

A Survey on the Integration of Machine Learning with Sampling-based Motion Planning

Other titles in Foundations and Trends® in Robotics

Adoption of Robots for Disasters: Lessons from the Response to COVID-19

Robin R. Murphy, Vignesh B.M. Gandudi, Justin Adams, Angela Clendenin and Jason Moats

ISBN: 978-1-68083-862-6

Cybersecurity in Robotics: Challenges, Quantitative Modeling, and Practice

Quanyan Zhu, Stefan Rass, Bernhard Dieber and Víctor Mayoral Vilches

ISBN: 978-1-68083-860-2

A Roadmap for US Robotics – From Internet to Robotics 2020 Edition

Henrik Christensen, Nancy Amato, Holly Yanco, Maja Mataric, Howie Choset, Ann Drobni, Ken Goldberg, Jessy Grizzle, Gregory Hager, John Hollerbach, Seth Hutchinson, Venkat Krovi, Daniel Lee, Bill Smart, Jeff Trinkle and Gaurav Sukhatme

ISBN: 978-1-68083-858-9

The State of Industrial Robotics: Emerging Technologies, Challenges, and Key Research Directions

Lindsay Sanneman, Christopher Fourie and Julie A. Shah

ISBN: 978-1-68083-800-8

Semantics for Robotic Mapping, Perception and Interaction: A Survey

Sourav Garg, Niko Sünderhauf, Feras Dayoub, Douglas Morrison, Akansel Cosgun, Gustavo Carneiro, Qi Wu, Tat-Jun Chin, Ian Reid, Stephen Gould, Peter Corke and Michael Milford

ISBN: 978-1-68083-768-1

A Survey on the Integration of Machine Learning with Sampling-based Motion Planning

Troy McMahon

Rutgers University
tm799@rutgers.edu

Aravind Sivaramakrishnan

Rutgers University
as2578@rutgers.edu

Edgar Granados

Rutgers University
eg585@rutgers.edu

Kostas E. Bekris

Rutgers University
kb572@rutgers.edu

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Robotics

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

T. McMahan *et al.*. *A Survey on the Integration of Machine Learning with Sampling-based Motion Planning*. Foundations and Trends[®] in Robotics, vol. 9, no. 4, pp. 266–327, 2022.

ISBN: 978-1-63828-135-1
© 2022 T. McMahan *et al.*

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

Foundations and Trends[®] in Robotics
Volume 9, Issue 4, 2022
Editorial Board

Editors-in-Chief

Julie Shah

Massachusetts Institute of Technology

Honorary Editors

Henrik Christensen

University of California, San Diego

Roland Siegwart

ETH Zurich

Editors

Minoru Asada

Osaka University

Antonio Bicchi

University of Pisa

Aude Billard

EPFL

Cynthia Breazeal

*Massachusetts Institute of
Technology*

Oliver Brock

TU Berlin

Wolfram Burgard

University of Freiburg

Udo Frese

University of Bremen

Ken Goldberg

*University of California,
Berkeley*

Hiroshi Ishiguro

Osaka University

Makoto Kaneko

Osaka University

Danica Kragic

KTH Stockholm

Vijay Kumar

University of Pennsylvania

Simon Lacroix

LAAS

Christian Laugier

INRIA

Steve LaValle

*University of Illinois at
Urbana-Champaign*

Yoshihiko Nakamura

The University of Tokyo

Brad Nelson

ETH Zurich

Paul Newman

University of Oxford

Daniela Rus

*Massachusetts Institute of
Technology*

Giulio Sandini

University of Genova

Sebastian Thrun

Stanford University

Manuela Veloso

*Carnegie Mellon
University*

Markus Vincze

Vienna University

Alex Zelinsky

DSTG

Editorial Scope

Topics

Foundations and Trends[®] in Robotics publishes survey and tutorial articles in the following topics:

- Mathematical modelling
- Kinematics
- Dynamics
- Estimation Methods
- Robot Control
- Planning
- Artificial Intelligence in Robotics
- Software Systems and Architectures
- Mechanisms and Actuators
- Sensors and Estimation
- Planning and Control
- Human-Robot Interaction
- Industrial Robotics
- Service Robotics

Information for Librarians

Foundations and Trends[®] in Robotics, 2022, Volume 9, 4 issues. ISSN paper version 1935-8253. ISSN online version 1935-8261. Also available as a combined paper and online subscription.

Contents

1	Introduction	3
2	Sampling-based Motion Planning	8
3	Learning Primitives of Sampling-based Motion Planning	13
3.1	Sampling Sequences	13
3.2	Collision Checking	18
3.3	Distance Metrics, Heuristics and Constraints	23
3.4	Steering Functions and Local Planners	29
3.5	Adaptive Selection of Primitives and Meta-Reasoning	31
4	Learning-based Pipelines	36
4.1	Retrieve and Repair Methods	36
4.2	Neural Motion Planning	37
4.3	Learning a lower dimensional planning space	39
5	SBMP with Learned Models	42
5.1	Modeling Robot Uncertainty	43
5.2	Modeling Uncertainty in Obstacle Dynamics	44

6 Discussion	46
6.1 Computational Efficiency	47
6.2 Path Quality	48
6.3 Ease of Use	48
6.4 Limitations of Existing Methods	49
6.5 Potential Future Work	50
6.6 Alternative Planning Frameworks	51
References	52

A Survey on the Integration of Machine Learning with Sampling-based Motion Planning

Troy McMahon^{1,*}, Aravind Sivaramakrishnan^{1,*}, Edgar Granados¹ and Kostas E. Bekris¹

¹*Department of Computer Science, Rutgers University, NJ, USA; kb572@rutgers.edu, as2578@rutgers.edu*

**Equal contribution.*

ABSTRACT

Sampling-based methods are widely adopted solutions for robot motion planning. The methods are straightforward to implement, effective in practice for many robotic systems. It is often possible to prove that they have desirable properties, such as probabilistic completeness and asymptotic optimality. Nevertheless, they still face challenges as the complexity of the underlying planning problem increases, especially under tight computation time constraints, which impact the quality of returned solutions or given inaccurate models. This has motivated machine learning to improve the computational efficiency and applicability of Sampling-Based Motion Planners (SBMPs). This survey reviews such integrative efforts and aims to provide a classification of the alternative directions that have been explored in the literature. It first discusses how learning has been used to enhance key components of SBMPs, such as node sampling, collision detection, distance or nearest neighbor computation, local

Troy McMahon, Aravind Sivaramakrishnan, Edgar Granados and Kostas E. Bekris (2022), “A Survey on the Integration of Machine Learning with Sampling-based Motion Planning”, *Foundations and Trends® in Robotics*: Vol. 9, No. 4, pp 266–327. DOI: 10.1561/23000000063.

©2022 T. McMahon *et al.*

planning, and termination conditions. Then, it highlights planners that use learning to adaptively select between different implementations of such primitives in response to the underlying problem's features. It also covers emerging methods, which build complete machine learning pipelines that reflect the traditional structure of **SBMPs**. It also discusses how machine learning has been used to provide data-driven models of robots, which can then be used by a **SBMP**. Finally, it provides a comparative discussion of the advantages and disadvantages of the approaches covered, and insights on possible future directions of research. An online version of this survey can be found at: <https://prx-kinodynamic.github.io/>

1

Introduction

Motion planning is the problem of finding valid paths, expressed as sequences of configurations, or trajectories, expressed as sequences of controls, which move a robot from a given start state to a desired goal state while avoiding obstacles. It has applications in problems ranging from mobile robotics (Barraquand and Latombe, 1991), manipulation planning (Oriolo and Mongillo, 2005), graphics and animation (Kallmann *et al.*, 2008), protein folding (Song and Amato, 2001) to crowd simulation (Bayazit *et al.*, 2002; Sud *et al.*, 2008; Toll *et al.*, 2012) and multi-robot applications (Svestka and Overmars, 1995; Clark and Rock, 2001). Variations of the motion planning problem can include dynamic constraints, which can be important in autonomous driving and aerial vehicles (Figure 1.1).

The motion planning problem is PSPACE-Complete (Reif, 1979; Canny, 1988; Latombe, 1991), and its complexity depends exponentially on the number of degrees of freedom of the robotic system. This makes traditional, complete methods difficult to apply for problems with more than 4 or 5 degrees of freedom. This has motivated work on developing sampling-based methods, which often use random samples to explore the underlying configuration space. Examples of popular




Figure 1.1: Example domains where machine learning has been integrated with sampling-based motion planning: a) grasping objects with robotic arms (image from Chamzas *et al.*, 2019), b) dexterous manipulation with adaptive hands (image from Sintov *et al.*, 2020), c) robot navigation of ground robots (image from Granados *et al.*, 2022), d) aerial vehicles (image from Faust *et al.*, 2017).

Sampling-Based Motion Planners (SBMPs) include the Probabilistic Roadmap (PRM) (Kavraki *et al.*, 1996), the Rapidly-exploring Random Tree (RRT) (Lavalle, 1998) and the Expansive Spaces (Hsu *et al.*, 1997) algorithms. Sampling-based motion planners give up on the traditional notion of completeness and instead aim for probabilistic completeness, which means that they are guaranteed to eventually discover a solution if one exists but cannot confirm solution non-existence. Progress in the field has also allowed the development of methods, such as RRT* and PRM* (Karaman and Frazzoli, 2011), which are also asymptotically optimal, i.e., they guarantee convergence to an optimal solution if one exists.

Beyond their algorithmic properties, these methods have proven quite effective in finding solutions for relatively high-dimensional challenges, where the traditional approaches do not scale. Their popularity also stems from the fact that they provide flexible frameworks, which are rather straightforward to implement and adapt for a large variety of robotic systems. Nevertheless, they may still face challenges as the complexity of the underlying planning problem increases:

1. Some of the challenges relate to *computational efficiency*, which may be hindered from the exploration of the underlying configuration space via sampling, especially when the key primitives of these planners, such as collision checking or forward propagation of the system's dynamics, are computationally expensive.
2. Other issues relate to *path quality*. Despite the progress in understanding the conditions for asymptotic optimality, convergence to high-quality solutions may be hindered in practice when naive exploration primitives are employed, such as the random sampling of controls.
3. Furthermore, SBMPs like most motion planning methods, typically assume the *availability of an accurate, complete model*. Traditional, engineered models may be inaccurate or unable to express all critical physical aspects of the problem or not predict how a dynamic environment may evolve. Furthermore, sensing constraints may introduce partial observability and uncertainty about the environment. These factors can limit the applicability of SBMPs.

These challenges motivate the use of machine learning to improve the computational efficiency of SBMPs, accelerate their practical convergence to high-quality solutions, and provide access to accurate-enough, data-driven models, which adapt to varying environmental conditions and sensing input.

Machine learning enables the autonomous derivation of solutions to problems based on prior experience and data. It promises to constantly improve performance by incorporating new data and identifying solutions that engineered approaches may not be able to achieve. This

makes machine learning especially useful for an application like robotics, where a robotic agent must contend with an endless variety of tasks and environments. There are also many scenarios where learned agents can approximate costly computations. In these cases, the learned agent can be trained to model the computations in a pre-processing step, then used during run-time in place of the computation or as a heuristic for it. This is especially useful in robotics, where the robot must act and react to situations in real-time.

Fundamentally, machine learning algorithms operate by building a model of observed data, that can predict and generalize to new examples. This data can come from various sources: an existing dataset, through human training or demonstration, or it can be accumulated from the results of previous predictions that the model has made. For model-based agents, learning is done by fitting the model's parameters to data, which can be done using methods such as regression or reinforcement learning. During query time, the model is queried to give predictions based on patterns observed in the data.

There is a large body of literature on the application of machine learning algorithms to improve the efficiency of robotic systems in general (Kober *et al.*, 2013; Kroemer *et al.*, 2019). Recently, there has been a lot of attention on the progress of deep learning methods, which has resulted in many efforts to utilize the corresponding tools in robotics (Sünderhauf *et al.*, 2018). This survey focuses specifically on integrating machine learning tools to improve the efficiency, convergence, and applicability of SBMPs.

This survey covers a wide breadth of robotic applications, including, but not limited to, mobile robot navigation, manipulation planning, and planning for systems with dynamic constraints. In particular, this monograph first reviews the attempts to use machine learning to improve the performance of individual primitives used by SBMPs (Section 3). It also studies a series of planners that use machine learning to adaptively select from a set of motion planning primitives. It then proceeds to study a series of integrated architectures that learn an end-to-end mapping of sensor inputs to robot trajectories or controls (Section 4). Finally, it studies how SBMPs can operate over learned models of robotic system that account for noise and uncertainty (Section 5).

The survey concludes with a comparative discussion of the different approaches covered in Sections 3 - 5. It evaluates these approaches in terms of their impact on computational efficiency of the planner, quality of the computed paths, and their overall applicability. It then outlines the broad difficulties and limitations of these methods, as well as potential directions of future work.

References

- Aoude, G., B. Luders, J. Joseph, N. Roy, and J. How. (2013). “Probabilistically safe motion planning to avoid dynamic obstacles with uncertain motion patterns”. *Autonomous Robots*. 35: 51–76.
- Arslan, O. and P. Tsiotras. (2015). “Machine learning guided exploration for sampling-based motion planning algorithms”. In: *IROS*.
- Arslan, O. and P. Tsiotras. (2013). “Use of relaxation methods in sampling-based algorithms for optimal motion planning”. In: *ICRA*.
- Bähnemann, R., M. Burri, E. Galceran, R. Siegwart, and J. Nieto. (2017). “Sampling-based motion planning for active multirotor system identification”. In: *ICRA*. 3931–3938.
- Baldwin, I. and P. Newman. (2010). “Non-parametric learning for natural plan generation”. In: *IROS*.
- Barraquand, J. and J.-C. Latombe. (1991). “Nonholonomic multibody mobile robots: controllability and motion planning in the presence of obstacles”. In: *ICRA*.
- Bayazit, O. B., J.-M. Lien, and N. M. Amato. (2002). “Better Group Behaviors in Complex Environments Using Global Roadmaps”. In: *Proceedings of the Eighth International Conference on Artificial Life*.
- Bekris, K. E. and R. Shome. (2020). “Asymptotically Optimal Sampling-Based Planners”. In: *Encyclopedia of Robotics*. Ed. by M. H. Ang, O. Khatib, and B. Siciliano. 1–12.

- Berenson, D., P. Abbeel, and K. Goldberg. (2012). “A robot path planning framework that learns from experience”. In: *ICRA*.
- Bharatheesha, M., W. Caarls, W. J. Wolfslag, and M. Wisse. (2014). “Distance metric approximation for state-space RRTs using supervised learning”. In: *IROS*.
- Bhardwaj, M., S. Choudhury, and S. Scherer. (2017). “Learning Heuristic Search via Imitation”. In: *CoRL*.
- Bhardwaj, M., S. Chowdhury, B. Boots, and S. Srinivasa. (2019). “Leveraging Experience in Lazy Search”. In: *RSS*.
- Bialkowski, J., M. Otte, and E. Frazzoli. (2013). “Free-configuration biased sampling for motion planning”. In: *IROS*.
- Bialkowski, J., M. Otte, S. Karaman, and E. Frazzoli. (2016). “Efficient collision checking in sampling-based motion planning via safety certificates”. *IJRR*. 35: 767–796.
- Bohlin, R. and L. E. Kavraki. (2000). “Path planning using lazy PRM”. In: *ICRA*.
- Bruce, J., K. Caluwaerts, A. Iscen, A. P. Sabelhaus, and V. SunSpiral. (2014). “Design and evolution of a modular tensegrity robot platform”. In: *ICRA*.
- Burns, B. and O. Brock. (2005). “Sampling-Based Motion Planning Using Predictive Models”. In: *ICRA*.
- Burri, M., M. Bloesch, Z. Taylor, R. Siegwart, and J. Nieto. (2018). “A framework for maximum likelihood parameter identification applied on MAVs”. *Journal of Field Robotics*. 35(1): 5–22.
- Canny, J. (1988). *The complexity of robot motion planning*. MIT press.
- Chamzas, C., A. Cullen, A. Shrivastava, and L. E. Kavraki. (2022). “Learning to Retrieve Relevant Experiences for Motion Planning”. *ICRA*.
- Chamzas, C., Z. Kingston, C. Quintero-Peña, A. Shrivastava, and L. E. Kavraki. (2021a). “Learning Sampling Distributions Using Local 3D Workspace Decompositions for Motion Planning in High Dimensions”. *ICRA*.
- Chamzas, C., C. Quintero-Pena, Z. Kingston, A. Orthey, D. Rakita, M. Gleicher, M. Toussaint, and L. E. Kavraki. (2021b). “MOTION-BENCHMARKER: A Tool to Generate and Benchmark Motion Planning Datasets”. *RA-L*. 7(2): 882–889.

- Chamzas, C., A. Shrivastava, and L. E. Kavraki. (2019). “Using local experiences for global motion planning”. In: *ICRA*.
- Chen, B., B. Dai, Q. Lin, G. Ye, H. Liu, and L. Song. (2019). “Learning to Plan in High Dimensions via Neural Exploration-Exploitation Trees”. In: *ICLR*.
- Chiang, H.-T. L., J. Hsu, M. Fiser, L. Tapia, and A. Faust. (2019). “RL-RRT: Kinodynamic Motion Planning via Learning Reachability Estimators From RL Policies”. *RA-L*. 4(4): 4298–4305.
- Chiang, L., A. Faust, S. Sugaya, and L. Tapia. (2018). “Fast Swept Volume Estimation with Deep Learning”. In: *WAFR*.
- Choudhury, S., S. Arora, and S. Scherer. (2015). “The planner ensemble: Motion planning by executing diverse algorithms”. *ICRA*.
- Choudhury, S., S. Javdani, S. Srinivasa, and S. Scherer. (2017). “Near-Optimal Edge Evaluation in Explicit Generalized Binomial Graphs”. In: *NeurIPS*.
- Choudhury, S., S. Srinivasa, and S. Scherer. (2018). “Bayesian Active Edge Evaluation on Expensive Graphs”. *IJCAI*.
- Clark, C. and S. Rock. (2001). “Randomized Motion Planning for Groups of Nonholonomic Robots”. In: *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*.
- Coleman, D., I. Sucas, M. Moll, K. Okada, and N. Correll. (2014). “Experience-Based Planning with Sparse Roadmap Spanners”. *ICRA*.
- Curtis, A., X. Fang, L. P. Kaelbling, T. Lozano-Pérez, and C. R. Garrett. (2022). “Long-horizon manipulation of unknown objects via task and motion planning with estimated affordances”. In: *ICRA*.
- Das, N. and M. Yip. (2020). “Learning-Based Proxy Collision Detection for Robot Motion Planning Applications”. *T-RO*.
- Ekenna, C., S. A. Jacobs, S. Thomas, and N. M. Amato. (2013). “Adaptive neighbor connection for PRMs: A natural fit for heterogeneous environments and parallelism”. In: *IROS*.
- Ekenna, C., S. Thomas, and N. M. Amato. (2015). “Adaptive local learning in sampling based motion planning for protein folding”. In: *BIBM*.

- Faust, A., K. Oslund, O. Ramirez, A. Francis, L. Tapia, M. Fiser, and J. Davidson. (2018). “PRM-RL: Long-range Robotic Navigation Tasks by Combining Reinforcement Learning and Sampling-Based Planning”. In: *ICRA*.
- Faust, A., I. Palunko, P. Cruz, R. Fierro, and L. Tapia. (2017). “Automated aerial suspended cargo delivery through reinforcement learning”. *AI*. 247: 381–398.
- Faust, A., P. Ruymgaart, M. Salman, R. Fierro, and L. Tapia. (2014). “Continuous Action Reinforcement Learning for Control-Affine Systems with Unknown Dynamics”. *IEEE/CAA Journal of Automatica Sinica*. 1(3): 323–336.
- Fulgenzi, C., A. Spalanzani, C. Laugier, and C. Tay. (2010). “Risk based motion planning and navigation in uncertain dynamic environment”. *Research Report*. 14. URL: <https://hal.inria.fr/inria-00526601>.
- Fulgenzi, C., C. Tay, A. Spalanzani, and C. Laugier. (2008). “Probabilistic navigation in dynamic environment using rapidly-exploring random trees and gaussian processes”. In: *IROS*.
- Gammell, J. D., T. D. Barfoot, and S. S. Srinivasa. (2020). “Batch Informed Trees (BIT*): Informed asymptotically optimal anytime search”. *IJRR*. 39(5): 543–567.
- Gammell, J. D. and M. P. Strub. (2021). “Asymptotically Optimal Sampling-Based Motion Planning Methods”. *Annual Review of Control, Robotics, & Autonomous Systems*. 4: 295–318.
- Granados, E., A. Boularias, K. Bekris, and M. Aanjaneya. (2022). “Model Identification and Control of a Low-Cost Wheeled Mobile Robot Using Differentiable Physics”. *ICRA*.
- Granados, E., A. Sivaramakrishnan, T. McMahan, Z. Littlefield, and K. E. Bekris. (2021–2021). “Machine Learning for Kinodynamic Planning (ML4KP)”. URL: <https://github.com/PRX-Kinodynamic/ML4KP>.
- Hauser, K. (2015). “Lazy collision checking in asymptotically-optimal motion planning”. In: *ICRA*.
- Hauser, K., T. Bretl, and J. Latombe. (2005). “Learning-Assisted Multi-Step Planning”. In: *ICRA*.

- Hauser, K. and Y. Zhou. (2016). “Asymptotically Optimal Planning by Feasible Kinodynamic Planning in a State–Cost Space”. *T-RO*. 32(6): 1431–1443.
- Hou, B., S. Choudhury, G. Lee, A. Mandalika, and S. S. Srinivasa. (2020). “Posterior sampling for anytime motion planning on graphs with expensive-to-evaluate edges”. In: *ICRA*.
- Hsu, D., G. Sanchez-Ante, and Zheng Sun. (2005). “Hybrid PRM Sampling with a Cost-Sensitive Adaptive Strategy”. In: *ICRA*.
- Hsu, D., J.-C. Latombe, and R. Motwani. (1997). “Path planning in expansive configuration spaces”. In: *ICRA*.
- Huh, J., V. Isler, and D. D. Lee. (2021). “Cost-to-Go Function Generating Networks for High Dimensional Motion Planning”. *ICRA*.
- Huh, J. and D. D. Lee. (2018). “Efficient sampling with Q-learning to guide rapidly exploring random trees”. *RA-L*. 3(4).
- Huh, J. and D. D. Lee. (2016). “Learning high-dimensional Mixture Models for fast collision detection in Rapidly-Exploring Random Trees”. In: *ICRA*.
- Ichter, B., J. Harrison, and M. Pavone. (2018). “Learning Sampling Distributions for Robot Motion Planning”. In: *ICRA*.
- Ichter, B., E. Schmerling, T. -. E. Lee, and A. Faust. (2020). “Learned Critical Probabilistic Roadmaps for Robotic Motion Planning”. In: *ICRA*.
- Ichter, B. and M. Pavone. (2019). “Robot motion planning in learned latent spaces”. *RA-L*. 4(3): 2407–2414.
- Ichter, B., P. Sermanet, and C. Lynch. (2021). “Broadly-Exploring, Local-Policy Trees for Long-Horizon Task Planning”. In: *5th Annual Conference on Robot Learning*.
- Janson, L., E. Schmerling, A. Clark, and M. Pavone. (2015). “Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions”. *IJRR*. 34(7): 883–921.
- Johnson, J. J., L. Li, F. Liu, A. H. Qureshi, and M. C. Yip. (2020). “Dynamically Constrained Motion Planning Networks for Non-Holonomic Robots”. In: *IROS*.
- Jurgenson, T. and A. Tamar. (2019). “Harnessing Reinforcement Learning for Neural Motion Planning”. In: *RSS*.

- Kalisiak, M. and M. van de Panne. (2006). “RRT-blossom: RRT with a local flood-fill behavior”. In: *ICRA*.
- Kalisiak, M. and M. van de Panne. (2007). “Faster Motion Planning Using Learned Local Viability Models”. In: *ICRA*.
- Kallmann, M., A. Aubel, T. Abaci, and D. Thalmann. (2008). “Planning Collision-Free Reaching Motions for Interactive Object Manipulation and Grasping”. In: *ACM SIGGRAPH 2008 Classes*.
- Karaman, S. and E. Frazzoli. (2011). “Sampling-based algorithms for optimal motion planning”. *IJRR*. 30(7): 846–894.
- Kavraki, L. E., P. Svestka, J. -. Latombe, and M. H. Overmars. (1996). “Probabilistic roadmaps for path planning in high-dimensional configuration spaces”. *T-RO*. 12(4): 566–580.
- Kew, J. C., B. Ichter, M. Bandari, T.-W. E. Lee, and A. Faust. (2020). “Neural Collision Clearance Estimator for Fast Robot Motion Planning”. In: *WAFR*.
- Kim, B., Z. Wang, L. P. Kaelbling, and T. Lozano-Pérez. (2019). “Learning to guide task and motion planning using score-space representation”. *IJRR*. 38(7): 793–812.
- Kimmel, A., A. Sintov, J. Tan, B. Wen, A. Boularias, and K. E. Bekris. (2019). “Belief-Space Planning using Learned Models with Application to Underactuated Hands”. *ISRR*.
- Kingston, Z., C. Chamzas, and L. E. Kavraki. (2021). “Using Experience to Improve Constrained Planning on Foliations for Multi-Modal Problems”. In: *IROS*.
- Kingston, Z., M. Moll, and L. E. Kavraki. (2019). “Exploring implicit spaces for constrained sampling-based planning”. *IJRR*. 38(10-11): 1151–1178.
- Kleinbort, M., E. Granados, K. Solovey, R. Bonalli, K. E. Bekris, and D. Halperin. (2020). “Refined Analysis of Asymptotically-Optimal Kinodynamic Planning in the State-Cost Space”. In: *ICRA*.
- Kleinbort, M., O. Salzman, and D. Halperin. (2016). “Collision Detection or Nearest-Neighbor Search? On the Computational Bottleneck in Sampling-based Motion Planning”. In: *WAFR*.
- Knobloch, A., N. Vahrenkamp, M. Wächter, and T. Asfour. (2018). “Distance-Aware Dynamically Weighted Roadmaps for Motion Planning in Unknown Environments”. *RA-L*. 3(3): 2016–2023.

- Kober, J., J. Bagnell, and J. Peters. (2013). “Reinforcement Learning in Robotics: A Survey”. *IJRR*. 32: 1238–1274.
- Kroemer, O., S. Niekum, and G. D. Konidaris. (2019). “A Review of Robot Learning for Manipulation: Challenges, Representations, and Algorithms”. *CoRR*. abs/1907.03146.
- Kuffner, J. J. and S. M. LaValle. (2000). “RRT-connect: An efficient approach to single-query path planning”. In: *ICRA*.
- Kumar, R., A. Mandalika, S. Choudhury, and S. Srinivasa. (2019). “LEGO: Leveraging Experience in Roadmap Generation for Sampling-Based Planning”. In: *IROS*.
- Latombe, J.-C. (1991). “Robot Motion Planning”.
- Lavalle, S. M. (1998). “Rapidly-Exploring Random Trees: A New Tool for Path Planning”. *Tech. rep.*
- Li, L., Y. Miao, A. H. Qureshi, and M. C. Yip. (2021). “MPC-MPNet: Model-Predictive Motion Planning Networks for Fast, Near-Optimal Planning under Kinodynamic Constraints”. *RA-L*. 6(3): 4496–4503.
- Li, S. and N. T. Dantam. (2021). “Learning Proofs of Motion Planning Infeasibility”. In: *RSS*.
- Li, Y. and K. E. Bekris. (2011). “Learning approximate cost-to-go metrics to improve sampling-based motion planning”. In: *ICRA*.
- Li, Y., R. Cui, Z. Li, and D. Xu. (2018). “Neural Network Approximation Based Near-Optimal Motion Planning With Kinodynamic Constraints Using RRT”. *IEEE Transactions on Industrial Electronics*. 65(11): 8718–8729.
- Li, Y., Z. Littlefield, and K. E. Bekris. (2016). “Asymptotically optimal sampling-based kinodynamic planning”. *IJRR*. 35(5): 528–564.
- Littlefield, Z. and K. Bekris. (2018). “Efficient and Asymptotically Optimal Kinodynamic Motion Planning via Dominance-Informed Regions”. In:
- Liu, K., M. Stadler, and N. Roy. (2020). “Learned Sampling Distributions for Efficient Planning in Hybrid Geometric and Object-Level Representations”. In: *ICRA*.
- Malone, N., B. Rohrer, L. Tapia, R. Lumia, and J. Wood. (2012). “Implementation of an embodied general reinforcement learner on a serial link manipulator”. In: *ICRA*.

- Mandalika, A., S. Choudhury, O. Salzman, and S. Srinivasa. (2019). “Generalized lazy search for robot motion planning: Interleaving search and edge evaluation via event-based toggles”. In: *Proceedings of the International Conference on Automated Planning and Scheduling*. Vol. 29. 745–753.
- McConachie, D., T. Power, P. Mitrano, and D. Berenson. (2020). “Learning When to Trust a Dynamics Model for Planning in Reduced State Spaces”. *RA-L*. 5(2): 3540–3547.
- Moll, M., C. Chamzas, Z. Kingston, and L. E. Kavraki. (2021). “HyperPlan: A Framework for Motion Planning Algorithm Selection and Parameter Optimization”. In: *IROS*.
- Morales, M., L. Tapia, R. Pearce, S. Rodriguez, and N. Amato. (2005). “A Machine Learning Approach for Feature-Sensitive Motion Planning”. In: *WAFR*.
- Mukadam, M., J. Dong, X. Yan, F. Dellaert, and B. Boots. (2018). “Continuous-time Gaussian process motion planning via probabilistic inference”. *IJRR*. 37(11): 1319–1340.
- Oriolo, G. and C. Mongillo. (2005). “Motion Planning for Mobile Manipulators along Given End-effector Paths”. In: *ICRA*.
- Palmieri, L. and K. O. Arras. (2015). “Distance metric learning for RRT-based motion planning with constant-time inference”. In: *ICRA*.
- Pan, J., S. Chitta, and D. Manocha. (2013). “Faster Sample-Based Motion Planning Using Instance-Based Learning”. *WAFR*.
- Phillips-Grafflin, C. and D. Berenson. (2015). “Reproducing Expert-Like Motion in Deformable Environments Using Active Learning and IOC”. In: *ISRR*.
- Qi, C. R., H. Su, K. Mo, and L. J. Guibas. (2017). “Pointnet: Deep learning on point sets for 3d classification and segmentation”. In: *CVPR*.
- Qureshi, A. H., A. Simeonov, M. J. Bency, and M. C. Yip. (2019). “Motion planning networks”. In: *ICRA*.
- Qureshi, A. H., J. Dong, A. Choe, and M. C. Yip. (2020). “Neural Manipulation Planning on Constraint Manifolds”. *RA-L*. 5(4): 6089–6096.
- Qureshi, A. H., J. Dong, A. Baig, and M. C. Yip. (2021). “Constrained Motion Planning Networks X”. *T-RO*: 1–19.

- Ratliff, N., M. Zucker, J. A. Bagnell, and S. Srinivasa. (2009). “CHOMP: Gradient optimization techniques for efficient motion planning”. In: *ICRA*.
- Reif, J. H. (1979). “Complexity of the mover’s problem and generalizations”. In: *20th Annual Symposium on Foundations of Computer Science (SFCS 1979)*. 421–427.
- Sintov, A., A. Kimmel, K. E. Bekris, and A. Boularias. (2020). “Motion Planning with Competency-Aware Transition Models for Underactuated Adaptive Hands”. In: *ICRA*.
- Sivaramakrishnan, A., E. Granados, S. Karten, T. McMahan, and K. E. Bekris. (2021). “Improving Kinodynamic Planners for Vehicular Navigation with Learned Goal-Reaching Controllers”. In: *IROS*.
- Song, G. and N. M. Amato. (2001). “Using Motion Planning to Study Protein Folding Pathways”. In: *Proceedings of the Fifth Annual International Conference on Computational Biology*. 287–296.
- Strudel, R., R. Garcia, J. Carpentier, J.-P. Laumond, I. Laptev, and C. Schmid. (2020). “Learning Obstacle Representations for Neural Motion Planning”. *CoRL*.
- Sud, A., R. Gayle, E. Andersen, S. Guy, M. Lin, and D. Manocha. (2008). “Real-Time Navigation of Independent Agents Using Adaptive Roadmaps”. In: *ACM SIGGRAPH 2008 Classes*.
- Sünderhauf, N., O. Brock, W. Scheirer, R. Hadsell, D. Fox, J. Leitner, B. Upcroft, P. Abbeel, W. Burgard, M. Milford, *et al.* (2018). “The limits and potentials of deep learning for robotics”. *IJRR*. 37(4-5): 405–420.
- Sung, Y., L. P. Kaelbling, and T. Lozano-Pérez. (2021). “Learning When to Quit: Meta-Reasoning for Motion Planning”. *IROS*.
- Sutanto, G., I. M. R. Fernández, P. Englert, R. K. Ramachandran, and G. S. Sukhatme. (2020). “Learning Equality Constraints for Motion Planning on Manifolds”. *CoRL*.
- Svestka, P. and M. Overmars. (1995). “Coordinated motion planning for multiple car-like robots using probabilistic roadmaps”. In: *ICRA*.
- Toll, W. van, A. IV, and R. Geraerts. (2012). “Real-Time Density-Based Crowd Simulation”. *Computer Animation and Virtual Worlds*. 23(Feb.): 59–69.

- Uwacu, D., C. Ekenna, S. L. Thomas, and N. Amato. (2016). “The Impact of Approximate Methods on Local Learning in Motion Planning”. In: *RLP*.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. (2017). “Attention is all you need”. *NeurIPS*.
- Vemula, A., Y. Oza, J. Bagnell, and M. Likhachev. (2020). “Planning and Execution using Inaccurate Models with Provable Guarantees”. In: *R:SS*.
- Wang, J., W. Chi, C. Li, C. Wang, and M. Q. .-. Meng. (2020). “Neural RRT*: Learning-Based Optimal Path Planning”. *T-ASE*.
- Wang, Z., C. R. Garrett, L. P. Kaelbling, and T. Lozano-Pérez. (2021). “Learning compositional models of robot skills for task and motion planning”. *IJRR*. 40(6-7): 866–894.
- Wells, A. M., N. T. Dantam, A. Shrivastava, and L. E. Kavraki. (2019). “Learning feasibility for task and motion planning in tabletop environments”. *IEEE RA-L*. 4(2): 1255–1262.
- Ye, G. and R. Alterovitz. (2017). “Demonstration-Guided Motion Planning”. In: *ISRR*.
- Yu, C. and S. Gao. (2021). “Reducing Collision Checking for Sampling-Based Motion Planning Using Graph Neural Networks”. In: *NeurIPS*.
- Yuan, Y., J. Liu, J. Wang, W. Chi, G. Chen, and L. Sun. (2021). “A Knowledge-Based Fast Motion Planning Method Through Online Environmental Feature Learning”. *ICRA*.
- Zhang, Y., J. Zhang, J. Zhang, J. Wang, K. Lu, and J. Hong. (2020). “A novel learning framework for sampling-based motion planning in autonomous driving”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. No. 01. 1202–1209.
- Zhao, L., J. Putman, W. Wang, and D. Balkcom. (2020). “PLRC*: A piecewise linear regression complex for approximating optimal robot motion”. In: *IROS*.
- Zheng, D. and P. Tsiotras. (2021). “Sampling-based Kinodynamic Motion Planning Using a Neural Network Controller”. In: *AIAA Scitech 2021 Forum*. 1754.
- Zhu, Q. (1991). “Hidden Markov model for dynamic obstacle avoidance of mobile robot navigation”. *T-RO*. 7(3): 390–397.

Zucker, M., J. Kuffner, and J. A. Bagnell. (2008). "Adaptive workspace biasing for sampling-based planners". In: *IROS*.