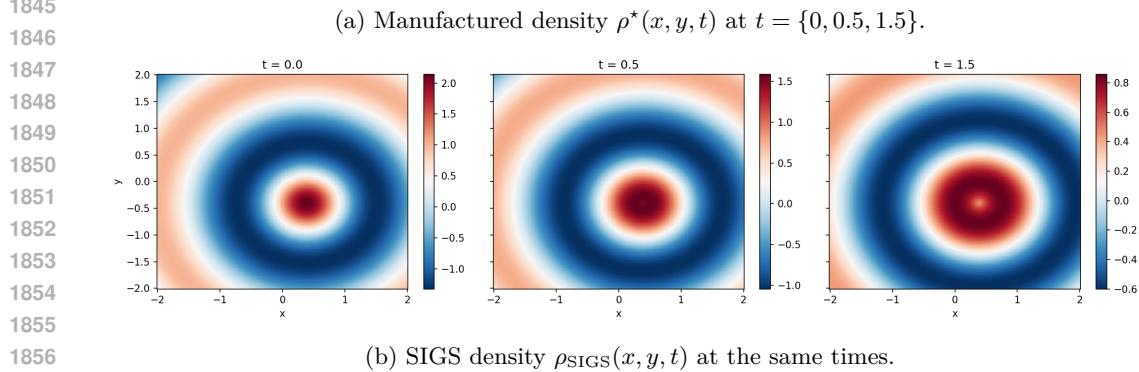
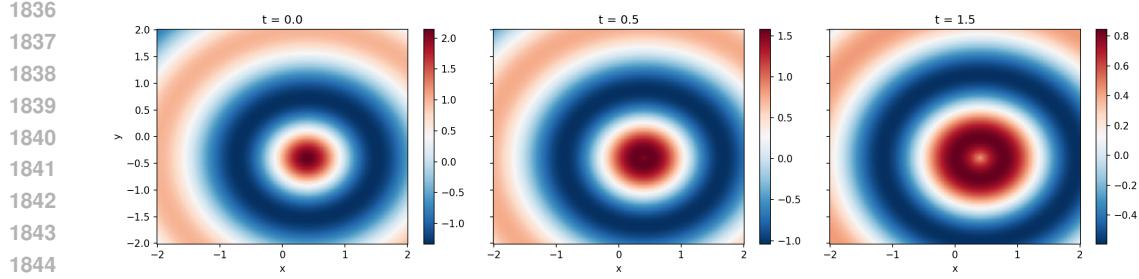


Figure 10: Shallow-water density fields. Top: manufactured solution from Eq. (SWE-MS); bottom: SIGS-refined solution. The three panels correspond to $t = 0, 0.5$, and 1.5 on the domain $(x, y) \in [-2, 2]^2$ with periodic boundary conditions.

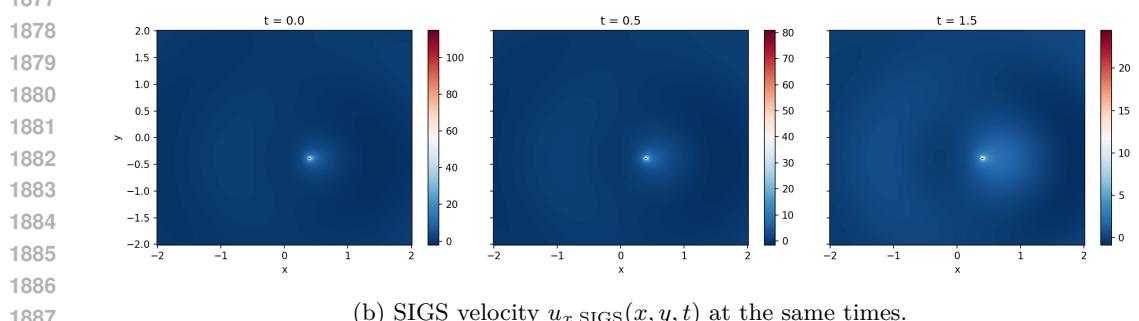
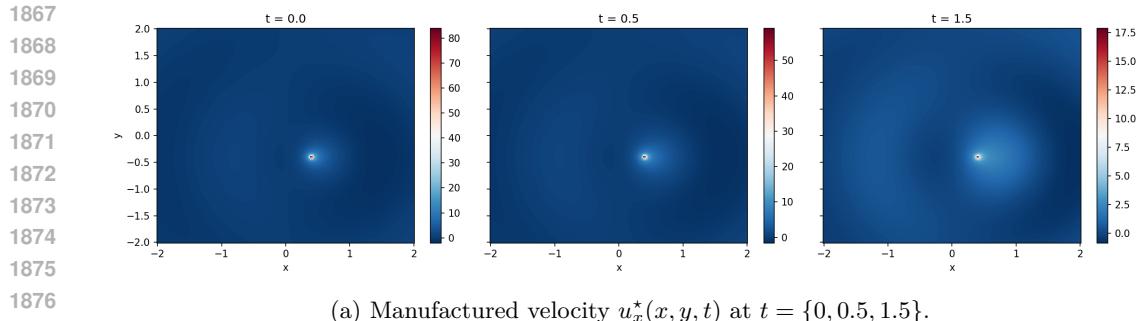


Figure 11: Shallow-water x -velocity fields. Top: manufactured solution; bottom: SIGS solution. Panels show $t = 0, 0.5$, and 1.5 .

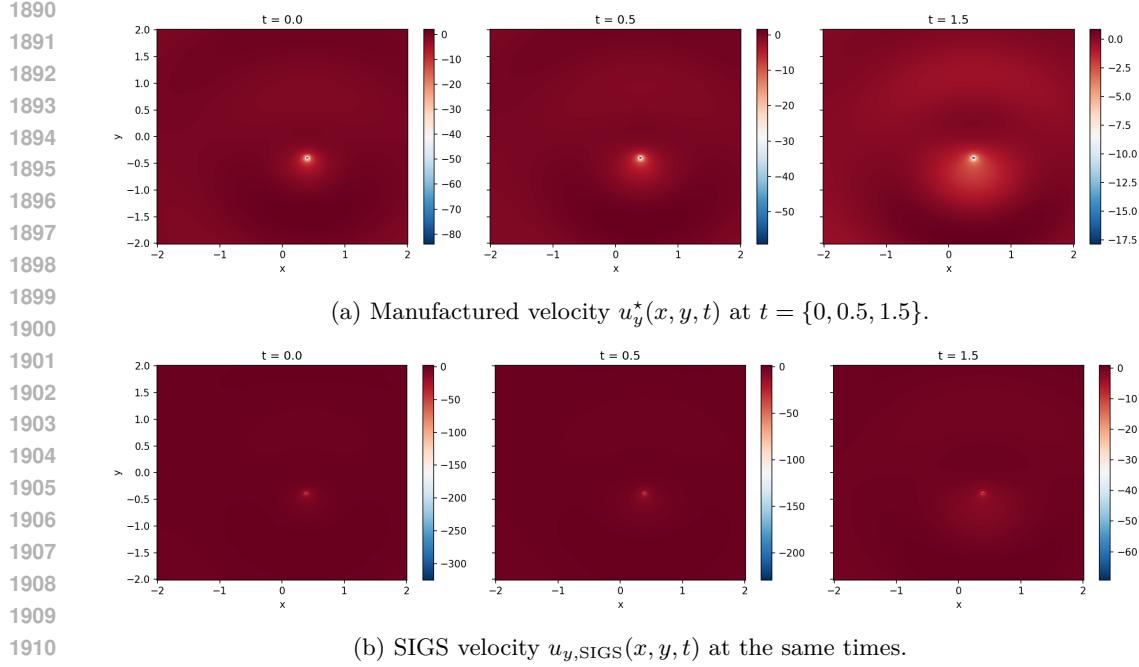


Figure 12: Shallow-water y -velocity fields. Top: manufactured solution; bottom: SIGS solution. Panels show $t = 0, 0.5$, and 1.5 .

C.8.3 Compressible Euler equations (CE)

We consider the two-dimensional compressible Euler equations (in steady state) on a periodic spatial domain. We use the method of manufactured solutions (MMS) to construct closed-form analytical baselines, and derive consistent forcing terms so that the manufactured fields satisfy the Euler system exactly. SIGS is then used to rediscover these fields by minimizing PDE residuals.

Let $\Omega = [0, 1]^2$ with $(x, y) \in \Omega$. We define the conservative state

$$U(x, y) = \begin{pmatrix} \rho(x, y) \\ \rho(x, y) u(x, y) \\ \rho(x, y) v(x, y) \\ E(x, y) \end{pmatrix}, \quad \rho > 0, \quad (u, v) \in \mathbb{R}^2,$$

where ρ is density, (u, v) are the velocity components, and E is the total energy density. In the MMS setting we use steady (time-independent) solutions, hence $\partial_t(\cdot) = 0$. The forced steady Euler system reads

$$\nabla \cdot F(U) = \begin{pmatrix} f_\rho \\ f_u \\ f_v \\ f_E \end{pmatrix} \quad \text{in } \Omega, \quad (12)$$

with flux

$$F(U) = \begin{pmatrix} \rho u & \rho v \\ \rho u^2 + p & \rho u v \\ \rho u v & \rho v^2 + p \\ u(E + p) & v(E + p) \end{pmatrix}.$$

1944 Equivalently, in component form,
 1945

$$\begin{aligned} \partial_x(\rho u) + \partial_y(\rho v) &= f_\rho, \\ \partial_x(\rho u^2 + p) + \partial_y(\rho u v) &= f_u, \\ \partial_x(\rho u v) + \partial_y(\rho v^2 + p) &= f_v, \\ \partial_x(u(E + p)) + \partial_y(v(E + p)) &= f_E. \end{aligned}$$

1951 We close the system with the ideal-gas equation of state
 1952

$$p = (\gamma - 1)(E - \frac{1}{2}\rho(u^2 + v^2)), \quad \gamma = 1.4. \quad (13)$$

1955 We impose periodic boundary conditions on all primitive variables: for $\phi \in \{\rho, u, v, p\}$,
 1956

$$\phi(0, y) = \phi(1, y), \quad \phi(x, 0) = \phi(x, 1).$$

1958 The Ansätze used by SIGS for each variable are sums of grammar atoms, with exponential
 1959 envelopes for ρ and p :
 1960

$$\begin{aligned} \rho(x, y) &= \exp\left(\sum_{i=1}^6 f_i(x, y)\right), \\ u(x, y) &= \sum_{i=1}^6 g_i(x, y), \\ v(x, y) &= \sum_{i=1}^6 h_i(x, y), \\ p(x, y) &= \exp\left(\sum_{i=1}^6 k_i(x, y)\right), \end{aligned}$$

1972 where f_i, g_i, h_i, k_i are spatial atoms generated by the grammar.
 1973

Field	Manufactured	Optimized	Rel. L^2 (%)
ρ	$\exp(-0.0887 \sin(\pi x) \sin(\pi y))$ + $0.504 \sin(\pi x) \sin(2\pi y)$ + $0.259 \sin(2\pi x) \sin(\pi y)$ + $0.140 \sin(2\pi x) \sin(2\pi y))$	$\exp(0.273 \sin(\pi x) \sin(2\pi y))$ + $0.217 \sin(2\pi x) \sin(\pi y)$ + $0.229 \sin(2\pi x) \sin(2\pi y)$ + $2.18 \times 10^{-3} \sin(2\pi y) \cos(\pi x))$	10.8
u	$\pi(-0.243 \sin(\pi x) \sin(\pi y))$ − $0.385 \sin(\pi x) \sin(2\pi y)$ − $0.494 \sin(2\pi x) \sin(\pi y)$ + $0.518 \sin(2\pi x) \sin(2\pi y))$	$-0.929 \sin(\pi x) \sin(\pi y)$ − $1 \sin(\pi x) \sin(2\pi y)$ + $0.0518 \sin(\pi x) \cos(3\pi y)$ − $1.50 \sin(2\pi x) \sin(\pi y)$ + $1.56 \sin(2\pi x) \sin(2\pi y)$	9.93
v	$\pi(0.0715 \sin(\pi x) \sin(\pi y))$ + $0.233 \sin(\pi x) \sin(2\pi y)$ − $0.536 \sin(2\pi x) \sin(\pi y)$ + $0.665 \sin(2\pi x) \sin(2\pi y))$	$0.718 \sin(\pi x) \sin(2\pi y)$ + $2.05 \sin(2\pi x) \sin(2\pi y)$ − $0.307 \sin(4\pi x) \sin(\pi y)$ − $1.22 \sin(\pi y) \cos(\pi x)$ + $1 \sin(\pi y) \cos(3\pi x)$	9.84
p	$\exp(0.235 \sin(\pi x) \sin(\pi y))$ − $0.322 \sin(\pi x) \sin(2\pi y)$ − $0.356 \sin(2\pi x) \sin(\pi y)$ − $0.448 \sin(2\pi x) \sin(2\pi y))$	$\exp(-0.389 \sin(\pi x) \sin(2\pi y))$ − $0.354 \sin(2\pi x) \sin(\pi y)$ − $0.410 \sin(2\pi x) \sin(2\pi y))$	12.1

1991 Table 22: Manufactured vs. optimized fields with coefficients truncated to three significant
 1992 figures.
 1993

1994 C.8.4 Results from classical methods 1995

1996 We assess the difficulty of the PDEs considered in this manuscript by trying to solve them
 1997 using different automated procedures, such as the state-of-the-art Computer Algebra System
 Mathematica Wolfram Research, Inc. (2024). Since manual solving might also lead to success

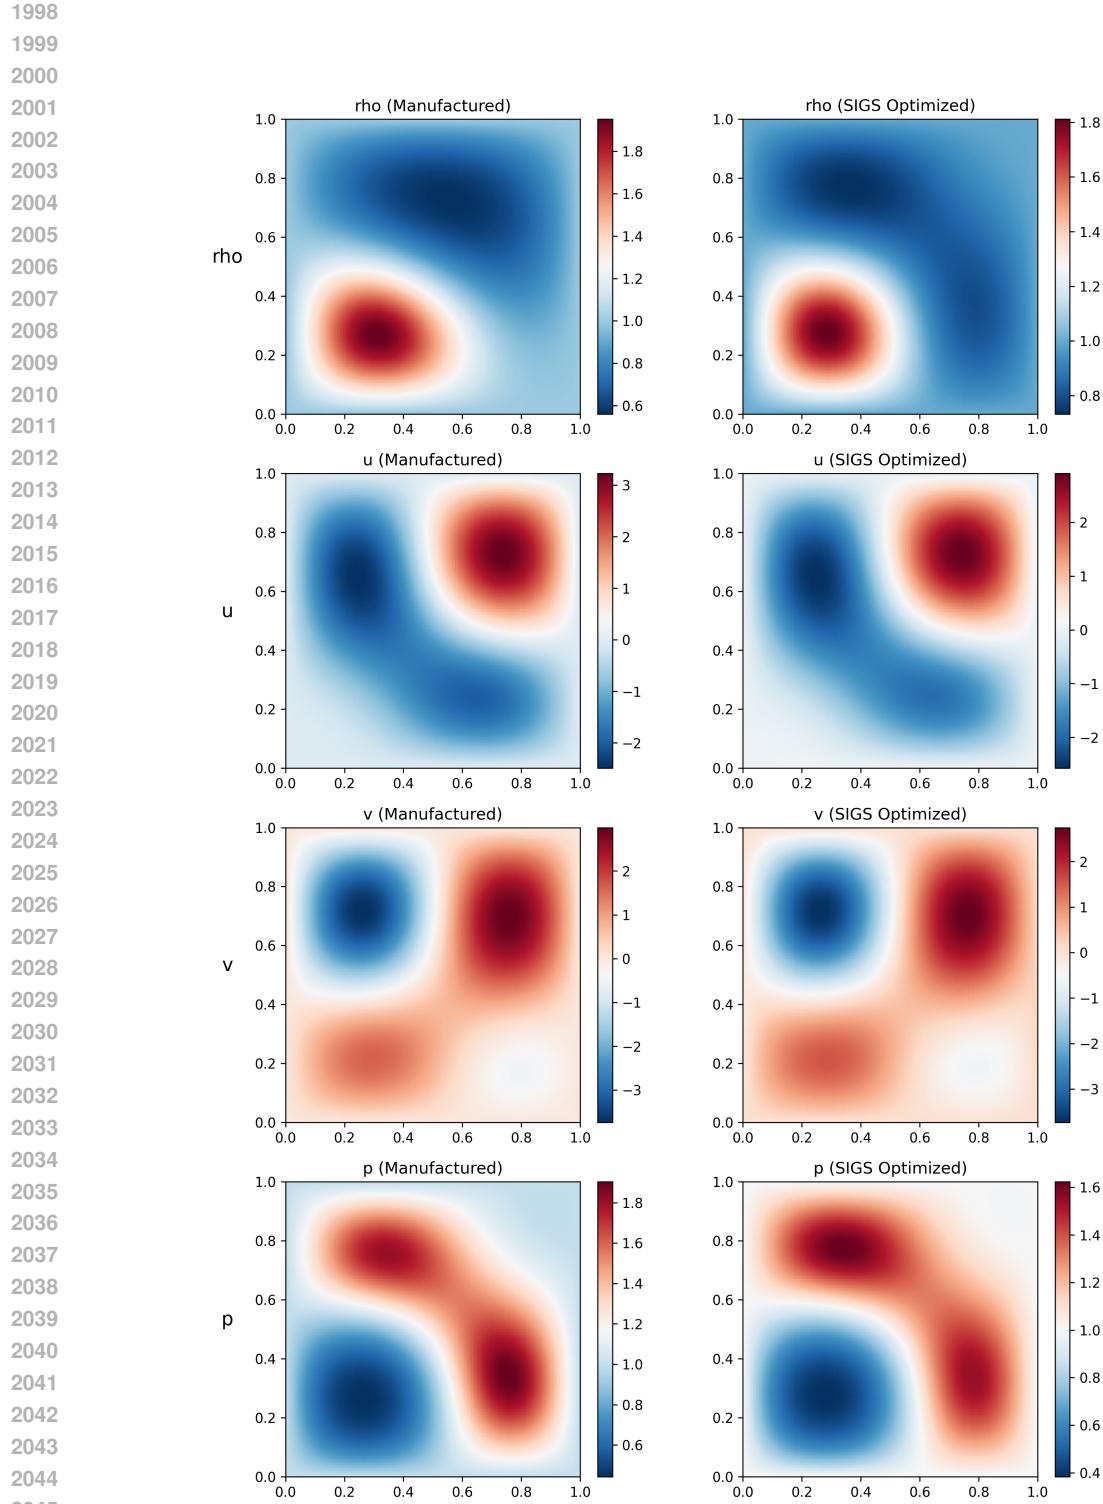


Figure 13: Compressible Euler manufactured vs. SIGS-optimized fields. Each row corresponds to one state variable (ρ, u, v, p), and columns show the manufactured reference (left) and the SIGS-refined solution (right) evaluated on the same grid. SIGS recovers the spatial structure and magnitude of all four coupled fields simultaneously, yielding a near-indistinguishable match to the manufactured solution.

in different cases, we ask ChatGPT 5.1 (Extended Reasoning) to solve the problems as a proxy for human ingenuity. We use the DSolveValue function of Mathematica on the same suite of problems that we tested SIGS on. We report the results in Table 23. The following solutions were found:

Mathematica: Dirichlet heat equation. Mathematica returns the standard Fourier–sine series solution of the Dirichlet heat problem:

$$u(x, t) = \sum_{n=1}^{\infty} b_n \exp\left(-D\left(\frac{n\pi}{L}\right)^2 t\right) \sin\left(\frac{n\pi x}{L}\right), \quad (14)$$

where the coefficients b_n are the sine–series coefficients of the initial condition (computed symbolically or numerically by DSolveValue):

$$b_n = \frac{2}{L} \int_0^L u_0(\xi) \sin\left(\frac{n\pi\xi}{L}\right) d\xi. \quad (15)$$

The infinite series solution found is written as:

$$\begin{aligned} u(x, t) = & \sum_{K=1}^{\infty} \frac{e^{-3.52tK^2} \sin(2.24K)}{1.69 \times 10^{63}K^6 - 5.94 \times 10^{64}K^4 + 4.40 \times 10^{65}K^2 - 3.82 \times 10^{65}} \\ & + \frac{(-1.28 \times 10^{64}K^4 + 2.57 \times 10^{64}K^2 - 8.38 \times 10^{65}) \sin(3.14K)}{1.69 \times 10^{63}K^6 - 5.94 \times 10^{64}K^4 + 4.40 \times 10^{65}K^2 - 3.82 \times 10^{65}} \\ & + \frac{K(-1.68 \times 10^{49}K^4 + 2.78 \times 10^{50}K^2 - 3.62 \times 10^{50}) \cos(3.14K)}{1.69 \times 10^{63}K^6 - 5.94 \times 10^{64}K^4 + 4.40 \times 10^{65}K^2 - 3.82 \times 10^{65}}. \end{aligned} \quad (16)$$

where K the mode index n in a sine expansion in, meaning $\sin(\pi nx)$, matching Dirichlet BC in x .

Mathematica: Poisson-Gauss (2 centers). Mathematica returns an eigenfunction/Green’s function representation based on the sine basis in x and the corresponding 1D Green’s function in y :

$$u(x, y) = 2 \sum_{n=1}^{\infty} \sin(n\pi x) \int_0^1 G_n(y, \eta) \left(\int_0^1 \sin(n\pi\xi) f(\xi, \eta) d\xi \right) d\eta, \quad (17)$$

where, for each Fourier mode $n \geq 1$, $G_n(y, \eta)$ is the 1D Green’s function for the operator $\partial_{yy} - (n\pi)^2$ on $y \in (0, 1)$ with homogeneous Dirichlet boundary conditions:

$$G_n(y, \eta) = \frac{1}{n\pi \sinh(n\pi)} \begin{cases} \sinh(n\pi y) \sinh(n\pi(1-\eta)), & 0 \leq y \leq \eta \leq 1, \\ \sinh(n\pi\eta) \sinh(n\pi(1-y)), & 0 \leq \eta \leq y \leq 1. \end{cases} \quad (18)$$

This is exactly the Green’s-function representation

$$u(x, y) = \iint_{(0,1)^2} G(x, y; \xi, \eta) f(\xi, \eta) d\xi d\eta, \quad (19)$$

with G expanded in the sine basis in x .

ChatGPT: Poisson-Gauss (2 centers). ChatGPT returned the following solution after a reasoning time of 8 m 16 s:

$$\begin{aligned} u_{ChatGPT}(x, y) = & -0.01 \left(\log(\sqrt{(x-0.3)^2 + (y-0.5)^2}) \right. \\ & - 0.5 \operatorname{Ei}\left(-\frac{(x-0.3)^2 + (y-0.5)^2}{0.02}\right) \\ & + \log(\sqrt{(x-0.7)^2 + (y-0.2)^2}) \\ & \left. - 0.5 \operatorname{Ei}\left(-\frac{(x-0.7)^2 + (y-0.2)^2}{0.02}\right) \right), \end{aligned} \quad (20)$$

shown in Fig. 14 and corresponding to a relative L^2 error of 1.576e+02%. Upon questioning the result, it admits: 'Exactly: the closed-form solution I gave you did not enforce the Dirichlet BC=0, u=0'.

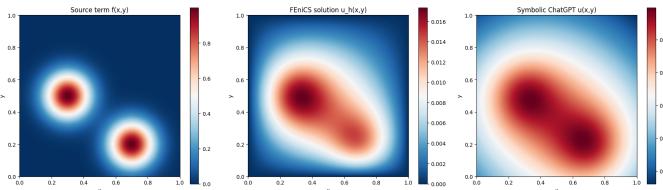


Figure 14: Performance of ChatGPT on the PG-2 problem: Initial condition, FEniCS solution, and approximation returned by ChatGPT.

ChatGPT: Burgers. ChatGPT correctly finds the analytical solution

$$u_{ChatGPT}(x, y) = 0.86 - 0.6 \tanh(30x - 25.8t - 9.9). \quad (21)$$

Here we need to clarify that ChatGPT actually "cheats" in the sense that it doesn't really solve the PDEs with manufactured solutions, but it uses the initial and boundary conditions to reverse engineer the manufactured solution.

Table 23: Success of classical solution methods on a selection of problems.

Problem	Mathematica	ChatGPT
Burgers	No	Yes
Damped Wave	No	
Diffusion equation	Yes (Infinite Series)	
Poisson–Gauss (2 centers)	Yes (Infinite Series)	Approximation
KdV equation	No	

Table 24: Performance metrics for ChatGPT for two mathematical problems.

Problem	Relative L^2 Error	Reasoning Time
Burgers	0.0%	7 m 25 s
Poisson–Gauss (2 centers)	1.576e+02%	8 m 16 s

C.9 Additions to the computational performance assessment

In Table 3, the relative L^2 error between the SIGS and FEniCS solutions was computed on the native 100×100 grid. We add the residuals $R(u)$ of both methods on different grids for the solutions by SIGS and FEniCS in Table 25. In general, FEniCS has a lower residual error, justifying to consider it as the reference in the relative error computation. For finer meshes, we see the FEniCS residual decreasing, while it stays constant for SIGS. An extended convergence study of $R(u)$ for FEniCS can be found in Table 17.

Table 25: Comparison of residuals $R(u)$ between SIGS and FEniCS for problems PG-2, PG-3 and PG-4, evaluated on two mesh resolutions.

Mesh	Model residual	PG-2	PG-3	PG-4
128×128	$R(u_{\text{FEniCS}})$	2.924×10^{-4}	3.663×10^{-4}	4.070×10^{-4}
128×128	$R(u_{\text{SIGS}})$	4.491×10^{-2}	4.476×10^{-2}	3.617×10^{-2}
256×256	$R(u_{\text{FEniCS}})$	8.048×10^{-5}	1.034×10^{-4}	1.133×10^{-4}
256×256	$R(u_{\text{SIGS}})$	4.493×10^{-2}	4.490×10^{-2}	3.618×10^{-2}

2160 In terms of runtimes, the model runs in Table 4 arrive at different levels of accuracy. In
 2161 order to facilitate the comparison, we stop the SIGS optimization at errors comparable to
 2162 the FEniCS solution in Table 2 and report this cropped SIGS runtime in Table 17.
 2163

2164 Table 26: Updated SIGS runtimes: time to discover an analytical solution at comparable
 2165 relative L^2 error $\approx 10^{-2}$ – 10^{-3} .

PDE problem	Rel. L^2 (SIGS)	Total SIGS time (s)
Burgers	4.85×10^{-3}	11.62
Diffusion	2.59×10^{-3}	14.67
Damped wave	1.44×10^{-2}	8.95

2172
 2173
 2174
 2175
 2176
 2177
 2178
 2179
 2180
 2181
 2182
 2183
 2184
 2185
 2186
 2187
 2188
 2189
 2190
 2191
 2192
 2193
 2194
 2195
 2196
 2197
 2198
 2199
 2200
 2201
 2202
 2203
 2204
 2205
 2206
 2207
 2208
 2209
 2210
 2211
 2212
 2213

2214 LLMs usage in the manuscript. The authors used LLMs to polish grammar and spelling
2215 through Overleaf tools and independent LLMs services.
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267