

Probabilistic Homotopy Optimization for Dynamic Motion Planning

Shayan Pardis*¹ Matthew Chignoli*² and Sangbae Kim²

Abstract—We present a homotopic approach to solving challenging, optimization-based motion planning problems. The approach uses Homotopy Optimization, which, unlike standard continuation methods for solving homotopy problems, solves a sequence of constrained optimization problems rather than a sequence of nonlinear systems of equations. The insight behind our proposed algorithm is formulating the discovery of this sequence of optimization problems as a search problem in a multidimensional homotopy parameter space. Our proposed algorithm, the Probabilistic Homotopy Optimization algorithm, switches between solve and sample phases, using solutions to easy problems as initial guesses to more challenging problems. We analyze how our algorithm performs in the presence of common challenges to homotopy methods, such as bifurcation, folding, and disconnectedness of the homotopy solution manifold. Finally, we demonstrate its utility via a case study on two dynamic motion planning problems: the cart-pole and the MIT Humanoid.

I. INTRODUCTION

The ability of optimization-based motion planners to generalize to unseen situations and reason about challenging constraints makes them powerful tools. Highly efficient motion planners can run in real-time for MPC [1], [2], while less efficient motion planners are still valuable as sources of high-quality reference data for training reinforcement learning policies [3]–[5]. Highly dynamic motions such as the back flip shown in Fig. 1 often have a highly non-convex cost landscape [6], which makes it difficult to converge to a solution reliably. This work introduces a systematic approach to solving such challenging nonlinear trajectory optimization problems. Our approach involves defining a homotopy between simple and challenging motion planning problems and then probabilistically searching the potentially multidimensional homotopy parameter space in a manner robust to common pitfalls such as bifurcation and disjoint solution manifolds. We demonstrate our method’s ability to solve challenging nonlinear optimization problems that are not solvable with off-the-shelf globalization strategies.

Homotopy methods [7] are a common approach to solving challenging Nonlinear Programs such as those arising from robotic motion planning. The Interior Point Method upon which many modern Nonlinear Program (NLP) solvers are built [8], [9] involves a homotopy over the nonlinear system of equations corresponding to the optimality conditions, where the homotopy parameter encodes a relaxation of the problem constraints. While this method has proven to be very effective for solving a wide variety of challenging

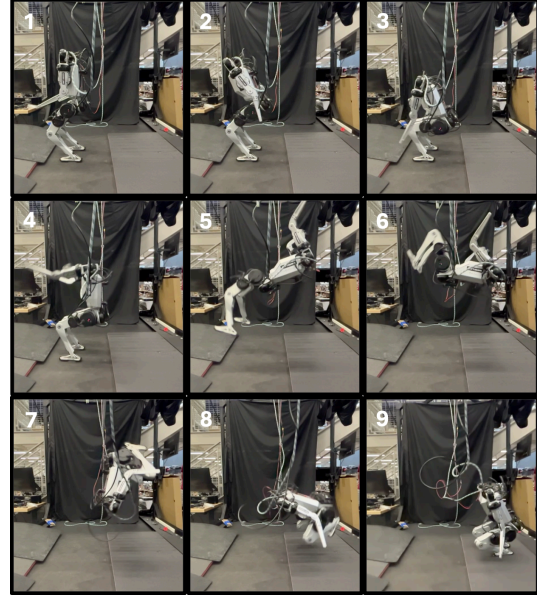


Fig. 1. Framewise animation of the back flip motion planned using Probabilistic Homotopy Optimization.

Nonlinear Programs [8], [9], restricting the homotopy to the optimality relaxation parameter does not allow domain knowledge about the problem or its structure to be leveraged. The Homotopy Optimization Method [10] can exploit such domain knowledge because it involves solving a sequence of constrained minimization problems rather than systems of nonlinear equations. Versions of the Homotopy Optimization Method have been employed for robotic motion planning [11]–[14], but have always pre-specified the full sequence of optimizations to solve. In contrast, our proposed algorithm eliminates the need for pre-specification of the sequence.

While the Homotopy Optimization Method is more general than homotopy over optimality relaxation, it can lead to a more challenging homotopic landscape due to bifurcations, folds, abbreviated paths, and disconnected branches [15]. Adding homotopy parameters can alleviate the folding problem [16]. However, adding homotopy parameters necessitates advanced methods for tracing a solution in the now multidimensional homotopy parameter space [17], [18]. These multidimensional tracing methods require gradients of the homotopy map with respect to the problem variables and homotopy parameters, which are not readily available in Homotopy Optimization.

In this work, we propose Probabilistic Homotopy Optimization (P-HO) algorithm for solving challenging NLPs

*Equal contribution.

¹Department of Electrical Engineering and Computer Science and
²Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA: chignoli@mit.edu

corresponding to dynamic motion planning for robotic systems. Like the stochastic extension of the Homotopy Optimization Method [10], our approach involves stochastically selecting and then solving a sequence of minimization problems ranging from easy to difficult. However, we offer a generalized formulation that (i) enables more flexible homotopy definitions, (ii) requires less hyperparameter tuning, and (iii) improves explorations to find better local minima. Additionally, our approach does not require gradients of the homotopy map to trace a solution. The crucial features of our homotopy method, and thus the main contributions of the paper, are

- Formulating the traversal of the homotopy’s Solution Manifold as a search problem
- A novel sampling-based algorithm to efficiently traverse the Solution Manifold
- Support for multidimensional homotopy parameterizations

P-HO’s method of growing the tree eliminates the need to tune hyperparameters related to the pre-specification the homotopy path. Furthermore, for each parameterization in homotopy space, P-HO solves the problem multiple times with different initial guesses, so it can explore multiple local minimas and choose the best among them.

Lastly, the ability to define multidimensional homotopy parameters allows the algorithm, rather than the user, to find the rate at which the difficulty of different aspects of the problem is increased. We demonstrate the capabilities of our proposed homotopy method by implementing it to solve a cart pole swing-up problem and multiple highly dynamic motion planning problems involving the MIT Humanoid robot.

The rest of the paper is organized as follows. Background information on relevant mathematical concepts and notation is provided in Sec. II. Sec. III formalizes our problem statement, describes the P-HO algorithm, and analyzes its properties. For comparison, two alternatives to the P-HO algorithm are presented in Sec. IV. Sec. V presents the results from case studies using the proposed algorithms to solve dynamic motion planning problems, and Sec. VI summarizes and concludes the work.

II. BACKGROUND

A. Trajectory Optimization

Nonlinear programming involves solving a problem of the form:

$$\begin{aligned} \mathcal{P}_\theta = \min_{\mathbf{x} \in \mathcal{X}} \quad & f(\mathbf{x}, \theta) \\ \text{s.t.} \quad & \mathbf{c}(\mathbf{x}, \theta) = \mathbf{0} \\ & \mathbf{x} \geq \mathbf{0} \end{aligned} \quad (1)$$

where $\mathbf{x} \in \mathcal{X}$ are the decision variables, $f : \mathcal{X} \rightarrow \mathbb{R}$ is an objective function, $\mathbf{c} : \mathcal{X} \rightarrow \mathbb{R}^m$ is a set of equality constraints, and $\theta \in \mathcal{Q}$ is a set of parameters whose value is fixed prior to solving. The feasible set of decision variables $\mathcal{C} \subset \mathcal{X}$ contains all decision variables satisfying $\mathbf{c}(\mathbf{x}, \theta) = \mathbf{0}$

and $\mathbf{x} \geq \mathbf{0}$. Any nonlinear program can be converted to the form of (1) through slack variables [19].

In trajectory optimization, the goal is to find the state and control trajectories that optimally accomplish a desired task. Thus, the decision variables encode these state and control trajectories, the cost encodes what it means to “optimally” accomplish the task, and the constraints encode the dynamics of the system as well as any other task-specific path or boundary constraints.

Analytical solutions for such trajectory optimization problems are typically intractable. Thus, due to the ready availability of local gradient information, gradient-based numerical methods such as Sequential Quadratic Programming [7], Interior Point Methods [7], and Differential Dynamic Programming [20] are used to solve these problems. We define the following operator to represent solving (1) with a numerical solver:

$$\text{solve}(\mathcal{P}_\theta, \mathbf{x}) : \mathcal{X} \rightarrow \mathcal{X}, \quad (2)$$

where \mathcal{P}_θ is the NLP being solved and \mathbf{x} is the initial point used by the solver. The solver will return a locally optimal solution $\mathbf{x}^* \in \mathcal{C}$ if and only if the initial point is in the basin of attraction of \mathbf{x}^* . Otherwise, it will return an unusable, infeasible solution not in \mathcal{C} .

Definition 1 (Basin of Attraction). *The set of all initial guesses \mathbf{x} for which the numerical solver finds a locally optimal solution \mathbf{x}^* to the problem \mathcal{P}_θ .*

$$\mathcal{B}_\theta(\mathbf{x}^*) = \{\mathbf{x} : \text{solve}(\mathcal{P}_\theta, \mathbf{x}) = \mathbf{x}^*\}.$$

Note that the behavior of $\text{solve}(\mathcal{P}_\theta, \cdot)$ and the Basin of Attraction will depend on the solver used. Our proposed approach will work for any solver, but results will vary.

B. Homotopy Methods

Homotopy methods are numerical techniques that aim to solve a challenging nonlinear problem $F(\mathbf{x})$ by gradually transforming an easy problem $F_{\text{easy}}(\mathbf{x})$ into $F(\mathbf{x})$ via a homotopy [7]. The homotopy defines a continuous relationship between $F_{\text{easy}}(\mathbf{x})$ and $F(\mathbf{x})$ as a function of the homotopy parameter $\lambda \in \mathcal{H}_d$, where

$$\mathcal{H}_d = \left\{ \lambda = \begin{bmatrix} \lambda_1 \\ \vdots \\ \lambda_d \end{bmatrix} : 0 \leq \lambda_i \leq 1 \forall i = 1, \dots, d \right\}.$$

Typically, $d = 1$, but the homotopy parameter can be multidimensional, as will be the case in this work. A homotopy map $H(\mathbf{x}, \lambda)$ is defined such that when $\lambda = \mathbf{0}_d$, the solution to $H(\mathbf{x}, \mathbf{0}_d)$ is easily obtained solution to $F_{\text{easy}}(\mathbf{x})$, and the solution to $H(\mathbf{x}, \mathbf{1}_d)$ is the solution to $F(\mathbf{x})$. Here, $\mathbf{0}_d$ and $\mathbf{1}_d$ are the d -dimensional vectors of zeroes and ones, respectively.

Homotopy methods are most often used to solve systems of nonlinear equations. However, they can also be used to solve optimization problems [10]. This technique is referred to as “Homotopy Optimization.” In Homotopy Optimization, $F_{\text{easy}}(\mathbf{x})$ and $F(\mathbf{x})$ encode NLPs rather than systems

of equations. How difficult an NLP is to solve depends on the problem parameters θ . For example, $\theta = \theta_0$ might lead to simple constraints that permit a large basin of attraction, while $\theta = \theta^*$ might encode challenging constraints with a small basin of attraction. To relate problem difficulty to homotopy, we define a mapping between homotopy parameters and optimization problem parameters,

$$\Gamma(\lambda) : \mathcal{H}_d \rightarrow \mathcal{Q},$$

such that $\lambda = \mathbf{0}_d$ maps to the “easy” problem θ_0 and $\lambda = \mathbf{1}_d$ maps to the “difficult” goal problem θ^* .

Homotopy optimization methods aim to find a solution to $\lambda = \mathbf{1}_d$ using intermediate points on the Solution Manifold.

Definition 2 (Solution Manifold). *The Solution Manifold is the manifold formed in $\mathcal{X} \times \mathcal{H}_d$ by the union of all solvable local minima \mathbf{x}^* of $\mathcal{P}_{\Gamma(\lambda)}$.*

$$\mathcal{M}_S = \{(\mathbf{x}^*, \lambda) : \exists \mathbf{x} \in \mathcal{B}_{\Gamma(\lambda)}(\mathbf{x}^*)\}.$$

III. PROBABILISTIC HOMOTOPY OPTIMIZATION

In this work, we propose P-HO to solve challenging trajectory optimization problems. The approach involves searching in the space $\mathcal{X} \times \mathcal{H}_d$. We sample homotopy parameters $\lambda \in \mathcal{H}_d$ and then use previous solutions to find as many points on the Solution Manifold (i.e., local minima) as possible for each sampled λ . This sample and solve process is repeated until a solution to the goal NLP is found. This section will formally define our problem, detail our proposed search algorithm, and briefly analyze its properties.

A. Problem Statement

Given a problem \mathcal{P}_{θ^*} and a parameter map $\Gamma(\cdot)$, the goal of a homotopy optimization method is to find a homotopy path, represented with the sequence of parameters $\Lambda = \{\lambda_0, \dots, \lambda_K\}$, such that for a trivial initial guess \mathbf{x}_0

$$\begin{cases} \mathbf{x}_0^* = \text{solve}(\mathcal{P}_{\Gamma(\lambda_0)}, \mathbf{x}_0) \\ \mathbf{x}_k^* = \text{solve}(\mathcal{P}_{\Gamma(\lambda_k)}, \mathbf{x}_{k-1}^*) \quad \forall k = 1, \dots, K \\ \mathbf{x}_{k-1}^* \in \mathcal{B}_{\Gamma(\lambda_k)}(\mathbf{x}_k^*) \quad \forall k = 1, \dots, K \\ \lambda_0 = \mathbf{0}_d, \quad \lambda_K = \mathbf{1}_d. \end{cases} \quad (3)$$

The behavior of a numerical solver is, in general, intractable to characterize. We, therefore, treat the numerical solver as a black box, making no assumptions about the structure of the Solution Manifold or the Basins of Attraction for any of the problems \mathcal{P}_{θ} . Conventional path-tracing homotopy methods [7] require access to gradient information and are limited to following a single connected component of the Solution Manifold. Furthermore, the homotopy parameter’s multidimensionality makes it difficult to prescribe the relative rates of change of its components. As such, the problem of finding the optimal sequence Λ is well-suited to be solved with a sampling-based approach. Our sampled-based approach involves growing an “Optimization Tree” that captures information about which NLPs have been solved.

Definition 3 (Optimization Tree). *An Optimization Tree τ is a rooted tree data structure that keeps track of the*

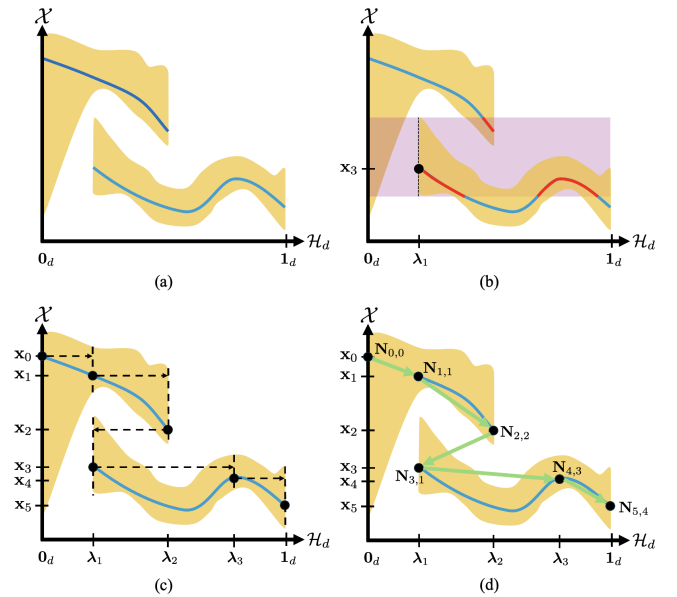


Fig. 2. Illustration of the P-HO algorithm: (a) Solution Manifold \mathcal{M}_S (blue) and the associated Basin of Attraction $\mathcal{B}_{\theta}(\mathbf{x})$ for each point on \mathcal{M}_S (yellow). (b) Highlights the range of \mathcal{X} in $\mathcal{B}_{\Gamma(\lambda_1)}(\mathbf{x}_3^*)$ (purple) and the points on the homotopy curve from which $\mathbf{N}_{3,1}$ is visible (red). (c) A sequence of nodes on \mathcal{M}_S and their projections on the Basins of attraction of the next node in the sequence (dashed line). (d) Potential solution path found by P-HO satisfying (3).

optimization problems that have been solved. Nodes $\mathbf{N}_{i,j}$ of the tree correspond to pairs of local minima and the problem parameters that lead to them: $\mathbf{N}_{i,j} = (\mathbf{x}_i^*, \lambda_j)$. An edge from $\mathbf{N}_{i,j}$ to $\mathbf{N}_{k,\ell}$ denotes that $\mathbf{x}_i^* \in \mathcal{B}_{\Gamma(\lambda_\ell)}(\mathbf{x}_k^*)$.

The root node of the tree is $\mathbf{N}_{0,0} = (\mathbf{x}_0^*, \mathbf{0}_d)$ and is assumed to have a Basin of attraction such that it is easy to initialize the tree with a trivial initial guess \mathbf{x}_0 . Thus, finding a sequence of parameters Λ satisfying (3) is equivalent to finding a path from the root $\mathbf{N}_{0,0}$ to a node $\mathbf{N}_{i,K}$ where $\lambda_K = \mathbf{1}_d$. Different branches of the tree allow for searching the parameter space in different homotopic directions.

B. Probabilistic Homotopy Optimization

While nodes of the tree are elements of $\mathcal{X} \times \mathcal{H}_d$, P-HO only samples in \mathcal{H}_d . For a sampled $\lambda \in \mathcal{H}_d$, the $\mathbf{x}^* \in \mathcal{X}$ component of the node comes from solving the problem encoded by $\Gamma(\lambda)$. This is a crucial feature of P-HO since it allows the same problem $\mathcal{P}_{\Gamma(\lambda)}$ to be solved multiple times, but with different initial guesses, as shown in Fig. 2. This leads to finding diverse solutions and better exploration of the trajectory space.

A summary of the P-HO algorithm is given by Alg. 1 and 2. Note that Alg. 1, as well as the other algorithms, use the operator $\text{sample}(\cdot)$ to randomly select an element from the set provided as an argument. P-HO begins by using the provided initial guess to solve the easy problem $\mathcal{P}_{\Gamma(\mathbf{0}_d)}$. If this problem cannot be solved, P-HO fails. Thus, the easy problem should permit a large basin of attraction. Next, the homotopy parameter set \mathcal{S}_λ and the solution attempts set \mathcal{S}_A

are initialized. The homotopy parameter set keeps track of the candidate NLPs, and the attempts set keeps track of which initial guesses we have attempted for each candidate NLP. The remainder of the algorithm involves swapping between the *sampling* and *solving* phases, depending on the solve/sample ratio.

The solve/sample ratio is computed via $\frac{|\mathcal{S}_A|}{|\tau| \times |\mathcal{S}_\lambda|}$, where $|\cdot|$ gives the number of elements in a set. Qualitatively, this number represents the fraction of possible unique solve() operations that have been attempted for the candidate NLPs. Thus, when $\frac{|\mathcal{S}_A|}{|\tau| \times |\mathcal{S}_\lambda|} = 1$, every candidate problem in \mathcal{S}_λ has been attempted to be solved with every possible initial guess from τ . When the solve/sample ratio is below the threshold ρ_A , P-HO executes the solve phase, where it tries to add a node to τ by solving an existing problem in \mathcal{S}_λ using an initial guess from τ that it has not yet tried on that problem. If the result from the solver \mathbf{x}_{new}^* is feasible and no similar trajectory exists, it will be added to τ . Regardless of the success of the solver, the attempt to solve $\mathcal{P}_{\Gamma(\lambda)}$ will be added to \mathcal{S}_A . Once enough solutions have been attempted to raise the ratio above the threshold, the sample phase is executed, where a new candidate parameterization is added to \mathcal{S}_λ .

Algorithm 1 Probabilistic Homotopy Optimization

Require: Initial guess \mathbf{x}_0

Hyperparams: Iteration limit q , solve/sample threshold ρ_A

```

1:  $\tau \leftarrow \{(\text{solve}(\mathcal{P}_{\Gamma(\mathbf{0}_d)}, \mathbf{x}_0), \mathbf{0}_d)\}$ 
2:  $\mathcal{S}_\lambda \leftarrow \{\mathbf{1}_d, \mathbf{0}_d\}$ ,  $\mathcal{S}_A \leftarrow \{\}$ 
3: for  $q$  iterations do
4:   if  $\frac{|\mathcal{S}_A|}{|\tau| \times |\mathcal{S}_\lambda|} \leq \rho_A$  then
5:      $(\mathbf{N}_{i,j}, \lambda_\ell) = \text{sampleAttempt}(\tau, \mathcal{S}_\lambda, \mathcal{S}_A)$ 
6:      $\mathcal{S}_A \leftarrow \mathcal{S}_A \cup (\mathbf{N}_{i,j}, \lambda_\ell)$ 
7:      $\mathbf{x}_{new}^* \leftarrow \text{solve}(\mathcal{P}_{\Gamma(\lambda_\ell)}, \mathbf{x}_i^*)$ 
8:     if  $\mathbf{x}_{new}^* \in \mathcal{C}$  AND  $\mathbf{x}_{new}^* \neq \mathbf{x}_j^* \forall \mathbf{N}_{j,\cdot} \in \tau$  then
9:        $\tau \leftarrow \tau \cup (\mathbf{x}_{new}^*, \lambda_\ell)$ 
10:    end if
11:  else
12:     $\mathcal{S}_\lambda \leftarrow \mathcal{S}_\lambda \cup \text{sample}(\mathcal{H}_d)$ 
13:  end if
14: end for
15: return  $\{\mathbf{x}_i^* : \mathbf{N}_{i,j} \in \tau, \lambda_j = \mathbf{1}_d\}$ 

```

Algorithm 2 sampleAttempt

Require: Optimization Tree τ , Parameters \mathcal{S}_λ , Attempts \mathcal{S}_A

Hyperparams: Probability of sampling goal P_g

```

1: if with probability  $P_g$  then
2:    $\mathbf{N} = \text{sample}(\{\mathbf{N}_{i,j} \in \tau : (\mathbf{N}_{i,j}, \mathbf{1}_d) \notin \mathcal{S}_A\})$ 
3:   return  $(\mathbf{N}, \mathbf{1}_d)$ 
4: else
5:   return  $\text{sample}(\{(\mathbf{N} \in \tau, \lambda \in \mathcal{S}_\lambda) : (\mathbf{N}, \lambda) \notin \mathcal{S}_A\})$ 
6: end if

```

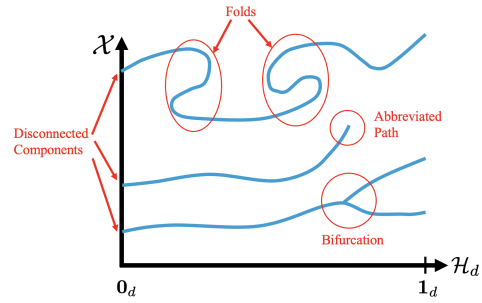


Fig. 3. Visualization of common pitfalls encountered by Homotopy Methods.

C. Algorithm Analysis

Our algorithm can find solutions in the presence of common pitfalls encountered in homotopy methods. These pitfalls are illustrated in Fig. 3. While path tracing methods are restricted to continuously following a single connected component of the Solution Manifold, P-HO can take discrete jumps, potentially large ones, in the homotopy parameter space. These jumps can allow the P-HO to skip over folds, move between disjoint components of the Solution Manifold, and continue progressing after an abbreviated path. Furthermore, P-HO supports storing multiple local minima for a given λ . This allows the algorithm to traverse all branches emerging from a bifurcation and track multiple disjoint components of the Solution Manifold over the same region of the homotopy space.

If a sequence Λ satisfying (3) exists, P-HO will eventually sample a set of homotopy parameters that are close enough to Λ . Since P-HO tries every found solution as the initial guess to every candidate problem, the algorithm will eventually find a sequence of parameters to solve $\mathcal{P}_{\Gamma(\mathbf{1}_d)}$. Note that this is not the same as guaranteeing that if a feasible solution $\mathcal{P}_{\Gamma(\mathbf{1}_d)}$ exists, P-HO will find it. The choice of parameters θ for (1) and the homotopy map $\Gamma(\cdot)$ influence the number of iterations P-HO needs to find a solution to $\mathcal{P}_{\Gamma(\mathbf{1}_d)}$.

IV. LINEAR INTERPOLATION AND RRT HOMOTOPY OPTIMIZATION

For the sake of comparison, we discuss alternative methods for expanding nodes in the Optimization Tree to find a sequence Λ that satisfies (3). In this section, we discuss Linear Interpolation Homotopy Optimization (LI-HO) and RRT Homotopy Optimization (R-HO) and compare them to P-HO.

A. Linear Interpolation Homotopy Optimization

While P-HO supports λ of arbitrary dimension, LI-HO requires a scalar homotopy parameter $\lambda \in \mathcal{H}_1$ and uses the parameter mapping

$$\Gamma(\lambda) = (1 - \lambda)\theta_0 + \lambda\theta^*.$$

This scalar parameterization produces an algorithm similar to the original Homotopy Optimization Method proposed in [10]. They are similar in that both use a one-dimensional

homotopy parameter to transform a simple optimization (rather than a system of equations) into a difficult one. However, LI-HO uses a dynamic step adjusting strategy for λ , while the original Homotopy Optimization Method uses a predetermined, fixed step size. Specifically, LI-HO uses the following strategy

- After each k_1 consecutive rounds of optimization being solvable $\Delta\lambda$ is multiplied by a constant $c_1 > 1$
- after k_2 consecutive failures $\Delta\lambda$ is multiplied by $c_2 < 1$
- Terminate algorithm if $\Delta\lambda < \epsilon$ occurs

where $k_1, c_1, k_2, c_2, \epsilon$, and the initial $\Delta\lambda$ are all hyperparameters. This strategy is useful as it avoids predetermining the steps and allows the algorithm to increase the step size quickly in "easy" regimes. A summary of the LI-HO algorithm is shown in Algorithm 3.

Algorithm 3 Linear Interpolation Homotopy Optimization

Require: Initial guess \mathbf{x}_0

Hyperparams: $k_1, c_1, k_2, c_2, \epsilon, \Delta\lambda_0$

```

1:  $\mathbf{x}_{cur}^* \leftarrow \text{solve}(\mathcal{P}_{\Gamma(\mathbf{0}_d)}, \mathbf{x}_0)$ 
2:  $\lambda \leftarrow \mathbf{0}_d, \Delta\lambda \leftarrow \Delta\lambda_0$ 
3: while  $\lambda \neq \mathbf{1}_d$  do
4:    $\lambda \leftarrow \lambda + \Delta\lambda$ 
5:    $\mathbf{x}_{new}^* = \text{solve}(\mathcal{P}_{\Gamma(\lambda)}, \mathbf{x}_{cur}^*)$ 
6:   if  $\mathbf{x}_{new}^* \in \mathcal{C}$  then
7:      $\mathbf{x}_{cur}^* \leftarrow \mathbf{x}_{new}^*$ 
8:     Re-adjust  $\Delta\lambda$ 
9:   else
10:    Re-adjust  $\Delta\lambda$  or return NULL
11:   end if
12: end while
13: return  $\mathbf{x}_{cur}^*$ 

```

B. RRT Homotopy Optimization

R-HO is inspired by another probabilistic search algorithm, Rapidly Exploring Random Tree (RRT) [21]. The key distinction holds: R-HO samples only in the homotopy parameter space \mathcal{H}_d , not the full configuration space associated with the nodes of the tree. R-HO grows the Optimization Tree by sampling a point $\lambda_{new} \in \mathcal{H}_d$, finding the node $\mathbf{N}_{close} \in \tau$ with the closest homotopy parameter to the sampled point, and then attempting to connect the sampled point to \mathbf{N}_{close} by solving $\mathcal{P}_{\Gamma(\lambda_{new})}$ using \mathbf{x}_{close}^* for an initial guess. A summary of the R-HO algorithm is shown in Algorithm 4.

Important implementation details for R-HO include the sampling distribution and the distance metric used to determine the closest node in the tree. We chose the sampling distribution to be uniform over \mathcal{H}_d (Line 6) and we use Euclidean distance metric in parameter space (Line 8).

C. Comparison with P-HO

Although we observed R-HO, and to a lesser extent LI-HO, to be effective algorithms, we also observed that both of them are susceptible to converge to lower quality local minima compared to P-HO. An analysis of the Solution Manifolds can explain this. Analyzing the Solution Manifolds of

Algorithm 4 RRT Homotopy Optimization

Require: Initial guess \mathbf{x}_0

Hyperparams: Goal probability P_g , iteration limit q

```

1:  $\tau \leftarrow \{(\text{solve}(\mathcal{P}_{\Gamma(\mathbf{0}_d)}, \mathbf{x}_0), \mathbf{0}_d)\}$ 
2: for  $q$  iterations do
3:   if with probability  $P_g$  then
4:      $\lambda_{new} \leftarrow \mathbf{1}_d$ 
5:   else
6:      $\lambda_{new} \leftarrow \text{sample}(\mathcal{H}_d)$ 
7:   end if
8:    $\mathbf{N}_{close} \leftarrow \arg \min_{\mathbf{N}_{i,j} \in \tau} \|\lambda_{new} - \lambda_j\|$ 
9:    $\mathbf{x}_{new}^* \leftarrow \text{solve}(\mathcal{P}_{\Gamma(\lambda_{new})}, \mathbf{x}_{close}^*)$ 
10:  if  $\mathbf{x}_{new}^* \in \mathcal{C}$  then
11:     $\tau \leftarrow \tau \cup (\mathbf{x}_{new}^*, \lambda_{new})$ 
12:    if  $\lambda_{new} = \mathbf{1}_d$  then
13:      return  $\mathbf{x}_{new}^*$ 
14:    end if
15:  end if
16: end for
17: return NULL

```

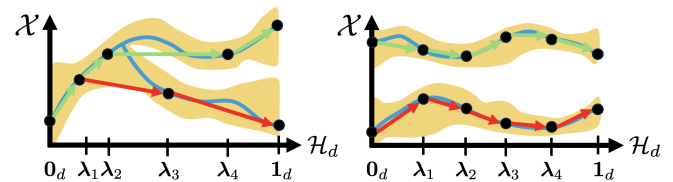


Fig. 4. Illustration of how R-HO and P-HO deal with bifurcation (left) and multiple local minima (right). After the green path has been found, only P-HO can also find the red path.

real-world systems, even toy systems like the cart-pole, is intractable. Therefore, we first discuss how the structure of Solution Manifolds impacts our presented algorithms using visuals of hypothetical Solution Manifolds. Later in Sec. V, we will present empirical results supporting this theoretical analysis.

1) *Bifurcation:* In case of bifurcation of the Solution Manifold, R-HO and LI-HO can only follow one of the branches. This is because once a node on one branch is chosen, that node will be the closest neighbor any time a parameter in the bifurcated region is sampled. Because P-HO allows multiple nodes to have the same parameter and does not expand the tree based on the proximity of sampled nodes, it can explore multiple branches. This distinction between the algorithms is illustrated in Fig. 4.

2) *Distance Bias:* R-HO assumes that if λ_i is the closest to λ_j , then the solution found for $\mathcal{P}_{\Gamma(\lambda_i)}$ will have the best chance of being in a Basin of Attraction associated with $\mathcal{P}_{\Gamma(\lambda_j)}$. This assumption, however, is fundamentally flawed. First of all, it is predicated on the definition of "closeness" between two points in a space where distance can be ill-defined. P-HO requires no such definition. Furthermore,

consider the case where a local minima near the goal is found and is not in the basin of attraction of the goal solution. R-HO may fail to converge since all samples in the goal region will find this node as the closest node.

3) *Diversity of Solutions*: Since P-HO maintains multiple different solutions for each λ , it can explore different local minima for a single λ . This results in a diverse set of solutions for both the end result and the intermediary nodes. Consider, for example, Fig. 4, where the Solution Manifold consists of multiple disjoint components. Because R-HO cannot store multiple solutions \mathbf{x}^* for a given λ , it is restricted to following either the top component of the manifold or the bottom component. Since the top and bottom components find different local minima for the goal problem, the quality of the solution R-HO returns is stochastic. P-HO can store multiple solutions per λ , so it will find the local minima for both the top and bottom components.

V. RESULTS

We supplement the theoretical analysis of P-HO from Sec. IV-C with empirical results from case studies involving dynamic motion planning for two systems: the underactuated Cart-Pole and the MIT Humanoid robot. The hyperparameters used for all algorithms for both case studies are given by Table I, unless noted otherwise.

TABLE I

HYPERPARAMETERS FOR HOMOTOPY OPTIMIZATION ALGORITHMS

P_g	ρ_A	q	k_1	k_2	c_1	c_2	ϵ	$\Delta\lambda_0$
0.3	1.0	10^4	2	1	1.5	0.3	10^{-9}	10^{-2}

A. Cart-Pole

Cart-Pole is a simple yet dynamically rich system often studied in robotics. The motion planning task is to swing up the pole by only applying horizontal force to the cart. We benchmark P-HO, R-HO, LI-HO algorithms on this task and compare them to directly optimizing from $\mathbf{N}_{0,0}$ as an initial guess (referred to as one-depth method).

There homotopy parameters are: mass of the cart (m_{cart}), mass of the pole (m_{pole}), force limit (F_{max}), length of the pole (l_{pole}), and horizontal movement range of the cart (x_{max}). The time horizon is 5 seconds. Figure 6 shows an example of how the trajectory evolves as parameters change.

To compare the performance of different algorithms, 1000 samples of $\theta_i \in \mathcal{Q}$ are uniformly selected from a "difficult" region in parameter space. Starting from the same root node $\mathbf{N}_{0,0}$, each algorithm attempts to solve for $\theta^* = \theta_i$ using no more than 200 queries of the NLP solver. The results in Fig. 5 show R-HO and P-HO outperform LI-HO and the one-depth method. Additionally, LI-HO is extremely sensitive to small changes in the goal parameters as we observe the set of solvable goal parameters using LI-HO is extremely sparse. Such sparsity may seem non-intuitive, but it can be explained by the potential presence of bifurcations, discontinuities, and folds in the Solution Manifold. In general, the path that LI-HO tries to traverse may contain regions of physical

infeasibility. R-HO and P-HO avoid these by finding a path in higher dimensions.

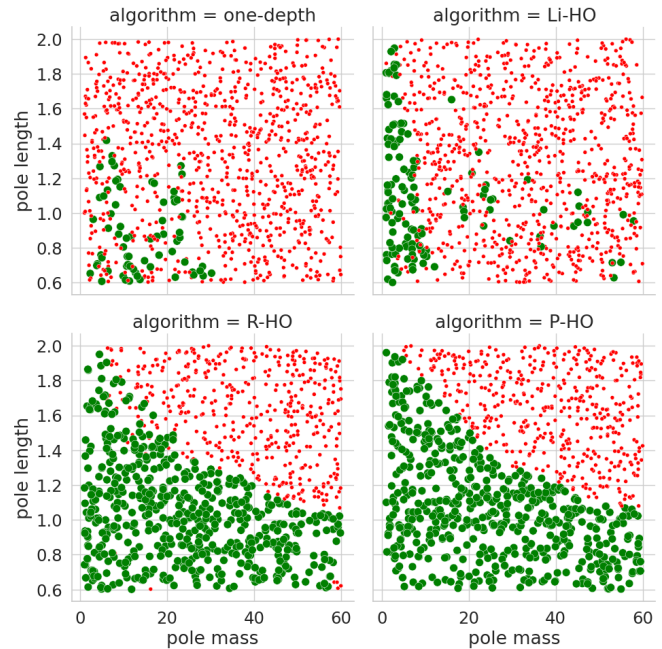


Fig. 5. Result of solving Cart-Pole swing up. For each parameter θ , a green dot represents the algorithm solving \mathcal{P}_θ successfully within the iteration limit, and a red dot represents an unsuccessful attempt. Parameters are uniformly drawn from $1kg \leq m_{pole} \leq 60kg$, $0.6m \leq l_{pole} \leq 2m$, $x_{max} = 1.6m$, $F_{max} = 100N$, $m_{cart} = 20kg$ (SI units). With an initial guess of zero, none of the \mathcal{P}_θ in this range are solvable.



Fig. 6. The trajectories \mathbf{x}_0^* , \mathbf{x}_K^* found by P-HO shown from left to right. As we go from θ_0 to θ_K , m_{pole} increases from $1kg$ to $60kg$ and F_{max} goes from $200N$ to $100N$ thus more swings are needed.

B. MIT Humanoid

Next, we will look at dynamic motion planning for the MIT Humanoid. Details about the baseline formulation can be found in [22]. The baseline motion planning problem was solved using KNITRO's interior point algorithm [9] but required substantial cost function tuning and manual initial-guess crafting. For the current work, we modify the baseline formulation in ways that make the original, manual solution method intractable, thus necessitating the algorithm proposed in this work. The key features of the modified formulation are:

- Modeling the actuator level dynamics of the robot, including the kinematic loops [23] that arise,
- Constraining the power consumed by the robot's motors via a quadratic regression of their efficiencies,

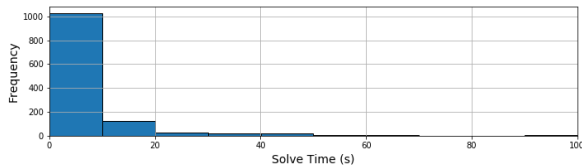


Fig. 7. Histogram of solve times for the MIT Humanoid back/front flip trajectory optimizations solved over the course of LI-HO, R-HO, and P-HO.

- Planning up to the point of landing (rather than the point of takeoff).

The typical solve times for this trajectory optimization problem are shown in Fig. 7.

We compare the performance of the LI-HO, R-HO, and P-HO with the following experiment. For the sake of fairness, we set a total solve-time limit T_{max} on each algorithm. For each algorithm, if the desired motion planning problem is solved prior to T_{max} , the objective function value of the solution corresponding to θ^* and the total solve time required are recorded. Thus, we compare the algorithms based on two criteria: (1) average cost when solved within T_{max} and (2) likelihood to find a solution within T_{max} , referred to as “success rate.” LI-HO is deterministic, its success rate within a given time interval is either 0% or 100%, and the cost associated with the solution it finds every time is fixed. R-HO and P-HO, however are probabilistic. Therefore, these algorithms were repeated and averaged over 30 runs per motion planning problem. The results of the comparisons for a back flip (Fig. 1) and front flip planning problem are shown in Fig. 8 and 9, respectively. The parameters that define these motion planning homotopies, as well as their values, are shown in Table II.

TABLE II
HOMOTOPY PARAMETERS FOR HUMANOID MOTION PLANNING

	Back Flip		Front Flip	
	θ_0	θ^*	θ_0	θ^*
Dynamics Timestep (s)	0.03	0.01	0.03	0.01
Mid-Flight Body Height (m)	0.4	1.2	0.4	1.2
Final Pitch Rotation ($^\circ$)	0	-360	0	360
Self-Collision Relaxation	50	1	50	1

We first note that Homotopy Optimization is required for these motion planning problems. Neither can be solved with a trivial initial guess, such as all zeros or a constant joint position at every timestep. For both motions, LI-HO is the fastest at finding a valid solution to the desired problem in the sense that LI-HO always finds solutions within 100s, while R-HO and P-HO take at least 700s to reach success rates of 80% or better. However, also, in both cases, the solution LI-HO finds has a significantly higher objective function value than the solutions returned by R-HO and P-HO. For these problems, the objective function value corresponds to control effort, so finding high-quality (i.e., low-cost) local minima is crucial. Thus, if the user is willing to sacrifice time efficiency

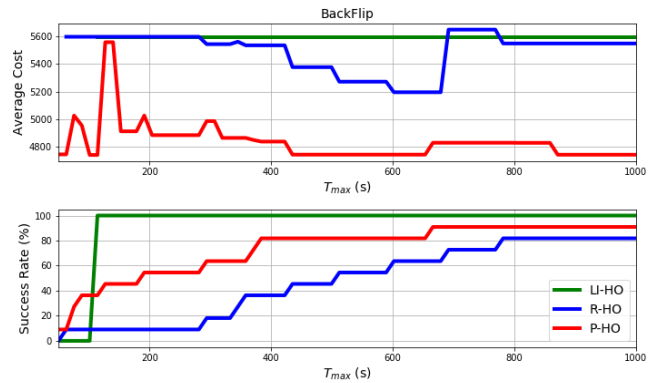


Fig. 8. Average cost and success rate for planning a back flip of the MIT Humanoid robot within T_{max} .

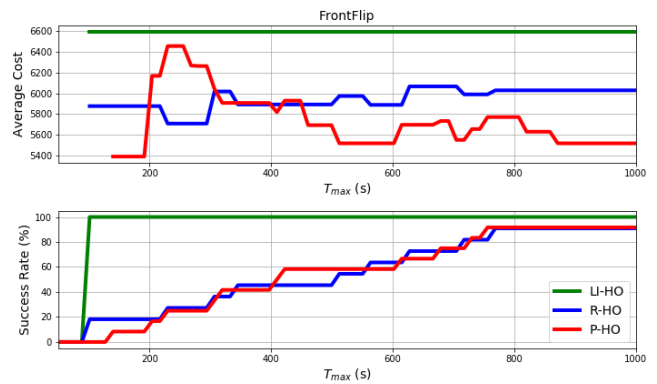


Fig. 9. Average cost and success rate for planning a front flip of the MIT Humanoid robot within T_{max} .

for solution quality, R-HO and P-HO are better options.

Lastly, P-HO has the advantage over R-HO by constantly improving its solution until T_{max} is reached. Therefore, as time passes, P-HO is likely to find better local minima, even the global minimum, if such a path exists. R-HO, on the other hand, terminates once a solution is found. While “anytime” variants such as RRT* exist [24], they rely on rewiring the tree as new nodes are added based on the edge costs. However, our search problem does not have edge costs. The cost for our problem is the objective function value of the solved goal NLP, which is only accrued once the goal node has been reached. This incompatibility with “anytime” variants can explain the unintuitive increase over time shown by the average cost of R-HO. Analyzing the structure of the Solution Manifold for such a complex problem is outside the scope of this work. However, we hypothesize that a bifurcation in the Solution Manifold could explain the increasing cost. While P-HO could follow all branches emerging from the bifurcation, R-HO can only follow one branch, and the choice of the branch can be stochastic. One branch may efficiently lead to a low-cost solution while the other leads slowly to a high-cost solution.

VI. CONCLUSION

This work presented the P-HO algorithm, a homotopic approach for solving challenging, optimization-based motion planning problems. The algorithm emerges from formulating the problem of multidimensional Homotopy Optimization as a search problem in the space of homotopy parameters. Given this formulation, we described the P-HO algorithm and analyzed its properties, such as its ability to deal with bifurcated, disconnected, and folded Solution Manifolds. P-HO was then compared to alternative Homotopy Optimization methods, the deterministic LI-HO and the stochastic R-HO. Two case studies were presented. The Cart-Pole study showed that P-HO and R-HO methods reliably solve much more challenging optimization tasks. The MIT Humanoid study showed that P-HO is consistently able to find higher-quality local minima if some additional time is allowed for the algorithm to run.

Future work will focus on leveraging knowledge of the parameter space to make search more efficient, as well as investigating alternative parameterizations. For example, finding a meaningful distance metric in the parameter space can improve the notion of “closest node” in allow R-HO and lead to more efficient tree growth. For P-HO, non-uniform parameter sampling strategies, such as Gaussian Sampling in Probabilistic Roadmap (PRM) [25], can lead to a more efficient search. Finally, we will extend the algorithm to enable discrete parameterizations that will allow, for example, constraints to be altogether turned on and off rather than simply softened.

ACKNOWLEDGMENTS

This work was supported by Naver Labs, LG Electronics, and the National Science Foundation (NSF) Graduate Research Fellowship Program under Grant No. 4000092301.

REFERENCES

- [1] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter, “Perceptive locomotion through nonlinear model-predictive control,” *IEEE Transactions on Robotics*, 2023.
- [2] G. Kim, D. Kang, J.-H. Kim, S. Hong, and H.-W. Park, “Contact-implicit mpc: Controlling diverse quadruped motions without pre-planned contact modes or trajectories,” *arXiv preprint arXiv:2312.08961*, 2023.
- [3] F. Jenelten, J. He, F. Farshidian, and M. Hutter, “Dtc: Deep tracking control,” *Science Robotics*, vol. 9, no. 86, p. eadh5401, 2024.
- [4] P. Brakel, S. Bohez, L. Hasenclever, N. Heess, and K. Bousmalis, “Learning coordinated terrain-adaptive locomotion by imitating a centroidal dynamics planner,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 10335–10342, IEEE, 2022.
- [5] T. Marcucci, M. Petersen, D. von Wrangel, and R. Tedrake, “Motion planning around obstacles with convex optimization,” *Science robotics*, vol. 8, no. 84, p. eadf7843, 2023.
- [6] G. Bledt and S. Kim, “Extracting legged locomotion heuristics with regularized predictive control,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 406–412, IEEE, 2020.
- [7] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 1999.
- [8] 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, pp. 25–57, 2006.
- [9] R. H. Byrd, J. Nocedal, and R. A. Waltz, “Knitro: An integrated package for nonlinear optimization,” *Large-scale nonlinear optimization*, pp. 35–59, 2006.
- [10] D. M. Dunlavy and D. P. O’Leary, “Homotopy optimization methods for global optimization,” tech. rep., Sandia National Laboratories (SNL), Albuquerque, NM, and Livermore, CA, 2005.
- [11] N. Rosa and K. M. Lynch, “Extending equilibria to periodic orbits for walkers using continuation methods,” in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3661–3667, IEEE, 2014.
- [12] M. Raff, N. Rosa, and C. D. Remy, “Generating families of optimally actuated gaits from a legged system’s energetically conservative dynamics,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8866–8872, IEEE, 2022.
- [13] F. Jenelten, R. Grandia, F. Farshidian, and M. Hutter, “Tamols: Terrain-aware motion optimization for legged systems,” *IEEE Transactions on Robotics*, vol. 38, no. 6, pp. 3395–3413, 2022.
- [14] M. R. Turski, J. Norby, and A. M. Johnson, “Staged contact optimization: Combining contact-implicit and multi-phase hybrid trajectory optimization,” *arXiv preprint arXiv:2304.04923*, 2023.
- [15] L. T. Watson, “Globally convergent homotopy methods: a tutorial,” *Applied Mathematics and Computation*, vol. 31, pp. 369–396, 1989.
- [16] D. M. Wolf and S. R. Sanders, “Multiparameter homotopy methods for finding dc operating points of nonlinear circuits,” *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 43, no. 10, pp. 824–838, 1996.
- [17] H. Vazquez-Leal, R. Castaneda-Sheissa, F. Rabago-Bernal, L. Hernandez-Martinez, A. Sarmiento-Reyes, and U. Filobello-Nino, “Powering multiparameter homotopy-based simulation with a fast path-following technique,” *International Scholarly Research Notices*, vol. 2011, 2011.
- [18] C. Grenat, S. Baguet, C.-H. Lamarque, and R. Dufour, “A multiparametric recursive continuation method for nonlinear dynamical systems,” *Mechanical Systems and Signal Processing*, vol. 127, pp. 276–289, 2019.
- [19] R. Fletcher, *Practical methods of optimization*. John Wiley & Sons, 2000.
- [20] D. Mayne, “A second-order gradient method for determining optimal trajectories of non-linear discrete-time systems,” *International Journal of Control*, vol. 3, no. 1, pp. 85–95, 1966.
- [21] S. M. LaValle, J. J. Kuffner, B. Donald, *et al.*, “Rapidly-exploring random trees: Progress and prospects,” *Algorithmic and computational robotics: new directions*, vol. 5, pp. 293–308, 2001.
- [22] M. Chignoli, D. Kim, E. Stanger-Jones, and S. Kim, “The mit humanoid robot: Design, motion planning, and control for acrobatic behaviors,” in *IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*, 2021.
- [23] M. Chignoli, N. Adrian, S. Kim, and P. M. Wensing, “Recursive rigid-body dynamics algorithms for systems with kinematic loops,” *arXiv preprint arXiv:2311.13732*, 2023.
- [24] S. Karaman, M. R. Walter, A. Perez, E. Frazzoli, and S. Teller, “Anytime motion planning using the rrt,” in *2011 IEEE international conference on robotics and automation*, pp. 1478–1483, IEEE, 2011.
- [25] V. Boor, M. Overmars, and A. van der Stappen, “The gaussian sampling strategy for probabilistic roadmap planners,” in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, vol. 2, pp. 1018–1023 vol.2, 1999.