Genetic and agentic symbolic regression of distributed rate-and-state friction models

Luigi Romano
Department of Electrical Engineering
Linköping University
Linköping, Sweden
Email: luigi.romano@liu.se

Alessandro Lucantonio
Department of Mechanical
and Production Engineering
Aarhus University
Aarhus, Denmark
Email: a.lucantonio@mpe.au.dk

Marco Virgolin
Unlayer AI
Amsterdam, the Netherlands
Email: marco@unlayer.ai

***Abstract*—Friction plays a fundamental role in a wide range of engineering systems, yet its accurate modeling remains a challenging problem, particularly in the presence of distributed dynamics. This paper proposes symbolic regression methods, based on genetic programming and large language models, to identify analytical expressions for the friction coefficient within the distributed LuGre and Dahl formulations. The resulting models are obtained by minimizing the discrepancy between experimentally measured tire friction forces and those predicted by the corresponding analytical representations. Unlike traditional black-box approaches, the proposed framework yields interpretable, physics-consistent expressions that can be directly embedded into control-oriented formulations. Experimental results demonstrate the effectiveness of the approach and highlight its potential for advancing data-driven yet explainable friction modeling in complex dynamical systems.**

I. Introduction

Modern society relies on the dependable operation of several tribological systems, ranging from engines, turbines, and bearings to hard-drive heads and miniature electromechanical devices, where friction-related inefficiencies carry an enormous economic toll. In particular, estimates suggest that energy losses, wear, and maintenance costs arising from poorly controlled friction account for roughly 3-6% of the gross national product in several developed countries [1], representing hundreds of billions of euros lost annually through wasted energy, increased fuel consumption, and premature component failure [2], [3]. Friction is also a critical resource for control, stability, and interaction with the physical world: walking, braking, manipulating objects, and nearly every form of mechanical actuation rely on frictional interactions [1], [4]. In many applications, such as automotive safety, robotic grasping, or haptic interfaces, performance improves not by eliminating friction, but by sapiently modulating it [5]–[10]. Enhancing our ability to predict, reduce, and control friction is therefore not merely a scientific endeavor but a pressing economic, environmental, and societal necessity [1].

Accurate and realistic friction models are essential for enhancing the performance of estimators and controllers aimed at minimizing friction losses and optimally regulating frictional systems. In most cases, friction is described through analytical, human-interpretable laws expressing the friction coefficient as a function of the relative velocity between bodies in sliding or rolling contact. Whilst such physically motivated models provide valuable insight into the underlying mechanisms, they often lack accuracy due to the intrinsic complexity of friction, which spans multiple scales, from macroscopic interactions to microscopic and atomic phenomena. Conversely, empirical models are typically more accurate but difficult to interpret, as their parameters lack clear physical meaning. In this context, hybrid approaches that combine physically grounded descriptions with data-driven augmentation offer a promising alternative, reconciling interpretability with improved agreement with experimental observations.

Building on this perspective, the present work proposes to integrate distributed dynamic friction models with symbolic regression in order to derive friction laws that are physically and thermodynamically consistent, interpretable, and suitable for control design. Furthermore, two approaches are considered: *genetic programming* (GP) – suited for inexpensive goodness of fit calculations – and agents based on *large language models* (LLMs) – suited for expensive goodness of fit calculations. Sections I-A and I-B review rate-and-state friction models and recent advances in AI-augmented descriptions.

Data and code to reproduce this paper is available at: https://github.com/Unlayer-AI/friction-modeling-symreg2026.

A. Rate-and-state distributed friction models

Static friction laws typically describe the friction force as an instantaneous function of the relative velocity, thereby neglecting the internal dynamics of the contact interface. Whilst this approach enables simple and computationally efficient representations, it fails to reproduce several experimentally observed phenomena, including memory effects, rate dependence, and transient behaviors during the onset of motion. As a result, static models are inadequate when an accurate description of transient regimes is essential [11].

Conversely, dynamic friction models – formulated in terms of *ordinary differential equations* (ODEs) – address these limitations by introducing internal state variables that encode the evolution of the contact interface. In particular, in rate-and-state formulations, friction is no longer solely determined by the instantaneous relative velocity, but also by additional state

variables that capture, for instance, the average deformation of surface asperities or the history of contact conditions. This structure enables the description of key features such as pre-sliding displacement, hysteresis loops, and velocity-dependent friction transitions, whilst retaining a level of interpretability that is compatible with physical reasoning [11]–[13].

A further refinement is achieved by considering distributed formulations, in which the internal state is described by spatially distributed variables rather than lumped quantities. This approach is particularly relevant in applications such as tire–road or wheel–rail interactions, where the contact patch exhibits non-uniform stress and deformation fields [12], [13]. Distributed friction models naturally arise from the coupling of macroscopic system dynamics with the local behavior of the contact interface, leading to infinite-dimensional descriptions that take the form of *partial differential equations* (PDEs).

Within this class of models, the FrBD framework [14], [15] provides a systematic way to derive distributed extensions of classical rate-and-state formulations, including LuGre and Dahl. By embedding friction dynamics into a continuum description of the contact interface, the framework captures the spatial variability of internal states and allows for a consistent treatment of rolling and sliding conditions. This results in models that retain the physical interpretability of lumped formulations whilst offering increased fidelity in scenarios where spatial effects play a dominant role. Additionally, the FrBD models are particularly well-suited for control-oriented applications, as they automatically ensure dissipation via Lyapunov-like inequalities.

B. Related work on data-driven discovery of friction models

Multiple research works exist on discovering or improving models of ordinary and partial differential equations from data. Symbolic regression based on genetic programming has proven capable of retrieving physical laws from data [17], but, to the best of the authors' knowledge, it has not been applied to friction modeling specifically. Instead, neural approaches such as *physics-informed neural networks* have already shown success in friction modeling: Kolmogorov–Arnold networks were employed in [20] to learn static friction characteristics, whilst [21] proposed a PINN-based approach for rate-and-state dynamic friction models, including the LuGre formulation.

Regarding LLM-based approaches, recent works have seen their adoption to SR – again, not specifically to friction modeling. Such LLM-based approaches range from pre-training transformers from scratch on synthetic data for "translating" tokens that represent the input dataset into tokens representing mathematical expressions [22], [23], to leveraging general LLMs that are iteratively prompted to improve a program or symbolic formula [18]. For example, *AlphaEvolve* and *ShinkaEvolve* orchestrate LLM calls within a population-based, evolutionary framework [16], [19].

This paper is different from the cited works in several ways. From the modeling perspective, it considers a distributed parameter representation of friction dynamics, instead of relying on static nonlinearities or ODEs. From the approach perspective, both classic genetic programming and the use of an LLM-based agent is considered. The latter is different from the aforementioned LLM-based approaches, which orchestrate LLM calls in a pre-configured manner. Instead, similar to Karpathy's *autoresearch* idea [24], this paper relies on a modern agent to plan its actions and attempts autonomously, using only a minimal set of guidelines and starting code.

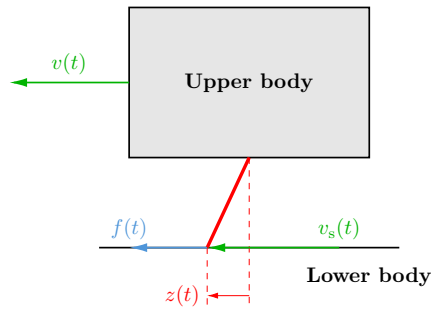

Fig. 1. A schematic representation of the friction models.

II. Description of the friction model

In the following, Section II-A details the governing equations of the distributed LuGre and Dahl models, whereas a possible parametrization procedure – explicitly leveraging the model decomposition into a rheological and frictional component – is elucidated in Section II-B.

A. Distributed LuGre and Dahl friction models

A schematic representation of the LuGre and Dhal friction models is illustrated in Figure 1, where a deformable body travels with rigid velocity $v(t) \in \mathbb{R}$ with respect to a rigid substrate. The deformation of the upper body varies locally within the contact area, and is undergone by bristle-like elements $z(t) \in \mathbb{R}$. The sliding velocity of the bristle tips is given by $\mathbb{R} \ni v_s(t) = v(t) + \dot{z}(t)$.

Considering a distributed configuration in steady-state conditions, where $z(t) = z(\xi, t)$, the total derivative of the bristle deflection becomes

$$\dot{z}(\xi, t) = \frac{\mathrm{d}z(\xi, t)}{\mathrm{d}t} = V \frac{\partial z(\xi)}{\partial \xi}, \tag{1}$$

and the governing equations of the models may be deduced as

$$V \frac{\partial z(\xi)}{\partial \xi} = -\frac{\bar{k}_0(F_z)|v|_\varepsilon}{\theta(F_z)\bar{\mu}(v)} z(\xi) - v, \quad \xi \in (0, L), \tag{2a}$$

$$z(0) = 0, \tag{2b}$$

where $z(\xi) \in \mathbb{R}$ denotes the spatially-varying bristle deflection, $v \in \mathbb{R}$ the relative velocity between the contacting bodies, V the transport velocity, $\bar{k}_0(F_z) \in \mathbb{R}_{>0}$ the normalized bristle stiffness, $\bar{\mu} : \mathbb{R} \mapsto [\bar{\mu}_{\text{min}}, \infty)$, with $\bar{\mu}_{\text{min}} \in \mathbb{R}_{>0}$ the normalized friction coefficient, $\theta(F_z) \in \mathbb{R}_{>0}$ a scaling factor that accounts for normal load variations, $L = L(F_z)$ the contact length, and $F_z \in \mathbb{R}_{>0}$ the applied normal load. Finally, the function

$|\cdot|_\varepsilon \in C^0(\mathbb{R};\mathbb{R}_{\geq 0})$, with $\varepsilon \in \mathbb{R}_{\geq 0}$, is a regularization of the absolute value $|\cdot|$ for $\varepsilon \in \mathbb{R}_{>0}$, converging uniformly to $|\cdot|$ in $C^0(\mathbb{R};\mathbb{R}_{\geq 0})$ for $\varepsilon \to 0$ (e.g., $|v|_\varepsilon = \sqrt{v^2+\varepsilon}$), and with $|\cdot|_\varepsilon \in C^1(\mathbb{R};\mathbb{R}_{\geq 0})$ for $\varepsilon \in \mathbb{R}_{>0}$ [14], [15].

The solution to (2) is given by

$$z(\xi) = -\left[1 - \exp\left(-\frac{\bar{k}_0 |v|_\varepsilon}{V\theta\bar{\mu}(v)}\xi\right)\right]\frac{\theta\bar{\mu}(v)}{\bar{k}_0}\operatorname{sgn}_\varepsilon(v), \quad \xi \in [0, L], \tag{3}$$

where $\operatorname{sgn}_\varepsilon(v) \triangleq \frac{v}{|v|_\varepsilon}$ and the functional dependency on F_z has been omitted for compactness. Starting with (3), and neglecting internal damping effects, the bristle force, which must oppose friction, is calculated as

$$f(\xi) = \bar{k}_0 z(\xi). \tag{4}$$

Moreover, assuming a constant pressure distribution inside the contact area, the tangential force may be computed as

$$F = \int_0^L \frac{F_z}{L(F_z)} f(\xi)\,\mathrm{d}\xi. \tag{5}$$

In particular, the following functional dependencies are usually postulated:

$$\bar{L}(F_z) = a_1 + a_2\sqrt{F_z}, \tag{6a}$$
$$\bar{k}_0(F_z) = a_3 - a_4 F_z, \tag{6b}$$
$$\theta(F_z) = a_5 - a_6 F_z, \tag{6c}$$

with $a_1, a_3, a_5 \in \mathbb{R}_{>0}$ and $a_2, a_4, a_6 \in \mathbb{R}_{\geq 0}$.

Finally, to completely characterize the friction models, it is necessary to specify an analytical expression for the normalized coefficient $\bar{\mu}(v)$. Common choices are the Stribeck function

$$\bar{\mu}(v) = \bar{\mu}_\mathrm{d} + (\bar{\mu}_\mathrm{s} - \bar{\mu}_\mathrm{d})\exp\left(-\left(\frac{|v|}{v_\mathrm{S}}\right)^{\delta_\mathrm{S}}\right), \tag{7}$$

and the simpler Savkoor law

$$\bar{\mu}(v) = \frac{\bar{\mu}_\mathrm{s}}{1 + \alpha_\mathrm{S}|v|^{\beta_\mathrm{S}}}, \tag{8}$$

where $\bar{\mu}_\mathrm{d} \in \mathbb{R}_{>0}$ is the normalized dynamic friction coefficient, $\bar{\mu}_\mathrm{s} > \mu_\mathrm{d}$ is the normalized static friction coefficient, $v_\mathrm{S} \in \mathbb{R}_{>0}$ the Stribeck velocity, $\delta_\mathrm{S} \in \mathbb{R}_{\geq 0}$ the Stribeck exponent, $\alpha_\mathrm{S} \in \mathbb{R}_{\geq 0}$ the Savkoor coefficient, and $\beta_\mathrm{S} \in \mathbb{R}_{\geq 0}$ the Savkoor exponent. In both (7) and (8), the absolute value may be replaced by the smooth function $|\cdot|_\varepsilon$ to ensure global differentiability.

B. Model parametrization

The parametrization of the model described by (2)–(7) can be simplified by noting that the parameters $L(F_z)$, $\bar{k}_0(F_z)$, and $\theta(F_z)$ are primarily rheological and, as such, largely independent of frictional effects. For a given normal force F_z, these coefficients can therefore be identified by considering operating conditions in which frictional contributions are negligible, as it typically happens for low relative velocities v.

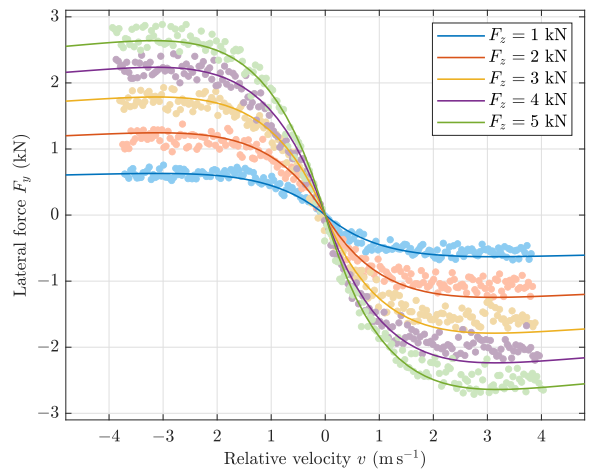

Fig. 2. Comparison between the optimized nominal friction model with Stribeck friction law (7) (solid lines) and real tire data (cloud points) for different values of the normal load F_z.

Once these coefficients have been determined, the parameters governing the normalized friction coefficient can be estimated separately, using data collected in sliding-dominated regimes at higher values of v.

This enables a two-stage parametrization procedure: first, the rheological properties of the model are identified independently of sliding effects; second, the friction coefficient is estimated based on the characteristics of the contact pair and the tribological properties of the interface.

This approach was employed to identify the rheological and tribological parameters entering the normalized friction coefficient $\bar{\mu}(v)$, as defined in (7), using tire data collected at different levels of vertical load, namely $F_z = 1$, 2, 3, 4, and 5 kN. The force predicted by the distributed model described in Section II-A is compared against the corresponding experimental measurements in Figure 2. The values of the parameters used to produce Figure 2, and obtained using a standard genetic particle swarm optimization algorithm, are $a_1 = 0.067$, $a_2 = 1 \cdot 10^{-4}$, $a_3 = 360.49$, $a_4 = 0.023$, $a_5 = 0.646$, $a_6 = 3.07 \cdot 10^{-5}$, $\bar{\mu}_\mathrm{s} = 1.72$, $v_\mathrm{S} = 3.7$, and $\delta_\mathrm{S} = 1.37$.

It is worth emphasizing that the expression for the normalized friction coefficient, postulated according to (7), must be specified a priori. In contrast, it is of interest to infer such an expression directly from data as the frictional behavior evolves. This is motivated by the fact that the friction coefficient may vary over time as the tribological properties of the contact interface change due to, for instance, lubrication effects or temperature variations. The objective is therefore to learn a symbolic representation of $\bar{\mu}(v)$ directly from data, without prescribing a specific functional form in advance.

III. Discovery of Symbolic Friction Models

To identify a symbolic expression for the normalized friction coefficient $\bar{\mu}(v)$, two approaches are considered: a GP-based

method and an LLM agent-based method. The experimental data are partitioned into five subsets corresponding to distinct values of the normal load F_z. The subsets for $F_z = 1$, 2, and 3 kN are used for training, whilst those for $F_z = 4$ kN and 5 kN are reserved for validation and holdout testing, respectively. This split is primarily motivated by the higher noise levels observed in the data associated with larger vertical loads.

A. GP-based discovery

Following a GP-based approach to symbolic regression, a model for $\bar{\mu}(v)$ was discovered by evolving a population of candidate expressions represented as trees to maximize their *fitness*. The Python-based evolutionary framework `Flex` [25] was used. Each tree incorporated a random number of numerical constants and combined operators taken from the set $\{+, -, *, /, \sin, \exp, \mathrm{pow}\}$. The search space was explored using a distributed island model to maintain genetic diversity and prevent premature convergence. The population was partitioned into 4 independent islands, each containing 800 individuals, and evolved over 50 generations to lead to a similar time compared to LLM agent runs (approx. 7 minutes).

The fitness f of an individual was evaluated using a cost function balancing accuracy with simplicity (parsimony pressure): $f = -(\mathrm{MSE} + \lambda N)$, where MSE is the mean square error between the tangential force computed from (5)-(6) and the corresponding training data labels, N is the number of nodes of the expression tree and $\lambda = 10^{-5}$ serves as the penalty coefficient. The numerical parameters in the expressions were tuned via Particle Swarm Optimization (20 generations, swarm size 50) to minimize the MSE. To ensure the mathematical and physical validity of the results, the algorithm discarded singular expressions and those containing nested sine functions.

For each generation, parent individuals to be subject to either one-point cross-over or mutation (probabilities: 20%/80%) were selected via stochastic universal sampling to ensure a zero-bias selection process whilst maintaining high pressure toward fitter candidates. The mutation operator was randomly selected from three distinct types: one-point, random subtree, or shrink operations. Survival selection was implemented by truncating the combined parents and offspring, keeping the highest-performing individuals.

Random search (RS) was also implemented, specifically by continuous sampling of random individuals, with the same runtime budget of GP. To account for stochasticity, GP and RS runs were repeated 5 times.

B. Agent-based discovery

An LLM agent, specifically the GitHub Copilot harness with the `GPT-5.4` model, was instructed to discover an analytical expression for the function $\bar{\mu}(\cdot)$, subject to the constraint $\bar{\mu} \in C^1(\mathbb{R}; \mathbb{R}_{>0})$. The agent was provided with the Stribeck law in (7) (with the parameter values reported in Section II-B) as a starting point to improve upon, and was configured, with an `AGENTS.md` file, to reason about how the law could be modified, edit the respective python file (`mu_func.py`) and then run the loss calculation (`tire_model_loss.py`), for 20 iterations. The agent was instructed to consider the training and validation loss to decide what to attempt next. As in Section III-A, the loss was calculated as the MSE between the tangential force from (5)-(6) and the corresponding training data labels. A total of 5 runs was performed (median runtime of approx. 7 minutes).

C. Results and discussion

The main results of the GP- (RS-) and LLM agent-based approaches are collected in Section III-C1, where performance metrics are also reported, whereas Section III-C2 discusses qualitatively their main differences.

1) Results: The best model discovered by GP was

$$\bar{\mu}(v) = \bar{\mu}_{\mathrm{GP}}(v) = 1.02^v - \frac{1.1}{v - \exp(v)}, \tag{9}$$

whereas the formulae discovered by the RS and LLM agent read, respectively,

$$\bar{\mu}(v) = \bar{\mu}_{\mathrm{RS}}(v) = \left(\frac{4.52 - v + \exp(v)}{0.2 + \exp(v)}\right)^{0.37\exp(0.5v)} \tag{10}$$

and

$$\begin{aligned}\bar{\mu}(v) = \bar{\mu}_{\mathrm{LLM}}(v) = 1 + 0.53\exp\left(-\left(\frac{|v|_\varepsilon}{1.45}\right)^{1.2}\right) \\ + 0.2\exp\left(-\left(\frac{|v|_\varepsilon}{4.8}\right)^{2.2}\right) + \frac{0.03}{1 + \left(\frac{|v|_\varepsilon}{9}\right)^6}.\end{aligned} \tag{11}$$

The expressions are all well-behaved and free of singularities.

The training and validation losses for GP, RS, and the LLM-agent are plotted against the number of generations and attempts in Figure 3, where the medians are shown together with the minimum and maximum values. Performance metrics in terms of RMSE and R^2 are reported in Table I. The results show that GP achieves better results than the LLM agent. RS performs worse than GP but better than the LLM agent. However, when considering RS with the same budget of attempts used by the LLM agent (20 attempts, Figure 3 central panel), then outcome flips: the LLM agent performs better. Note that GP with a random population of size 20 would be equivalent to RS. When considering the R^2 values, a similar level of explanatory accuracy is reached by all methods.

Figure 4 compares the experimental data normalized with respect to the vertical load and aggregated considering all the subsets with the baseline model, using the Stribeck friction law (7), and the GP-, RS-, and LLM-augmented formulations. Interestingly, the friction characteristics predicted by the augmented models coincide almost completely with that of the baseline formulation for $v \in [-1, 1]$ m s^{-1}, where nonlinear effects induced by friction are negligible and the model response is dominated by the rheological component encoded in the parameters $L(F_z)$, $\bar{k}_0(F_z)$, and $\theta(F_z)$. On the other hand, a marked discrepancy may be observed for larger values of v, and particularly after the friction peaks, corresponding to $v \approx \pm 2.5$ m s^{-1}.

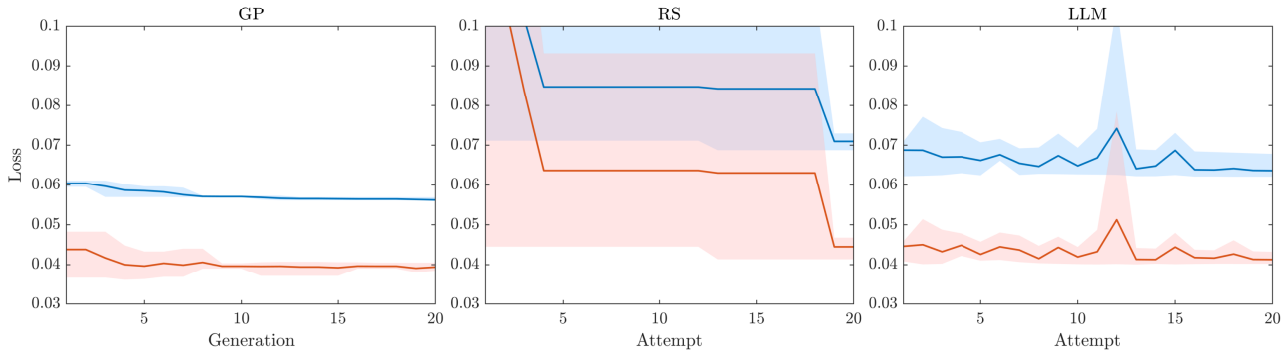

Fig. 3. Model losses (MSE) over the generations/attempts. For GP, the first 20 generations are shown (further progress is minimal), but it should be noted that a population of 800 individuals is employed. For RS, the first 20 attempts are displayed, to favor a direct comparison can be made with the 20 attempts of the LLM agent. (GP and RS display only best-so-far, LLM agent reports also worsening attempts.)

TABLE I
RMSE AND R^2 STATISTICS FOR THE MODELS FOUND BY GP, RS, AND LLM AGENT

Statistics		**GP**			**RS**			**LLM agent**		
		Train	*Valid.*	*Test*	*Train*	*Valid.*	*Test*	*Train*	*Valid.*	*Test*
RMSE	*Median*	0.0560	0.0381	0.0390	0.0578	0.0408	0.0439	0.0621	0.0408	0.0394
	Best	0.0560	0.0380	0.0386	0.0573	0.0384	0.0396	0.0621	0.0402	0.0368
R^2	*Median*	0.9883	0.9933	0.9926	0.9876	0.9924	0.9906	0.9856	0.9923	0.9924
	Best	0.9883	0.9934	0.9927	0.9878	0.9932	0.9923	0.9857	0.9926	0.9934

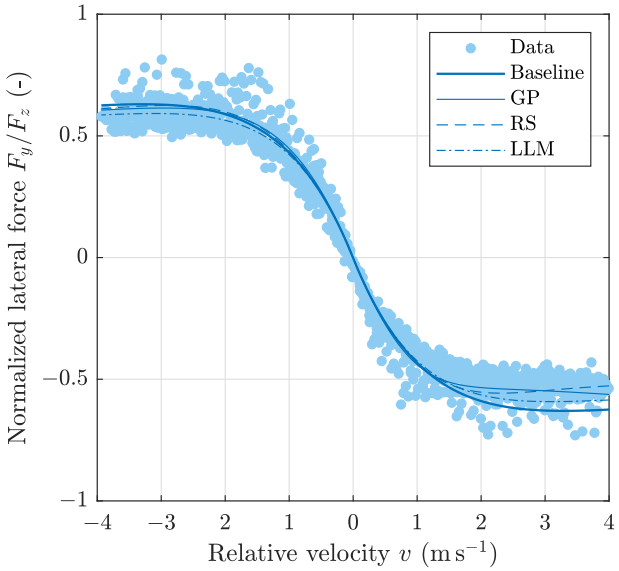

Fig. 4. Best-found SR model predictions from GP, RS, and LLM agent, versus the baseline model and experimental data.

2) Discussion: The comparison between GP (or RS) and the LLM agent symbolic discovery approaches should be interpreted in light of the fundamentally different cost structures underlying the two methods. GP operates with extremely low-cost mutations, enabling the exploration of a very large number of candidate expressions (e.g., hundreds or thousands of individuals over many generations). As a result, when the evaluation of the fitness function is relatively inexpensive, GP can efficiently exploit its large search budget and is expected to outperform more structured approaches. In contrast, the LLM agent incurs a significantly higher cost per candidate generation, as each modification of the friction law requires a nontrivial reasoning step. This limits the total number of attempts that can be performed within a given computational budget. However, this asymmetry becomes advantageous in scenarios where the fitness evaluation is computationally expensive (e.g., requiring high-fidelity simulations which take minutes or hours). In such cases, the total cost is dominated by the evaluation rather than the generation of candidates, and methods that can produce more informed proposals with fewer iterations – such as LLM-based approaches – become competitive or even preferable. Therefore, the two methods should be viewed as complementary: GP is well-suited for problems with cheap evaluations and large search budgets, whereas LLM-based approaches are more appropriate when evaluations are costly and the number of feasible iterations is severely constrained.

Other interesting differences between the two approaches emerge when comparing the expressions for the normalized friction coefficient in (9), (10), and (11). In particular, whilst (9) and (10) capture the asymmetry in the friction force for positive and negative relative velocities observed in the experimental data, (11) yields a symmetric trend with respect to the vertical axis. In fact, (11) can be interpreted as the superposition of two regularized Stribeck laws (7) and the Savkoor one (8) (despite (8) not being provided as input to the LLM agent). More generally, this behavior was representative of the functions identified by the LLM agent, which tended to be more conservative and composed of exponential-like terms,

whereas GP and RS were capable of producing expressions that deviated significantly from classical Stribeck-type friction laws.

Finally, some qualitative examples of descriptions of attempts made by the LLM agent in the run producing its best model are reported below:

- Attempt 3: *Exponential core with rational shoulder and ultra-broad tail*,
- Attempt 7: *Three-scale stretched exponential decay emphasizing fast core and broad shoulder*,
- Attempt 11: *Broader core and narrower shoulder around the early winning family*,
- Attempt 14: *Sharper core with very broad high-order exponential shoulder*,
- Attempt 20: *Selected final candidate: steeper core with weak rational tail*.

IV. Conclusion

This paper proposed a symbolic regression framework to identify analytical expressions for the normalized friction coefficient within distributed rate-and-state friction models. By combining the FrBD formulation with data-driven discovery, the approach enabled the derivation of interpretable and physically consistent friction laws that accurately reproduce experimental observations.

Complementary strategies were investigated, namely genetic programming (GP), a random search baseline (RS), and a large language model (LLM) agent. All methods successfully identified compact and well-behaved expressions with comparable predictive performance, although the GP-based approach could produce far more attempts within the same runtime, and achieved slightly lower errors. Their comparison highlights a fundamental trade-off: GP efficiently explores large search spaces when fitness evaluations are inexpensive, whereas LLM-based approaches generate more informed candidates and become advantageous when evaluations are costly.

The results suggest that these paradigms are best viewed as complementary rather than competing, and could be effectively combined in hybrid frameworks. Future work will focus on extending the approach to more general distributed friction models and on incorporating structural properties, such as stability and passivity, directly into the symbolic discovery process.

Acknowledgment

LR acknowledges financial support from the project FASTEST (Reg. no. 2023-06511), funded by the Swedish Research Council. The work of AL is supported by the European Union (European Research Council (ERC), ALPS, 101039481). Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the ERC Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

[1] Persson BNJ. Sliding friction: physical principles and applications. 1st ed., Springer Science & Business Media, 2013.

[2] Holmberg K, Andersson P, Erdemir A. Global energy consumption due to friction in passenger cars. Tribology International 47:221-234, 2012.

[3] Holmberg K, Erdemir A. Influence of tribology on global energy consumption, costs and emissions. Friction 5(3):263-284, 2017.

[4] Marques F, Flores P, Pimenta Claro JC et al. A survey and comparison of several friction force models for dynamic analysis of multibody mechanical systems. Nonlinear Dyn 86:1407-1443, 2016.

[5] Yao LLB, Wang Q, Chen Z. Adaptive robust control of linear motors with dynamic friction compensation using modified LuGre model. Automatica 45(12):2890-2896, 2009.

[6] Yao J, Deng W, Jiao Z. Adaptive Control of Hydraulic Actuators With LuGre Model-Based Friction Compensation. IEEE Trans. Industr. Electro. 62(10):6469-6477, 2015.

[7] Zeng H, Sepehri N. Tracking Control of Hydraulic Actuators Using a LuGre Friction Model Compensation. ASME. J. Dyn. Sys. Meas. Control 130(1):014502, 2008.

[8] Freidovich L, Robertsson A, Shiriaev A, et al. LuGre-Model-Based Friction Compensation. IEEE Trans. Contr. Syst. Tech. 18(1):194-200, 2010.

[9] Bittencourt AC, Wernholt E, Sander-Tavallaey S, Brogårdh T. An extended friction model to capture load and temperature effects in robot joints. 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 6161-6167, Taipei, Taiwan, 2010.

[10] Bittencourt AC, Gunnarsson S. Static Friction in a Robot Joint—Modeling and Identification of Load and Temperature Effects. ASME. J. Dyn Sys Meas Control 134(5): 051013, 2012.

[11] Canudas-de-Wit C, Olsson H, Åström KJ, Lischinsky P. A new model for control of systems with friction. IEEE Trans Autom Control 40(3):419–425 (1995).

[12] Canudas-de-Wit C, Tsiotras P, Velenis E, et al. Dynamic friction models for road/tire longitudinal interaction. Veh Syst Dyn 39(3):189–226 (2003).

[13] Deur J, Asgari J, Hrovat D. A 3D brush-type dynamic tire friction model. Veh Syst Dyn 42(3):133–173 (2004).

[14] Romano L. Two-dimensional FrBD friction models for rolling contact. Nonlinear Dynamics 114, 444 (2026).

[15] Romano L. Two-dimensional FrBD friction models for rolling contact: extension to linear viscoelasticity. Tribology International 220, 111953 (2026).

[16] Lange RT, Imajuku Y, Cetin E. ShinkaEvolve: Towards open-ended and sample-efficient program evolution. arXiv preprint arXiv:2509.19349, 2025.

[17] Schmidt M, Lipson H. Distilling free-form natural laws from experimental data. Science 324.5923: 81-85, 2009.

[18] Merler M, Haitsiukevich K, Dainese N, Marttinen P. In-context symbolic regression: Leveraging large language models for function discovery. Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop), 2024.

[19] Novikov A, Vũ N, Eisenberger M, Dupont E, Huang PS, Wagner AZ, Shirobokov S, Kozlovskii B, Ruiz FJ, Mehrabian A, Kumar MP. Alphaevolve: A coding agent for scientific and algorithmic discovery. arXiv preprint arXiv:2506.13131, 2025.

[20] Wang Y, Rabczuk T, Liu Y. Physics-informed Machine Learning for Static Friction Modeling in Robotic Manipulators Based on Kolmogorov-Arnold Networks. arXiv:2511.10079v1, 2025.

[21] Ozmen A, Hespanha JP, Byl K. Learning Transferable Friction Models and LuGre Identification via Physics Informed Neural Networks. arXiv:2504.12441v1, 2025.

[22] Kamienny PA, d'Ascoli S, Lample G, Charton F. End-to-end symbolic regression with transformers. Advances in Neural Information Processing Systems, 2022.

[23] Kamienny PA, Lample G, Lamprier S, Virgolin M. Deep generative symbolic regression with Monte-Carlo-tree-search. International Conference on Machine Learning, 2023

[24] Karpathy A. Autoresearch. Github repository. https://github.com/karpathy/autoresearch

[25] Manti S, Lucantonio A. Discovering interpretable physical model using symbolic regression and discrete exterior calculus. Machine Learning: Science and Technology 5, 015005, 2024.