

Bull, Adam D.

Convergence rates of efficient global optimization algorithms. (English) Zbl 1280.90094

J. Mach. Learn. Res. 12, 2879-2904 (2011).

Summary: In the efficient global optimization problem, we minimize an unknown function f , using as few observations $f(x)$ as possible. It can be considered a continuum-armed-bandit problem, with noiseless data, and simple regret. Expected-improvement algorithms are perhaps the most popular methods for solving the problem; in this paper, we provide theoretical results on their asymptotic behaviour. Implementing these algorithms requires a choice of Gaussian-process prior, which determines an associated space of functions, its reproducing-kernel Hilbert space (RKHS). When the prior is fixed, expected improvement is known to converge on the minimum of any function in its RKHS. We provide convergence rates for this procedure, optimal for functions of low smoothness, and describe a modified algorithm attaining optimal rates for smoother functions. In practice, however, priors are typically estimated sequentially from the data. For standard estimators, we show this procedure may never find the minimum of f . We then propose alternative estimators, chosen to minimize the constants in the rate of convergence, and show these estimators retain the convergence rates of a fixed prior.

MSC:

90C26 Nonconvex programming, global optimization

60G15 Gaussian processes

62G20 Asymptotic properties of nonparametric inference

Cited in **25** Documents

Keywords:

convergence rates; efficient global optimization; expected improvement; continuum-armed bandit; Bayesian optimization

Full Text: [Link arXiv](#)