

Approximate Topological Optimization using Multi-Mode Estimation for Robot Motion Planning

Andreas Orthey¹ and Florian T. Pokorny² and Marc Toussaint^{1,3}




Fig. 1: Initial test of multi-mode estimation on different state spaces. Problem is to move a point robot from a given start state (green) to a goal state (red) while optimizing a minimum-length cost functional. Local modes found after convergence are shown in red. **Left:** Sphere with two punctured holes (in grey) and four modes (in red). **Middle Left:** Punctured Mobius strip with three modes. **Middle Right:** Punctured Klein bottle with three modes. **Right:** Punctured torus with four modes.

Abstract—In this extended abstract, we report on ongoing work towards an approximate multimodal optimization algorithm with asymptotic guarantees. Multimodal optimization is the problem of finding all local optimal solutions (modes) to a path optimization problem. This is important to compress path databases, as contingencies for replanning and as source of symbolic representations. Following ideas from Morse theory, we define modes as paths invariant under optimization of a cost functional. We develop a multi-mode estimation algorithm which approximately finds all modes of a given motion optimization problem and asymptotically converges. This is made possible by integrating sparse roadmaps with an existing single-mode optimization algorithm. Initial evaluation results show the multi-mode estimation algorithm as a promising direction to study path spaces from a topological point of view.

I. INTRODUCTION

We develop a new multi-mode estimation algorithm to find modes. Modes are paths invariant under optimization of a cost functional. Estimating modes is important to provide completeness guarantees to optimization algorithms, as contingencies for rapid replanning [35, 18, 14], and to provide admissible heuristics for more complex problems [33, 15]. This is useful in multi-robot navigation [14, 8], to investigate long-horizon planning problems [4], to sparsify path databases [19], or as symbols for high-level planning [30].

However, the robotics community has thus far concentrated almost exclusively on single-mode optimization algorithms. Methods include CHOMP [36], TrajOpt [23], STOMP [6],

KOMO [28, 29], Kernel Projection [7], Bayesian Optimization [32] or Gaussian Process Planning [12]. While those optimization methods find a *single mode* from an existing solution, we like to leverage them for *multi-mode* estimation.

Multi-mode estimation is closely related to topological optimization. Topological concepts like homotopy [2, 1], homology [19, 20] or braids [9, 10] have been successfully applied in robotics. Our work differs by explicitly assuming that a single-mode optimization algorithm is given. Motivated by Morse theory [11], we concentrate on the problem of enumerating modes. While finding modes has been studied extensively in (evolutionary) optimization [21], in terms of visualization [27] and in terms of mode-relationships [13, 31], there are just a few works applying this concept to path planning [5, 22, 17, 33]. Our work is complementary, in that we do multi-mode estimation but provide asymptotic guarantees.

While we previously tackled the same problem [16, 14], we significantly improve upon this work. In particular, our multi-mode estimation algorithm can enumerate the modes of a given optimal planning problem while (i) asymptotically converges to all modes (see proof in Appendix A), (ii) being anytime, and (iii) can use iterative optimizers. This is achieved by having an integrated 2-stage process combining sparse roadmaps with a single-mode path optimization method.

II. APPROXIMATE TOPOLOGICAL OPTIMIZATION

We consider optimal motion planning problems of the form (X, x_I, x_G, c, σ) whereby X is the state space, x_I is an initial state, x_G a goal state, c is a cost functional of the form

$$c(p) = \int_0^1 L(x, p(x), p'(x)) dx \quad (1)$$

¹Max Planck Institute for Intelligent Systems, Stuttgart, Germany. Marc Toussaint thanks the MPI-IS for the Max Planck Fellowship.

²KTH Royal Institute of Technology, Stockholm, Sweden

³Technical University of Berlin, Berlin, Germany. Emails: {aorthey}@is.mpg.de, {ftoussaint}@tu-berlin.de, {fpokorny}@kth.se,




Fig. 2: Overview of multi-mode estimation. As input, we use a state space on which we grow a sparse roadmap [3]. New paths in the sparse roadmap are proposed to the multi-mode database for storage and optimization. We terminate if both the sparse roadmap and all database paths have converged.

whereby $p : I \rightarrow X$ is a path on the state space from x_I to x_G and σ is a path optimizer. We assume that σ is deterministic, does not increase path cost and asymptotically converges to a fixed point path (See Appendix A for more details). Our goal is to find the modes, i.e. the set of paths which are invariant under optimizer σ over the cost functional Eq. (1).

A. Overview of Multi-Mode Estimation

Our general method is an integrated 2-stage process depicted in Fig. 2. We first grow a sparse roadmap [3] on the state space. Whenever we add a new edge to the roadmap, we check if this edge creates a new path from start to goal. This path is then added to the multi-mode database. Our algorithm terminates if no samples are added to the sparse roadmap for M subsequent iterations *and* if all paths in the database have asymptotically converged under optimizer σ . This algorithm is anytime, i.e. it returns all current paths on premature termination.

B. Sparse Roadmaps

In the first step of our algorithm we grow a sparse roadmap on the state space. This method follows closely previous algorithms [25, 3] and uses the parameter Δ as the visibility radius to reject samples. Whenever we add a new edge to the sparse graph, we check if this edge adds a path from start to goal. This is done by first checking if the source vertex of the edge is in the same component as the start and the goal vertex of the graph. We then construct a path from start to the source vertex and another path from source vertex to the goal vertex. Both paths are then concatenated and send to the multi-mode estimation database. We say that the sparse roadmap is converged if M subsequent infeasible samples have been drawn [25].

C. Multi-mode database

The multi-mode database gets as input a stream of paths and adds those paths either to the database or uses them to update an existing path. In the first step, we apply the optimizer σ for one iteration. We then iterate through all database paths and check for path equivalence. Path equivalence is tested using

the Hausdorff metric which is defined as

$$d_H(p, p') = \sup_{s \in I} \inf_{t \in I} \{d(p(s), p'(t)) \mid t \in I\}, \quad (2)$$

whereby \sup is the supremum (least upper bound), \inf is the infimum (greatest lower bound) and d is a metric on X . We say that two paths p, p' are equivalent if $d_H(p, p') < \epsilon$, whereby ϵ is a resolution parameter specific to the optimizer. If two paths are equivalent, we either update the existing path if its cost is lower or remove it. If no paths are equivalent, we add the new path to the database. After every iteration, we randomly pick a non-converged path, optimize it for one iteration, and check for path equivalence. This stage is said to be converged if all paths have converged.

III. DEMONSTRATIONS

As an initial test of the multi-mode estimation algorithm, we use four 2-d examples. In particular, we use punctured versions of the sphere, the Mobius strip, the Klein bottle and the torus (Fig. 1). Each of those state spaces has been implemented in the open motion planning library (OMPL) [26], whereby we implement uniform sampling using curvature-based rejection sampling [34] and implement interpolation functions which take the gluing of the state space into account.

Using parameters as detailed in Appendix B, we let the multi-mode estimation run 10 times and report on t , the time until convergence in seconds, m , the average number of minima found, and n , the average number of nodes in the sparse roadmap after convergence. For the Mobius strip, we have $t = 8.75 \pm 1.24$, $m = 3$ and $n = 78.3$. For the sphere, the results are $t = 17.76 \pm 4.32$, $m = 4$, $n = 42.1$, for the torus $t = 16.39 \pm 1.52$, $m = 5$, $n = 116.0$ and for the Klein bottle $t = 62.41 \pm 35.82$, $m = 3.5$, $n = 97.2$. The results indicate that the multi-modal estimation algorithm robustly converges while having a low number of nodes in the sparse roadmap. We observe that in the Klein bottle scenario, the optimizer converges to 3 modes in half of the cases, while converging to 4 in the rest of them. This is due to numerical instabilities in the optimizer, which needs to be addressed in future work. However, our multi-mode estimation algorithm is conservative, i.e. it reliably finds all existing modes in all cases.

IV. CONCLUSION

This extended abstract described a method to estimate modes of a robot motion planning and optimization problem. To apply this methodology to more complex problems, we further need to solve two problems. First, we need to scale the algorithm up to high-dimensional systems. This can be achieved by aggregating modes using multilevel abstractions [14, 15]. Second, while we provided proof for asymptotic convergence in the Appendix, it is often difficult to achieve this in practice due to jumping behavior in the optimizer. This could potentially be alleviated by using more powerful optimizers with iteration step guarantees. However, we believe multi-mode estimation to be a useful and promising tool to study high-dimensional path spaces from a topological point of view.

REFERENCES

- [1] Subhrajit Bhattacharya and Robert Ghrist. Path homotopy invariants and their application to optimal trajectory planning. *Annals of Mathematics and Artificial Intelligence*, 84(3-4):139–160, 2018.
- [2] Subhrajit Bhattacharya, Maxim Likhachev, and Vijay Kumar. Topological constraints in search-based robot path planning. *Autonomous Robots*, 33(3), 2012.
- [3] Andrew Dobson and Kostas E. Bekris. Sparse roadmap spanners for asymptotically near-optimal motion planning. *International Journal of Robotics Research*, 33(1): 18–47, 2014.
- [4] Valentin N. Hartmann, Andreas Orthey, Danny Driess, Ozgur S. Oguz, and Marc Toussaint. Long-horizon multi-robot rearrangement planning for construction, 2021. arXiv:2106.02489 [cs.RO].
- [5] Léonard Jaillet and Thierry Siméon. Path deformation roadmaps: Compact graphs with useful cycles for motion planning. *International Journal of Robotics Research*, 2008.
- [6] Mrinal Kalakrishnan, Sachin Chitta, Evangelos Theodorou, Peter Pastor, and Stefan Schaal. Stomp: Stochastic trajectory optimization for motion planning. In *IEEE International Conference on Robotics and Automation*, 2011.
- [7] Zita Marinho, Byron Boots, Anca Dragan, Arunkumar Byravan, Geoffrey J. Gordon, and Siddhartha Srinivasa. Functional gradient motion planning in reproducing kernel hilbert spaces. In *Robotics: Science and Systems*, Ann Arbor, Michigan, June 2016. doi: 10.15607/RSS.2016.XII.046.
- [8] Christoforos Mavrogiannis, Jonathan A. DeCastro, and Siddhartha S. Srinivasa. Implicit multiagent coordination at unsignalized intersections via multimodal inference enabled by topological braids, 2020.
- [9] Christoforos I Mavrogiannis and Ross A Knepper. Decentralized multi-agent navigation planning with braids. In *Workshop on the Algorithmic Foundations of Robotics*, 2016.
- [10] Christoforos I Mavrogiannis and Ross A Knepper. Multi-agent path topology in support of socially competent navigation planning. *International Journal of Robotics Research*, 38(2-3):338–356, 2019.
- [11] Marston Morse. *The calculus of variations in the large*, volume 18 of *Colloquium Publications*. American Mathematical Society, 1934.
- [12] Mustafa Mukadam, Jing Dong, Xinyan Yan, Frank Dellaert, and Byron Boots. Continuous-time gaussian process motion planning via probabilistic inference. *International Journal of Robotics Research*, 37(11):1319–1340, 2018.
- [13] Gabriela Ochoa, Sébastien Verel, Fabio Daolio, and Marco Tomassini. Local optima networks: A new model of combinatorial fitness landscapes. In *Recent advances in the theory and application of fitness landscapes*, pages 233–262. Springer, 2014.
- [14] Andreas Orthey and Marc Toussaint. Visualizing local minima in multi-robot motion planning using multilevel morse theory. *Workshop on the Algorithmic Foundations of Robotics*, 2020.
- [15] Andreas Orthey, Sohaib Akbar, and Marc Toussaint. Multilevel motion planning: A fiber bundle formulation, 2020. arXiv:2007.09435 [cs.RO].
- [16] Andreas Orthey, Benjamin Frész, and Marc Toussaint. Motion planning explorer: Visualizing local minima using a local-minima tree. *Robotics and Automation Letters*, 5(2):346–353, April 2020.
- [17] Takayuki Osa. Multimodal trajectory optimization for motion planning. *The International Journal of Robotics Research*, 39(8):983–1001, 2020.
- [18] Előd Páll, Arne Sieverling, and Oliver Brock. Contingent contact-based motion planning. In *IEEE International Conference on Intelligent Robots and Systems*, pages 6615–6621. IEEE, 2018.
- [19] Florian T Pokorny, Majd Hawasly, and Subramanian Ramamoorthy. Topological trajectory classification with filtrations of simplicial complexes and persistent homology. *International Journal of Robotics Research*, 35(1-3): 204–223, 2016.
- [20] Florian T Pokorny, Danica Kragic, Lydia E Kavraki, and Ken Goldberg. High-dimensional winding-augmented motion planning with 2d topological task projections and persistent homology. In *IEEE International Conference on Robotics and Automation*, pages 24–31. IEEE, 2016.
- [21] Mike Preuss. *Multimodal optimization by means of evolutionary algorithms*. Springer, 2015.
- [22] Christoph Rösmann, Frank Hoffmann, and Torsten Bertram. Integrated online trajectory planning and optimization in distinctive topologies. *Robotics and Autonomous Systems*, 88:142–153, 2017.
- [23] John Schulman, Yan Duan, Jonathan Ho, Alex Lee, Ibrahim Awwal, Henry Bradlow, Jia Pan, Sachin Patil, Ken Goldberg, and Pieter Abbeel. Motion planning with sequential convex optimization and convex collision checking. *The International Journal of Robotics Research*, 33(9):1251–1270, 2014.
- [24] Sepanta Sekhavat, Petr Svestka, Jean-Paul Laumond, and Mark H Overmars. Multilevel path planning for nonholonomic robots using semiholonomic subsystems. *International Journal of Robotics Research*, 17(8):840–857, 1998.
- [25] Thierry Siméon, J-P Laumond, and Carole Nissoux. Visibility-based probabilistic roadmaps for motion planning. *Advanced Robotics*, 14(6):477–493, 2000.
- [26] Ioan A Şucan, Mark Moll, and Lydia Kavraki. The open motion planning library. *Robotics and Automation Magazine*, 19(4):72–82, 2012.
- [27] Julien Tierny, Guillaume Favelier, Joshua A Levine, Charles Gueunet, and Michael Michaux. The topology toolkit. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):832–842, 2017.

- [28] Marc Toussaint. Robot trajectory optimization using approximate inference. In *Proceedings of the 26th annual international conference on machine learning*, pages 1049–1056, 2009.
- [29] Marc Toussaint. Newton methods for k-order markov constrained motion problems. *arXiv preprint arXiv:1407.0414*, 2014.
- [30] Marc Toussaint, Kelsey R Allen, Kevin A Smith, and Josh B Tenenbaum. Differentiable physics and stable modes for tool-use and manipulation planning. In *Robotics: Science and Systems*, 2018.
- [31] German Treimun-Costa, Elizabeth Montero, Gabriela Ochoa, and Nicolás Rojas-Morales. Modelling parameter configuration spaces with local optima networks. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 751–759, 2020.
- [32] Ngo Anh Vien, Heiko Zimmermann, and Marc Toussaint. Bayesian functional optimization. In *AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [33] Vojtěch Vonásek, Robert Pěnička, and Barbora Kozl(i)ková. Searching multiple approximate solutions in configuration space to guide sampling-based motion planning. *Journal of Intelligent & Robotic Systems*, pages 1–17, 2020.
- [34] JF Williamson. Random selection of points distributed on curved surfaces. *Physics in Medicine & Biology*, 32(10):1311, 1987.
- [35] Yuandong Yang and Oliver Brock. Elastic roadmaps—motion generation for autonomous mobile manipulation. *Autonomous Robots*, 28(1):113–130, 2010.
- [36] Matthew Zucker, Nathan Ratliff, Anca Dragan, Mihail Pivtoraiko, Matthew Klingensmith, Christopher Dellin, J. Andrew (Drew) Bagnell, and Siddhartha Srinivasa. CHOMP: Covariant Hamiltonian Optimization for Motion Planning. *International Journal of Robotics Research*, 2013.

APPENDIX A PROOF OF ASYMPTOTIC CONVERGENCE

In this appendix, we prove that the multi-mode estimation algorithm converges to all modes of a given optimization problem, i.e. that it asymptotically converges. This result depends on the parameters M and Δ of the sparse roadmap. In particular, we prove that multi-mode estimation converges with probability 1 as M approaches infinity to all modes with a basin of attraction (neighborhood of paths converging to the mode) of radius at least Δ .

A. Notations and Assumptions

Let X be the state space, a compact manifold (closed, bounded, locally euclidean), $x_I \in X$ be the start and $x_G \in X$ be the goal state. Let us further assume that we are given a distance function $d : X \times X \rightarrow \mathbb{R}$ and a constraint function

$\phi : X \rightarrow \{0, 1\}$, which evaluates to zero if a state is constraint-free and to one otherwise (this function encapsulates e.g. joint limits, self-collisions, robot-environment-collisions or robot-robot-collisions). The constraint function implicitly defines the free state space $X_{\text{free}} = \{\phi(x) = 0 \mid x \in X\}$. We then define the path space P to be $P = X_{\text{free}}^I \Big|_{x_I}^{x_G}$, i.e. the set of paths $p : I \rightarrow X_{\text{free}}$ which start at x_I and end at x_G .

Given a path space, we like to estimate modes. Recall that a mode is a path invariant under *optimization* of a *cost functional*. We assume a given cost functional $c(p)$ as introduced in Eq. (1) which assigns a cost value to each path in our path space and we assume a path optimizer $\sigma : P \rightarrow P$, which is a mapping taking as input a path and returning as output another path. We make additional assumptions on this optimizer.

- **Non-increasing cost.** If $c(p)$ is the cost of p , then $c(\sigma(p)) \leq c(p)$.
- **Deterministic.** Any two applications of $\sigma(p)$ will yield the same unique deterministic path.
- **Converges Asymptotically to a Fixed Point.** The path optimizer will eventually converge to a path, i.e. there exists an $N > 0$ such that for any $n > N$ we have $\sigma^n(p) = \sigma^{n-1}(p)$ whereby σ^n is the n times repeated application of the mapping σ .

Note that we do not make additional assumptions such as continuity, i.e. that input and output are continuously deformable into each other (homotopic), or idempotence, i.e. that the optimizer converges after one iteration. While this slightly increases the complexity of our algorithm, it vastly increases the scope of single-mode optimizers we can use, such as shortcutting algorithms [24] or iterative gradient descent methods [29].

B. Basin of Attraction

Let p be a mode of Eq. (1). We define its *basin of attraction* as the set of paths $p' \in P$ such that there exists an $N > 0$ such that for all $n > N$ we have $d_H(p, \sigma^n(p')) \leq \epsilon$, i.e. p is a fixed point for the optimizer σ when applied to path p' . In this case, the mode is also called an attractor.

For the theoretical treatment, it is important to quantify the size of the basin of attraction. We introduce $r_\sigma(p)$, which we define as the *radius of the basin of attraction* of mode p , which is the largest real number such that all paths p' of distance $d_H(p, p') \leq r_\sigma(p)$ are inside the basin of attraction for optimizer σ .

In practice, we might have to deal with optimizers which exhibit jumping-behavior, where it looks like the optimizer converged (i.e. we stay below ϵ for one iteration), while in reality the optimizer might have only made small steps and has in fact not yet converged. This can be dealt with by defining a parameter N_E , which is the number of steps during which the optimizer needs to stay below ϵ for us to consider the path to be converged to a fixed point.

C. Asymptotic Convergence

We are ready to state our two theorems. The first theorem establishes that for every mode p , there exists (in the limit) a path on the sparse roadmap which has a Hausdorff distance to the mode which is upper bounded by the visibility radius Δ . In the second theorem, we then establish that all modes with a radius of basin of attraction larger than the constant will eventually be found with probability $1 - \frac{1}{M}$. Before we state those theorem, we restate a Lemma by Dobson and Bekris [3], which will come in handy in our proofs.

Lemma 1 (Dobson and Bekris [3]). *Let p be a path. Then there exists a connected series of vertices $V = \{v_1, \dots, v_M\}$ on the sparse graph, with probability approaching 1 as M goes to infinity, such that every point $p(s)$ on the image of p lies in the visibility region of at least one vertex in V .*

This is basically a slight restatement of Lemma 1 in Dobson and Bekris [3]. We use this Lemma to prove that paths are upper bounded in the Hausdorff distance.

Theorem 1. *Let p be a mode. Then there exists a path p' on the sparse roadmap, with probability approaching 1 as M goes to infinity, such that $d_H(p, p')$ is upper bounded by Δ for d_H being the Hausdorff distance.*

Proof: By Lemma 1, p is covered by a path p' consisting of a connected sequence of vertices. Let $p(s)$ be an element on the image of the mode. Then $p(s)$ lies in its visibility radius of some vertex v on p' and its distance to the path is therefore upper bounded by $d(p(s), v) \leq \Delta$. Since this is true for every point, the Hausdorff distance can therefore be upper bounded by Δ . ■

Theorem 1 establishes that there exists, for every mode, in the limit, a path on the sparse roadmap which is inside a Δ -neighborhood of the mode. We use this result to establish that every mode with a sufficiently bounded basin of attraction will be found in the limit.

Theorem 2. *Let p be a mode with radius of basin of attraction*

of $r_\sigma(p)$. If $r_\sigma(p) \geq \Delta$, then multi-mode estimation with optimizer σ will eventually converge to p with probability approaching 1 as M goes to infinity.

Proof: By Theorem 1, there exists, in the limit as M goes to infinity, a path p' on the sparse roadmap with Hausdorff distance upper bounded by Δ . Since this path is by assumption below the radius of the basin of attraction $r_\sigma(p)$, this path will, eventually, converge to p by repeated application of the mapping σ . ■

Those two theorems show that we will find all modes for optimizer σ with probability approaching 1 as M goes to infinity and which have a basin of attraction of at least Δ . Modes which belong to tight narrow passages or where the cost functional is pathological will often have small basin of attraction and are therefore not guaranteed to be found. However, we can tune both M to be arbitrarily large and Δ to be arbitrarily small, thereby accommodating even ill-behaved modes and ill-behaved cost functionals.

APPENDIX B PARAMETERS

In practice, we need to choose additional parameters for the multi-mode estimation, such that (1) we converge to all (desired) modes and (2) the runtime does not blow up. By trial and error, we converged to the set of parameters shown in Table I, both for the sparse roadmap and for the optimizer.

	Parameter	Description	Value
Sparse Roadmap	Δ	Visibility radius of a node	0.1μ
	M	Number of subsequent failures	5000
	t	Stretch factor	3
Optimizer	ϵ	Accuracy for path equivalence	0.3
	ϵ_E	Threshold for convergence error	$1e-2$
	N_E	Number of sub-threshold iterations to declare convergence	10

TABLE I: Parameters used in the multi-mode estimation algorithm. The constant μ is the measure of the state space.