



LIPS - Learning Industrial Physical Simulation benchmark suite

Milad Leyli-Abadi, David Danan, Mouadh Yagoubi, Seif-Eddine Attoui,
Antoine Marot, J Picault, B Donnot, P Dimitrov, C Etienam, A Farjallah

► To cite this version:

Milad Leyli-Abadi, David Danan, Mouadh Yagoubi, Seif-Eddine Attoui, Antoine Marot, et al.. LIPS - Learning Industrial Physical Simulation benchmark suite. NeurIPS - Data & Benchmark Track, Nov 2022, New Orleans, United States. hal-03780816

HAL Id: hal-03780816

<https://hal.science/hal-03780816>

Submitted on 19 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

LIPS - Learning Industrial Physical Simulation benchmark suite

M. Leyli-Abadi, D. Danan, M. Yagoubi, S. Attoui
IRT SystemX, Palaiseau, France

A. Marot, J. Picault, B. Donnot
RTE France, Paris, France

P. Dimitrov, C. Etienam, A. Farjallah
Nvidia

Abstract

Physical simulations are at the core of many critical industrial systems. However, today's physical simulators have some limitations such as computation time, dealing with missing or uncertain data, or even non-convergence for some feasible cases. Recently, the use of data-driven approaches to learn complex physical simulations has been considered as a promising approach to address those issues. However, this comes often at the cost of some accuracy which may hinder the industrial use. To drive this new research topic towards a better real-world applicability, we propose a new benchmark suite "Learning Industrial Physical Simulations"(LIPS) to meet the need of developing efficient, industrial application-oriented, augmented simulators. To define how to assess such benchmark performance, we propose a set of four generic categories of criteria. The proposed benchmark suite is a modular and configurable framework that can deal with different physical problems. To demonstrate this ability, we propose in this paper to investigate two distinct use-cases with different physical simulations, namely: the power grid and the pneumatic. For each use case, several benchmarks are described and assessed with existing models. None of the models perform well under all expected criteria, inviting the community to develop new industry-applicable solutions and possibly showcase their performance publicly upon online LIPS instance on Codabench.

1 Introduction

Physical simulations constitute today a key enabler for real-world complex industrial systems (power grid management, rail infrastructure, aeronautics, pneumatic, gas production plants, thermal comfort, etc.), and are used at several critical stages of the system life-cycle (system design, solutions exploration, system V&V, etc) to enhance decision making. The main drawback of using numerical simulations in general, is their high computational cost to reach satisfactory solutions. It can become prohibitive for complex systems requiring large number of simulations. To tackle this issue, several techniques have been explored in the literature to design simplified physical models [1, 2, 3, 4], dimension reduction, or considering simplified assumptions to linearize the problem. In recent years, there has been a growing interest in using machine learning techniques to solve physical problems [5] for which conventional modeling approaches are very expensive to compute. The main goal is to accelerate the computation time while maintaining an acceptable accuracy of simulation predictions under some specified tasks. Going even further to reach the best trade-off, Deep Neural Networks (DNN) have recently led to promising results in various domains (see e.g., [6, 7, 8, 9, 10]), allowing an important speed-up of simulations by substituting some computational bricks with data-driven numerical models.

These approaches emulate often existing simulators by learning from them in a supervised fashion and are generally used to complement them. Other approaches also aim at developing new kinds of differential solvers [11, 12, 13] in an unsupervised manner, and could possibly directly fit into existing simulator core. They often fall in the class of Physics-informed machine learning [14], where the learning is performed through a residual loss function and then physical constraints are verified on the learned model to validate the obtained solution. They could lead to stronger convergence and generalization than emulators. As automated learning of complex physical simulations is still considered as a new field of research, there exists a lack of common benchmarking pipeline, starting from available datasets, across various applications and finally common evaluation criteria as reviewed in section 3. This may allow to rigorously compare these methods and drive further advances into real-world applications, in particular when considering industrial use-cases.

In this paper, we propose a new benchmark suite "Learning Industrial Physical Simulations (LIPS)" to facilitate the use and the assessment of augmented physical systems, when applied on real-world applications. Depending on the application scope, the set of required physical variables to be considered may be different. The trade-off between computation acceleration and accuracy, as well as the expected generalization capability, may be specific to each industrial domain and the considered application. The compliance to physical laws of the learnt simulations may also be very important to validate them and consequently increase the user trust toward these augmented simulators.

To develop the LIPS benchmark suite over several physical domains, we use a bottom-up approach by investigating two use-case described in section 2 with distinct physics: power grid and pneumatic. Those 2 industrial domains both contribute in tackling ongoing real-world challenges, such as Climate Change, by transforming our energy system through electricity decarbonization and gains in transportation energy efficiency, or improving the decision-making efficiency regarding industrial products. These 2 use cases also allow, thanks to their heterogeneity in terms of physics lying behind the modeling, a better assessment of our proposed benchmark. Preliminary ML-models to benchmark also exist in the respective literatures. Our contributions described in section 4 hence lie in:

1. defining application-oriented benchmark tasks for industry use cases as opposed to general-purpose simulation tasks;
2. proposing four categories of evaluation criteria that generalize to several physical, industrial and application domains and challenges beyond usual ML-only evaluation metrics;
3. sharing as open-source LIPS benchmarking suite framework with associated datasets;
4. opening a publicly available Codabench [15] thread providing a shared result table for user's submission and a fully automated and comparable evaluation.

Baseline experiments to demonstrate the usefulness of these benchmarks are run with existing state-of-the-art methods in section 5 and further discussed, highlighting the relevance of our benchmark.

2 Use-cases

2.1 The power grid case

Industrial context Power System Operators are in charge of managing the security of large critical power grids (thousands electrical lines and substations that can be reconfigured) in real time and co-ordinate the supply of and demand for electricity in a manner that avoids fluctuations in frequency or interruptions of supply. Above all, the grid needs to be robust to blackouts at any time, which means in particular avoiding powerline overflows that can lead to a cascading failure (Figure 1, left). Operators have to face unexpected events (loosing a line for example due to weather constraints) or to anticipate events such as variation of production during the day or as equipment's maintenance. They do so by assessing the risks and leveraging grid flexibility through simulations, carefully choosing sets of remedial actions which act on the grid topology or on the production levels.

Applications Near real-time operations of a power grid can be classified into three steps with different expected speed and accuracy simulation trade-offs (Table 1):

- 83 1. *Risk assessment*, i.e. identifying problematic contingencies over a large possible number while
- 84 assessing their severity (anticipating for instance lines overloads, maintenance operations...);
- 85 2. *Remedial action search*, i.e. exploring for solutions to find a set of remedial actions on the grid
- 86 such as topology change, to solve a local problem and assess its overall impact;
- 87 3. *Decision making*, i.e. selection and validation of one of the best solutions before implementation.

88 **Physical Simulations** The computation of the grid state involves a set of physical laws (see
 89 appendix C.1) such as Kirchhoff's law or Joule effect. More specifically, the physical resolution of
 90 the problem is derived from a set of powerflow equations [16] described at any node k of the grid.
 91 The power injected at a node of the network s_k is the sum of active (p_k) and reactive powers (q_k):
 92 $s_k = p_k + jq_k$. From Kirchhoff energy conservation law, the relation between voltage angle and
 93 magnitude can be formulated for node k and neighboring nodes m as follows:

$$\begin{cases} 0 = -p_k + \sum_{m=1}^K |v_k||v_m|(g_{k,m} \cdot \cos(\theta_k - \theta_m) + b_{k,m} \sin(\theta_k - \theta_m)) & \text{Active power} \\ 0 = q_k + \sum_{m=1}^K |v_k||v_m|(g_{k,m} \cdot \sin(\theta_k - \theta_m) - b_{k,m} \cos(\theta_k - \theta_m)) & \text{Reactive power} \end{cases} \quad (1)$$

94 where: phasors θ_k are unknown for all node k ; either voltage $|v_k|$ or reactive power q_k are known input at any
 95 given node k ; active power p_k is a known input and $g_{k,m}$, $b_{k,m}$ known line characteristics for all nodes. For
 96 each line l , active p^l and reactive q^l powerflows or the current a^l can further be derived with Ohm's law.

97 **Significance** The problem 1 is *non linear* and *non convex*. To estimate these variables, a Newton-
 98 Raphson power flow solver such as LightSim2grid[17] can be used. Over the past years, the amount
 99 of simulation required has drastically increased due to emerging trends [18] – mainly driven by
 100 Energy Transition initiatives, with increasing renewable energy share as well as stronger exchanges
 101 with neighboring countries over the whole European grid, both leading to a greater stochasticity. In
 102 this context, the computation time of physical solvers becomes an obstacle for upgraded decision
 103 support [19]. An acceleration by several order of magnitudes is now expected.




Figure 1: Left, small IEEE14 grid with 2 line overflows and highlighted topology changes (node splitting). Right, cross section of a simplified tire.

104 2.2 The pneumatic domain use case

105 **Industrial Context** In order to assess pneumatic performances, several physical considerations have
 106 to be taken into account such as durability, ground adhesion or robustness for instance, depending on
 107 the numerous applications. To do so, it is required to model the tire behaviour during either rolling
 108 cycles or configurations where it is subject to crushing forces. Besides, as the forces arising from
 109 the the tire/ground interface are located on the contact area, it is also critical to accurately estimate
 110 these quantities to predict the global behaviour. Only then, by being as close as possible to the actual
 111 configuration (including considerations about the vehicle velocity with respect to the ground, the
 112 pressure, the friction, the material behaviour...) we can optimize and improve our understanding of
 113 the processes involved and find the best compromise in terms of performances.

114 **Applications** Real-world tests performed on tires involve in particular two classical configurations
 115 depending on the accuracy/computational time simulation trade-offs (Table 1):

- 116 1. *Wheel sustaining*: assess whether the wheel is able to properly sustain the weight of a vehicle;
- 117 2. *Rolling cycle*: assess the behaviour of the tire during the rolling phase.

118 **Physical Simulations** The computation of the tire state, that is to say the displacement and the
 119 contact stress arising from the contact conditions on the discretized domain (i.e. the mesh), is
 120 performed based on the resolution of a set of Partial Differential Equations (PDE) (see appendix C.2
 121 for more details) through Finite Element (FE) formalism. As such, the solution is evaluated at any
 122 node of the domain’s mesh, for the displacement, at the nodes on the contact boundary for the contact
 123 stress. In these PDEs, several physical considerations are involved such as: the behaviour law, the
 124 relation between the stress acting on a body and the displacement, the law of motion, the unilateral
 125 contact conditions (equivalent to assuming the ground is perfectly rigid), the Coulomb’s law of dry
 126 friction. For more details about Contact Mechanics, we refer to [20, 21, 22]. Note that, while the
 127 displacements and contact stress at the contact boundary are the actual unknowns of the problem, it
 128 may not be enough depending on the usecase. For instance, some physical quantities relevant to a
 129 given application can be computed as a post-processing of the unknowns.

130 Figure 1, right, depicts the cross section of a simplified tire. We consider an idealized straight rolling
 131 on an undeformable ground at constant speed.

132 **Significance** Such a problem is strongly *nonlinear* due to the non linear behavior law, the large
 133 deformation framework and the frictional contact conditions. In order to estimate the displacement
 134 and the stress, the FE solver "Getfem" [23] is used. In practical applications, rolling simulations in
 135 particular provides a lot of useful information, such as for instance the contact area, forces, contact
 136 pressure and moments. Classical methods exists [24, 25], however, because of the problem inherent
 137 complexity, the computation time is prohibitively expensive. Running over a day sometimes, it
 138 limits its use in industrial applications compared to simpler surrogate model. An order of magnitude
 139 acceleration with acceptable accuracy would democratize its usage.

Table 1: Grid and pneumatic apps: speed vs accuracy and physical law compliance trade-offs

		Application	Variables to predict	Accuracy & PL compliance	Speed
Use cases	Grid	(1) Risk assessment	a^ℓ	+	+++
		(2) Action Search	a^ℓ, p^ℓ, v_k	++	++
		(3) Decision Making	$a^\ell, p^\ell, v_k, q^\ell, \theta_k$	+++	+
	Tire	(1) Wheel sustaining	u_Ω	+++	++
		(2) Rolling cycle	u_Ω, λ_c	++	+++

140 2.3 Added value of ML

141 Generally speaking, ML-model can provide more direct and faster predictions than a Newton-Raphson
 142 resolution over the non-linearities of both use cases. It can leverage a learning memory of any given
 143 grid, mesh or last rolling cycle iteration of interest without restarting the resolution from scratch as if
 144 it was a new system or problem. Also, as only some variables such as the flows or contact forces
 145 are sometimes of interest, in some of our benchmark tasks we will only require machine learning
 146 models to predict them instead of predicting all variables as the physical solvers. ML-models could
 147 finally provide more factorized computation such as for varying grid topologies (varying number of
 148 electrical nodes) as existing physical solvers does not offer factorization over such dimension.

149 3 Related works and novelty

150 **Simulations and benchmarks in power grids.** Although simulation time and convergence have
 151 improved over decades thanks to benchmarks based on shared power grid cases and some contests
 152 [26, 27], it remains too slow to compute large volume of simulations. In addition, existing simulators
 153 are general purpose and not application-oriented which we would like to address with this work. Some

	paper reference	evaluation criteria categories				impact meaningful thresholding	readability standard comparative viz	environment setup			
		ML-related	industrial readiness	ood generalization	PL compliance			dataset	ref Physical simulator	ML model repository	baselines
ML-model papers	LeapNet [51]	yes	no	yes	no	no	no	large, simple prod distributions	fast - Hades2 (proprietary)	maintained	yes & diverse
	GNS/DSS [13,31,34]	yes	no	no	partial	no	no	data generation shared, simple distributions	slow - PandaPower (open source)	unmaintained	no
	Fast contingency analysis [32]	yes	partial	no	no	no	no	not shared, realistic prod but simple topo	slow - PandaPower (open source)	no access	yes & diverse
	Physics-informed GNN [35]	yes	no	no	no	no	no	not shared, realistic prod but simple topo	none	no access	yes but uniform
	Gridwarm [36]	yes	partial	no	partial	no	no	data generation shared, simple distributions	none	unmaintained	yes & diverse
Benchmark suites	LIPS	yes	yes	yes	comprehensive	yes	yes	large & doc & realistic prod + topo distributions	fast - LightSimgrid (open source)	maintained	yes & diverse
	SimBench [28]	no (optimization & heuristics)	yes	yes	no	no	yes	medium & doc & realistic prod but simple topo	slow - PandaPower (open source)	no ML	yes & diverse
	PowerGridLib [29]	no (optimization)	partial	yes	partial	no	no	small & doc & realistic prod but static topo	med-speed - PowerModel (open source)	no ML	yes but uniform

Figure 2: Comparative table between LIPS and related work for the power grid case. This highlights that our framework offers a comprehensive evaluation setup that was often barely covered by others.

application-oriented simulation-related benchmarks emerged lately in the power system community (SimBench [28], Power Grid Lib [29]). However they are mostly designed to drive advances in operational research algorithms. In comparison, our benchmark: *a)* stresses the importance of considering the complexity of varying grid topologies for industrial applications; *b)* unlocks the creation of data-driven models by providing comprehensive data distributions to train them, similarly to [30] for other power grid related applications. *c)* define specific metrics to evaluate them such as physics compliance, out-of-distribution generalization over unseen topologies or industrial readiness considering available data volume and scalability. It eventually allows a fair comparison of pre-existing ML-models [8, 31, 32, 33, 34, 35] over all necessary dimensions as summarized in Figure 2. [36] also concurrently advocates for such evaluation over defined categories as a first step towards proper benchmarks.

PDEs simulations and benchmarks for pneumatic. In the last few years, the success of deep learning techniques has encouraged researchers to investigate their capability to solve PDE problems. Several works were proposed to hybridize PDE-based physical problems with Neural Networks (NN), from black-box resolution on unstructured meshes with graphs NN [37], to more interpretable approaches like the physics informed NN [14]. Some other works have focused on using un-supervised learning techniques to avoid the mesh construction (mesh-free methods)[38, 39]. Regarding pneumatic domain in particular, several attempts to use these techniques have already been made so far: the first tire/pavement contact-stress model based on artificial NN in [40] using a Neuro-Patch Model, tire modeling was investigated in [41] relying on a feedforward back propagation algorithm and [42] proposed a Structure-Preserving NN to predict the stress field within the tire. While providing promising results, none of these works attempt to compare fairly the performances of several ML-models with respect to a set of significant application-based criteria and we propose to fill that gap. To our knowledge, this is the first ML-friendly benchmark for pneumatic.

Benchmark for Learning to simulate physics. Learning to simulate benchmarks started to emerge recently. [7] studied extensively a neural network architecture performance over several simulators based on different physics. Unfortunately no resulting benchmark has been made available yet. It mainly relies on qualitative visual analysis, while more quantitative metrics as well as physical law verification could help for better comparison as we later formalize. The authors eventually indicate that scalability to large systems remains an issue for now, as well as proper generalization in regions with high variability, highlighting the need for further advances. [43] also proposed lately a new benchmark over four PDEs canonical physic systems to drive forward the development of data-driven time integration solutions. Both focus primarily on scientific needs, with limited evaluation criteria categories, as opposed to industrial needs and applications.

Identified research question (RQ). Thus, compared to these various works, the research questions we address in this work can be described as follows: 1) There has been ongoing ML research for physical simulations for several years now. Are current evaluation setup comprehensive enough to actually provide applicable models in the industry ? If not, what is missing ? 2) Can we define an homogeneous evaluation framework, with generic and comprehensive categories of criteria, for different industrial domains that could systematize the creation of such benchmarks and possibly drive cross-domain advances? 3) How can we represent an exhaustive set of benchmark results in an interpretable way ?

We also set open research questions (ORQ) yet to be addressed that should be of interest for ML research: 1) What kind of inductive biases could help enforce ood generalization and physical consistency without sacrificing speed? 2) Is there a one-size fits all simulation model that performs best for all applications in a given domain or should it be more tailored to achieve better application-specific trade-off? 3) Could we foster the emergence of foundational models accross domains?

Regarding RQ1 and considering the power grid use case as a reference one since ML research has been ongoing for more than a decade now, we can answer that current evaluation setup are very heterogeneous from Figure 2, each with distinct attributes. Hence none looks comprehensive enough and a standardization is needed with meaningful categories and targets to drive research towards industrial impact which has yet to be seen.

4 Benchmark suite design

4.1 Comprehensive evaluation criteria for benchmarking industrial physical simulations

The first step towards LIPS benchmark is a generic and yet comprehensive categories of evaluation criteria that standardizes it, allowing for comparison within and across physical domains, while being expressive enough to represent industrial needs and expectations. ML-related only metrics are not sufficient in that regard. Thus, we introduce four categories of criteria of importance for industrial applications and illustrate their applicability and usability on our 2 use cases in section 5.

ML-related performance Among classical ML metrics, we focus on the trade-offs of typical model accuracy metrics such as Mean Absolute Error (MAE) vs computation time (optimal ML inference time without batch size consideration as opposed to application time later).

Industrial Readiness When deploying a model in real-world applications, it should consider the real data availability and scale-up to large systems. We hence consider: 1) *Scalability*: the computational complexity of a surrogate method should scale well depending on the problem size, e.g. number of nodes on power grid, mesh refinement level in pneumatics; 2) *Application Time*: as we are looking for a tailored model to an application, we aim at measuring the computation time in this context. We consider a proper finite batch size for that application which can affect the speed-up.

Application-based out-of-distribution Generalization For industrial physical simulation, there is always some expectation to extrapolate over minimal variations of the problem geometry depending on the application. We hence consider ood geometry evaluation such as unseen power grid topology or unseen pneumatic mesh variations.

Physics compliance Physical laws compliance is decisive when simulation results are used to make consistent real-world decisions. Depending on the expected level of criticality of the benchmark, this criterion aims at determining the type and number of physical laws that should be satisfied.

4.2 Power grid application-oriented benchmarking task descriptions and datasets

From applications in table 1, we define two application-oriented benchmarks. Note that the third application will be considered in a future benchmark. The Benchmark datasets all depart from the same published datasets of realistic production and consumption distributions [44], [45] over two widely studied grids (IEEE 14 and IEEE 118 bus-systems) in the power system literature [46]. The benchmark datasets however each differs from the application specific grid topology variations which

are applied using Grid2Op [47] framework. Ground truth of physical variables are further computed using LightSim2Grid [17] physical solver with industrial-like performance on the selected grids.

1. **Benchmark 1 - Risk assessment through contingency screening.** The problem is to anticipate near real-time potential threats on the power grid and warn the operators accordingly [48]. It simulates incidents (aka contingencies) involving various elements of the grid (such as the disconnection of a line), one by one. For each contingency, a risk is identified when overloads on lines are detected. On a real grid, this scenario means running hundred of thousands of simulations, thereby, computation time is critical, especially since this risk assessment is refreshed every few minutes. We consider large simulation batches and the main physical variable is the line electric current a^ℓ , because an overload occurs when it exceeds the line capacity.

Dataset specificity: It presents grid snapshots including all possible line disconnections (N-1) for few different reference grid topologies. An ood topology test set containing N-2 line disconnections (2 line disconnections combined) is also attached to test for such generalization.

2. **Benchmark 2 - Remedial action search.** We need to explore possible solutions (aka "remedial actions") to identified risks for recommendation to the grid operator as in [49]. A solution consists in predefined topological change on the grid that alleviates the previous overflow without generating any new problem. Those changes such as node splitting (see Figure 1) bring more non-linearity than line disconnections in benchmark1, making the distributions more complex. We here target medium-sized batches. Additional physical variables are predicted: active power flows p^ℓ and voltages v_k . Level of compliance with more related physical laws is expected. This allows the operator to better assess the system state in a difficult situation with some consistency.

Dataset specificity: It presents grid snapshots over single substation topological reconfiguration among a set of specified ones. It also considers some possible line contingencies that could cause overloads. An ood topology test set containing combination of 2 topological unitary actions is also attached to test for such generalization.

For more details about dataset input and output variables and dimensions for both industrial use cases, please refer to appendix C with notation description in appendix B. Our "Datasheet for dataset" in appendix A will also give you additional information about our datasets creation and content.

For the power grid use case, you might also want to look at the Grid2op documentation¹ and take a look at the baseline architecture with variable names and dimensions in E.2 for a visual illustration.

4.3 Pneumatic application-oriented benchmarking task descriptions and datasets

We focus on the mechanics of tires on hard surfaces, one of the basic types of problem in the mechanics of tires the vehicle engineers are concerned with. As in table 1, we define two application-oriented benchmarks addressed in the literature, for instance in [42] for the rolling. To generate the datasets, we rely on the tire and experiment configuration described in [50]. Both the reference physical solution and the physical criteria of interest are computed by using the FE physical solver Getfem [23] and used as ground truth. Thus, the fields predicted by an augmented simulator are given to the physical solver in order to compare the physical criteria prediction to the ground truth of these very quantities. Note that, the pure mechanical criteria computation part is performed by the solver for convenience, as their calculation rely on the underlying physical model at hand.

1. **Benchmark 1 - Static vertical stiffness.** One of the basic function of a pneumatic tire is to support the vehicle weight. When a normal load is applied to a tire, it deflects as the load increases. Then, using the vertical load–deflection curves, we can estimate the so-called static vertical stiffness of tires. Such a criteria is known to have significant impacts on riding comfort, steering stability, and driving performance. Experimentally, this scenario implies running several simulations where different loads are applied on the wheel (inputs) to observe the resulting displacement of the structure (output). To be more specific, the physical variable we are interested in is the displacement u_Ω .

¹<https://grid2op.readthedocs.io/en/latest/observation.html#main-observation-attributes>

Dataset specificity: It presents displacement snapshots for different forces applied on the tire. Each displacement field arise from the simulation of a different static problem on a fixed axisymmetric mesh for the same physic.

2. **Benchmark 2 - Design testing during a rolling cycle.** Next, the aim is to assess the behaviour of the tire under the action of displacement-enforced rolling. Rather than the actual value of the criteria, we are interested about the relevancy of the design, i.e. whether the criteria value are within an acceptable range. Unlike the first scenario, this is a quasi-static configuration. Instead of running several static simulations, a single quasi-static problem is run for several time instants within a time interval over several rolling cycles. The physical variables we are interested in are the displacement u_Ω and the contact stress on the contact boundary λ_c .

Dataset specificity: It presents displacement and contact stresses snapshots evaluated at different instant during the rolling process. The idea is to train the model during $[0, t_1]$ and then evaluate the model for $t > t_1$; as such it is a pure out of distribution example. Unlike the first case, it involves a single quasi-static problem on a fixed non-axisymmetric mesh with time as input variable.

4.4 Configurable benchmark suite architecture & ressources

These benchmarks are implemented within a extensible platform that consists of three modules combining data management, surrogate models integration and evaluation metrics. The developed platform is flexible and allow to integrate more benchmarks from other similar domains.

A Benchmark is instantiated by a combination of a dataset, an augmented simulator and an evaluation objects as in Figure 3. Each of them is defined in a generic configuration file, making it easy to adapt it for different setups. They further comply with simple interfaces, making it modular to add new evaluation metric or new physical domains.

Benchmark resources The benchmark implementation and data are provided in open-source via a github link ² alongside a starting kit aiming to facilitate the use of main functionalities. In addition, we make LIPS available on Codabench³ – an open, public platform that allows to submit easily surrogate models, and compare fairly submissions, under the same settings and in a fully automated way. Results of participants are summarized on a result table, which facilitates progress monitoring. We strongly encourage anyone working on ML for real physical problems to use this platform to submit and evaluate their algorithms against the benchmarks we described above. Public results could be highlighted and discussed at NeurIPS 2022.

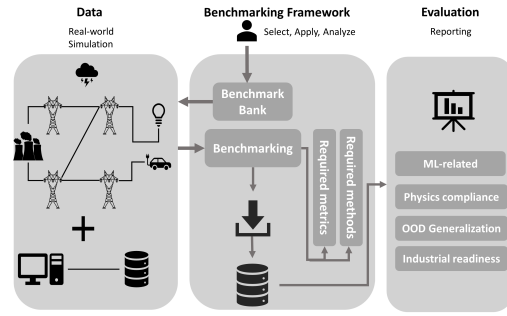


Figure 3: Benchmarking framework

5 Experiments

This section presents the evaluation results of baseline methods for each scenario of both use cases, alongside the experimental configurations used to obtain them.

5.1 Experimental setup

Regarding the stochastic nature of the optimisation methods based on gradient descent, 5 trials with different seed has been executed and the performances based on mean and standard deviation of different runs. All the experiments in the following sections are performed using a server equipped with AMD EPYC 7502P 32-Core Processor, NVIDIA RTX A6000 GPU and 128 GB of RAM. All computation time evaluation are run on the CPU with time measured per simulation or prediction.

²<https://github.com/Mleyliabadi/LIPS>

³<http://htmlpreview.github.io/?https://github.com/Mleyliabadi/LIPS/blob/main/codabench/codabench.html>

Power grid - Our reference simulator LightSim2Grid has comparable speed on those IEEE grid cases than the proprietary RTE solver Hades 2, and is faster than PandaPower [51], the physical simulator used in SimBench, [28] by at least a factor 30 (see [17]) and also faster than the one used in Power Grid Lib [29] by at least a factor 5 on the hardware setup in section 5.1. This reference simulator choice hence makes our benchmark quite challenging. We have looked at a first baseline with differently tuned reference simulator, with a maximum of 1 solver iteration to see the maximum speed we can get with approximate results. We never go beyond a factor 5 speed-up and hence don't consider it further as it is far from expected speedup. But it definitely set a lower bound to outperform. We have then considered three different baselines for evaluation: a physics based simplification of power flow calculus which is DC approximation [52] and two augmented simulators which are Fully Connected (FC) architecture and a state-of-the art LEAP net [53], where contrary to FC, the topology intervenes in the latent space and demonstrate better combinatorial generalisation capabilities. Note that we have conducted automated grid search to find the best performing network hyperparameters for both architectures (see appendix E.4). Through this benchmark suite, we encourage the community to contribute and to suggest approaches aiming to improve the performances of the existing baselines.

Pneumatic - We have considered two types of augmented simulators within the first benchmark: a FC architecture and a Unet [54] architecture. For the latter, the numerical solution evaluated by the physical solver on an unstructured mesh is projected on a 128×128 grid then, after the evaluation by the augmented solver, it is projected back to the mesh. For the second benchmark, two FC architecture are used: one to predict the displacement and one to predict the contact stress on the contact boundary.

5.2 Benchmark results and experiments

Table 2 summarizes the benchmark results for both use cases and their specific scenarios. In order to enhance the readability, we have sometimes used three qualitative levels from "not acceptable" to "great", relying on application-relevant threshold values reported in appendix C (tables 3 and 6). The full quantitative table from which this table is derived is also provided in section J.1 of appendix.

Power grid As it can be seen in this table, the ML based models (FC and LeapNet) shows better accuracy for target variables than the baseline DC approximation. However, their performance on out-of-distribution dataset is still challenging and not acceptable. While the LeapNet shows a bit better generalization performance, the accuracy is still above 6% error, compared to the reported performance in [53]. Maybe surprisingly, quantitative OOD results on the small (inner small circles) grid are worse than the larger grid. It could be explained by the fact that the small network is subject of more variations when introducing some changes such as N-1 actions but easier to track and understand the variations. The only possible physical law is also verified for this benchmark. Looking at a more challenging benchmark 2, we can observe that more variables should be predicted and more laws should be verified. The DC approximation respect most of the laws as it is based on physical solver, however it comes with some costs from the accuracy point of view. A high order speed-up when using ML models can be observed in comparison to a very optimized solver. We emphasize that such speed-up time is dependant on the application context which needs to be considered. For more detailed comparison concerning the physics-based criteria, the readers may refer to appendix J.2.

Pneumatic Likewise, regarding the pneumatic use case, it seems ML models perform relatively well in the first benchmark with FC architecture for the prediction of the displacement field, despite questionable results regarding the physics. For the Unet, while the results obtained are clearly not acceptable at all for this benchmark, they can be explained by the small dataset considered. Such an architecture requires further investigations to assess whether it is actually suitable for this benchmark. For the second benchmark, despite the fact that it is a pure out-of-distribution case, the ML-model behaves surprisingly well: the prediction is quite accurate for the contact stress. However, given the results obtained for the displacement, we can reasonably assume the scaler proposed to handle the out-of-distribution displacement prediction is not working as expected, which prevent an acceptable convergence of the optimizer associated to the augmented simulator and thus acceptable results.

Table 2: Benchmark result table for the two use cases under 4 categories of evaluation criteria. The performances are reported using three colors computed on the basis of two thresholds. Colors and symbol meaning: ● Not acceptable ● Acceptable ● Great \odot two problem scales reported (in that case, speed-up for smaller scale is in parenthesis). The number of circles corresponds to the number of variables or laws that are evaluated. For quantitative values from which this table is derived, please refer to section J.1 of appendix.

			Criteria category						
			Methods	ML-related		Readiness	OOD Gen.	Physics	
				Quality	Speed-up	Speed-up	Quality	Domain laws	
Use cases	Power Grid	Bench1	DC	<u>A</u> 	NA	19 (7)	<u>A</u> 	<u>P1</u> 	
			FC		19 (22)	17 (20)			
			LeapNet		17 (19)	14 (17)			
		Bench2	DC	<u>A</u> <u>P</u> <u>V</u> 	NA	5 (3)	<u>A</u> <u>P</u> <u>V</u> 	<u>P1</u> <u>P2</u> <u>P3</u> <u>P4</u> <u>P5</u> <u>P6</u> <u>P7</u> <u>P8</u> 	
			FC		99 (157)	57 (27)			
			LeapNet		90 (140)	54 (24)			
	Pneumatic	Bench1	FC	<u>u_Ω</u> 	18	NA	NA	<u>P1</u> <u>P2</u> <u>P3</u> 	
			UNet		18	NA	NA		
		Bench2	FC	<u>u_Ω</u> <u>λ_c</u> 	11	NA	<u>u_Ω</u> <u>λ_c</u> 	<u>P1</u> <u>P2</u> <u>P3</u> 	

In all the cases investigated, the speed-up observed is at least one order of magnitude for both benchmarks compared to the physical solver, which was precisely our aim for the rolling case in the first place. However, given the accuracy for the displacement field, it is far from satisfactory and as such only partially met with our requirements.

We have shown with two very distinct industrial and physical domains that we can systematize the creation of comprehensive and yet homogeneous benchmarks for the use of physical simulation in industry, hence answering our RQ2. Our result table also displays a lot of benchmark results, yet in a compact and readable way through the use of meaningful thresholding, colors and symbols: this answers our RQ3 and is an original benchmark result representation in the ML community to the best of our knowledge.

6 Conclusion and perspectives

This paper has investigated the definition and the implementation of a new benchmark suite, called LIPS (Learning Industrial Physical Simulations). The aim is to address simulation-based industrial use-cases augmented with machine learning techniques. Two distinct industrial use cases (with different physics) have been considered to illustrate the proposed framework, with several application-oriented benchmarks. Experiments have shown how several comparative studies based on proposed categories of criteria are conducted. The obtained results have also clarified the remaining challenges of existing state-of-the-art augmented simulator to emulate the behavior of a physical simulator in an industrial context. Although, they are much faster for providing the appropriate results, their interesting but yet insufficient out-of-generalization properties and vulnerability vis-à-vis the physics compliance highlights the requirement for further improvements. This benchmark opens the door for designing more robust and reliable augmented simulators that will find better real-world applicability. Future work to extend the suite will focus on integrating new industrial use cases related to other physical (e.g. aeronautics, transport,...) which would help to improve the generalization of LIPS.

Acknowledgments and Disclosure of Funding

This research was supported by IRT SystemX, RTE France, Michelin in the context of HSA project (hybridization of simulation and ML). Thanks to Reynaldo Gomez and Benoit Bastien, it was also supported by NVIDIA through an open collaboration to develop benchmarking framework and baselines for industrial physical simulation. We thank all contributors and colleagues from RTE, Michelin, IRT SystemX and NVIDIA who provided insight and expertise that greatly assisted the research. We are finally very thankful to EDF Exaion for sponsoring through their infrastructure these developments by making GPU servers available for training all baselines as well as for running the Codabench public evaluation thread.

References

- [1] Michel Bergmann, Andrea Ferrero, Angelo Iollo, Edoardo Lombardi, Angela Scardigli, and Haysam Telib. A zonal galerkin-free pod model for incompressible flows. *Journal of Computational Physics*, 352, 10 2017.
- [2] Fabien Casenave, Nissrine Akkari, Felipe Bordeu, Christian Rey, and D. Ryckelynck. A nonintrusive distributed reduced order modeling framework for nonlinear structural mechanics – application to elastoviscoplastic computations. *International Journal for Numerical Methods in Engineering*, 121, 08 2019.
- [3] J. Kutz, Steven Brunton, Bingni Brunton, and Joshua Proctor. *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. 11 2016.
- [4] Francisco Chinesta, Elías CUETO, Emmanuelle Abisset-Chavanne, Jean Duval, and Fouad Khaldi. Virtual, digital and hybrid twins: A new paradigm in data-based engineering and engineered data. *Archives of Computational Methods in Engineering*, 27, 11 2018.
- [5] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Reviews of Modern Physics*, 91(4):045002, 2019.
- [6] J. Tompson, K. Schlachter, P. Sprechmann, and K. Perlin. Accelerating eulerian fluid simulation with convolutional networks. *ArXiv: 1607.03597*, 2016.
- [7] MF Kasim, D Watson-Parris, L Deaconu, S Oliver, P Hatfield, DH Froula, G Gregori, M Jarvis, S Khatiwala, J Korenaga, et al. Building high accuracy emulators for scientific simulations with deep neural architecture search. *Machine Learning: Science and Technology*, 3(1):015013, 2021.
- [8] Benjamin Donnot, Balthazar Donon, Isabelle Guyon, Zhengying Liu, Antoine Marot, Patrick Panciatici, and Marc Schoenauer. Leap nets for power grid perturbations. *arXiv preprint arXiv:1908.08314*, 2019.
- [9] S. Rasp, M. S. Pritchard, and P. Gentine. Deep learning to represent sub-grid processes in climate models. *PNAS*, 2018.
- [10] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International Conference on Machine Learning*, pages 8459–8468. PMLR, 2020.
- [11] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. *arXiv preprint arXiv:1806.07366*, 2018.
- [12] Samuel J Greydanus, Misko Dzumba, and Jason Yosinski. Hamiltonian neural networks. 2019.
- [13] Balthazar Donon, Zhengying Liu, Wenzhuo Liu, Isabelle Guyon, Antoine Marot, and Marc Schoenauer. Deep statistical solvers. In *NeurIPS*, 2020.

- [14] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [15] Zhen Xu, Huan Zhao, Wei-Wei Tu, Magali Richard, Sergio Escalera, and Isabelle Guyon. Codabench: Flexible, easy-to-use and reproducible benchmarking for everyone. *arXiv preprint arXiv:2110.05802*, 2021.
- [16] Daniel K Molzahn, Ian A Hiskens, et al. A survey of relaxations and approximations of the power flow equations. 2019.
- [17] Lightsim2grid. <https://github.com/BDonnot/lightsim2grid>.
- [18] Antoine Marot, Adrian Kelly, Matija Naglic, Vincent Barbesant, Jochen Cremer, Alexandru Stefanov, and Jan Viebahn. Perspectives on future power system control centers for energy transition. *Journal of Modern Power Systems and Clean Energy*, 10(2):328–344, 2022.
- [19] Antoine Marot, Alexandre Rozier, Matthieu Dussartre, Laure Crochepierre, and Benjamin Donnot. Towards an ai assistant for human grid operators. *arXiv preprint arXiv:2012.02026*, 2020.
- [20] Peter Wriggers. *Computational Contact Mechanics*. 01 2006.
- [21] Patrick Tallec. Numerical methods for nonlinear three dimensional elasticity. *Handbook of Numerical Analysis*, 3, 12 1994.
- [22] Tod Laursen. *Computational Contact and Impact Mechanics: Fundamentals of Modeling Interfacial Phenomena in Nonlinear Finite Element Analysis*. 01 2003.
- [23] Yves Renard and Konstantinos Poullos. Getfem: Automated fe modeling of multiphysics problems based on a generic weak form language. *ACM Transactions on Mathematical Software*, 47:1–31, 12 2020.
- [24] Yves Renard. Generalized newton’s methods for the approximation and resolution of frictional contact problems in elasticity. *Computer Methods in Applied Mechanics and Engineering*, 256:38–55, 04 2013.
- [25] Peter Wriggers. *Computational Contact Mechanics*. 01 2006.
- [26] Ray Zimmerman, Carlos Murillo-Sanchez, and Defang Gan. Matpower-a matlab power system simulation package: User’s manual. 12 1997.
- [27] Zhuo Li, Raju Balasubramanian, Frank Liu, and Sani Nassif. 2011 tau power grid simulation contest: Benchmark suite and results. In *2011 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 478–481, 2011.
- [28] Steffen Meinecke, Džanan Sarajlić, Simon Drauz, Annika Klettke, Lars-Peter Lauen, Christian Rehtanz, Albert Moser, and Martin Braun. Simbench—a benchmark dataset of electric power systems to compare innovative solutions based on power flow analysis. *Energies*, 13:3290, 06 2020.
- [29] Sogol Babaeinejadsarookolae, Adam Birchfield, Richard D. Christie, Carleton Coffrin, Christopher DeMarco, Ruisheng Diao, Michael Ferris, Stephane Fliscounakis, Scott Greene, Renke Huang, Cedric Jozs, Roman Korab, Bernard Lesieutre, Jean Maeght, Terrence W. K. Mak, Daniel K. Molzahn, Thomas J. Overbye, Patrick Panciatici, Byungkwon Park, Jonathan Snodgrass, Ahmad Tbaileh, Pascal Van Hentenryck, and Ray Zimmerman. The power grid library for benchmarking ac optimal power flow algorithms, 2021.

- 487 [30] Xiangtian Zheng, Nan Xu, Loc Trinh, Dongqi Wu, Tong Huang, S Sivaranjani, Yan Liu, and
488 Le Xie. Psml: A multi-scale time-series dataset for machine learning in decarbonized energy
489 grids. *arXiv preprint arXiv:2110.06324*, 2021.
- 490 [31] Balthazar Donon, Benjamin Donnot, Isabelle Guyon, and Antoine Marot. Graph neural solver
491 for power systems. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages
492 1–8. IEEE, 2019.
- 493 [32] Florian Schäfer, Jan-Hendrik Menke, and Martin Braun. Evaluating machine learning models
494 for the fast identification of contingency cases. *Applied AI Letters*, 1(2):e19, 2020.
- 495 [33] Valentin Bolz, Johannes Rueß, and Andreas Zell. Power flow approximation based on graph
496 convolutional networks. In *2019 18th IEEE International Conference On Machine Learning
497 And Applications (ICMLA)*, pages 1679–1686. IEEE, 2019.
- 498 [34] Balthazar Donon, Rémy Clément, Benjamin Donnot, Antoine Marot, Isabelle Guyon, and Marc
499 Schoenauer. Neural networks for power flow: Graph neural solver. *Electric Power Systems
500 Research*, 189:106547, 2020.
- 501 [35] Laurent Pagnier and Michael Chertkov. Physics-informed graphical neural network for param-
502 eter & state estimations in power systems. *arXiv preprint arXiv:2102.06349*, 2021.
- 503 [36] Shimiao Li, Amritanshu Pandey, and Larry Pileggi. Gridwarm: Towards practical physics-
504 informed ml design and evaluation for power grid. *arXiv preprint arXiv:2205.03673*, 2022.
- 505 [37] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W Battaglia. Learning
506 mesh-based simulation with graph networks. *arXiv preprint arXiv:2010.03409*, 2020.
- 507 [38] Jens Berg and Kaj Nyström. A unified deep artificial neural network approach to partial
508 differential equations in complex geometries. *Neurocomputing*, 317:28–41, 2018.
- 509 [39] Justin Sirignano and Konstantinos Spiliopoulos. Dgm: A deep learning algorithm for solving
510 partial differential equations. *Journal of computational physics*, 375:1339–1364, 2018.
- 511 [40] *Development of a Tire/Pavement Contact-Stress Model Based on an Artificial Neural Network*,
512 volume Innovations in Vehicle Design and Development of ASME International Mechanical
513 Engineering Congress and Exposition, 11 1999.
- 514 [41] Hoyong Kim and Paul Ro. A tire side force model by artificial neural network. 02 1995.
- 515 [42] Quercus Hernández, Alberto Badias, David González, Francisco Chinesta, and Elías Cueto.
516 Deep learning of thermodynamics-aware reduced-order models from data. *Computer Methods
517 in Applied Mechanics and Engineering*, 379:113763, 06 2021.
- 518 [43] Karl Otness, Arvi Gjoka, Joan Bruna, Daniele Panozzo, Benjamin Peherstorfer, Teseo Schneider,
519 and Denis Zorin. An extensible benchmark suite for learning to simulate physical systems.
520 2021.
- 521 [44] Antoine Marot, Benjamin Donnot, Camilo Romero, Balthazar Donon, Marvin Lerousseau, Luca
522 Veyrin-Forrer, and Isabelle Guyon. Learning to run a power network challenge for training
523 topology controllers. *Electric Power Systems Research*, 189:106635, 2020.
- 524 [45] Antoine Marot, Benjamin Donnot, Gabriel Dulac-Arnold, Adrian Kelly, Aïdan O’Sullivan, Jan
525 Viebahn, Mariette Awad, Isabelle Guyon, Patrick Panciatici, and Camilo Romero. Learning
526 to run a power network challenge: a retrospective analysis. *arXiv preprint arXiv:2103.03104*,
527 2021.
- 528 [46] Ray Daniel Zimmerman, Carlos Edmundo Murillo-Sánchez, and Robert John Thomas. Mat-
529 power: Steady-state operations, planning, and analysis tools for power systems research and
530 education. *IEEE Transactions on power systems*, 26(1):12–19, 2010.

- 531 [47] Grid2op. <https://github.com/rte-france/Grid2Op>.
- 532 [48] Benjamin Donnot, Isabelle Guyon, Marc Schoenauer, Antoine Marot, and Patrick Panciatici.
533 Fast power system security analysis with guided dropout. *arXiv preprint arXiv:1801.09870*,
534 2018.
- 535 [49] A Marot, B Donnot, S Tazi, and P Panciatici. Expert system for topological remedial action
536 discovery in smart grids. 2018.
- 537 [50] Julien Gillard. *An Efficient Partitioned Coupling Scheme for Tire Hydroplaning Analysis*. 05
538 2018.
- 539 [51] Leon Thurner, Alexander Scheidler, Florian Schäfer, Jan-Hendrik Menke, Julian Dollichon,
540 Friederike Meier, Steffen Meinecke, and Martin Braun. pandapower—an open-source python
541 tool for convenient modeling, analysis, and optimization of electric power systems. *IEEE*
542 *Transactions on Power Systems*, 33(6):6510–6521, 2018.
- 543 [52] Brian Stott, Jorge Jardim, and Ongun Alsac. Dc power flow revisited. *IEEE Transactions on*
544 *Power Systems*, 24(3):1290–1300, 2009.
- 545 [53] Balthazar Donon, Benjamin Donnot, Isabelle Guyon, Zhengying Liu, Antoine Marot, Patrick
546 Panciatici, and Marc Schoenauer. Leap nets for system identification and application to power
547 systems. *Neurocomputing*, 416:316–327, 2020.
- 548 [54] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for
549 biomedical image segmentation. volume 9351, pages 234–241, 10 2015.

550 Checklist

- 551 1. For all authors...
- 552 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
553 contributions and scope? **[Yes]**
- 554 (b) Did you describe the limitations of your work? **[Yes] The conclusion in section 6**
555 **includes a discussion about the limitations of our work**
- 556 (c) Did you discuss any potential negative societal impacts of your work? **[N/A] We do**
557 **not expect any potential negative societal impacts of this work.**
- 558 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
559 them? **[Yes]**
- 560 2. If you are including theoretical results...
- 561 (a) Did you state the full set of assumptions of all theoretical results? **[N/A] No theoretical**
562 **results**
- 563 (b) Did you include complete proofs of all theoretical results? **[N/A] No theoretical**
564 **results**
- 565 3. If you ran experiments (e.g. for benchmarks)...
- 566 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
567 mental results (either in the supplemental material or as a URL)? **[Yes] Yes, via the**
568 **following github link: <https://github.com/Mleyliabadi/LIPS>**
- 569 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
570 were chosen)? **[Yes] We provide the details about the experimental setup in the**
571 **main paper and in the supplemental material. We also make available some**
572 **notebooks to reproduce the results. The platform includes a module for hyper**
573 **parameter tuning which is used for the presented models and is described in**
574 **supplemental material.**

- 575 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
 576 ments multiple times)? **[Yes] yes, evaluation criteria are reported using the mean**
 577 **and standard deviation over various runs of models**
- 578 (d) Did you include the total amount of compute and the type of resources used (e.g., type
 579 of GPUs, internal cluster, or cloud provider)? **[Yes] This information is provided in**
 580 **section 5.1**
- 581 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 582 (a) If your work uses existing assets, did you cite the creators? **[Yes] We have included**
 583 **the references to the core software libraries and the input dataset used in the**
 584 **benchmarks**
- 585 (b) Did you mention the license of the assets? **[Yes] We are only using open-source**
 586 **assets. The license of the assets we are using is provided at the URL we provide as**
 587 **a reference to the assets.**
- 588 (c) Did you include any new assets either in the supplemental material or as a URL? **[Yes]**
 589 **The license (MPL 2.0) is provided with the code of our benchmark framework.**
- 590 (d) Did you discuss whether and how consent was obtained from people whose data you're
 591 using/curating? **[N/A] No human data used or curated.**
- 592 (e) Did you discuss whether the data you are using/curating contains personally identi-
 593 fiable information or offensive content? **[N/A] Our data only include the result of**
 594 **numerical simulations.**
- 595 5. If you used crowdsourcing or conducted research with human subjects...
- 596 (a) Did you include the full text of instructions given to participants and screenshots, if
 597 applicable? **[N/A] No human participants**
- 598 (b) Did you describe any potential participant risks, with links to Institutional Review
 599 Board (IRB) approvals, if applicable? **[N/A] No human participants**
- 600 (c) Did you include the estimated hourly wage paid to participants and the total amount
 601 spent on participant compensation? **[N/A] No human participants**