Hessisches Kompetenzzentrum
für Hochleistungsrechnen

Stochastic Search for Applications in Robotics



Project Manager
Dr. Oleg Arenz

Researchers
Philipp Becker and Maximilian Hüttenrauch

Principal Investigator
Prof. Jan Peters (PhD)

Project Term
2019 - 2020

Clusters
Lichtenberg Cluster Darmstadt

Software
TensorFlow

Institute
Intelligent Autonomous Systems

University
Technische Universität Darmstadt

Introduction

Stochastic-search algorithms are problem independent algorithms well-suited for black-box optimization of an objective function. They only require function evaluations and are used when the objective function cannot be modeled analytically and no gradient information is available. This is often the case for real world problems such as robotics, medical applications, or forensic identification. While the typical stochastic search objective is to find a point estimate for the optimal solution, in certain problem domains such as variational inference, density estimation, or reinforcement learning, it is important to not only learn a single parameter vector that performs well. Instead, we want to capture a meaningful distribution over parameters that maximizes the expected fitness. We re-introduce Model-based Relative Entropy Search (MORE) a versatile, general purpose stochastic-search optimization algorithm, based on information-theoretic principles, which incorporates an entropy objective and is, thus, suitable for learning such distributions over parameters. Based on experimental findings, we aim to improve MORE in terms of convergence speed and sample complexity by applying a (block) coordinate descent strategy on the updates for the mean and covariance of the search distribution. Additionally, we present a more advanced entropy regulation strategy based on an evolution path, similar to the one already employed in CMA-ES. To evaluate our improved version of MORE, we run extensive experiments on simulated robotics tasks, as well as a set of benchmark optimization functions.

Methods

JUSTUS-LIEBIG-UNIVERSITÄT GIESSEN
TECHNISCHE UNIVERSITÄT DARMSTADT
UNIKASSEL VERSITÄT
Philipps Universität Marburg

We investigate several changes to the original MORE algorihms. MORE limits the change of the search distribution during every update in terms of its Kullback-Leibler divergence. We propose to decouple the mean and covariance update by employing a block coordinate ascent strategy for the mean and the covariance matrix which allows for setting different bounds on each of these components. We introduce a scheme for adapting the entropy during optimization, which is similar to the evolution paths used by related stochastic-serach algorithms. Additionally, we propose an adaptive iterative normalization and clipping scheme for preprocessing the data at every iteration, which is based on the excess kurtosis of the target values. We evaluate these changes for episodic reinforcement learning, and on typical benchmark suits for black-box optimization.

Results

We first notice that our modified MORE algorithm (CAS-MORE) has a significant advantage over the original formulation of MORE on almost all test functions. Furthermore, we can see that CAS-MORE is competitive to or improving state of the art results by competitors like CMA-ES and XNES, especially on functions from the group of moderate and ill-conditioned functions. Some functions on the other hand, for example the Attractive Sector function, MORE has trouble optimizing. We found this is mainly due to difficulties estimating a good model for this objective. Other hard functions include multi modal functions such as the Rastrigin function, where MORE often focuses on a local optimum.

Discussion

We presented CAS-MORE, a new version of Model-based Relative Entropy Stochastic Search, based on a coordinate ascent strategy on the mean and covariance of the search distribution and an adaptive entropy schedule. We are competitive on benchmark functions and show the benefits of learning a model and optimizing the expectation of a distribution on reinforcement learning tasks that involve noise indicates an advantage over other approaches. Open research questions are improved model learning and sampling strategies and scaling to higher dimensions.

Last Update: 2022-03-28 15:42