



Offer #2025-08814

PhD Position F/M PhD position F/M Data-driven methods for vision-based robotic motion control

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Fonction : PhD Position

Context

The Willow team will be a resourceful environment to carry out this project, as it is recognized for its contributions in the fields of computer vision and robotics. The research led during this PhD will also benefit from the regional AI ecosystem via the Parisian AI institute PR[AI]RIE-PSAI.

Non-discrimination, openness and transparency: Partners of PR[AI]RIE-PSAI are committed to supporting and promoting equality, diversity, and inclusion within their communities. We encourage applications from diverse backgrounds, which we will ensure are selected through an open and transparent recruitment process. (

Français : Non-discrimination, ouverture et transparence : L'ensemble des partenaires de PR[AI]RIE-PSAI s'engagent à soutenir et promouvoir l'égalité, la diversité et l'inclusion au sein de ses communautés. Nous encourageons les candidatures issues de profils variés, que nous veillerons à sélectionner via un processus de recrutement ouvert et transparent.)

Assignment

Context: Motion control in robotics is subject to Moravec's paradox: robots can execute physically impressive motions, yet they fail at seemingly simple tasks. The

most impressive humanoid robot today, the Atlas from BostonDynamics, came to fame by performing athletic back-flips; yet it would not be able to get up from lying in bed, unless specialized engineers work on implementing that new behavior. All its motions are executed under a set of assumptions about the robot's environment, for instance: the robot has its torso tilted less than 90 degrees from gravity, has its feet on a flat floor or a moderately-tilted terrain, is facing stairs with a specific number of steps, etc.

This approach has led to a *decoupling* between perception and locomotion. On the one hand, perception experts work on geometric *exteroception* problems such as detecting walls and steppable surfaces; on the other hand, locomotion experts implement control strategies that assume knowledge of the environment and focus on *proprioception* (torque and force measurements, inertial measurements, etc.). The main drawback of this decoupling, however, is that it makes both the vision and locomotion problems harder than they would be if they were addressed jointly. For instance, in a study of a quadruped robot walking in a forest, Miki *et al.* [1] observed that surface reconstruction methods would frequently fail, in which case the locomotion behavior essentially downgraded to blind locomotion. (Needless to say, walking blindly in a forest is high-risk, even for humans.)

The current paradigm to make robots walk outside of controlled lab environments relies on deep reinforcement learning from massively-parallel simulations [2, 3]. It does not revisit the extero-proprioceptive decoupling: rather, locomotion policies are made robust against defective perception via domain randomization. Including vision involves a major update to this paradigm, as the ability to see objects from afar fundamentally creates a synchronization bottleneck that breaks simulation parallelism. In this thesis, our plan is to explore a line of ideas orthogonal to the simulation-intensive approach.

Scientific objectives: This project explores questions that arise when relaxing assumptions about the structure of the world that are at the core of the extero-proprioceptive decoupling. What if locomotion is allowed to decide motions from implicit rather than explicit representations? What if vision contributed to locomotory decisions, and not only the other way round?

This thesis is articulated around three main axes: *perception*, *control*, and *learning* from real-robot data. Scientifically, our key idea is to leverage physical contacts on legged robots to establish ground-truth validation between vision and proprioception. We will focus on real-robot data rather than massively-parallel simulation, using low-cost open-source robots for rapid prototyping and data collection. Our first application will be in perception, where we will study the question of contact estimation using both visual and motor data. We will then consider the broader topic of including visual inputs to extend model predictive control into interpretable motion policies. Our overall objective throughout the work will be to evaluate how learning from limited real-robot data, but including visual inputs, can be applied on real-robots to solve challenging tasks such as agile locomotion.

Application process: Applications will only be considered if they are submitted online from the Inria website. The deadline for submitting an application is *May 15th, 2025*. After this deadline, a screening process will take place and results will

be communicated in two stages:

- Pre-screening until May 30th, 2025, at 1:00pm CEST.
- Final selection by a PRAIRIE-PSAI committee before June 15th, 2025.

Each application should include:

- An up-to-date CV
- A one-page motivation letter covering (1) the candidate's ambitions for this topic and (2) the candidacy's fit to the PhD topic described below.
- Scans of the latest diplomas.

Main activities

Research plan

Axis 1: Connecting vision and proprioception through contact estimation

One goal of this project will be to handle visual and proprioceptive data jointly when learning new motion policies. We will consider contact estimation as an approach to validate the connection between the two. For limbed robots, contact estimation is the problem of determining whether any part of a limb, for instance the sole of a foot on a humanoid leg, is in contact with the environment. When it comes to estimating contact, vision validates the no-contact hypothesis (if one sees space between two bodies, they are not touching) whereas proprioception validates the in-contact hypothesis (if one feels resistance below their feet, they are on the ground). Contact estimation is commonly solved using prior models, such as in contact-aided invariant Kalman filtering [4] or probabilistic contact estimation [5, 6]. These methods rely on priors due to data scarcity, as it is expensive to collect data from large expensive robots operated by specialized technicians [4, 7, 8].

Our proposal is to collect a joint visual and proprioceptive dataset from real-robot data, from which we may learn visual representations and motion control simultaneously. Technically, we will be able to collect larger datasets on open-source wheeled-legged robots available at Inria Paris, which are easier to operate and cheaper to maintain than large-scale legged robots, yet have the same challenging properties for locomotion (underactuated dynamics, importance of collision avoidance, ...) Scientifically, we will follow up on the idea of encoding visual inputs to a latent space and learning latent-space dynamics [9, 10], with the novelty of taking contact constraints into account. For instance, if the robot is making contact with a wall in front of it, the learned dynamics will be trained to predict that trying to go forward will result in increased proprioceptive contact forces and marginal visual changes (and conversely, visual motion and marginal force increase in the absence of contact).

Our goal will thus be to lay the foundations for an implicit representation shared between vision and locomotion components, attacking the problem through the well-

defined question of contact estimation, where we will have existing baselines to compare against, and an original angle in terms of methods, with data-based machine learning rather than model-based state estimation.

Axis 2: Model predictive control with visual inputs

Controlling physical robots means dealing with complex and various sensory inputs such as vision, velocity, acceleration and force measurements, etc. In agile robot locomotion, optimal control has stood out as a relevant paradigm to derive effective controllers, whether it is via model predictive control and online numerical optimization [11, 7] or via reinforcement learning [12, 3]. The main driver behind this adoption is that optimal control represents and adapts to the physics underlying the problem at hand. Yet, optimal control requires a model of forward dynamics $\dot{x} = f(x, u)$, which is discretized as $x_{t+1} = f_d(x_t, u_t)$ and unrolled either directly in model predictive control optimizations or indirectly in a simulator training a parameterized policy by reinforcement learning. This pipeline has worked successfully for low-dimensional proprioceptive inputs that map nicely to states x_t , yet lead to blind policies. How can we deal with visual inputs to train perceptive policies?

In this axis, we will explore machine learning of maps from visual inputs to not only system dynamics (as in Axis 1) but full-fledged optimal control problems. Advances on this topic have been made possible thanks to recent works on the differentiation of convex optimization problems [13] and convex optimal control [14, 10]. We will focus on model predictive control, where optimal control problems are solved repeatedly over a receding horizon. The receding horizon consists of two parts: near-future, where dynamics and constraints are fully taken into account while optimizing an objective function, and post-horizon, where an approximation of the value function for the terminal state is used to approximate infinite-time optimization. We will explore how visual inputs can map to both of these regimes. In the near-future regime, via (i) initial state, (ii) objective function, (iii) system dynamics and (iv) constraints, and in the post-horizon regime, via (v) value function approximation. Preliminary results [14, 15] showed that nonlinear tasks could be approximated by a convex model predictive controller using a vision-trained map to the objective function (ii). In this axis, we will consider the whole spectrum (i)–(v) where vision can enrich model predictive control to solve more complex perceptive tasks.

Unlike reinforcement learning of black-box function approximators, the problems we will predict from vision will also be interpretable. For instance, a collision avoidance task mapped to (iv) constraints will produce a polytope (hence a volume of space that we can visualize) that the model predictive controller will certifiably avoid. We will include in our study questions not only of optimizing task performance, but also of trading it off with model interpretability with applications to user feedback.

Axis 3: Sim-to-real transfer of vision-based policies

Experimental validation in simulation is an important step to assess the robustness

and capabilities of a proposed solution. In robotics, the sim-to-real gap is particularly dominant: mathematical models of the robot, the environment and their interactions are usually simplified, prompting the need to validate solutions on real hardware as often as possible. In this project, we will work with hardware and real-robot data distributions right from the start. The data necessary to train our controllers will be gathered from wheeled-biped robots built and maintained at Inria.

We will consider tasks that are challenging to transfer to real robots, such as stair climbing. Stair climbing is challenging even for seasoned roller skaters, and has not been demonstrated dynamically yet on wheeled bipeds. It also encompasses the salient components of visual predictive control: stairs are only visible from afar, prompting practical confrontation with a split receding horizon. As a first challenge, we will consider the task of continuous stair climbing by always keeping both wheels on the ground. This approach will require fine contact estimation, as the robot will need to be able to discriminate between vertical and horizontal forces exerted on its wheels, allowing us to evaluate the effectiveness of the representations built in Axis 1. We will then further consider the question of dynamic stair climbing, the richer behavior where the robot is allowed to lift its legs and cannot stop mid-step. The dynamic version of the task is achievable, as demonstrated by seasoned roller skaters, yet has not been demonstrated on any wheeled-biped robot so far.

References

1. T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning robust perceptive locomotion for quadrupedal robots in the wild,” *Science Robotics*, vol. 7, no. 62, 2022.
2. A. Kumar, Z. Fu, D. Pathak, and J. Malik, “Rma: Rapid motor adaptation for legged robots,” in *Robotics: Science and Systems*, 2021.
3. N. Rudin, D. Hoeller, P. Reist, and M. Hutter, “Learning to walk in minutes using massively parallel deep reinforcement learning,” in *Proceedings of the 5th Conference on Robot Learning*, vol. 164 of *Proceedings of Machine Learning Research*, pp. 91–100, PMLR, 08–11 Nov 2022.
4. R. Hartley, M. G. Jadidi, J. Grizzle, and R. M. Eustice, “Contact-aided invariant extended kalman filtering for legged robot state estimation,” in *Proceedings of Robotics: Science and Systems*, (Pittsburgh, Pennsylvania), June 2018.
5. J. Hwangbo, C. D. Bellicoso, P. Fankhauser, and M. Hutter, “Probabilistic foot contact estimation by fusing information from dynamics and differential/forward kinematics,” in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3872–3878, IEEE, 2016.
6. U. B. Gökbakan, F. D. ?umbgen, and S. Caron, “A Data-driven Contact Estimation Method for Wheeled-Biped Robots,” in *IEEE International Conference on Robotics and Automation*, May 2025.
7. S. Caron, A. Kheddar, and O. Tempier, “Stair climbing stabilization of the hrp-4 humanoid robot using whole-body admittance control,” in *2019 International conference on robotics and automation (ICRA)*, pp. 277–283, IEEE, 2019.
8. J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning quadrupedal locomotion over challenging terrain,” *Science robotics*, vol. 5,

- no. 47, 2020.
9. N. Hansen, X. Wang, and H. Su, “Temporal difference learning for model predictive control,” in ICML, 2022.
 10. O. Bounou, J. Ponce, and J. Carpentier, “Learning system dynamics from sensory input under optimal control principles,” in CDC 2024 Conference on Decision and Control, 2024.
 11. J. Di Carlo, P. M. Wensing, B. Katz, G. Bledt, and S. Kim, “Dynamic locomotion in the mit cheetah 3 through convex model-predictivecontrol,” in 2018 IEEE/RSJ international conference on intelligent robots and systems (IROS), pp. 1–9, IEEE, 2018.
 12. J. Tan, T. Zhang, E. Coumans, A. Iscen, Y. Bai, D. Hafner, S. Bohez, and V. Vanhoucke, “Sim-to-real: Learning agile locomotion for quadruped robots,” arXiv preprint arXiv:1804.10332, 2018.
 13. A. Bambade, F. Schramm, A. Taylor, and J. Carpentier, “Qplayer: efficient differentiation of convex quadratic optimization,” 2023.
 14. A. Meduri, H. Zhu, A. Jordana, and L. Righetti, “Mpc with sensor-based online cost adaptation,” in 2023 IEEE International Conference on Robotics and Automation (ICRA), pp. 996–1002, IEEE, 2023.
 15. V. Tordjman-Levavasseur and S. Caron, “Collision avoidance from monocular vision trained with novel view synthesis.”, pre-print, Mar. 2025.

Skills

- Skills: robotics (M2), computer vision (M2), machine learning (M2), Python (advanced)
- Language: English (French is a plus)
- Additional skills (not required but appreciated): convex optimization, C++, Linux, Git

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Leave: 7 weeks of annual leave + 10 extra days off due to RTT (statutory reduction in working hours) + possibility of exceptional leave (sick children, moving home, etc.)
- Possibility of teleworking and flexible organization of working hours
- Professional equipment available (videoconferencing, loan of computer equipment, etc.)
- Social, cultural and sports events and activities
- Access to vocational training
- Social security coverage

General Information

- **Theme/Domain** : Vision, perception and multimedia interpretation
Scientific computing (BAP E)
- **Town/city** : Paris
- **Inria Center** : [Centre Inria de Paris](#)
- **Starting date** : 2025-09-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2025-06-30

Contacts

- **Inria Team** : [WILLOW](#)
- **PhD Supervisor** :
Caron Stephane / stephane.caron@inria.fr

About Inria

Inria is the French national research institute dedicated to digital science and technology. It employs 2,600 people. Its 200 agile project teams, generally run jointly with academic partners, include more than 3,500 scientists and engineers working to meet the challenges of digital technology, often at the interface with other disciplines. The Institute also employs numerous talents in over forty different professions. 900 research support staff contribute to the preparation and development of scientific and entrepreneurial projects that have a worldwide impact.

The keys to success

Carrying out this PhD project will require in particular:

- Integrating into a dynamic scientific environment: having an analytical mind, but also a taste for learning, listening and sharing thoughts will be essential.
- Past studies in machine learning or robotics at the M2 level, including motion planning, kinematics and dynamics modeling, as well as computer vision at the M2 level.
- Previous experience of scientific research during an M2 internship.
- Appetite for experimenting on real robots.

Warning : you must enter your e-mail address in order to save your application to Inria. Applications must be submitted online on the Inria website. Processing of applications sent from other channels is not guaranteed.

Instruction to apply

Defence Security :

This position is likely to be situated in a restricted area (ZRR), as defined in Decree No. 2011-1425 relating to the protection of national scientific and technical potential (PPST). Authorisation to enter an area is granted by the director of the unit, following a favourable Ministerial decision, as defined in the decree of 3 July 2012 relating to the PPST. An unfavourable Ministerial decision in respect of a position situated in a ZRR would result in the cancellation of the appointment.

Recruitment Policy :

As part of its diversity policy, all Inria positions are accessible to people with disabilities.