

Caring versus Sharing: How to Maintain Engagement and Diversity in Coevolving Populations

John Cartlidge and Seth Bullock

School of Computing, University of Leeds, Leeds, LS2 9JT, UK

Abstract. Coevolutionary optimisation suffers from a series of problems that interfere with the progressive escalating arms races that are hoped might solve difficult classes of optimisation problem. Here we explore the extent to which encouraging moderation in one coevolving population (termed parasites) can alleviate the problem of coevolutionary disengagement. Results suggest that, under these conditions, disengagement is avoided through maintaining variation in relative fitness scores. In order to explore whether standard diversity maintenance techniques such as resource sharing could achieve the same effects, we compare moderating virulence with resource sharing in a simple matching game. We demonstrate that moderating parasite virulence differs significantly from resource sharing, and that its tendency to prevent disengagement can also reduce the likelihood of coevolutionary optimisation halting at mediocre stable states.

1 Introduction

Offering an attractive alternative to standard evolutionary approaches—by removing the difficult but necessary task of defining an adequate fitness function—competitive coevolution has been successfully utilised for optimisation in several domains (for e.g., [1,2]) However, as an optimisation technique, competitive coevolution suffers from some difficult problems stemming from the relative nature of fitness assessment—individuals receive a fitness based upon their success against contemporary opponents.

In general, coevolutionary systems are difficult to *direct*. Individuals may *over-fit* their contemporary competitors, resulting in potentially *brittle* solutions [3]. Rather than enter a progressive *arms-race*, competing populations may stabilise into a sub-optimal equilibrium, or *mediocre stable state* [2,4] As individuals are only rewarded for outperforming their *contemporary* opponents, it is possible for earlier adaptations to be lost, potentially leading to cycling [4,5]. Finally, if one population outperforms the other to the extent that every opponent is beaten, the gradient for selection disappears and the populations *disengage* and drift [3,6]. As drift during disengagement is *random* rather than *neutral*, near-optimal populations are likely to degenerate.

Although there are methods for counter-acting particular coevolutionary problems, (e.g., fitness sharing and the “hall of fame” [7]), few of these address the problem of disengagement (one exception is the domain-specific approach presented in [4]). In this paper we present a biologically inspired method for combating coevolutionary disengagement by *moderating parasite virulence*. Continuing upon previous work [6], we demonstrate that reducing parasite virulence can reduce the effects of disengagement in

the *Counting Ones* domain. The possibility that disengagement could also be combated via existing techniques for maintaining population diversity is considered. In order to explore this, a simple *Matching Game* is used to compare the effects of moderating virulence with those of *resource sharing*, a common diversity maintenance technique. We demonstrate that, whilst fitness sharing encourages phenotypic diversity via niching, which is prone to mediocre stability, moderating virulence tackles disengagement through increasing diversity in *relative fitness scores*, which actively resists mediocre stability in the scenarios that we consider. Fundamentally, moderating parasite virulence is *not* just an apparatus for maintaining population diversity.

2 Moderating Parasite Virulence

Artificial coevolutionary systems are often described as analogous to natural predator-prey or host-parasite systems. Given that most coevolutionary algorithms employ only two populations, the host-parasite analogy is probably closer [8]—one population is typically considered to pose problems for the other resulting in a series of adaptations and counter-adaptations that may result in an escalating arms-race.

Coevolutionary algorithms typically differ from natural systems in the way that they deal with parasite virulence. In order to ensure survival long enough to reproduce, it is not always in the best interests of a natural parasite to be as virulent as possible [9]. As a result, virulence varies dramatically between natural host-parasite systems (compare, for instance, cholera and the common cold), and over time within a particular system (e.g., the history of the myxoma virus in Australian rabbit populations [10]). However, when parasites are used in artificial coevolution, they are generally encoded to be maximally virulent—their fitness varies inversely with the success of the hosts that they compete against. Might coevolutionary algorithms benefit from treating parasite virulence more naturally?

Coevolutionary disengagement occurs when one population secures a significant advantage over the other, such that each competitor from the advantaged population beats each of their opponents in competition. In this way, individuals in both populations become indiscriminable in terms of fitness, all scoring maximally in one population and minimally in the other. Without intra-population fitness diversity there can be no selective forces—each individual is *as likely* to reproduce as any other—resulting in coevolutionary drift. Disengagement continues until the populations re-engage by chance, by which time there may have been a dramatic reduction in the objective quality (*absolute* fitness) of both populations—disengagement hinders optimisation [6].

Often, coevolutionary systems are asymmetric—hosts and parasites may differ genetically (in terms of encoding) or behaviourally (in terms of goal strategy). Such asymmetry may result in an inherent advantage for one population. When coevolving pursuers and evaders, for example, it is often much easier, at least initially, to be a successful evader [5]. Given that disengagement results from one population out-performing the other, it is intuitive that an inherent asymmetrical advantage toward a particular population will encourage the likelihood of coevolutionary disengagement.

However, if parasite virulence were moderated—favouring parasites that achieve less than 100% success against opponents—this trend could be reversed. By reward-

ing parasites capable of discriminating hosts—those that occasionally lose—with *more* offspring, the asymmetrical advantage will be reduced. Moderating virulence may thus reduce the likelihood of disengagement. Critically, preventing disengagement will improve coevolutionary optimisation if a reduction in periods of degrading coevolutionary drift can be achieved without sacrificing the selection pressure that ensures progress.

Implementing Moderated Virulence. Canonically, parasites receive fitness proportionally to their ability to defeat the hosts they compete against. In order to moderate parasite virulence it is necessary to change this relationship. Throughout this paper we use the term *score* to refer to the ability of a parasite to defeat the hosts it is pitted against. Parasite scores are normalised with respect to the maximum score achieved that generation such that the best current parasite always achieves a score of 1. We define parasite fitness as a function of score, x , and virulence, λ ($0.5 \leq \lambda \leq 1.0$), such that: $f(x, \lambda) = \frac{2x}{\lambda} - \frac{x^2}{\lambda^2}$. Thus, a parasite achieves optimum fitness by winning a proportion of contests equal to a fraction λ of that achieved by the best parasite. By varying λ , parasites can be encouraged to be more, or less, virulent. Although there is a continuum of possible curves, throughout this paper, we use only three values of λ . These are labelled as *maximum* virulence ($\lambda = 1.0$) where parasites are encouraged to beat as many hosts as possible, *moderate* virulence ($\lambda = 0.75$) where parasites are encouraged to achieve a win-rate three-quarters that of the highest scoring current parasite, and *null* virulence ($\lambda = 0.5$), where the fittest parasites achieve half the win-rate of their highest-scoring conspecifics. Notice that a value of lambda lower than 0.5 would encourage cooperation between parasites and hosts.

3 Study One: Counting Ones

In order to introduce the concept of coevolutionary disengagement, Watson and Pollack [3] used a minimal substrate to highlight the effects of disengagement in the *Counting Ones* domain. In this section we utilise an adaptation of the Counting Ones domain to demonstrate the effect that moderating virulence has upon disengagement.

Two reproductively isolated populations of size 25 are coevolved. Individuals in each population consist of binary strings containing 100 bits, with each bit initialised to 0 in generation 0. The *aim* of the Counting Ones problem is to evolve strings containing as many ones as possible. Of course, in this *toy* example, as observers we can assess the absolute fitness or objective quality of each individual by counting its 1-alleles. This allows us, as experimenters, a useful way of measuring progress. However, the coevolutionary algorithm does not make use of this absolute measure, only having access to the relative fitness measure described below.

Members of one population are selected to play a set of *pair-wise* contests against a random sample of 5 opponents from the competing population. During each contest, the individual with the genotype containing the greatest number of 1-alleles receives a fitness point. Each opponent receives half a fitness point for contests resulting in a draw. Individuals in both populations reproduce asexually with parents chosen through tournament selection (tournament size 5; winner reproduces). Offspring have a small probability of mutation, m . Unless specified otherwise, $m = 0.03$.




Fig. 1. Typical results from the Counting Ones domain ($B_p = 0.75$; same random seed each run).

An asymmetry was introduced by varying mutation bias B_p ($0 \leq B_p \leq 1$) in favour of one of the two coevolving populations, henceforth classified as the *parasite* population. Given mutation at a particular parasite locus, the substitution of a 1 or 0 occurs with probability B_p and $(1 - B_p)$, respectively. In contrast, the coevolving *host* population substitutes a 0 or 1 with *equal likelihood* whenever mutation occurs. We thus see that if $B_p > 0.5$, then there is a bias in favour of evolving parasites with more ones—an asymmetry that favours the parasite population.

This type of problem asymmetry is common in coevolutionary problems. It is often the case that one side of the coevolutionary contest enjoys some (perhaps temporary) advantage over the other in terms of the ease with which successful counter-adaptations are discovered. For example, at the outset of coevolving list-sorting algorithms it is easier to find a list that is difficult to sort than an algorithm that is difficult to defeat [1].

Several parasite mutation bias values and three parasite virulence levels were tested over a series of runs; $B_p = [0.50, 0.99]$ and $\lambda = 0.50, 0.75, 1.0$. Unless otherwise stated, the value of λ remained constant throughout each run.

Results. Fig. 1 displays three stereotypical runs, each using a parasite mutation bias, B_p , of 0.75. When using maximally virulent parasites (Fig. 1, left) the populations have a tendency to disengage. This can be observed between generations 150 – 175 and again between 250 – 500. During these periods of disengagement the populations drift back to their relative baseline performance, equal to the mutation bias, $B_p = 0.75$ and $B_{host} = 0.5$. Only once the populations re-engage by chance is there an improvement in absolute fitness. Repeating the run with the same random seed, the second period of disengagement depicted in Fig. 1 is prevented if moderate virulence is introduced at generation 250 (middle). Notice that the left and middle graphs are identical until generation 250—the point at which parasite virulence is changed to moderate. In contrast to maximum virulence, when moderate parasites are used (right), the populations remain engaged throughout the entire run, achieving a continuously high level of performance.

These results are sensitive to variation in both population size and the number of opponents played by each individual. As either parameter increases, the probability of disengagement decreases due to the increased frequency of meeting varied opponents. However, the results observed in this section are qualitatively robust to mutation rate ($m = [0.005, 0.05]$) and tournament size ($tourney = [2, 15]$).

Fig. 1 clearly demonstrates that moderating parasite virulence in asymmetric coevolution can reduce the effects of disengagement. Further results (not shown¹) also suggest that the greater the inherent asymmetry, the greater the effect moderating virulence has upon results—i.e., in order to reduce disengagement, virulence should be reduced as engagement increases. The asymmetry imposed in this model gave the biased parasite population a great advantage over the coevolving host population. Purely by stochastic effects one would expect all the individuals from parasite populations to contain more ones than individuals in host populations. This is observed in Fig. 1. The mutation bias alone pushes the populations toward the expected ratio of ones to zeros, i.e., 0.50 for hosts and 0.75 for the parasites. The difference in fitness acceleration between populations, forced by the mutation bias, ensures that once disengagement occurs, the populations quickly diverge to different equilibrium levels. Both populations will remain disengaged until a very large, and thus very *unlikely*, mutation occurs allowing the gap to be, at least temporarily, bridged, for e.g., Fig. 1, left, generation 500.

The first non-zero parasite generation will on average contain many more ones than the host population. However, under moderated virulence, any parasites that beat all opponents are *less* fit than those parasites that lose a small percentage of contests. In this way, acceleration is decreased as the parasites resist their mutation bias. Moderate virulence parasites appear to actively prevent disengagement. Using the continuous selection pressure hosts evolve to a greater level than would otherwise be possible. It should not be overlooked, however, that moderate parasites gain from this relationship too, as both populations evolve to a greater standard than either would alone (Fig. 1, right). However, as parasite virulence is decreased there is a tendency for coevolution to stagnate at a sub-optimal but highly engaged fluid local optimum. In order to push populations to optimal solutions, stronger selection pressure is required (see below).

Diversity Maintenance. Disengagement occurs when intra-population fitness diversity reduces to zero. Moderating virulence counter-acts disengagement by selecting for reproduction parasites that are occasionally beaten. This preserves a selection gradient for hosts which, in turn, maintains relative fitness diversity in both populations.

A tendency toward reduced population diversity (and associated premature convergence) has long been a major concern of the evolutionary computation research community. As such, a suite of diversity maintenance techniques have been proposed, including for e.g., competitive fitness sharing [7], resource sharing [4], and spatial embedding [1]. These approaches are attempts to maintain genetic diversity on the assumption that a loss of diversity can be harmful to optimisation as it may restrict search to local optima.

Resource (or competitive fitness) sharing maintains genetic diversity in a population by encouraging niching—individuals are rewarded for being able to beat opponents that few others can. This idea has been extended to coevolutionary scenarios where opponents are treated as a commodity or resource. Rather than gain a fitness point for each victory against an opponent (simple fitness), one fitness point is shared among the competitors that beat a particular individual. Thus, individuals are rewarded less for *how many* opponents they beat and more for *who* they beat, rewarding genetic novelty and maintaining diversity.

¹ see [6], also <http://www.comp.leeds.ac.uk/johnc/publications/ecal03long.pdf>




Fig. 2. The probability, $p = \frac{1}{2} \left(1 + \tanh\left(\frac{x-50}{7}\right) \right)$, of parasites playing the 1-allele half of the matching game (left). Depending upon which half of the matching game the parasite plays, hosts try to match *either* 0– or 1-alleles (right).

Since disengagement is associated with a loss of diversity, could it be prevented by simple diversity maintenance approaches? Perhaps moderating virulence is only preventing disengagement by mimicking these existing techniques? If so, it is largely superfluous. The next study contrasts moderating virulence with resource sharing in order to explore whether they are effectively the same or different in some fundamental sense.

4 Study Two: Matching Game

In order to compare the influence of parasite (and host) resource sharing with that of moderating parasite virulence, we need to choose an appropriate and simple problem domain. Here we develop a simple matching game, in which hosts are rewarded for matching parasites, but parasites are punished. Games with this type of dynamic often suffer from coevolutionary cycling, as hosts chase parasites through the strategy space. Although a global optimum strategy exists, populations are easily diverted from it as they exploit the temporary idiosyncrasies of their opponents. Resource sharing is one way of discouraging this type of short-termist behaviour. By maintaining a diverse strategy-base in each population, the value of exploiting idiosyncrasies is reduced, encouraging generalists. Unfortunately, an alternative mediocre stable scenario is possible in which populations “speciate” such that they exhibit a number of different sub-optimal strategies that together form a stable combination. In this sense, the game is similar to any number of scenarios in which a generalist strategy is globally optimum, but difficult to evolve in practice—e.g., scissors-paper-stone, immune systems, etc.

Two genetically isolated populations of size 50 are coevolved—hosts and parasites. Individuals in each population consist of binary strings containing 100 bits, initialised *randomly* in generation 0. Each generation, members of the host population are selected to play a set of pair-wise contests against a random sample of 10 opponents from the parasite population. The aim for hosts is to match as many parasite alleles as possible. Antagonistically, parasites aim to mis-match host alleles. Both populations breed asexually, with each individual having a small probability of unbiased mutation per locus, $m = 0.03$. Tournament selection was used (tournament size 5) with the winner of each tournament always chosen to reproduce.

Not all loci are involved in this matching game. For parasites with many 1-alleles, the matching game tends to involve only those loci at which the parasite possesses 1-alleles.

For parasites with many 0-alleles, the game tends to involve only those loci at which the parasite possesses 0-alleles. Whether 1-allele loci or 0-allele loci are involved is determined probabilistically. The probability, p , of a game involving matching 1-alleles increases with the total number of 1-alleles. Once the game has been decided, a host wins by matching alleles in *at least* $T = 30$ loci, else the parasite wins (see Fig. 2).

Having several antagonistic points of attraction, the Matching Game domain is designed to exhibit interesting coevolutionary dynamics. Mutation bias attracts both populations toward genotypes containing 50% 1-alleles and 50% 0-alleles. However, given a host plays a parasite at the 1-allele (0-allele) half of the matching game, it is advantageous for the host to have *as many* 1s (or 0s) as possible. Thus, the host population is attracted toward homogeneous genotypes (all 1s or all 0s). The direction of attraction for hosts (toward either 100% 1- or 0-alleles) depends upon the frequency with which the parasite population plays either the 1-allele or 0-allele halves of the game. This occurs with increasing frequency the further parasite genotypes vary from 50% 1s. Thus, parasites are also attracted away from 50% 1s, but in the *opposite* direction to hosts. Parasites deviating too far from 50% 1s, however, become *too predictable*. In general, the most difficult parasites to match are those having approximately 50% 1s.

This matching game resembles the density classification task for 1-D cellular automata. The aim of the density classification task is to evolve a set of cellular automata (CA) rules which correctly classify the density of an initial condition (IC)—a binary string—as having less than, or more than, 50% 1s. Whilst coevolving CA rules, Juillé and Pollack found it necessary to moderate the virulence of ICs in order to stop disengagement, despite the use of resource sharing [4]. Although very successful, their approach is domain dependent, relying heavily upon domain-specific knowledge. As such, it is not transferable to arbitrary coevolutionary optimisation scenarios. Here we tease apart the contribution of two domain general approaches to improving coevolutionary optimisation, resource sharing and moderating parasite virulence.

As before, three λ values were tested over a series of runs; $\lambda = 1.0$ (maximum), $\lambda = 0.75$ (moderate) and $\lambda = 0.5$ (null). The value of λ remained constant throughout each run. Runs were performed under four conditions: maximum virulence without resource sharing (i.e., standard coevolution); maximum virulence with resource sharing; moderated virulence without resource sharing; both moderated virulence and resource sharing. Under each condition, the degree of niching or genotypic diversity within each population was calculated using a *linkage disequilibrium* measure, particularly sensitive to the effects of resource sharing.

Results. Fig. 3 displays four stereotypical graphs from the Matching Game domain, resulting from the four test conditions. Both resource sharing and moderating parasite virulence have clear effects on coevolutionary dynamics.

Under condition one—maximum virulence with no resource sharing, i.e., typical coevolutionary optimisation (top-left)—the system exhibits cycling. After the initial generations, hosts may begin to recruit more 1-alleles in order to defeat parasites playing the 1-allele half of the game. However, as parasites counter-adapt, by recruiting more 0-alleles, they increase the likelihood of playing the 0-allele half of the game. In response, hosts appear with a greater proportion of 0s, with the entire population

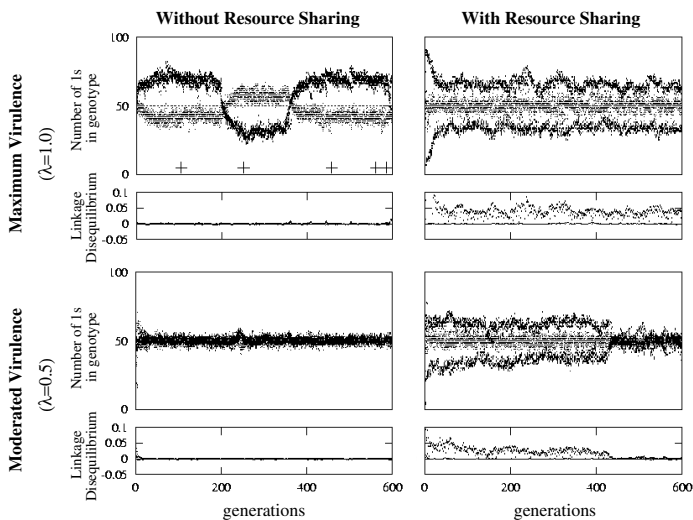


Fig. 3. Typical coevolution in the Matching Game domain.

eventually switching strategy, in order to concentrate on winning the 0-half of the game. Subsequently, parasites again regain the *upper-hand* by recruiting 1-alleles, and so on. Under these conditions, the Matching Game is inherently easier for the parasite population. Hosts find it difficult to be successful generalists—incapable of matching parasites along *both* dimensions—and are encouraged to become brittle *specialists*. As a result, maximally virulent parasites win the majority of competitions and occasionally win *all* competitions, resulting in disengagement (indicated by crosses).

Under condition two—maximum virulence with resource sharing (top-right)—the system reaches mediocre stability. At the beginning of the run, hosts immediately niche into two groups, each specialising on one half of the matching game. In order to be as unpredictable as possible, parasites tend toward 50% 1s—any deviation from this distribution will be punished by one of the specialist host niches. At this mediocre equilibrium the host population as a whole achieves roughly 50% victories over parasites, but each individual host is extremely vulnerable to parasites playing the opposite half of the game. In contrast, parasites tend to become maximally unpredictable and play either half of the matching game with roughly equal probability. Although a generalist strategy exists for hosts, they are unable to discover it.

Under condition three—moderated virulence ($\lambda = 0.5$) without resource sharing (bottom-left)—the system stabilises with generalist hosts. After the initial generations, the host population settles into generalist strategies, capable of matching some parasites whichever allele is triggered. Moderate virulence ensures that parasites are rewarded when occasionally matched, thus allowing hosts to succeed without having to concentrate on winning one half of the matching game. It should be noticed that moderating virulence *does not* result in host-parasite collusion, which would tend to result in homogeneous parasites—the simplest to match. Rather, parasites remain challenging and unpredictable. Any deviation from 50% 1s is quickly punished by the generalist hosts. As

such, both hosts and parasites engage in competition in the most *difficult* regions of space. This is equivalent to discovering the “play random” strategy in scissors-paper-stone, or a generalist immune system capable of defeating a wide range of intruders.

Under condition four—moderated virulence with resource sharing (bottom-right)—the system initially achieves mediocre stability, before encouraging hosts to become generalists, strongly engaged with parasites. Early in the run, resource sharing encourages the host population into two niches, each concentrating on one half of the matching game. In this way, the system reaches mediocre stability with hosts and parasites sharing victories. However, unlike condition two, mediocre stability does not persist. Recall that null virulence encourages parasites to achieve a win-rate half that of the highest scoring parasite. This scheme lures parasites away from the mediocre equilibrium at which they achieve a 50% win-rate. As parasites become more easily matched, they reduce the pressure upon hosts to concentrate on one half of the matching game. In this way hosts are steered toward a more generalist strategy of 50% 1s. Hosts engage parasites in a difficult region of space, unattainable without moderated parasite virulence.

Results clearly demonstrate that imposing moderate virulence on parasites alters coevolutionary dynamics in a fundamentally different way to that achieved by resource sharing. Whilst resource sharing encourages *within-population* genetic (and phenotypic) diversity, observable as niching in the host population, moderate virulence encourages diversity in *relative fitness*.

5 Conclusions

We have demonstrated the effects that resource sharing and moderating virulence have upon coevolutionary dynamics. Though both tools increase diversity each affects coevolution in a fundamentally different manner. Resource sharing encourages a population to diversify into separate niches, thus reducing the likelihood of *over-focusing* and coevolutionary cycling. However, niching may produce mediocre stability—reaching a sub-optimal equilibrium—whereby niches *share* success. In contrast, moderating virulence does not encourage intra-population diversity and rather encourages engagement—the extent to which coevolving populations *interact*.

Resource sharing adds a second layer of coupling between conspecifics. In addition to the standard competition that conspecifics experience—striving to beat more opponents than each other—they are forced to *share* their success with one another. This encourages individuals to beat different opponents—i.e., to be different from one another. Niching results from this additional intra-population coupling. In contrast, moderating parasite virulence increases *inter-population* coupling—it ensures that individuals in one population *care* about the success of individuals in the other. In particular, through attempting to achieve moderate success, parasites care about the variation in relative fitness achieved by their opponents—they are selected to cause a range of scores in their opponents. However, this is not achieved through niching, or genetic diversity *per se*. Rather, it is a direct consequence of the moderation that maintains engagement.

It is true that increased genetic diversity has some relationship with coevolutionary engagement. If genetic diversity reduces to zero, populations will disengage (individuals will achieve equivalent scores). However, the converse is not true. Genetic diversity *does*

not ensure engagement. Both populations may feature a diverse array of phenotypes, yet still suffer disengagement if each and every phenotype in one population defeats each and every phenotype in the other. Indeed, periods of disengagement often increase genetic diversity through random drift without necessarily increasing engagement. While this coevolutionary coupling (engagedness) is affected by genetic diversity (and noise, sampling error, etc.), it is not determined by it. These considerations ensure that moderating virulence and resource sharing are *complementary*, rather than *exclusive*, tools. It is not necessary to choose one over another. Indeed, the greatest success may result from using both [4].

Finally, it is important to re-iterate the fact that moderating virulence involves a *trade-off* between engagedness and optimisation. Whilst reducing λ increases engagement, optimisation may suffer as a result, since parasites are being encouraged to present a less-than perfect opponent. Balancing this trade-off—maintaining engagement while encouraging optimisation—proves to be quite difficult. In general, a parameter such as λ must be constantly altered over the course of coevolution to achieve this optimal balance. Currently we have no successful way of adaptively altering λ . Future work involving an experimental interface for *manually steering* virulence during coevolution [11], will aim to discover insights into disengagement that lead to a technique for automatically adapting virulence in the necessary manner.

In conclusion, this paper has demonstrated that moderating virulence is *not* an alternative diversity maintenance technique but a complementary, and novel domain-independent tool, capable of improving coevolutionary optimisation through reducing coevolutionary disengagement.

References

1. Hillis, W.D.: Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D* **42** (1990) 228–234
2. Pollack, J., Blair, A., Land, M.: Coevolution of a backgammon player. In Langton, C.G., Shimohara, T., eds.: *Artificial Life V*, MIT Press (1996)
3. Watson, R.A., Pollack, J.B.: Coevolutionary dynamics in a minimal substrate. In: Proc. of the Genetic and Evolutionary Computation Conference, IEEE Press (2001) 702–709
4. Juillé, H., Pollack, J.: Coevolutionary learning: a case study. In Shavlik, J.W., ed.: Proc. of the Fifteenth International Conference on Machine Learning, Morgan Kaufmann (1998)
5. Cliff, D., Miller, G.F.: Tracking the Red Queen: Measurements of adaptive progress in coevolutionary simulations. In Morán, F., Moreno, A., Merelo, J.J., Chacón, P., eds.: *Third European Conference on Artificial Life*, Springer (1995) 200–218
6. Cartlidge, J., Bullock, S.: Learning lessons from the common cold: How reducing parasite virulence improves coevolutionary optimization. In Fogel, D., ed.: *Congress on Evolutionary Computation*, IEEE Press (2002) 1420–1425
7. Rosin, C.D., Belew, R.K.: New methods for competitive coevolution. *Evolutionary Computation* **5** (1997) 1–29
8. Janzen, D.H.: When is it co-evolution? *Evolution* **34** (1980) 611–612
9. Futuyma, D.J., Slatkin, M.: *Coevolution*. Sinauer Associates, Sunderland, Mass. (1983)
10. Fenner, F., Ratcliffe, F.N.: *Myxomatosis*. Cambridge University Press, Cambridge (1965)
11. Bullock, S., Cartlidge, J., Thompson, M.: Prospects for computational steering of evolutionary computation. In Bilotta, E., ed.: *Artificial Life 8, Workshops*, MIT Press (2002) 131–137