

João Müller Carvalho

✉ joao@robot-learning.de 🌐 joao-carvalho.net 🐙 github.com/joaoamcarvalho

🏠 Frankfurt am Main, Germany 📞 +49 157 76439163

RESEARCH INTERESTS

In my research, I specialize in developing **machine learning** and **reinforcement learning** algorithms for **robotics**. My work focuses on **deep generative models** (e.g., diffusion, flow matching) for motion planning, grasping, and imitation learning, as well as **deep reinforcement learning** for solving contact-rich tasks and improving policy search.

EDUCATION

Ph.D. in Computer Science — TU Darmstadt, Germany **2020 - 2025**

- Focused on machine learning for robotics.
- Thesis: Enhancing Robot Manipulation Skills through Learning (advisor: [Prof. Jan Peters](#))
- Grade: 1.0 (magna cum laude) (5-1 scale)

M.Sc. in Computer Science — Universität Freiburg, Germany **2016 - 2019**

- Focused on machine learning and reinforcement learning
- Thesis: Nonparametric Off-Policy Policy Gradient (advisor: [Prof. Samuele Tosatto](#))
- Grade: 1.3 (5-1 scale)

M.Sc. & B.Sc. in Electrical and Computer Engineering — Instituto Superior Técnico, Portugal **2007 - 2012**

- Focused on energy systems and control
- Thesis: Electric Distribution Network Loss Calculation and Minimization (advisor: [Prof. Pedro Carvalho](#))
- Grade: 17 (0-20 scale)

Exchange student — TU Delft, The Netherlands **2011 - 2012**

- ERASMUS programme

WORK EXPERIENCE

Research Assistant — TU Darmstadt, Darmstadt, Germany **2019 - present**

- Conducted independent research in ML and robotics at the Intelligent Autonomous Systems (IAS) group.
- Developed novel algorithms for robot manipulation using deep generative models and reinforcement learning.
- Authored and co-authored 10+ works at top-tier ML and robotics conferences/journals (e.g., PAMI, ICLR, AISTATS, IROS).
- Set up and maintained the first GPU cluster computer of IAS, serving 100+ users.
- Prepared teaching materials, presented exercise lectures with 300+ students, and supervised 10+ theses.

Student Teaching Assistant — TU Darmstadt, Darmstadt, Germany **2019**

- Prepared teaching material for the Statistical Machine Learning lecture with 200+ students.

Software Developer (part-time) — Medizinische Planungssysteme GmbH, Freiburg, Germany **2017 - 2018**

- Developed software for infrastructure and deployment tools within the DevOps team.

Electrical Engineer — Energias de Portugal SA, Lisbon, Portugal **2012 - 2016**

- Developed home automation products using different communication protocols (Wi-Fi, ZigBee, PLC).
- Extracted and analyzed power consumption data and built electricity demand forecasting models.
- Developed software tools to automate analysis and reduce simulation times of electricity and gas forecasting models.

Student Research Assistant — Instituto Superior Técnico / INESC-ID, Lisbon, Portugal **2010 - 2012**

- Developed statistical algorithms to estimate and minimize power losses in low-voltage distribution networks. (supervisor: [Prof. Pedro Carvalho](#))
- Developed discrete optimal control algorithms to optimize a vehicle's fuel consumption in cruise control mode. (supervisor: [Prof. João Miranda Lemos](#))

TECHNICAL SKILLS & LANGUAGES

- **Programming:** Python, Numpy, PyTorch, C++, Linux, Git/GitHub, LaTeX, PyBullet, IsaacGym, ROS
- **Languages:** Portuguese (native), English (fluent), German (fluent), Spanish (intermediate)

AWARDS & GRANTS

- **ROBOSTRUCT project** — cooperation with Volkswagen AG within the Software Campus 2024-2025, sponsored by the German Federal Ministry of Education and Research (BMBF). Wrote the proposal and secured 115k Eur in research funding.
- **Best paper award** — R:SS 2024 Workshop Priors4Robots
- **Best interactive paper award finalist** — 2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)
- **Travel award** — 2022 Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)

OTHERS

Research projects

| Project | Funding agency | Role |
|-----------------------|--|-----------------------------|
| ROBOSTRUCT (01S23067) | German Federal Ministry of Education and Research (BMBF) | Researcher, project manager |
| IKIDA (1IS20045) | German Federal Ministry of Education and Research (BMBF) | Researcher |
| Kobo34 (16SV798) | German Federal Ministry of Education and Research (BMBF) | Researcher |
| SE2A (109580PRJ) | Portuguese Foundation for Science and Technology (FCT) | Student researcher |

Further education

Mediterranean Machine Learning Summer School 2021 Nordic Probabilistic AI School 2022

Scientific reviewing

Robotics and Automation Letters (RA-L) Conference on Robot Learning (CoRL)
International Conference on Intelligent Robots (IROS) International Conference on Robotics and Automation (ICRA)
Robotics: Science and Systems (R:SS) NeurIPS Workshop on Interpretable Inductive Biases

Teaching

- Teaching assistant at TU Darmstadt for the lectures:

Robot Learning Integrated Project (2022) Computational Engineering and Robotics (German) (2020, 2021)
Robot Learning (2020) Statistical Machine Learning (2019)

- Supervised several students' master's and bachelor's theses and research projects (see list below).

Invited Talks

- Motion Planning Diffusion, University of Mannheim, Prof. Dr. Leif Döring, 2023

PUBLICATIONS

Full list at <https://scholar.google.com/citations?user=xYUxMF0AAAAJ&hl=en>

Preprints

1. **Carvalho, J.**; Le, A.; Kicki, P.; Koert, D.; Peters, J. (submitted). Motion Planning Diffusion: Learning and Adapting Robot Motion Planning with Diffusion Models
2. **Carvalho, J.**; Le, A.; Jahr, P.; Sun, Q.; Urain, J.; Koert, D.; Peters, J. (submitted). Grasp Diffusion Network: Learning Grasp Generators from Partial Point Clouds with Diffusion Models in $SO(3)\times R^3$
3. Funk, N.; Urain, J.; **Carvalho, J.**; Prasad, V.; Chalvatzaki, G.; Peters, J. (submitted). ActionFlow: Equivariant, Accurate, and Efficient Policies with Spatially Symmetric Flow Matching.
4. Palenicek, D.; Lutter, M.; **Carvalho, J.**; Dennert, D.; Ahmad, F.; Peters, J. (submitted). Diminishing Return of Value Expansion Methods
5. Le, A.; Hansel, K.; **Carvalho, J.**; Watson, J.; Urain, J.; Biess, A.; Chalvatzaki, G.; Peters, J. (submitted). Global Tensor Motion Planning

Journal Papers

6. Tosatto, S.; **Carvalho, J.**; Peters, J. (2022). Batch Reinforcement Learning with a Nonparametric Off-Policy Policy Gradient, IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)

Conference Papers

7. **Carvalho, J.**; Le, A.; Baierl, M.; Koert, D.; Peters, J. (2023). Motion Planning Diffusion: Learning and Planning of Robot Motions with Diffusion Models, 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

8. Palenicek, D.; Lutter, M.; **Carvalho, J.**; Peters, J. (2023). Diminishing Return of Value Expansion Methods in Model-Based Reinforcement Learning, International Conference on Learning Representations (ICLR)
9. **Carvalho, J.**; Koert, D.; Daniv, M.; Peters, J. (2022). Adapting Object-Centric Probabilistic Movement Primitives with Residual Reinforcement Learning, 2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)
10. Vorndamme, J.; **Carvalho, J.**; Laha, Riddhiman; Koert, D.; Figueredo, L.; Peters, J.; Haddadin, S. (2022). Integrated Bi-Manual Motion Generation and Control shaped for Probabilistic Movement Primitives, 2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids) — **Best interactive paper award finalist**
11. **Carvalho, J.**; Tateo, D.; Muratore, F.; Peters, J. (2021). An Empirical Analysis of MeasureValued Derivatives for Policy Gradients, International Joint Conference on Neural Networks (IJCNN)
12. Tosatto, S.; **Carvalho, J.**; Abdulsamad, H.; Peters, J. (2020). A Nonparametric OffPolicy Policy Gradient, Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)

Workshop Papers

13. Le, A. T.; Hansel, K.; **Carvalho, J.**; Urain, J.; Biess, A.; Chalvatzaki, G.; Peters, J. (2024). Global Tensor Motion Planning, CoRL 2024 Workshop on Differentiable Optimization Everywhere
14. Funk, N., Urain, J., **Carvalho, J.**, Prasad, V., Chalvatzaki, G., Peters, J. (2024). ActionFlow: Equivariant, Accurate, and Efficient Manipulation Policies with Flow Matching, CoRL 2024 Workshop on Mastering Robot Manipulation in a World of Abundant Data
15. Funk, N., Urain, J., **Carvalho, J.**, Prasad, V., Chalvatzaki, G., Peters, J. (2024). ActionFlow: Efficient, Accurate, and Fast Policies with Spatially Symmetric Flow Matching, R:SS 2024 Workshop Priors4Robots — **Best paper award**
16. **Carvalho, J.**; Baierl, M.; Urain, J.; Peters, J. (2022). Conditioned Score-Based Models for Learning Collision-Free Trajectory Generation, NeurIPS 2022 Workshop on Score-Based Methods
17. **Carvalho, J.**; Peters, J. (2022). An Analysis of Measure-Valued Derivatives for Policy Gradients, Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM) — **Travel award**

Book Chapters

18. Watson, J., Urain, J., **Carvalho, J.**, Funk, N., Peters, J., Robot Learning: An Introduction, Robotics Goes MOOC, Springer International — In preparation

STUDENT SUPERVISION

Theses

1. Jahr, P. (2025). Residual Stable Vector Fields
2. Striebel, N. (2025). Bimanual Imitation Learning (w/ Niklas Funk, Michael Drolet)
3. Sun Q. (2024). Grasp Diffusion Network (w/ An Thai Le)
4. Kappes, N. (2023). Natural Gradient Optimistic Actor-Critic
5. Hilt, F. (2023). Statistical Model-Based Reinforcement Learning (w/ Joe Watson)
6. Keller, L. (2023). Context-Dependent Variable Impedance Control with Stability Guarantees (w/ Dorothea Koert)
7. Herrmann, P. (2023). 6DCenterPose: Multi-object RGB-D 6D pose tracking with synthetic training data (w/ Suman Pal)
8. Brosseit, J. (2023). The Principle of Value Equivalence for Policy Gradient Search
9. Baierl, M. (2023). Score-Based Generative Models as Trajectory Priors for Motion Planning (w/ Julen Urain, An Thai Le)
10. Hellwig, J. (2023). Residual Reinforcement Learning with Stable Priors
11. Xue, C. (2022). Robot Task Classification and Local Manipulation Controllers (w/ Suman Pal)
12. Zhao, P. (2021). Improving Gradient Directions for Episodic Policy Search
13. Kaemmerer, M. (2021). Measure-Valued Derivatives for Machine Learning
14. Daniv, M. (2022). Graph-Based Model Predictive Visual Imitation Learning

Research Projects

1. Striebel, N., Mulder, A. (2024). Reinforcement Learning of Insertion Tasks: A Comparison Between Policy Structures
2. Striebel, N., Mulder, A. (2023). Building a Framework to Solve Insertion Tasks with Residual Reinforcement Learning in the Real World
3. Meier, H. (2023). Model-Based Multi-Object 6D Pose Estimation (w/ Felix Kaiser, Arjun Vir Datta)
4. Kappes, N., Herrmann, P. (2021). Trust Region Optimistic Actor-Critic
5. Kappes, N., Herrmann, P. (2021). Second Order Extension of Optimistic Actor-Critic
6. Hellwig, J., Baierl, M. (2021). A Hierarchical Approach to Active Pose Estimation (w/ Julen Urain)
7. Hellwig, J., Baierl, M. (2021). Active Visual Search with Partially Observable Monte-Carlo Planning (w/ Julen Urain)
8. Hilt, F., Kolf, J., Weiland, C. (2021). Graph neural networks for robotic manipulation
9. Hilt, F., Kolf, J., Weiland, C. (2020). Balloon Estimators for Improving and Scaling the Nonparametric Off-Policy Policy Gradient (w/ Samuele Tosatto)
10. Musekamp, D., Rettig, M. (2020). Learning Robot Skills From Video Data (w/ Dorothea Koert)
11. Shi, J. (2023). Rapid Adaptation for Contact-Rich Tasks