



HAL
open science

Augmenting differentiable physics with randomized smoothing

Quentin Le Lidec, Louis Montaut, Cordelia Schmid, Ivan Laptev, Justin Carpentier

► **To cite this version:**

Quentin Le Lidec, Louis Montaut, Cordelia Schmid, Ivan Laptev, Justin Carpentier. Augmenting differentiable physics with randomized smoothing. RSS 2022 - Robotics Science and Systems, Workshop on Differentiable Simulation For Robotics, Jun 2022, New York, United States. hal-03703324

HAL Id: hal-03703324

<https://hal.science/hal-03703324>

Submitted on 23 Jun 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.

Augmenting differentiable physics with randomized smoothing

Quentin Le Lidec¹, Louis Montaut^{1,2}, Cordelia Schmid¹, Ivan Laptev¹, Justin Carpentier¹

¹ Inria and Département d’Informatique de l’École Normale Supérieure,

PSL Research University, Paris, France

firstname.lastname@inria.fr

² Czech Institute of Informatics, Robotics and Cybernetics,

Czech Technical University, Prague, Czech Republic

firstname.lastname@cvut.cz

Abstract—In the past few years, following the differentiable programming paradigm, there has been a growing interest in computing the gradient information of physical processes (e.g., physical simulation, image rendering). However, such processes may be non-differentiable or yield uninformative gradients (i.d., null almost everywhere). When faced with the former pitfalls, gradients estimated via analytical expression or numerical techniques such as automatic differentiation and finite differences, make classical optimization schemes converge towards poor quality solutions. Thus, relying only on the local information provided by these gradients is often not sufficient to solve advanced optimization problems involving such physical processes, notably when they are subject to non-smoothness and non-convexity issues. In this work, inspired by the field of zero-th order optimization, we leverage randomized smoothing to augment differentiable physics by estimating gradients in a neighborhood. Our experiments suggest that integrating this approach inside optimization algorithms may be fruitful for tasks as varied as mesh reconstruction from images or optimal control of robotic systems subject to contact and friction issues.

I. INTRODUCTION

A. Motivations

Various physical processes such as image rendering [11, 10] or physical simulation [16, 2], are subject to non-smooth phenomena. These result in badly conditioned gradients which may be uninformative (i.d., null almost everywhere) or even undefined. Typically, rendering [11, 10] is naturally non-differentiable and yields either null or undefined derivatives. Similarly, differentiable simulators exhibits null gradients over entire regions of the trajectory space [17]. In both cases, these physical gradients present some discontinuities or lack of regularity, which, in the end, may drastically hinder gradient-based optimization algorithms. In this presentation, we show how randomized smoothing appears as a natural solution to avoid the pitfalls induced by a lack of regularity originating from classic gradient estimators. We notably show its applications in the context of differentiable rendering [8] and differentiable simulation for control [9].

B. Background on randomized smoothing

Originally used in black-box optimization algorithms [12, 13], randomized smoothing was recently introduced in the

machine learning community [4, 1] in order to integrate non-differentiable operations inside neural networks.

In more details, a function g can be approximated by convolving it with a probability distribution μ :

$$g_\epsilon(x) = \mathbb{E}_{Z \sim \mu} [g(x + \epsilon Z)] \quad (1)$$

which corresponds to the randomly smoothed counterpart of g and can be estimated with a Monte-Carlo estimator as follows:

$$g_\epsilon(x) \approx \frac{1}{M} \sum_{i=0}^M g(x + \epsilon z^{(i)}) \quad (2)$$

where $\{z^{(1)}, \dots, z^{(M)}\}$ are i.i.d. samples and M is the number of samples. Intuitively, the convolution makes g_ϵ smoother than its original counterpart g and, thus, yields better conditioned gradients.

Using an integration by part yields a zero-th order estimator:

$$\nabla_x^{(0)} g_\epsilon(x) = \frac{1}{M} \sum_{i=0}^M -g(x + \epsilon z^{(i)}) \frac{\nabla \log \mu(z^{(i)})^\top}{\epsilon}. \quad (3)$$

This zero-th order estimator can be used even when g is non differentiable to obtain first-order information usable in a gradient-based optimization scheme, as in Sec. II.

Alternatively, a direct way to estimate the gradients whenever g is differentiable, is to use the first-order estimator:

$$\nabla_x^{(1)} g_\epsilon(x) = \frac{1}{M} \sum_{i=0}^M \nabla g(x + \epsilon z^{(i)}). \quad (4)$$

Such a first-order estimator averages gradients over a noise distribution and, thus, mechanically augments the quantity of captured information in the gradients of g even when they are uninformative e.g null in a region of the optimization space, as in Sec. III.

II. DIFFERENTIABLE RENDERING

An example of a non-differentiable process which can be turned into a differentiable one via zero-th order randomized smoothing is rendering. Rendering is the process which, given a 3D scene and a camera parameterization,

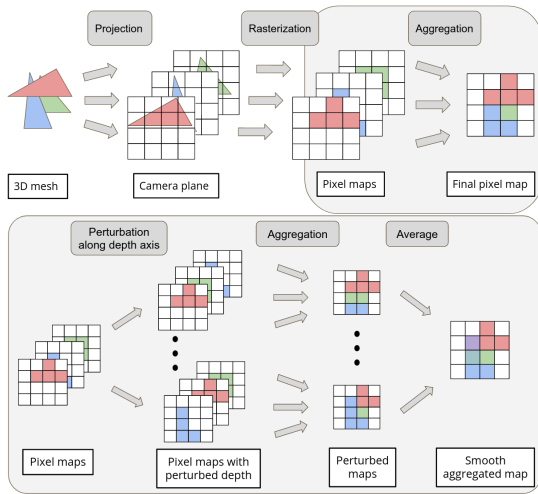


Fig. 1: **Top:** Overview of the rendering process: both rasterization and aggregation steps induce non-smoothness in the computational flow. **Bottom:** Illustration of the differentiable perturbed aggregation process. The rasterization step is made differentiable in a similar way.

outputs the corresponding RGB image. As illustrated in Fig. 1, rasterization-based renderers need to go through two discontinuous steps, called rasterization and aggregation. These discontinuities seem natural: if an object moves on a plane parallel to the camera, some pixels will immediately change color at the moment the object enters the camera view or becomes unoccluded by another object. Mathematically, the operations occurring during these steps can be written as *argmin* operators of Linear Programming problems:

$$y^*(\theta) = \underset{y \in C}{\operatorname{argmin}} \langle \theta, y \rangle \quad (5)$$

Such operators have null gradients almost everywhere and undefined otherwise, hence making them unusable for optimization.

Differentiable rendering approximates the underlying process by a smoother one by adding blur and transparency effects [10, 8, 6]. In particular, in [8], we propose to do it by replacing the original $y^*(\theta)$ operators from rasterization and aggregation by their randomly smoothed counterparts (Fig. 1). If randomized smoothing alleviates the burden of badly conditioned gradients by working with an approximate but smoother renderer, it is necessary to consider a way to reduce the smoothing during optimization in order to solve the original problem. In [8], we introduce a criterion allowing to automatically reduce the injected noise across iterations. This approach is applied to tackle pose optimization and mesh reconstruction tasks and reaches results competitive with the state of the art [8].

III. DIFFERENTIABLE SIMULATION FOR CONTROL

An example of a differentiable process for which first-order gradient estimation via randomized smoothing can be applied to enrich gradients' information is physical simulation. Simulating a physical system subject to contacts

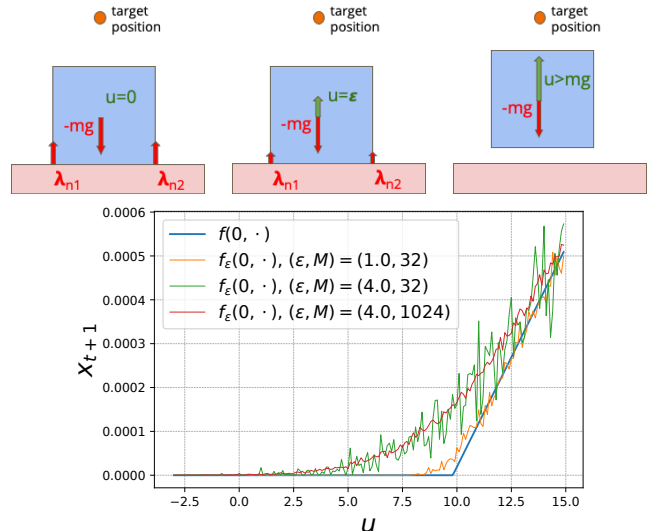


Fig. 2: **Top:** A slight vertical force is not able to break the unilateral contact, leaving cube on the floor and, thus, the state unchanged. **Bottom:** The non-smoothness of physics induces null gradients which results in the failure of classical optimization techniques

and frictions requires to solve a problem formed by the Euler-Lagrange equations of motion, the complementarity constraint accounting for contacts, the Maximum Dissipation Principle and Coulomb's law which governs friction. Solutions of this problem are typically obtained by solving an associated optimization problem [16, 3, 7] and, then, their gradients can be computed via implicit differentiation. However, as illustrated by Fig. 2 and detailed in [17], contacts and friction cause the simulation to be non-smooth and its gradients to remain null over large regions of the trajectory space.

In such a setting, randomized smoothing allows to approximate the underlying physics by a smoother one as shown in [9, 15] and illustrated by Fig. 2. Interestingly, this approach also allows to frame the stochasticity of exploration as a key element of the success of RL on solving physical problems with contacts and frictions [9]. If randomized smoothing allows to augment the quantity of information contained in the gradient estimator, it also introduces stochasticity and it is necessary to design accordingly an optimization procedure that preserve the convergence guarantees. To enforce the convergence towards an optimal (local) solution, it remains crucial to reduce the noise injected via the randomized smoothing across the iterations. In [9], we propose to decrease it in a way that adapts to the problem and avoids the smoothing being reduced too quickly, which would lead to performances similar to classical DDP, or too slowly, which would induce an unnecessary large number of iterations.

Finally, we apply our algorithm to solve a task on the Solo robot [5], an 18-DoF robotics system, with frictional contacts. Here, the goal is to reach a final target pose (lifting of the tip of one leg while keeping the three others on the ground),



Fig. 3: Randomized-DDP leverages the smooth dynamics to precisely schedule movements requiring to break contacts.




Fig. 4: Randomized DDP makes it possible to achieve complex tasks requiring breaking contacts with the floor on Solo robot such as lifting a leg (see Fig. 3)

which requires breaking some initial existing contacts, as illustrated in Fig. 3. Randomized smoothing coupled with Differential Dynamic Programming (R-DDP) allows to solve this task while classical DDP algorithm is not able to precisely apprehend contacts because of non-informative gradients of the dynamics (Fig. 4). Typically, the control obtained from DDP will only approach this pose while maintaining the feet on the ground which results in Solo bending its leg instead of lifting it.

IV. DISCUSSION

Gradients from differentiable physical processes should not be thought about separately from the actual problem and the associated optimization algorithm. It is often the case that exact gradients estimated from analytical expressions or automatic differentiation are likely to provide insufficient information for classical optimization to get satisfying solutions. Randomized smoothing proposes to capture additional information by using a noise distribution to evaluate gradients in the neighborhood of the current optimization stage. Doing so, it moves from a deterministic to a stochastic setting which may hinder performance of optimization algorithms (convergence and reproducibility) and thus requires some adjustments. We propose an approach which decreases the noise adaptively to improve robustness. However, detecting when to increase it is left as future work. Alternatively, analysing the respective stability properties of the zero-th and first order estimators represents another promising research direction [14].

ACKNOWLEDGMENTS

This work was supported in part by the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute), the European Regional Development Fund under the project IMPACT (reg. no. CZ.02.1.01/0.0/0.0/15 003/0000468), the Grant Agency of the Czech Technical University in Prague, grant No.

SGS21/178/OHK3/3T/17, and Louis Vuitton ENS Chair on Artificial Intelligence.

REFERENCES

- [1] Quentin Berthet, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis Bach. Learning with differentiable perturbed optimizers. *Advances in neural information processing systems*, 33:9508–9519, 2020.
- [2] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2019.
- [3] Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and J Zico Kolter. End-to-end differentiable physics for learning and control. *Advances in neural information processing systems*, 31, 2018.
- [4] John C Duchi, Peter L Bartlett, and Martin J Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.
- [5] Felix Grimminger, Avadesh Meduri, Majid Khadiv, Julian Viereck, Manuel Wüthrich, Maximilien Naveau, Vincent Berenz, Steve Heim, Felix Widmaier, Thomas Flayols, et al. An open torque-controlled modular robot architecture for legged locomotion research. *IEEE Robotics and Automation Letters*, 5(2):3650–3657, 2020.
- [6] Krishna Murthy Jatavallabhula, Miles Macklin, Florian Golemo, Vikram Voleti, Linda Petrini, Martin Weiss, Breandan Considine, Jerome Parent-Levesque, Kevin Xie, Kenny Erleben, Liam Paull, Florian Shkurti, Derek Nowrouzezahrai, and Sanja Fidler. gradsim: Differentiable simulation for system identification and visuomotor control. *International Conference on Learning Representations (ICLR)*, 2021. URL https://openreview.net/forum?id=c_E8kFWfhp0.
- [7] Quentin Le Lidec, Igor Kalevtykh, Ivan Laptev, Cordelia Schmid, and Justin Carpentier. Differentiable simulation for physical system identification. *IEEE Robotics and Automation Letters*, 6(2):3413–3420, 2021.
- [8] Quentin Le Lidec, Ivan Laptev, Cordelia Schmid, and Justin Carpentier. Differentiable rendering with perturbed optimizers. *Advances in Neural Information Processing Systems*, 34, 2021.
- [9] Quentin Le Lidec, Louis Montaut, Cordelia Schmid, Ivan Laptev, and Justin Carpentier. Leveraging randomized smoothing for optimal control of nonsmooth dynamical systems. *arXiv preprint arXiv:2203.03986*, 2022.
- [10] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2019.
- [11] Matthew Loper and Michael Black. Opendr: An approximate differentiable renderer. 09 2014. doi: 10.1007/978-3-319-10584-0_11.
- [12] J Matyas et al. Random optimization. *Automation and Remote control*, 26(2):246–253, 1965.
- [13] Boris Polyak. *Introduction to Optimization*. 07 2020.
- [14] HJ Suh, Max Simchowitz, Kaiqing Zhang, and Russ Tedrake. Do differentiable simulators give better policy gradients? *arXiv preprint arXiv:2202.00817*, 2022.
- [15] Hyung Ju Terry Suh, Tao Pang, and Russ Tedrake. Bundled gradients through contact via randomized smoothing. *IEEE Robotics and Automation Letters*, 7(2):4000–4007, 2022.
- [16] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pages 5026–5033. IEEE, 2012.
- [17] Keenon Werling, Dalton Omens, Jeongseok Lee, Ioannis Exarchos, and C. Karen Liu. Fast and Feature-Complete Differentiable Physics Engine for Articulated Rigid Bodies with Contact Constraints. In *Proceedings of Robotics: Science and Systems*, Virtual, July 2021. doi: 10.15607/RSS.2021.XVII.034.