

Neural-Fields in Numerical Optimization with L4CasADi

Tim Salzmann¹, Jon Arrizabalaga¹, Lukas Pries¹ and Markus Ryll¹

Abstract—While real-world problems are often challenging to analyze analytically, deep learning excels in modeling complex processes from data. Existing optimization frameworks like CasADi facilitate seamless usage of solvers but face challenges when integrating learned process models into numerical optimizations. To address this gap, we present the *Learning for CasADi* (L4CasADi) framework, enabling the seamless integration of PyTorch-learned models with CasADi for efficient and potentially hardware-accelerated numerical optimization. The applicability of *L4CasADi* is demonstrated with two robotic tutorial examples: First, we optimize a fish’s trajectory in a turbulent river for energy efficiency where the turbulent flow is represented by a PyTorch model. Second, we demonstrate how an implicit Neural Radiance Field environment representation can be easily leveraged for optimal control with *L4CasADi*.

L4CasADi is available under MIT license at
<https://github.com/Tim-Salzmann/l4casadi>

I. INTRODUCTION

Accurate mathematical problem formulation is at the core of every numerical optimization procedure. While many real-world problems are hard to formulate analytically, data-driven methods, especially deep learning methods, thrive in modeling complex processes from data. Efficiently integrating learned process models into numerical optimizations is challenging, as data-driven models and numerical optimization come with their respective tools and characteristics: Deep learning models, commonly constructed in PyTorch [1], TensorFlow [2], or JAX [3], leverage first-order optimization methods (backpropagation) in time-expensive offline training on a large amount of data. Solutions to complex problems can be learned from data without defining any prior structure on the task. Once trained, however, their inference is fast and can be further accelerated on dedicated hardware (GPUs). Numerical optimization techniques, in contrast, employ second-order methods, primarily interior point and sequential quadratic programming, to formulate optimization problems, potentially subject to constraints. Solving these problems necessitates formulating a specific problem structure and proficiently selecting a suitable solver algorithm (IPOPT, SNOPT, qpOASES, OSQP, ECOS, etc.[4]–[8]). Aiming to remove the overhead of adapting the syntax for each solver, CasADi [9] allows seamless usage across them.

Thus, both PyTorch and CasADi have established themselves as prominent tools within their respective research domains, owing to their comprehensive functionalities, user-friendliness, and adaptability. However, the increasing importance of data-driven approaches in optimization poses a

challenge for CasADi as it lacks native support for learned functions.

We seek to close this gap by presenting the *Learning for CasADi* (L4CasADi) framework, which enables the seamless integration of learned PyTorch models with the numerical optimization framework CasADi [9]. *L4CasADi* enables hardware acceleration for learned components in a CasADi optimization.

II. L4CASADI - SYNTAX AND USAGE

L4CasADi was designed with three key desiderata in mind: (I) Simplicity for the user, (II) generalizability across PyTorch model architectures, and (III) efficiency in runtime. We will provide an insight into how these desiderata manifest within the user experience of *L4CasADi*.

Similarly to PyTorch, *L4CasADi* models are constructed in Python. Defining an *L4CasADi* model in Python given a pre-defined PyTorch model is as easy as

```
import l4casadi as l4c
# Construct L4CasADi Model from PyTorch Model
l4casadi_model = l4c.L4CasADi(
    pyTorch_model,
    device='cuda',          # Accelerator Device
    name='l4casadi_f'      # Unique name
)
```

where the architecture of the PyTorch model is unrestricted and large models can be accelerated with dedicated hardware. Once an *L4CasADi* model is defined it can be employed in a variety of settings, as will be outlined in the following.

Python

Once defined, an *L4CasADi* model can be seamlessly integrated with a CasADi symbolic graph within CasADi’s Python interface. On the first call, *L4CasADi* will automatically generate C++ code and compile the *L4CasADi* model for runtime efficiency.

```
# Use L4CasADi Model in CasADi Symbolic Graph
y: casadi.MX = l4casadi_model(x: casadi.MX)
```

The resulting symbolic output variable can be included in any further CasADi operations. This seamless integration empowers the utilization of CasADi’s extensive toolkit to formulate and solve optimization problems involving *L4CasADi* models.

```
# Minimize the L4CasADi Model using IPOPT
nlp = {'x': x, 'f': y}
solver = casadi.nlpsol("solver", "ipopt", nlp)
sol = solver()
```

III. TUTORIAL EXAMPLES

To showcase the practical application of the proposed framework and how it opens up new research avenues at ease, we present two illustrative case studies. First, we formulate a trajectory generator that finds the minimum energy path for a fish swimming upstream in a turbulent river. The second case study demonstrates how *LACasADi* facilitates the incorporation of cutting-edge computer vision models into optimization problems. In this case, we optimize a collision-free trajectory through an implicit environment representation given as a Neural Radiance Field (NeRF). These two case studies are distinguished by their simplicity and ease of comprehension, and thereby, serve as excellent templates for users to adapt to more intricate scenarios.

A. Fish Navigation in Turbulent Flow

The example shows how to design an optimization-based trajectory generator to navigate a turbulent fluid flow. The flow, and thus the fish's dynamics, are modeled by a Neural Network in PyTorch, while the optimal trajectory optimization problem is formulated in CasADi. By using *LACasADi*, both can be combined and optimized jointly. When doing so, we seek to find the minimum energy trajectory that allows a fish to navigate from the starting point to the goal. For this purpose, the fish needs to swim upstream a river where a circular stone causes the flow to be turbulent.

1) *Problem Formulation and Implementation:* We model the fish as a planar point mass actuated by velocity commands under the influence of the river's velocity field:

$$\dot{\mathbf{p}}(t) = \mathbf{v}(t) + \mathbf{v}_{fl}(t, \mathbf{p}(t)), \quad (1)$$

where $\mathbf{p}(t), \mathbf{v}(t) \in R^2$ are the position and velocity of the fish at time t and $\mathbf{v}_{fl}(t, \mathbf{p}(t)) \in R^2$ is the velocity of the river at time t and position $\mathbf{p}(t)$. In simpler terms, the fish's movement results from a combination of its own effort, and the influence of the turbulent flow. Eq. (1) can be written as a standard nonlinear dynamic system f , defined as

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t), t) = \mathbf{u}(t) + \mathbf{v}_{fl}(t, \mathbf{x}(t)), \quad (2)$$

whose states and inputs are the fish's position $\mathbf{x}(t) = \mathbf{p}(t)$ and velocity command $\mathbf{u}(t) = \mathbf{v}(t)$. To add realism to the problem, we constrain the fish's actuation by imposing box constraints on the velocity commands $\mathbf{u} \in [\underline{\mathbf{u}}, \bar{\mathbf{u}}]$, limit the trajectory to lie within the rivers bounds $[\underline{\mathbf{p}}, \bar{\mathbf{p}}]$ and constrain the trajectory to dodge the stone generating the turbulent flow.

To compute the trajectory that minimizes the fish's effort to reach the goal while swimming upstream, we formulate the

following Nonlinear Program (NLP):

$$\min_{\substack{\mathbf{x}_0, \dots, \mathbf{x}_N, \\ \mathbf{u}_0, \dots, \mathbf{u}_{N-1}}} \sum_{k=0}^{N-2} \left\| \frac{\mathbf{u}_{k+1} - \mathbf{u}_k}{\Delta t} \right\|^2 \quad (3a)$$

$$\text{s.t. } \mathbf{x}_0 = \mathbf{p}_0, \mathbf{x}_N = \mathbf{p}_f, \quad (3b)$$

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta t \cdot f(\mathbf{x}_k, \mathbf{u}_k, t_k), \quad (3c)$$

$$\underline{\mathbf{u}} \leq \mathbf{u}_k \leq \bar{\mathbf{u}}, \quad k = 0, \dots, N-1, \quad (3d)$$

$$\underline{\mathbf{p}} \leq \mathbf{x}_k \leq \bar{\mathbf{p}}, \quad k = 0, \dots, N, \quad (3e)$$

$$\|\mathbf{x}_k\|^2 \geq r_{st}^2, \quad k = 0, \dots, N. \quad (3f)$$

where \mathbf{p}_0 and \mathbf{p}_f are the initial and goal positions, Δt is the time step and r_{st} is the radius of the stone located at the origin. The cost function (3a) minimizes the differences in subsequent control (velocity) inputs which effectively minimizes the energy introduced into the system, constraint (3b) sets the starting and ending positions, (3c) enforces the system dynamics defined in (2), (3d) bounds the input commands, (3e) guarantees that the fish remains within the river's bounds and (3f) ensures that the fish does not collide with the stone. Notice that, due to the spatial bounds in (3f) and yet-to-be-defined turbulent flow model $\mathbf{v}_{fl}(\cdot)$ within the dynamic model $f(\mathbf{x}, \mathbf{u})$ in (3c), the NLP (3) is nonlinear and non-convex.

LACasADi allows for implementing the NLP in (3) in CasADi, modeling the turbulent flow $\mathbf{v}_{fl}(t, \mathbf{p}(t))$ by a Neural Network that has previously been trained in PyTorch. Given the nonlinear and non-convex structure of (3), the state-of-the-art interior-point solver IPOPT [4] is chosen as a solver.

2) *Outcome and Visualization:* The trajectory obtained from solving the NLP (3) is illustrated in Fig. 1 through a series of image sequences. The upper segment of figure (A) displays the magenta trajectory of the fish juxtaposed with the vorticity of the turbulent flow. Meanwhile, the middle segment (B) portrays the velocity field of the turbulent flow. In both instances, the stone is represented by a gray circle. The third segment (C) shows the velocity commands necessary to trace this trajectory. For a better insight, we encourage readers to view the accompanying animations for this example¹.

Fig. 1 reveals that the computed trajectory strategically utilizes the velocity field of the turbulent flow to minimize the energy required for reaching the objective. To put it simply, the object behaves like it's riding/surfing the river's current. Examining Fig. 1 enables us to comprehend this phenomenon from three distinct perspectives: Firstly, in (A), it becomes evident that the fish exploits the positive vortices (blue) to ascend and the negative ones (red) to descend. This observation is similarly evident in (B), where the fish adeptly positions itself per the velocity field, enabling it to exert minimal effort in the y-axis direction throughout the majority of the navigation, as depicted by the green line in (C).

¹ https://github.com/Tim-Salzmann/l4casadi/blob/main/examples/fish_turbulent_flow/README.md.

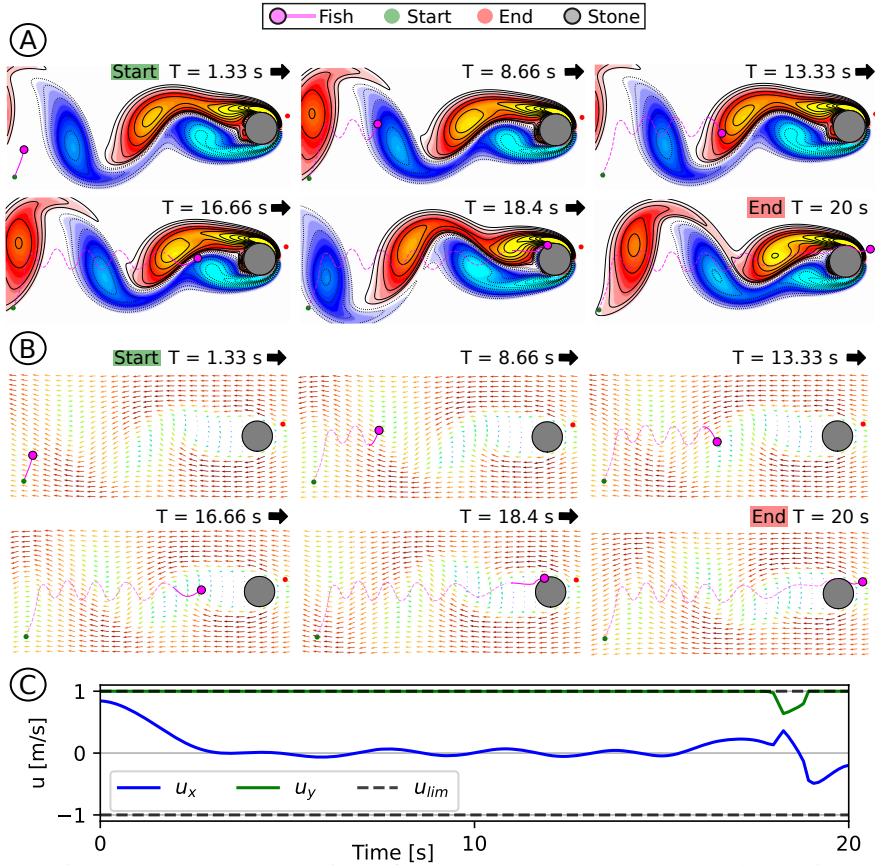


Fig. 1: Minimum energy trajectory (magenta) to navigate from the start (green) to the goal (red) in the presence of a turbulent flow represented by its vorticity (A) and velocity field (B). The trajectory generator is implemented in CasADi, while the flow is modeled by a Neural Network in PyTorch¹.

B. Trajectory Optimization in NeRFs

Neural Radiance Fields (NeRFs) are a powerful 3D representation technique that leverages deep learning to reconstruct 3D scenes from a collection of 2D images. NeRFs encode the scene’s geometry and appearance into a continuous function, allowing for the rendering of photorealistic images from arbitrary viewpoints. As a byproduct, the densities (translucencies) of the scene’s objects at any point in the 3D environment are implicitly captured in the network. In this example, we demonstrate that such state-of-the-art learned models from computer vision research can be easily incorporated into optimization procedures using *LACasADi*. To showcase this, we present the problem of finding a collision-free trajectory through the densities represented by a learned NeRF, where densities below a predefined threshold are deemed as unobstructed regions within the environment.

1) *Problem Formulation*: The trajectory to be planned is assumed to be given by a time-parametric polynomial of degree 9:

$$\mathbf{r}(\mathbf{c}, t) = \sum_{i=0}^9 \mathbf{c}_i t^i \quad (4)$$

where the parametric variable t is time and $\mathbf{c} \in \mathbb{R}^{9 \times 3}$ are the polynomial’s coefficients.

LACasADi enables the integration of a NeRF as an implicit environment representation into the optimization problem. For this purpose, we define the following function:

$$\rho = f_{\text{NeRF}}(\mathbf{p}) \quad (5)$$

where $\rho \in \mathbb{R}^+$ is the density of a location whose Euclidean coordinates are given by $\mathbf{p} \in \mathbb{R}^3$, i.e., it returns 0 for obstacle-free space while its value increases as the location becomes occupied.

Having defined the analytical expression of the trajectory in (4) and the NeRF-based environment representation in (5), we formulate a NLP that minimizes the curve’s snap, while ensuring that it remains collision-free:

$$\min_{\mathbf{c}} \sum_{k=0}^N \left\| \mathbf{r}^{(4)}(\mathbf{c}, t_k) \right\|^2 \quad (6a)$$

$$\text{s.t. } \mathbf{r}(\mathbf{c}, 0) = \mathbf{p}_0, \mathbf{r}^{(1)}(\mathbf{c}, 0) = \mathbf{0}, \mathbf{r}^{(2)}(\mathbf{c}, 0) = \mathbf{0}, \quad (6b)$$

$$\mathbf{r}(\mathbf{c}, T) = \mathbf{p}_f, \mathbf{r}^{(1)}(\mathbf{c}, T) = \mathbf{0}, \mathbf{r}^{(2)}(\mathbf{c}, T) = \mathbf{0}, \quad (6c)$$

$$\bar{\rho} > f_{\text{NeRF}}(\mathbf{r}(\mathbf{c}, t_k)), \quad k = 0, \dots, N, \quad (6d)$$

where $\mathbf{r}^{(n)}(\cdot)$ is the n -th time derivative of eq. (4), N is the number of evaluation points, T is the total time assigned to the trajectory. Constraints (6b) and (6c) define the starting and




Fig. 2: NeRF renders of three examples with a collision-free minimum snap optimized trajectory through the NeRF object. In each example, we vary the start point (blue) and goal point (green). Right: 2D slice with all three trajectories.

ending conditions and (6d) ensures that the NeRF density at all points along the trajectory is below a threshold $\bar{\rho}$ and thus collision-free.

This problem, in its original form without the NeRF, is well-studied in the planning and trajectory optimization community. Its mathematical structure allows it to be formulated as a quadratic programming (QP) problem, which can be solved efficiently. Being able to incorporate computer vision models in the problem formulation expands the scope of research possibilities but can simultaneously introduce greater complexity to the optimization task: The NeRF in (5) makes the resulting NLP highly non-convex due to the non-smooth density landscape as depicted in Section III-B.

2) *Implementation:* We utilize a pre-trained NeRF of a yellow Lego bulldozer, a commonly employed example in NeRF research. IPOPT [4] is again used as a solver. However, due to the complexity of the problem, we apply a two-stage approach. The first phase optimizes the trajectory to closely follow a sparse set of predetermined collision-free points, excluding the NeRF constraint in Eq. (6d). This results in a sub-optimal yet feasible trajectory. In the second phase, we initialize the solver with the first-phase solution and optimize for the full NLP in Eq. (6) including constraining all points to have a lower NeRF density than $\bar{\rho} = 1$. To ensure IPOPT utilizes the warmstarted trajectory in phase two, we initialize IPOPT with a small barrier parameter $\mu = 1e^{-4}$ (See Eq. (3a) in [4]). This compels IPOPT to remain within the region of feasible solutions from the beginning.

3) *Outcome and Visualization:* The optimized trajectory for three different configurations with varying start- and goal-points are shown in Fig. 2. For a better understanding of the problem space and the solution trajectories, we also visualize a 2D slice through the environment with the three optimal trajectories in Section III-B. The optimal minimal snap trajectories are clearly collision-free with the NeRF’s density representation. We encourage the reader to view the accompanying animations for this example².

IV. RELATED WORK

The increasing prevalence of data-driven models in optimization has spurred a surge of scholarly contributions, making it infeasible to comprehensively review the entire landscape of this domain within the scope of this discussion. Instead, we will concentrate on specific frameworks that integrate learned models with numerical optimization techniques, specifically PyTorch and CasADi.

a) *Learning and Numerical Optimization:* Multiple approaches have been proposed to bring advanced numerical optimization algorithms within learning frameworks like PyTorch and TensorFlow [10], [11]. However, these integrations often result in the re-implementation of individual solver algorithms within the learning framework, which may not exhibit the same robustness and maturity as the well-established solvers within CasADi. Recently, NeuroMANCEr [12] enables the formulation of optimization problems entirely in PyTorch. Naturally, learned components in PyTorch can be included, even optimized within the problem solution. However, because NeuroMANCEr is restricted to the PyTorch environment they are mostly limited to the first-order solvers within — not using second-order approaches such as IPOPT or SQP which have been proven to be efficient and robust.

b) *PyTorch and CasADi:* The demand for such a framework within the community is evident from prior endeavors aimed at uniting the two concepts of learning in PyTorch and optimization in CasADi. Multiple packages, *do-mpc* [13], *HILO-MPC* [14], and *ML-CasADi* [15], provide the capabilities to rebuild simple architectures of PyTorch models directly in CasADi by copying the learned weight tensors and formulate matrix multiplications and activation functions in CasADi. This approach, is restricted to models comprised of the limited CasADi function set which is not optimized for large matrix multiplications and can not use hardware acceleration. Additionally, [15] calculates function evaluations and sensitivities for the learned model separated from CasADi in the PyTorch framework and subsequently injects these results into the CasADi graph. The applicability of this approach, however, is limited to specific optimization algorithms and introduces inefficiencies in the form of excessive context switches and memory transfers between CasADi and PyTorch.

²https://github.com/Tim-Salzmann/l4casadi/blob/main/examples/nerf_trajectory_optimization/README.md.

REFERENCES

- [1] A. Paszke, S. Gross, F. Massa, *et al.*, “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” *Advances in Neural Information Processing Systems*, Dec. 2019.
- [2] M. Abadi, A. Agarwal, P. Barham, *et al.*, “TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems,” Mar. 2016.
- [3] J. Bradbury, R. Frostig, P. Hawkins, *et al.*, *JAX: composable transformations of Python+NumPy programs*, 2018. [Online]. Available: <http://github.com/google/jax>.
- [4] A. Wächter and L. T. Biegler, “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming,” *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, Mar. 2006, ISSN: 0025-5610. DOI: 10.1007/s10107-004-0559-y. [Online]. Available: <http://link.springer.com/10.1007/s10107-004-0559-y>.
- [5] P. E. Gill, W. Murray, and M. A. Saunders, “SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization,” *SIAM Review*, vol. 47, no. 1, pp. 99–131, Jan. 2005, ISSN: 0036-1445. DOI: 10.1137/S0036144504446096. [Online]. Available: <http://epubs.siam.org/doi/10.1137/S0036144504446096>.
- [6] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, “qpOASES: a parametric active-set algorithm for quadratic programming,” *Mathematical Programming Computation*, vol. 6, no. 4, pp. 327–363, Dec. 2014, ISSN: 1867-2949. DOI: 10.1007/s12532-014-0071-1. [Online]. Available: <http://link.springer.com/10.1007/s12532-014-0071-1>.
- [7] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, “OSQP: an operator splitting solver for quadratic programs,” *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, Dec. 2020, ISSN: 1867-2949. DOI: 10.1007/s12532-020-00179-2. [Online]. Available: <http://link.springer.com/10.1007/s12532-020-00179-2>.
- [8] A. Domahidi, E. Chu, and S. Boyd, “ECOS: An SOCP solver for embedded systems,” in *2013 European Control Conference (ECC)*, IEEE, Jul. 2013, pp. 3071–3076, ISBN: 978-3-033-03962-9. DOI: 10.23919/ECC.2013.6669541. [Online]. Available: <https://ieeexplore.ieee.org/document/6669541/>.
- [9] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, “CasADi: a software framework for nonlinear optimization and optimal control,” *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, Mar. 2019, ISSN: 1867-2949. DOI: 10.1007/s12532-018-0139-4. [Online]. Available: <http://link.springer.com/10.1007/s12532-018-0139-4>.
- [10] B. Amos, I. D. J. Rodriguez, J. Sacks, B. Boots, and J. Z. Kolter, “Differentiable MPC for End-to-End Planning and Control,” in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, ser. NIPS’18, Red Hook, NY, USA: Curran Associates Inc., 2018, pp. 8299–8310.
- [11] W. Wang, Y. Zhang, J. Gao, *et al.*, “GOPS: A general optimal control problem solver for autonomous driving and industrial control applications,” *Communications in Transportation Research*, vol. 3, p. 100096, Dec. 2023, ISSN: 27724247. DOI: 10.1016/j.commtr.2023.100096.
- [12] J. Drgona, A. Tuor, J. Koch, M. Shapiro, and D. Vrabie, “NeuroMANCEr: Neural Modules with Adaptive Nonlinear Constraints and Efficient Regularizations,” 2023. [Online]. Available: <https://github.com/pnnl/neuromancer>.
- [13] F. Fiedler, B. Karg, L. Lüken, *et al.*, “do-mpc: Towards FAIR nonlinear and robust model predictive control,” *Control Engineering Practice*, vol. 140, p. 105676, Nov. 2023, ISSN: 09670661. DOI: 10.1016/j.conengprac.2023.105676.
- [14] J. Pohlodek, B. Morabito, C. Schlauch, P. Zometa, and R. Findeisen, *Flexible development and evaluation of machine-learning-supported optimal control and estimation methods via HILO-MPC*, 2022. DOI: 10.48550/ARXIV.2203.13671.
- [15] T. Salzmann, E. Kaufmann, J. Arrizabalaga, M. Pavone, D. Scaramuzza, and M. Ryll, “Real-Time Neural MPC: Deep Learning Model Predictive Control for Quadrotors and Agile Robotic Platforms,” *IEEE Robotics and Automation Letters*, vol. 8, no. 4, pp. 2397–2404, Apr. 2023, ISSN: 2377-3766. DOI: 10.1109/LRA.2023.3246839. [Online]. Available: <https://ieeexplore.ieee.org/document/10049101/>.