THEORY OF EVOLUTIONARY COMPUTATION: ADVANCED TUTORIA

Kristian Lehre, Pietro S. Oliveto

Populations are at the heart of evolutionary algorithms (EAs). They provide the genetic variation which selection acts upon. A complete picture of EAs can only be obtained if we understand their population dynamics. Continuing from the first tutorial, this second tutorial focuses on techniques that facilitate runtime analysis of complex, population-based EAs.

We first give a brief overview of the population-based EAs that are covered by the techniques. We recall the common stochastic selection mechanisms and how to measure the selection pressure they induce. The main part of the tutorial covers in detail widely applicable techniques tailored to the analysis of populations such as advanced drift theorems and level-based analysis.

To illustrate the application of these techniques, we consider several fundamental questions: What are some conditions that prevent EAs from optimising a function efficiently? What is the appropriate balance between exploration and exploitation and how does this depend on relationships between mutation and selection rates? When are populations necessary for efficient optimisation with EAs? What determines an EA's tolerance for uncertainty, e.g. in the form of noisy or partially available fitness?