Learning from Failures in Evolutionary Computation

Workshop at GECCO 2009

Organized by Nicola Beume and Mike Preuss, TU Dortmund

Rez-de-Chaussee / Lobby Level, Verriere A
Wednesday, July 8, 2009, 14:00 - 18:00

Program and Abstracts

14:00 - 14:10	Opening by Organizers

14:10 - 15:00	*Failures as Stepping Stones to Success* *or "per aspera ad astra"* Hans-Paul Schwefel (Invited Speaker)

The implicit thesis of this talk's title will be underpinned with some examples from (my) real life. A first example leads back to the 1960s, when I simulated the (1+1)-ES with discrete mutations on a two-dimensional parabolic ridge by means of a Z23 computer. The result - getting stuck in certain search directions - led to making use of gaussian variations. The second example comes from experimental investigations to determine the shape of a hot water flashing nozzle, the water being really hot and not simulated on a computer. In search for a multimembered evolutionary algorithm with effective self-adaptation of the mutation strengths, a couple of failures occurred. These, however, rendered deep insight into basic prerequisites to achieve the goal. And finally, some theory will be re-presented about the optimal failure rate in two black-box situations.

15:00 – 15:25	*On the Limitations of Adaptive Resampling using the Student's t-test in Evolution Strategies* Johannes W. Kruisselbrink, Michael T.M. Emmerich, Thomas Bäck

Evolutionary Algorithms are believed to be relatively robust on noisy objective functions, but generally stagnate in the (later) stages of the evolution process when the population has zoomed in on a particular area of the search space when the noise ratio becomes too large compared to the differences in fitness. The occurrence of stagnation in the search process has been proven for Evolution Strategies using a constant number of repeated samples (resampling size) to evaluate individuals. To prevent stagnation and speed-up convergence, a straightforward and appealing idea is to use the Student's t-test for deciding on the number of individuals to sample in one generation. This paper seeks to study this strategy for $(1,\lambda)$-ES on the noisy sphere model. Besides showing gains achieved with such an adaptive approach in the early stage of runs we also show its limitations: Stagnation cannot be prevented in the long run. Additional studies aim to explain these results.

15:25 - 15:50	*Reinforcement Learning for Games: Failures and Successes* Wolfgang Konen, Thomas Bartz-Beielstein

We apply CMA-ES, an evolution strategy with covariance matrix adaptation, and TDL (Temporal Difference Learning) to reinforcement learning tasks. In both cases these algorithms seek to optimize a neural network which provides the policy for playing a simple game (TicTacToe). Our contribution is to study the effect of varying learning conditions on learning speed and quality. Certain initial failures with wrong fitness functions lead to the development of new fitness functions, which allow fast learning. These new fitness functions in combination with CMA-ES reduce the number of required games needed for training to the same order of magnitude as TDL. The selection of suitable features is also of critical importance for the learning success. It could be shown that using the raw board position as an input feature is not very effective – and it is orders of magnitudes slower than different feature sets which exploit the symmetry of the game. We develop a measure “feature set utility”, F_U, which allows to characterize a given feature set in advance. We show that the lower bound provided by F_U is largely in accordance with the results from our repeated experiments for very different learning algorithms, CMA-ES and TDL.

15:50 - 16:10	Break

16:10 - 17:00	*Compete to Win, Lose to Learn — Experiences from Game Competitions* Pier Luca Lanzi (Invited Speaker)

Sorry, no abstract yet.

17:00 - 17:25	*A Series of Failed and Partially Successful Fitness Functions for Evolving Spiking Neural Networks* J. David Schaffer, Heike Sichtig, Craig Laramee

One of the “black arts” of evolutionary computation is the design of effective fitness functions. For some tasks the appropriate function is easy to identify, but many times the “obvious” approach induces unforeseen failures of the evolutionary process to discover genomes with the desired properties. We present a series of fitness functions we have tried on the task of evolving spiking neural networks. The paradigm is to compare the output spike trains produced by evolving networks to provided target spike trains. The initial attempts failed dramatically, and subsequent versions revealed new failure modes until the third version which seems to be yielding better performance. We close with some speculations on possible limitations to this approach.

17:25 - 17:50	*Lessons Learned in Evolutionary Computation: 11 Steps to Success* Jörn Mehnen

Everybody makes mistakes – we all make one eventually if we just work hard enough! This is good news and bad news. We learn from mistakes but mistakes are also painful and could turn out to be costly in terms of money, reputation and credibility. One is prone to make mistakes particularly with new and complex techniques with unknown or not exactly known properties. This paper talks about some of my more unfortunate experiences with evolutionary computation. The paper covers design and application mistakes as well as misperceptions in academia and industry. You can make a lot of technical mistakes in evolutionary computation. However, technical errors can be detected and rectified. Algorithms are implemented, presented and analysed by humans who also discuss and measure the impact of algorithms from their very individual perspectives. A lot of ’bugs’ are actually not of a technical nature, but are human flaws. This text tries also to touch on these ’soft’ aspects of evolutionary computing.

17:50 - 18:00	General Discussion and Closing by Organizers