From GAs to Artificial Immune Systems: Improving Adaptation in Time Dependent Optimization

Alessio Gaspar
gaspar@nessie.essi.fr
http:www.essi.fr:g̃aspar

Philippe Collard
Philippe.Collard@i3s.unice.fr

I3S Laboratory
2000 Rte des Colles Sophia Antipolis
BIOT 06460 (France)

Abstract-

This paper explores Time Dependent Optimization (Tdo) as a measure of adaptiveness in artificial systems. We first discuss this choice and review classical Tdo models to propose a canonic benchmark. Then, we underline the central role of *diversity* in adaptive dynamics for biological and cybernetic systems and illustrate by a state of the art of evolutionary Tdo (Etdo).

A Simple Artificial Immune System (Sais) is then proposed and experimentally compared to Etdo. Encouraging results are explained by strong analogies between Sais and GAs as well as Sais's ability to manage stable heterogeneous populations as a model of *Idiotypic Networks*. We conclude on the relevance of Artificial Immune Systems as genuinely adaptive artificial systems.

1 Introduction

1.1 Are GAs Adaptive Systems ?

Directly inspired by the Darwinian model of natural evolution, the Genetic Algorithm (GA) has been introduced by John Holland in the sixties as an artificial adaptive system ([14]). Since then, its successful application *as* optimizer in non trivial combinatorial problems has progressively narrowed the research of the growing GA community to optimization related issues. Even if *GAs are not Function Optimizers* ([6, 12] they nevertheless share with them many features which make it difficult to decide whether they are innovative optimization techniques or genuinely adaptive artificial systems. Their success over traditional optimizers in a wide range of applications is seen as an evidence of their superior adaptiveness. But their efficiency as optimizers ends up contradicting some of the principles of natural evolution.

In particular, GAs converge and therefore lose progressively diversity as the optimum spread over their population. Such a behavior is fundamental for optimizers as they are valued according to the certitude and quickness of their convergence but for artificial adaptive systems, this uniforming is completely antagonistic with the polymorphism observed in populations undergoing natural evolution. Moreover, convergence inhibits progressively the effects of crossover and reduces the GA's exploration dynamics to mutation only.

Here is indeed the fundamental misconception of adaptation which leads many to consider Genetic Algorithms as adaptive systems. A natural system is alive as long as it is staying in viable states. Therefore, viable states must enable the system to either reach another viable state or stay in that one. By analogy, this rule adapts to artificial evolution where a GA ceases to be an adaptive system as soon as it reaches a state which inhibits any further adaptation. For instance, a GA converging toward a uniform population clearly reaches such a non viable state. This reveals a perception of adaptation as a time bounded dynamic where it should be a perpetual one. This vision is incompatible with Artificial Life purposes as well as the pragmatic requirements of Tdo.

In order to investigate adaptive behaviors in an optimization framework, this paper advocates the use of a particular class of optimization problems which requires a reactive rather than convergent dynamic from the system. The next section details the Time Dependent Optimization (Tdo) problem underlying its key role as well as the nature of the solutions we are going to consider in this paper.

1.2 A Rosetta Stone for Adaptiveness

A static fitness function associates to each individual a real valued fitness. Similarly, a Time Dependent Fitness Function associates *at time t* such a fitness to each individual. The Time Dependent Optimization Problem (Tdo), which consists in tracking the successive optima of such a non-stationary function, is known to be difficultly handled by classical GAs. As early underlined by Salomon and Eggenberger ([23]) Tdo is at the crossing road between research on evolutionary optimization and adaptive behaviors. In order to optimize efficiently such environments, GAs must abandon the convergent dynamic and maintain their ability to react to unpredictable environmental changes. Tdo can therefore be indifferently considered as (1) the problem of maintaining a viable state in a reactive GA, (2) the efficient optimization in dynamic environments or (3) the realization of adaptive vs convergent dynamics in artificial systems. Each can be studied with the convenience of an optimization framework i.e. numeric evaluation of progress, reproducibility of experiments and availability of benchmarks. For these reasons, we advocate to explore non-convergent optimization dynamics to measure GAs' adaptiveness improvements. Therefore, improving adaptiveness means counteracting convergence. We

will refer to such approaches as *Diversity Respectful Evolutionary Computation* as they rely on the maintenance of polymorphism in populations to prevent their homogenization.
The central role of diversity in these approaches is justified by biology and cybernetics. For biologists, the elitist nature of Darwin's *survival of the fittest* theory differs from the diversity observed in ecosystems where no specie dominates. Similarly, *redundancy* seems to be the keyword in structures like DNA, neural networks or immune systems which all dispatch information over thousands of elements and, for the two latter, implement efficiently robust associative memory features. W. Ross Ashby, one of the fathers of cybernetics, expressed this idea with the *principle of selective variety* : *"The larger the variety of configurations a system undergoes, the larger the probability that at least one of these configurations, will be selectively retained"*. Clearly, the larger the polymorphism the larger the probability that at least one individual will be selectively retained [when the optimum changes].

Whether it is spelled redundancy or variety, diversity is the keystone of adaptive dynamics and the foremost motivation of many evolutionary approaches to Tdo.

1.3 Scope of the Paper

This paper is organized as follows. The next section reviews classical Tdo tasks and discusses their parameters to propose a canonical benchmark. Section 3 presents four evolutionary approaches to Tdo underlining their strengths and weaknesses. Then, we introduce the motivations that lead us to consider Artificial Immune Systems as good candidates for Tdo and detail our model named Sais before to evaluate it in section 5 by comparison with other evolutionary approaches. We eventually discuss both Reactive and Memory capacities of Sais and present work under progress.

2 Tdo Models

We review classical Time Dependent Optimization Problems and their parameters before justifying the choice of *Pattern Tracking* as a canonical benchmark.

2.1 State of the Art

We restrict for now to simple models instead of industrial ones which will be investigated in a further work to evaluate the scalability of our approach.

Pattern Tracking

In this model, introduced by K. Pettit and E. Swigger [21], the fitness of a chromosome is its similitude with an arbitrary bit-string which represents the current optimum and which changes regularly. F. Vavak and T. Fogarty ([25]) expressed the advantage of Pt: "*the problem used in this study [...] does not represent any particular application and was chosen so that analysis of the results was easy and clear.*".

Dynamical Gaussians

Dynamical Gaussians have been first used by Helen Cobb and J.J. Grefenstette [3]. In their model, a multidimensional landscape is composed of randomly rising and flattening gaussians. We already proposed a simplified version of this environment ([4, 8]) which revealed to be similar to the one used by Vavak et al. ([26]). When optimizing a Mdg, a chromosome encodes d variables and tracks an optimum sliding along the diagonal of the search space. The fitness function is: $\frac{1}{n}\sum_{i=1}^{n} e^{-\pi*\left(\frac{x_i}{2^{(\lambda-1)}}-t\right)^2}$ Where t is the generation, λ the chromosome's length and π a bell-shape coefficient.

Dynamical Knapsack Problem

Given a set of n objects $\{o_i\}$ of weights $\{w_i\}$, the Knapsack Problem (Kp) consists in choosing a subset of objects fitting closest to the weight capacity of a container w_c. GAs naturally apply by evolving n gene long chromosomes where each bit x_i specifies whether the i^{th} object is ($x_i = 1$) or is not (x_i=0) taken. The fitness is $f(x) = \sum_{i=1}^{n}(w_i x_i)$ with respect to w_c. Goldberg and Smith [10] have defined the *Time Varying KP* (TVKP) at first by changing the interpretation of a particular bit and then by varying the w_i values over time. In both cases, the GA had to be able to switch between two different problems.N.Mori et al. ([19]) proposed a *Recurrently Varying Kp* (RVKP) where the weight limit cycles through three different values over time to test memory ability of adaptive systems.

2.2 Parameterizing a Tdo Problem

We propose a way to tune Tdo difficulty by parameterizing the *transitions* which alter the fitness landscape.

Continuousness of Transitions

Transitions in Pattern Tracking, TVKP switch from one fitness landscape to another. On the other hand, the Mdg model implements a progressive switch by respectively decreasing and increasing previous and next optimum fitness. This characterizes *Continuous* rather than *Alternative* Tdo problems.

Transition Period

The Transition Period T_p is the number of generations between each transition. If the environment features n different optima over g generations then $T_p = g/n$.

Transition Range

When only one optimum moves around, the Transition Range is defined as the hamming distance between two consecutive optima $T_r = \delta_H(O^t, O^{(t+1)})$. This already revealed to be a good measure of Tdo Difficulty ([4]) and more global measures, such as the average fitness variation, would require to consider (too) large parts of the search space.

Meta Dynamics

By Meta Dynamics we mean the rules determining the new optimum. This definition encompasses characteristics such as the motion of the optimum in the search space and the cyclic, aperiodic or random nature of the succession of optima.

2.3 Toward a Canonical Benchmark

Despite its simplicity, the Pattern Tracking model enables to tune all the parameters previously discussed and even the continuous transitions can be achieved fairly easily. In fact, dynamic difficulty can be qualitatively and quantitatively easily adjusted and other environments, such as Kp, only differ by an increased static complexity. Even if this paper will only deal with the dynamic aspects, it is worth noticing that our model of Pt includes an extra parameter allowing the tuning of its static complexity as well. Evaluation starts by an X-or between the evaluated chromosome and the current optimum. The resulting bit-string is then complemented and the fitness is computed by applying to it a onemax function. The extra parameter is the replacement of the onemax function by any other function defined over the unitation such as the well-known trap functions ([9]) which deceptiveness is easily tunable. Eventually, Pattern Tracking also allows an instinctive interpretation of results in terms of the search space's metrics where other environments alter the fitness landscape in a non trivial but consequently barely analyzable way. For these reasons we consider Pt as suitable for comparisons.

3 Diversity Respectful GAs

The success of GAs in Tdo relies on the preservation of diverse populations which can be achieved in two ways. By *Heterostasic* approaches which preserve diversity or by *heterogenesic* ones which re-introduce explicitly diversity. Five approaches belonging to these classes are discussed hereafter.

3.1 Alleles Preservation

This heterostasic approach has been investigated by Rowe and East ([22]) and Culberson ([5]) separately. Their respective *Direct Replacement* and *Gene Invariant* approaches are highly similar and share a common principle of preservation of the initial population's diversity. More specifically, the allele coverage of even small populations makes them holding the necessary genetic material to adapt to environmental changes. By enforcing the replacement of parents by their offsprings after crossover, the authors prevent alleles from disappearing at any locus. Without convergence, the initial diversity is preserved and eases further adaptations.

3.2 Diversity Re-introduction

Instead of preserving diversity, Heterogenesic approaches counteract convergence by an explicit re-introduction of random genetic material. Various methods have been proposed to date such as the *Hypermutation* and the *Random Immigrant* operators ([11, 3]). The first method monitors the population's best fitness and temporarily boosts the mutation rate (hypermutation rate) if a degradation is observed. The choice of the monitoring criterion can be critical as discussed by H. Cobb herself. A decrease of the best fitness from one generation to the next, is not able to detect changes in environments where the successive optima can feature increasing fitness values over time. On the other hand, triggering hypermutation whenever a convergence rate is reached boils down to restart periodically the GA (see below).

In this respect, the Random Immigrant technique is better as it replaces a given percentage of the population by random individuals at each generation to simulate immigration. The drawback is that this dynamic is not correlated to the environment and provides diversity regardless of real needs. This is therefore paradoxal to speak of enhanced adaptiveness.

3.3 Restart Approach

Rather than improving GAs' diversity management, some approaches accentuate their convergent dynamics to make them even faster. Krishnakumar framed such an approach in a *restart-when-needed* philosophy ([18]). His Micro GA (μGA) works with a minimal population (5). Once convergence is achieved, the best chromosome is saved and th μGA restarted. The μGA outperformed various GAs with as many fitness evaluations on a real world problem.

3.4 Local Search

F.Vavak, T.C. Fogarty and K. Jukes have been investigating the benefits of the use of a hybrid evolutionary approach embedding a local search operator triggered whenever a convergence rate is reached ([26]). This so-called *shift operator* adds or subtracts to the values encoded by the chromosome the value of a random *shift register* which size increases progressively. This method differs from hypermutation insofar that the amount of diversity introduced is growing progressively rather than suddenly. The method was successful over several real world applications but remains inadapted to combinatorial (e.g. Kp and Pt) rather than numerical Problems. An alternative solution is to randomly flip an increasing number of bits instead of adding / subtracting an increasing numerical value. This approach introduces diversity more progressively than any Random Immigrant or hypermutation.

4 A Simple Artificial Immune System

We propose a Simple Artificial Immune System (Sais) as an alternative artificial Adaptive Systems.

4.1 Biological Metaphor

All animals are protected from potentially harmful pathogens by humors, and other non specific defense systems classified as *Natural Immunity*. Vertebrates feature an additional *adaptive immunity* characterized by (1) a *specific* response destroying substances (*antigens*) recognized as non-self and (2) a *memory* enabling more efficient responses in case of further infection by a similar antigen.

Let us consider a brief overview of the main actors by which the immune responses are processed :

- *Lymphocytes*

  They are white blood cells generated in the bone marrow, featuring receptors matching the 3d shape of a specific antigen. Auto-immune diseases occur when Lymphocytes matching self cells are released in the body instead of being eliminated in the bone marrow by a process called *negative selection*. In the lymphoid organs, some lymphocytes specialize into either B-cells (released in the fluids –Humoral Immunity), or into T-Cells processed in the Thymus (Cellular Immunity).

- *T-Cells* (Thymus originated cells)

  They can be either *Helper T-Cells* which, along with B-Cells, build up the immune response (cf below), or Cytotoxic T-Cells which eliminate the antigens detected by the immune system.

- *B-Cells* (Bone marrow originated cells)

  They produce *antibodies*, a special group of blood proteins (immunoglobins) which bind to antigens to immobilize them, prevent them to enter sane cells, or catalyse a destructive reaction with a so-called *complement system*. Antibodies also produce *InterLeukins* to alert Helper T-Cells when outnumbered by antigens.

  B-Cells combine with helper T-Cells to transform into *Blast Cells* which divide and clone, *Plasma Cells* which massively produce the antibody or *Memory Cells* that ease next reaction to similar antigens.

The primary response of the immune system is originated by antibodies that produce an interleukin (Il1) which causes helper T-Cells to transform into *LymphoBlasts* which, in turn, trigger the immune response at various levels. At first, the production of blood cells (and therefore lymphocytes) in the bone marrow is increased, then the growth of killer T-Cells is promoted along with the *clonal selection* of B-Cells into Plasma Cells. The latter goes along with a phase of *Somatic Hypermutation* of the produced clones which allows to explore better matching possibilities with the antigen.

The secondary response, is due to the fact that B-Cells can also be activated by other B-Cells and not only antigens. Therefore, B-Cells can be preserved in the population and activated by being integrated in a so-called *idiotypic network* featuring a self reinforcement dynamics in a group of B-Cells that match each other in a circular manner. This reinforcement in the absence of any infection preserves the structures needed to respond to any further invasion of a particular antigen already encountered.

We propose a computer model loosely based on this complex structure with the hope of embedding principles helping the realization of more adaptive artificial systems. By analogy, the current optimum plays the role of an antigen which has to be matched by B-Cells. Sais is expected to feature both primary (reactive) and secondary (memory of previous infections) immune responses. Sais starts with an initial random population of so-called B-Cells, each able to detect a given antigen specified by a λ binary bits long string. Then, it applies at each generation three operators : Evaluation, Clonal Selection and Recruitment which are discussed hereafter.

4.2 Evaluation Phase

For the sake of Pattern Tracking, *Evaluation* consists in matching these strings against the current optimum. This results in a "fitness value" here referred to as exogenic activation and noted $\alpha_{ex}^{i} = \lambda - \delta_h(B_i, O^t)$ where B_i refers to the i^{th} B-Cell of the population P, O^t to the optimum at time t, and δ_h to the hamming distance between them. This evaluation will allow us to consider the group of the n B-Cells featuring the highest exo-activation (n being an heuristic value).

Because B-Cells, contrary to GAs' chromosomes, can also be activated according to the presence of similar B-Cells, we compute a so-called *endogenic activation* for the remaining ones. For all of these the endo-activation is computed as follows: $\alpha_{en}^{i} = ((1/N_s) - d_i) + 1$ where N_s is the number of different types of B-Cells in the population and d_i the density of the i^{th} B-Cell of the population. This method already integrated a sharing principle as it tends to enforce the convergence, under evolutionary dynamics, toward a set of equally represented categories. This model attempts to capture the memory capacity of natural immune systems. It is based on the *Idiotypic Network Hypothesis* ([17, 20]) which states that B-Cells can activate other B-Cells and thus form a network dynamics that preserves B-Cells able to suitably respond to previously encountered antigens.

4.3 Clonal Selection Phase

The *Clonal Selection* mimics natural immune systems' evolutionary dynamics ([13]) which clones all B-Cells activated by an antigen. We apply a K-Tournament on the B-Cells featuring the best over-average α_{ex} values to generate a population P_{ex}. Then, we simulate the effects of the *Somatic Hypermutation*[1] by applying a random number of one-point mutations to each member of P_{ex}. Mutants are only kept if improving their α_{ex} value.

Concerning endo-activated B-Cells, the clonal selection applies differently. They also generate by selection and replication an intermediate population P_{en} but do not undergo any mutations. This enforces an equilibrated distribution of all strings that have been useful. Their densities in the population will fade with time leading oldest B-Cells to disappear. But meanwhile, no disruptiveness due to mutation or crossover will, contrary to what happen with niches ([9]) in GAs, disturbs the efficiency of this memory model.

4.4 Recruitment Phase

At this point, P_{ex} and P_{en} respectively contain the offsprings of B-Cells exo or endogenously activated. The *recruitment* consists in reintroducing worthy B-Cells in P and discarding old and worthless ones. The mechanism used differ, once again, according to the category of B-Cell concerned. Elements from P_{en} replace directly their homologs in P. This is

---

[1] In natural immune systems SH mutates the DNA of B-Cells resulting from a clonal selection process ([13]).

due to the role of equilibration of the memory dynamics here implemented. Elements from P and P_{ex} undergo a selection process on the basis of their exo-activation. Practically we used a *K-Tournament* scheme here also putting each element of P in competitions with the offsprings of P_ex.

4.5 Remarks

Sais is an over-simplified model of the immune system designed to be minimal while still capturing the essential mechanisms by which Immune Systems realize primary and secondary responses to dynamically changing complex environments. Among all details keeping away Sais from its natural homolog let us cite the choice of binary strings, the fact that we assimilate more or less antibodies and B-Cells densities and implement the somatic hypermutation as a Lamarckian operator ([13]). We also rely on the idiotypic networks hypothesis which, if advocated by some theorists ([17, 20]) is still largely contested by experimental researchers of the field.

The following results will show that Sais nevertheless captures the essence of the Immune Adaptive Response and is thereforeworth being further studied.

5 Experimental Study

We compare various GAs on a pattern tracking problem involving bit-strings of length $\lambda = 40$. Experiments are conducted over 400 generations and the dynamic parameters are set to $T_p = 50$ generations and $T_r = 5$ for the first transition. The Transition range is then increased by 5 at each transition in order to make the changes more and more complex to adapt to. The curves plot the best fitness vs generation and are averaged over 200 experiments.