Open Internship in the ESA Advanced Concepts Team in 2017

on

The evolution of differentiable robotic behaviors

Stage topic Description

The information contained in differentials (i.e. derivatives) is of paramount importance in machine learning. In reinforcement learning, policy gradient methods seek, for example, to use this information to learn optimal policies. In supervised learning, back-propagation methods are based on error gradients, gradient boosting methods are based on the error function differential structure. In the Covariance Matrix Adaptation Evolutionary Strategy, the covariance matrix can be seen as an inverse Hessian approximate. All these examples, indicate clearly how crucial the use of first and second order derivatives is to successful algorithms in machine learning algorithms. But what about higher order differentials? The open source AuDi tool recently developed by Advanced Concepts Team scientists allows for the efficient computation of high order derivatives, opening the way to studies on the use of such an information (a.k.a. High Order Taylor maps) in learning and, in particular, in robotics. In this internship we will explore this research avenue with a particular attention to space robotics problems.

Candidate's tasks

- Get acquainted with AuDi and the computation of High Order Taylor Maps in dynamical systems
- Study the use of high order Taylor expansion of the reward (error) for deterministic policy learning
- Assemble a number of tests cases of interest to space robotics and make a comparison to evolutionary robotics approaches and to reinforcement learning approaches.

The ideal candidate

Mandatory: excellent python skills. excellent knowledge in machine learning

Desirable: knowledge on differential algebra or automated differentiation, knowledge of differential calculus.

References

- Dario Izzo, & Francesco Biscani. (2016). AuDi: First Release (http://doi.org/10.5281/zenodo.164628)