Analytical Integral Global Trajectory Optimization with Contact Dynamics

Davide De Martini
University of Trento
davide.demartini@unitn.it

Majid Khadiv
Technical University of Munich
majid.khadiv@tum.de

Andrea Del Prete
University of Trento
andrea.delprete@unitn.it

Abstract—**Trajectory Optimization (TO) involves designing trajectories by minimizing a cost function subject to constraints, often formulated as an Optimal Control Problem (OCP). Traditional methods, such as Gradient Descent or sampling-based approaches, can converge to poor local minima due to nonconvexity or high dimensionality. In this work, we explore Analytical Integral Global Optimization (AIGO), a novel algorithm inspired by Randomized Smoothing (RS), for TO. Unlike RS, which relies on sampling, AIGO computes the smoothed objective analytically using unitary kernels (e.g., hyperboxes) and optimizes over the integral of the function, progressively shrinking the integration domain until the original problem is recovered. We evaluate AIGO with a point-mass with collisions demonstrating its ability to reliably find feasible and efficient trajectories where traditional methods struggle.**

I. INTRODUCTION

Trajectory Optimization (TO) is a popular framework for solving certain classes of robotics problems, such as the offline generation of reference trajectories and the online control of different robotic systems. TO is especially appreciated for its versatility, as users can define an arbitrary cost function to be minimized, while satisfying arbitrary constraints, typically used to represent physical limits of the robot, such as its torque limits and collision avoidance.

Gradient-based methods [1] are commonly used for solving such problems, leveraging efficient computation of function derivatives to quickly find optimal solutions. However, these methods are inherently local and can converge to poor local minima on non-convex problems—common in robotics due to system dynamics and collision-avoidance constraints [2]—often converging to poor trajectories. This is especially the case if the robot must make/break contacts with the environment, as rigid contacts make the whole problem non-differentiable [3], preventing the use of gradient-based solvers. However, smoothed contact models have been used to approximate rigid contacts [4], recovering the ability to use gradient-based solvers, at the cost of introducing stiff ordinary differential equations that are more challenging to integrate.

An interesting alternative is to use gradient-free (zeroth-order) optimization, which avoids gradient computation by relying on sampling. However, sampling becomes inefficient in high-dimensional spaces, typical of robotics problems. To overcome these limitations, we explore Analytical Integral Global Optimization (AIGO) [5], a recently proposed global optimization algorithm.

AIGO takes its roots from Randomized Smoothing (RS) [6], a zeroth-order optimization technique that transforms a potentially irregular or non-convex objective function into a smoother surrogate function. This is achieved by convolving the original function with a probability distribution, typically Gaussian. The convolution process effectively averages the function values around each point, resulting in a smoother landscape that is easier to optimize. Gradients of this smoothed function can then be estimated through sampling, allowing optimization to proceed even when the original function is non-differentiable or highly non-convex.

Recent work applied the concept of RS to TO problems, achieving promising results [7], [8]. However, such methods compute approximately the convolution between the cost function and a smoothing kernel (typically Gaussian) using sampling. Therefore, they suffer from the curse of dimensionality and tend to be inefficient in high dimensions.

AIGO emulates RS in an analytical manner. By using unitary kernels (i.e., uniform probability on a hyperbox), the convolution can be computed in closed-form as an analytical integral, eliminating the challenges associated with sampling. In this setting, the minimization objective becomes the integral of the function rather than the original function itself. The integration set is then progressively shrunk in successive steps until the original optimization problem is recovered.

When applying AIGO to TO problems, several adjustments are required. The current version of AIGO does not handle constraints; therefore, the OCP must be formulated as an unconstrained problem. Moreover, due to the need for analytical integration of the cost function, polynomials were selected as they enable efficient closed-form computation, while retaining the ability to approximate a large class of functions to arbitrary accuracy.

This requirement significantly impacts the overall formulation of the problem, including the choice of transcription method used to discretize the OCP.

The main contribution of this work is the development of a practical formulation that, based on AIGO principles, enables the solution of TO problems involving contact dynamics.

We present preliminary results on a one-dimensional point-mass interacting with a wall. The results demonstrate that the proposed method consistently identifies improved solutions compared to classical gradient-based approaches across different initial conditions and goal configurations.

II. BACKGROUND

A. Problem Formulation

We consider the class of optimal control problems formulated as:

$$\begin{aligned} \min_{x(\cdot),\, u(\cdot),\, \lambda(\cdot)} \quad & \int_0^T \ell(x(t), u(t), \lambda(t))\; dt + \ell_T(x(T)) \\ \text{s.t.} \quad & M(q)\ddot{q} + h(q, \dot{q}) = S^\top u + J_c(q)^\top \lambda, \\ & 0 \le \phi_i(q) \;\perp\; \lambda_i \ge 0, \quad i = 1, \dots, m, \\ & x(0) = x_0. \end{aligned} \tag{1}$$

The state is defined as $x = [q, \dot{q}]$, where $q \in \mathbb{R}^n$ denotes the generalized coordinates and $\dot{q}$ their time derivatives. The control input $u \in \mathbb{R}^{n_u}$ is given by the joint torques.

The matrix $M(q) \in \mathbb{R}^{n\times n}$ is the mass matrix, while $h(q, \dot{q}) \in \mathbb{R}^n$ collects the Coriolis, centrifugal, and gravitational forces. The matrix $S \in \mathbb{R}^{n_u \times n}$ maps actuator torques to generalized forces, accounting for possible underactuation. The matrix $J_c(q) \in \mathbb{R}^{m\times n}$ denotes the contact Jacobian, and $\lambda \in \mathbb{R}^m$ represents the contact forces.

The function $\phi(q) \in \mathbb{R}^m$ denotes the signed distance between the contact points and the environment. The complementarity conditions

$$0 \le \phi_i(q) \;\perp\; \lambda_i \ge 0, \quad i = 1, \dots, m,$$

impose that contact forces can only be nonzero when the corresponding contact is active. For sake of simplicity, we do not consider friction limits in this work.

B. Analytical Integral Global Optimization

Analytical Integral Global Optimization (AIGO) [5] is a recent algorithm that aims to leverage the intuitions from Gaussian Smoothing in an analytical manner, eliminating the need for sampling and mitigating the curse of dimensionality.

Its focus is solving unconstrained optimization problems of the form

$$x^\star = \operatorname*{argmin}_x \; f(x)$$

where $f(\cdot) : \mathbb{R}^n \to \mathbb{R}$ is a continuous function, $n \in \mathbb{N}_+$ is the dimension of the preimage of f, and $x^\star \in \mathbb{R}^n$ is the global minimizer of f.

One-dimensional case. Let us introduce first the one-dimensional case, and then expand it to the *n*-dimensional case. AIGO's key idea is to minimize the integral of f over an interval with fixed width $2w$ and variable center c.

$$\operatorname*{minimize}_c \underbrace{\int_{c-w}^{c+w} f(x)\; \mathrm{d}x}_{I(c,w)} .$$

The problem is first solved for a large value of w, at each iteration this half width is decreased and the previous solution is used for warm-starting the solver. The algorithm converges when $w \approx 0$ is reached.

***N*-dimensional case.** In the N-dimensional case, the minimization must be performed over both c and w, since the objective depends on both the position and the shape of the hyper-rectangle $A(c, w)$. Geometrically, this is because different choices of w can produce regions with the same volume but different extents along each dimension, thereby including or excluding different portions of the integrand. As a result, the value of the integral changes not only with the location c, but also with how the region is distributed across the coordinate axes. Therefore, the minimization problem becomes:

$$\operatorname*{minimize}_{c,w} \; I(c, w) \qquad \text{subject to} \qquad S(w) = s$$

where the integral of the cost function $I(c, w)$ and the size of the integration set $S(w)$ are defined by

$$I(c, w) = \int_{A(c,w)} f(x)\; \mathrm{d}^n x, \quad \text{and} \quad S(w) = \prod_{i=1}^{n} (2w_i).$$

In [5], a formal proof of global optimality is established for the one-dimensional case, whereas in higher dimensions the global behavior is supported empirically.

It is worth highlighting that, for the algorithm to be efficient, the cost function must be analytically integrable, as the computation of $I(c, w)$ and its gradient relies on closed-form expressions of these integrals.

III. MAIN CONTRIBUTION

This section presents the main contributions of this work, namely a novel formulation of trajectory optimization problems within the AIGO framework. Due to the unconstrained nature of AIGO, several modifications to the standard optimal control problem formulation are required. In particular, we discuss the design of suitable penalty functions for handling constraints, the choice of transcription method, and the resulting algorithmic structure. These elements collectively enable the application of AIGO to trajectory optimization problems with contact dynamics.

A. Penalty formulation

Constraints in trajectory optimization are often incorporated into the objective function via penalty methods. Equality constraints can be translated to quadratic penalties. Inequality constraints require additional structure and are commonly handled using techniques such as barrier functions or ReLU-squared penalties. Barrier functions penalize solutions that approach the boundary of the feasible set, effectively preventing constraint violation by making the cost grow rapidly near infeasible regions. In contrast, ReLU-squared penalties apply no cost when the constraint is satisfied and introduce a quadratic penalty only when violations occur.

B. Transcription method

This section describes the process that drove us to choosing the best transcription method.

Single Shooting. Single shooting is a discretization method where only the control inputs $u(t)$ are optimized, while the state trajectory is obtained by propagating the system

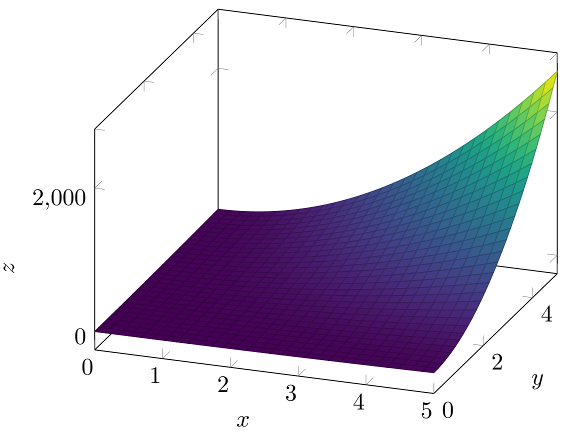

Fig. 1. Plot of the penalty function $(((x+w)^3-(x-w)^3)\cdot((y+w)^3-(y-w)^3))$, in this case w is set to 1. The structure of this function pushes the solution towards zero.

dynamics from the initial state. As a result, any cost or constraint involving the state must be rewritten as a function of the control sequence, e.g., $x_f = f_{\text{term}}(x_0, \mathbf{u})$. This leads to complex expressions due to the recursive dependence of states on inputs, which can become difficult or intractable to integrate analytically.

Collocation. Collocation is a discretization method in which both the state and input trajectories are treated as decision variables, improving stability and avoiding the need to express states as functions of the control sequence. In this framework, additional variables such as contact forces can also be included, and the main constraints become the system dynamics and complementarity conditions.

However, issues can arise in the formulation. In this case, the complementarity constraint $\phi(q)\cdot\lambda = 0$, when converted into a squared penalty and integrated, leads to a cost function whose minimum occurs at zero, as illustrated in Fig. 1. This causes the optimizer to converge to the trivial solution where all variables are zero, which is physically infeasible and compromises the validity of the optimization.

Collocation without forces. Collocation allows explicit inclusion of the system state, overcoming single shooting limitations, but introduces challenges from the complementarity constraint. An alternative is to retain collocation without enforcing it explicitly, using the implicit relation from the relaxed complementarity:

$$\phi(q)\cdot\lambda = \varepsilon \quad \Rightarrow \quad \lambda = \frac{\varepsilon}{\phi(q)},$$

where $\varepsilon > 0$ is a user-defined relaxation parameter. This removes λ from the decision variables and the complementarity penalty from the cost function, enabling any trajectory optimization problem to be solved with AIGO, provided all cost components are analytically integrable.

C. The algorithm

The implemented algorithm largely follows the structure of AIGO [5]. The principal difference lies in the optimization strategy: rather than employing a Block Coordinate Descent approach, updates are computed using the full gradient of the problem. This modification is motivated by an empirical evaluation of the original AIGO algorithm applied to trajectory optimization problems, which demonstrated suboptimal performance, with the optimizer frequently converging to local minima.

This behavior can be attributed to the fact that the centers and half-widths of the hyperboxes must be optimized jointly, as they mutually influence one another. The resulting procedure is outlined in Algorithm 1: s_{thr} denotes the convergence threshold, β and γ are hyperparameters introduced to ensure the monotonic decrease of the set size, and ζ represents the penalty associated with the soft constraint $S(w) = s$.

Algorithm 1: AIGO

```
1 Function aigo($f, c, w, s_{thr}, \beta, \gamma, \zeta$):
2     $s \leftarrow S(w)$;
3     while $s > s_{thr}$ do
4         $c, w \leftarrow \text{minimize}_{c,w}\ I_p(c, w, s, \zeta)$;
5         $w \leftarrow w - \beta|\dot{c}|$;
6         $s \leftarrow \min(S(w), s/\gamma)$;
7     return $c$;
```

The algorithm proceeds as follows. At line 2, the current size (hypervolume) of the integration set is computed as $s = S(w)$. The loop in line 3 continues until the size falls below a prescribed threshold s_{thr}. At each iteration (line 4), the objective $I_p(c, w, s, \zeta)$—which includes both the cost function and a penalty term on the set size—is minimized with respect to the centers c and half-widths w.

Subsequently, the half-widths are updated (line 5) proportionally to the variation of the centers, inducing a contraction of the integration domain. In line 6, the size is updated as $s = \min(S(w), s/\gamma)$, ensuring a controlled and monotonic shrinking independent of the optimization step. The procedure iterates until the size of the integration set is sufficiently small, at which point the final centers c are returned.

IV. RESULTS

To evaluate the trajectory optimization capabilities of AIGO, we consider a simple system: a one-dimensional point mass interacting with an idealized wall.

The resulting optimal control problem is

$$\begin{aligned} \underset{x(\cdot),u(\cdot)}{\text{minimize}} \quad & \int_{t_0}^{t_f} w_u\, u(t)^2\, dt + w_f\,(x(t_f) - x_{\text{des}})^2 \\ \text{subject to} \quad & \dot{x}(t) = f\big(x(t), u(t), \tfrac{\varepsilon}{x(t)}\big), \quad \forall t \in [t_0, t_f], \\ & x(0) = x_{\text{init}}. \end{aligned}$$

The cost penalizes control effort and the terminal deviation from the desired state.

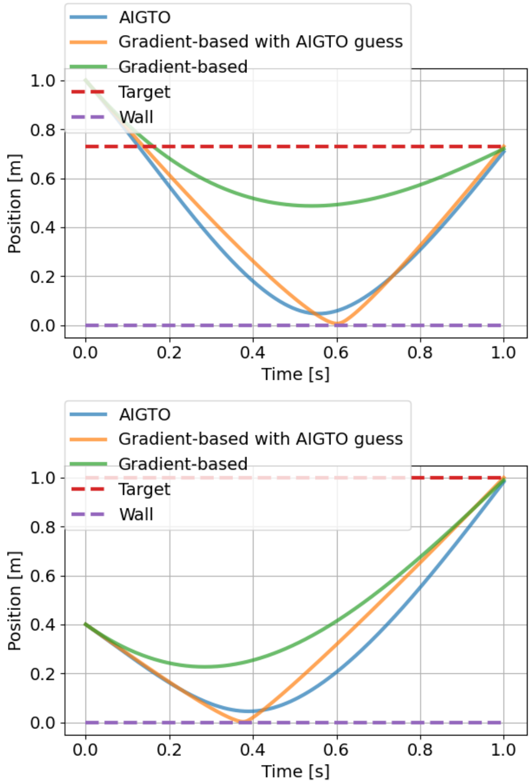

Fig. 2. In light blue the solution retrieved by AIGTO. In orange the solution found by the gradient solver with the initial guess provided by the AIGTO solution. Lastly in green the solution retrieved by a classic gradient based solver with a random initial guess.

The system is modeled as a double integrator with an additional contact force defined via a relaxed complementarity formulation. The discrete-time equations of motion are given by

$$\begin{cases} q_{k+1} = q_k + dt\,\dot{q}_k \\ \dot{q}_{k+1} = \dot{q}_k + dt\,(u_k + \lambda_k) \\ \lambda_k = \dfrac{\varepsilon}{q_k + \delta}, \end{cases}$$

where $q_k \in \mathbb{R}$ is the generalized position, $\dot{q}_k \in \mathbb{R}$ is the generalized velocity, $u_k \in \mathbb{R}$ is the control input, and $\lambda_k \in \mathbb{R}$ is the contact force. The parameter $\varepsilon > 0$ is a relaxation coefficient, and $\delta > 0$ is a small regularization term introduced for numerical stability. The subscript k denotes the discrete time index. The integration time step dt is equivalent to 5 ms.

Fig. 2 shows two representative results. In both cases, the gradient-based solver fails to exploit contact with the wall, leading to a deceleration–acceleration strategy and final costs of $1.19 \cdot 10^{-2}$ and $1.13 \cdot 10^{-2}$.

In contrast, AIGTO exploits contact effectively, minimizing control effort. When used to initialize the gradient-based solver, it yields significantly improved solutions, with final costs of $3.65 \cdot 10^{-5}$ and $8.7 \cdot 10^{-5}$, respectively.

Fig. 3 shows other reults achieved with different configurations. Also in these cases our method was able to retrieve better solutions compared to the gradient based approaches. Table I wraps up all the results presented displaying the different configurations and the final costs achieved.

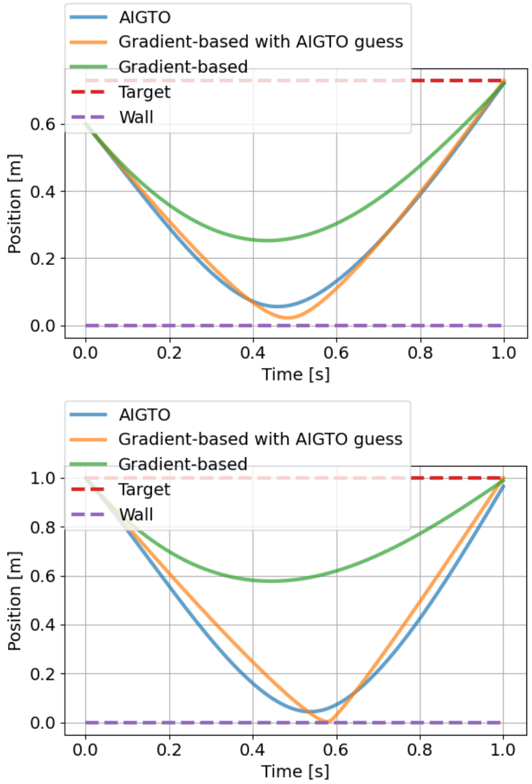

Fig. 3. Other configurations with different initial conditions, final conditions and relaxation variable ε.

(p_0, v_0)	p_f	ε	Grad. cost	AIGTO cost
$(1.0, -2.0)$	0.73	0.1	1.19×10^{-2}	3.65×10^{-5}
$(0.4, -1.2)$	1.0	0.1	1.13×10^{-2}	8.7×10^{-5}
$(0.6, -1.5)$	1.0	0.05	7.1×10^{-2}	5.42×10^{-5}
$(1.0, -2.0)$	1.0	0.02	1.13×10^{-2}	3.67×10^{-7}

TABLE I
COMPARISON OF GRADIENT-BASED OPTIMIZATION AND AIGTO ACROSS DIFFERENT INITIAL CONDITIONS.

V. CONCLUSIONS AND DISCUSSIONS

This work investigates the application of a recent global optimization algorithm (AIGO) to trajectory optimization (TO) problems with contact dynamics. Results show that AIGO effectively handles high-dimensional nonconvex objectives, outperforming classical methods in convergence speed while achieving near state-of-the-art solutions. In TO, AIGO yields solutions with objective values up to three orders of magnitude lower than gradient-based methods.

Future work will focus on evaluating robustness on more complex systems and integrating AIGO into a model predictive control (MPC) framework to provide high-quality initializations for nonconvex problems.

REFERENCES

- [1] T. A. Howell, S. Le Cleac'H, S. Singh, P. Florence, Z. Manchester, and V. Sindhwani, "Trajectory optimization with optimization-based dynamics," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, p. 6750 – 6757, 2022, cited by: 23.
- [2] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter, "Perceptive locomotion through nonlinear model-predictive control," *IEEE Transactions on Robotics*, vol. 39, no. 5, p. 3402 – 3421, 2023, cited by: 200; All Open Access, Green Open Access.
- [3] P. M. Wensing, M. Posa, Y. Hu, A. Escande, N. Mansard, and A. Del Prete, "Optimization-based control for dynamic legged robots," *IEEE Transactions on Robotics*, vol. 40, pp. 43–63, 2023.
- [4] E. Todorov, "A convex, smooth and invertible contact model for trajectory optimization," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 1071–1076.
- [5] S. Labbe and A. Del Prete, "Analytical integral global optimization," *Proceedings of Machine Learning Research vol*, vol. 283, pp. 1–12, 2025.
- [6] K. Gao and O. Sener, "Generalizing gaussian smoothing for random search," in *International Conference on Machine Learning*. PMLR, 2022, pp. 7077–7101.
- [7] Q. Le Lidec, F. Schramm, L. Montaut, C. Schmid, I. Laptev, and J. Carpentier, "Leveraging randomized smoothing for optimal control of nonsmooth dynamical systems," *Nonlinear Analysis: Hybrid Systems*, vol. 52, p. 101468, 2024.
- [8] H. J. T. Suh, T. Pang, and R. Tedrake, "Bundled gradients through contact via randomized smoothing," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4000–4007, 2022.