

Ruggedness and evolvability — an evolution’s-eye view

Lionel Barnett

Centre for Computational Neuroscience and Robotics, University of Sussex
lionelb@sussex.ac.uk

In both artificial and biological evolution, autocorrelation is commonly cited as a statistic which speaks to the “ruggedness” — and by implication evolvability — of a fitness landscape. But while the standard definition of autocorrelation involves *uniform* sampling of genotypes, it is a truism that evolution most decidedly does *not* sample a landscape uniformly. This is of particular significance in difficult artificial evolution problems, or indeed in natural evolution, where the vast majority of genotypes tend to be of poor or lethal fitness. On such landscapes uniform sampling is effectively biased towards precisely those (poor quality) genotypes which, from an evolutionary perspective, are of limited interest. To address this problem we suggest instead to take an “evolution’s-eye” view of autocorrelation: that is, we let evolution itself do the sampling.

How are we to go about this? We note first of all that autocorrelation may be considered naturally in terms of *mutation*. Indeed, the significance of autocorrelation to evolutionary dynamics lies precisely in the (statistical) relationship between the fitness of parents and their mutant offspring. We thus propose that a more cogent and useful statistic is just the correlation between parent/mutant fitnesses *as sampled over the ensemble of evolutionary histories*. We argue that this alternative autocorrelation is both conceptually compelling and also practicable, in the sense of being amenable to finite sampling.

We note that our new statistic is no longer “evolutionarily agnostic”; rather, it is tightly bound to the dynamics of a particular evolutionary scenario. This, however, we regard as a strength. We can imagine, for example, that the same fitness landscape might “appear smoother” to one evolutionary algorithm than to another, thus providing insight into the suitability of a particular evolutionary algorithm to a particular problem in artificial evolution.

We also demonstrate how autocorrelation may be derived from the *mutant fitness distribution* — a finer-grained statistic — and we introduce the notion of *linear regressive* fitness landscapes. We illustrate our ideas with generalised NK landscapes, which are particularly tractable to analysis.