

A Study of Spatial Distribution and Evolution

Peter A. Whigham¹ & Grant Dick²

¹Spatial Information Research Centre
University of Otago, Dunedin, New Zealand
Phone: +64 3 479-7391 Fax: +64 3 479-8311
Email: pwhigham@infoscience.otago.ac.nz

²Spatial Information Research Centre
University of Otago, Dunedin, New Zealand
Email:gdick@infoscience.otago.ac.nz

**Presented at SIRC 2002 - the 14th Annual Colloquium of the Spatial Information Research Centre
Victoria University of Wellington, Wellington, New Zealand
December 3-5th 2002**

ABSTRACT

The use of space for supporting evolution has been previously studied in the context of distributed computation and artificial life. However, these systems use a predefined structure for defining sub-populations and require a number of adjustable parameters to be set prior to evolution. This paper describes a self-organising mechanism for evolution using an unbounded spatial representation, along with the concept of a visibility for each population member. The resulting system is independent of spatial scale and visibility, and automatically adjusts the distribution of individuals based on the current fitness pressure in the population.

Keywords and phrases: spatial selection, distributed evolution, random model

1 INTRODUCTION

Evolutionary computation (EC) is a computational form of population-based search that uses a basic framework of evolution to discover near-optimal solutions to a diverse range of problems. Following the basic approach of natural systems, EC is composed of two main concepts: heritable variation and selection. Heritable variation in a population is a measure of the diversity of that population, resulting in potentially different behaviour (phenotypic expression) between each population member and the environment. The heritable nature of variation occurs across generations, either through single parent to child mutation, or through combining different parent genetics to make a child; selection drives the system by allowing 'fitter' individuals to survive and pass on their genetic description. During this process of heredity, variation to the genotype can occur that may result in changes to the phenotypic behaviour of the progeny. The process of selection repeats itself for each breeding event, ensuring that over time those individuals that are proportionally fitter tend to dominant in the population. Variation in the reproductive process is the source of change at the genetic level, which may translate into new innovation at the phenotypic level. Selection serves as a filtering mechanism to ensure that individuals of high fitness are maintained, and those of low fitness are removed along the way. Evolution, then, is the coupling of these two processes over time.

Evolution is a process that occurs in natural systems at a variety of spatial scales, from the sub-micron bacterial scale to the human scale. What is perhaps most important about the interaction of species and the spatial scales that they observe is that space is a limiting factor in the interaction of species and the possible exposure to other genetic material during their lifetime. Selection can only act on individuals that are localised in space and time (Hull 1980).

The use of space in EC has been limited to one of two main approaches: spatial models for use in distributed computation; and explicit spatial positioning of individuals when studying co-evolution. The use of a continuous space as a strict method for selection, which defines dynamic sub-populations, has not been previously considered in detail.

This paper is structured as follows: §2 gives the background to previous work and the motivation for the current study; §3 describes the basic random model of population and space used throughout this study, and presents results with no selection pressure for a variety of population sizes and spatial dimensions; §3.5 describes the behaviour of the population when a simple fitness pressure is applied to the selection of individuals; finally a discussion and conclusion regarding the current work is presented.

2 BACKGROUND AND MOTIVATION

The application of an explicit spatial model for selection and interaction has been previously applied in two main contexts. The use of space to limit interactions between individuals has been used to impose an environment for each population member, and to constrain the interaction between individuals. Such examples include 'Tierra' (Ray 1992), where organisms representing simple computer programs exist in a computer memory and compete for the resources of the machine. Extensions of this approach include 'Avida' (Adami 1998), the Echo Model (Forrest & Jones 1994, Hrabér & Milne 1997) and predator-prey models (Haynes & Sen 1996, Cliff & Miller 1996, Rosin & Belew 1997). Although these models all use space to explicitly limit the selection behaviour of individuals, the goal of these systems is to study coevolution and the dynamics of individuals that are all competing against each other. This differs from the current study where the population is not competing against each other, but rather is used to search for a single maximum (or minimum) in a fitness landscape.

The second context of using space involves frameworks for distributed computation of evolution. This is motivated by the embarrassing parallelism with EC that can be exploited to reduce the overall computation time for discovering solutions. The two main models in distributed EC are cellular-based and island-based.

Cellular or fine-grained models of EC contain a single population that is distributed over a grid (Robertson 1987). The shape of the grid is often a two-dimensional torus, although other topologies have been explored (Schwehm 1992). Interactions between individuals are based on the neighbourhood defined by the grid. For example, with a square lattice each individual has eight neighbours. The configuration of the space and the neighbourhood for each individual is fixed throughout the evolution.

Island models of EC take several sequential systems and execute them in parallel (Grefenstette 1981). Selection in this model is constrained by limiting parent selection to individuals that are within the same subpopulation. This leads to several subpopulations that tend to be small and therefore converge quite rapidly. To avoid this convergence, individuals are allowed to migrate between subpopulations so that the overall diversity of the system is maintained. This setup, although easy to construct in a distributed environment, requires a number of parameters (such as subpopulation size, migration rate and topology of the subpopulations) that require tuning to obtain reasonable performance.

2.1 Previous Results




Figure 1: Average fitness results with various spatial configurations.

The spatial framework examined in this paper has been previously applied to the Travelling Salesman Problem (TSP) (Dick & Whigham 2002). The system parameters are shown in Table 1. The tests were run with visibility settings of 5%, 25%, 50% and 100%, where this was interpreted as a percentage of the maximum required visibility to see all individuals after the initialisation of the population. The tests were run at each visibility setting 100 times and the final results were averaged out over these runs. The tests were run using a steady-state genetic algorithm.

Parameter	Value
Population Size	625
Visibility	5%, 25%, 50% and 100%
Base Selection	6-Tournament Selection
Replacement	Steady-State
Iterations	62500
Cities	30

Table 1: The Parameters for TSP with a steady-state Genetic Algorithm.

Figure 1 shows the average fitness of the individuals against the number of completed iterations. The “control” is the plot of the results produced by a genetic algorithm that does not incorporate spatially constrained selection. A lower fitness value indicates that the individuals of the population have been more successful in solving the problem. As can be seen, the population evolves fitter solutions when spatially constrained selection is used and converges more rapidly to a fit solution than the non-spatial counterpart. Examining Figure 1 it can be seen that after approximately 1/3 of the iterations for the non-spatial approach, the average fitness using spatial selection have reached to within 95% of the final result. This statistically significant improvement for convergence to a solution highlights the benefits of spatial selection EC.

The purpose of this paper is to explore the basic properties of the spatial selection method that has been used for this problem. In particular, the goal is to demonstrate that the behaviour of subpopulation construction can be automatically performed by the system, and is independent of any particular initial configuration.

3 A SIMPLE MODEL OF SPACE AND EVOLUTION

This section will describe the basic setup of the simple model of selection and replacement for a population-based computation model, and define the basic properties of the model. Commencing with a random model in 2 dimensions, the work will first explore the basic properties of the model, extend these to 1 and 3 dimensions, and then consider the behaviour when an explicit fitness selection pressure is introduced.

3.1 Simple Spatial Population Model

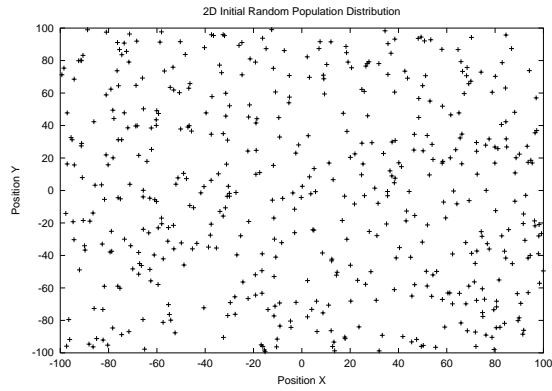
The simple random spatial model for two-dimensions, with population size N , and visibility v , may be defined as follows.

1. Create an initial *randomly distributed* population $P = \{p_1, p_2, \dots, p_N\}$ of N individuals over a 2-dimensional euclidean plane, bounded by $[x_{min}, y_{min}] \rightarrow [x_{max}, y_{max}]$.
2. **While not finished Do**
 - (a) Randomly select an individual $p_i \in P$.
 - (b) Randomly select an individual $p_j \in P$.
 - (c) Create a new individual p_k , with position randomly assigned within the circle of radius v from p_i .
 - (d) $P = P \cup \{p_k\} - \{p_j\}$


The purpose of this simple model is to allow sub-populations, based on the visibility v , to be automatically created. Here a sub-population implies the set of individuals that are visible from one population member. As an example, Figure 2 shows an initial population, bounded by $[-100, -100] \rightarrow [100, 100]$ and the subsequent distribution of the population after 100000 iterations for $v = 50, 500$ and 1000. Note that although the absolute dimensions of the final spatial extent increases with increasing v , the overall pattern of distribution is *similar*.

3.2 Influence of v on Spatial Distribution


The first property of this system to describe is the influence on the spatial distribution of individuals for different visibility settings. Figure 3 shows the average subpopulation size (i.e. the average number of individuals visible for each iteration), averaged over 100 runs. The population size is 500, initial spatial extent is $[-100, -100] \rightarrow [100, 100]$, and the number of iterations is 100000. Figure 3 shows that independent of the visibility, the population rapidly converges around a single mean value for the average visible members of each population. Figure 2.1 shows the final 20000 iterations for 2 extreme values of v . Note that there is no significant difference in the mean values of the subpopulation size at this point, even though after 500 iterations the average subpopulation sizes are 6.3 ± 0.3 for $v = 10$, and 391.7 ± 9.7 for $v = 1000$.



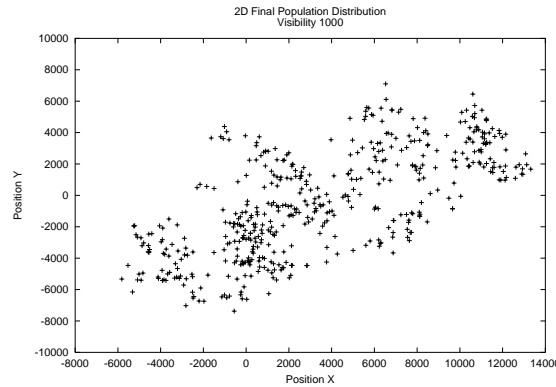
(a) Initial Population Distribution



(b) Final Population Distribution $v = 50$



(c) Final Population Distribution $v = 500$



(d) Final Population Distribution $v = 1000$

Figure 2: Random Spatial Model Behaviour after 100000 iterations.

3.3 Influence of Initial Bounding box on Spatial Distribution

The second property of this system to describe is the influence on the spatial distribution over time of individuals for different initial spatial configurations. For these experiments, the population was placed in a random bounding box of width and height 0.5, 1, 10 and 10000. Figure 4 shows the average number of individuals visible for each iteration, averaged over 100 runs. The population size is 500, $v = 1$, and the number of iterations is 100000. Note that independent of the initial size of the space, for a constant v , the average number of visible individuals for each member of the population converges to a mean value around 15. Hence the behaviour of the population appears to be independent of the initial size of the space. Since v was held constant for different size spaces the initial percentage of visible population ranged from 1 (i.e. the only individual visible was itself) through to 100% of the population.

3.4 Spatial Dimensions

The third property to describe is the influence on spatial dimensions to the distribution. Clearly there is nothing "special" about two-dimensions, however in terms of the final goal of producing a framework for a distributed computation system for EC, there are some arguments for using two dimensions (Dick & Whigham 2002). Figure 5 shows the behaviour for the system when 1, 2 and 3 dimensions are used, with a bounding line, box and cube of size edge length 100, $v = 100$, iterations = 100000, with the results averaged over 100 runs. Note that the visibility from the parent p_j for 1-D was a line segment of length $2 * v$ centred on p_j , for 2-D was a circle with radius v , centred on p_j , and for 3-D was a sphere of radius v , centred on p_j .

Figure 5 shows that similar behaviour is observed for 1, 2 and 3 dimensional spaces. However, as the number of dimensions increases, the amount of space available to the population increases. Therefore, the 1-dimensional

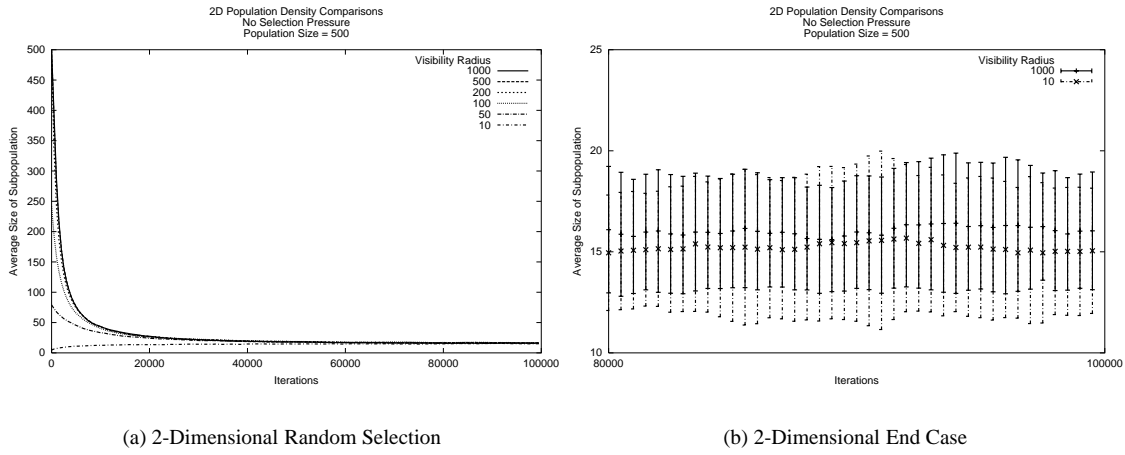


Figure 3: Average Subpopulation Size with No Selection Pressure.

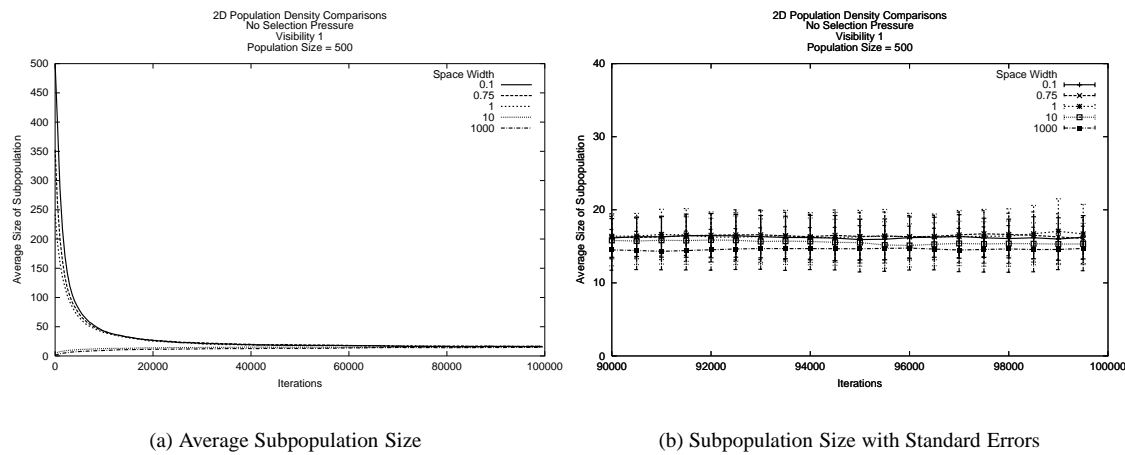


Figure 4: Random Spatial Model Behaviour for different initial Spatial Scales.

case converges to an average number of visible individuals that is greater than the 2-dimensional case, and so on for subsequent dimensions. However, a similar pattern of convergence can be observed for each dimension, independent of v .

3.5 Behaviour of the Model with Selection Pressure

Imposing selection pressure implies that there is a bias towards selecting certain individuals in the population based on their phenotypic performance. In EC there are a large variety of fitness-based selection techniques (Yao 1999), such as proportional-based fitness (roulette wheel), tournament, rank-based and elitist. Although each of these approaches have different characteristics, they all ensure that fitter individuals are more likely to be selected on average over many iterations of the model. For the purposes of this paper a simple M-Tournament selection scheme will be used.

M-Tournament selection, in its most basic form, is a selection technique that selects at random M individuals, and from that group selects the best (or worst) based on fitness. Clearly, as M increases the selection pressure increases. To explore the model behaviour once selection pressure is introduced a random initial fitness will be assigned to each population member, with mean 0 and standard deviation 1. Subsequent assignments of fitness for new individuals will have a fitness with mean 1 and standard deviation 0.5 added to the parent fitness. Hence, the child will most likely be fitter than the parent. A maximum fitness is defined as 100. Using M-tournament selection the spatial model is extended as follows:

1. Create an initial *randomly distributed* population $P = \{p_1, p_2, \dots, p_N\}$ of N individuals over a 2-dimensional

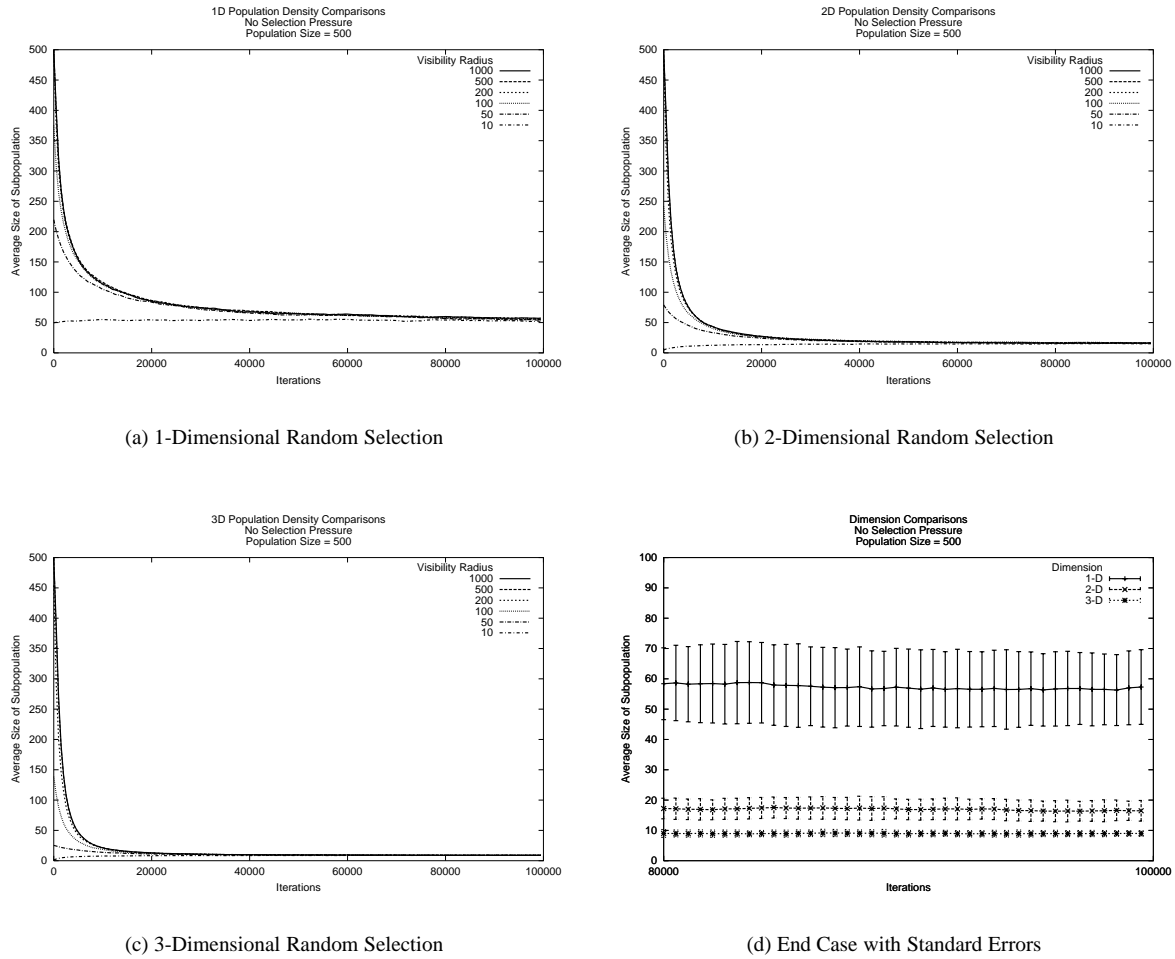


Figure 5: Average Visible Population for 1, 2 and 3 dimensions.

euclidean plane, bounded by $[x_{min}, y_{min}] \rightarrow [x_{max}, y_{max}]$.

Assign a random fitness value $f_{p_i} = N(0, 1)$, where $N(0, 1)$ is a normally distributed random number with mean 0 and standard deviation 1, to each $p_i \in P$.

2. While not finished Do

- (a) M-Tournament(best) select an individual $p_i \in P$.
- (b) M-Tournament(worst) select an individual $p_j \in P$.
- (c) Create a new individual p_k , with position randomly assigned within the circle of radius v from p_i .
Assign the fitness $f_{p_k} = \min(100, f_{p_i} + N(0.2, 0.1))$.
- (d) $P = P \cup \{p_k\} - \{p_j\}$

The results comparing a typical run using 6-Tournament selection and random selection are shown in Figure 6. The important behaviour to note is that when fitness pressure is applied to the population, there is an increase in the average size of the visible subpopulation while a fitness pressure is maintained. Note that after approximately 20000 iterations the population has converged to the maximum fitness (100), after which there is no longer any fitness pressure. At this point the population begins to behave in an equivalent fashion to the random model. Hence, the population is configuring itself based on fitness pressure, altering the number of individuals in a subpopulation. This is an attractive property for the system, since it is adjusting the spatial structure of the population due to fitness pressure, which has implications for using this approach as an automatic method for structuring sub-populations for selection. This will be further explored in §3.6.

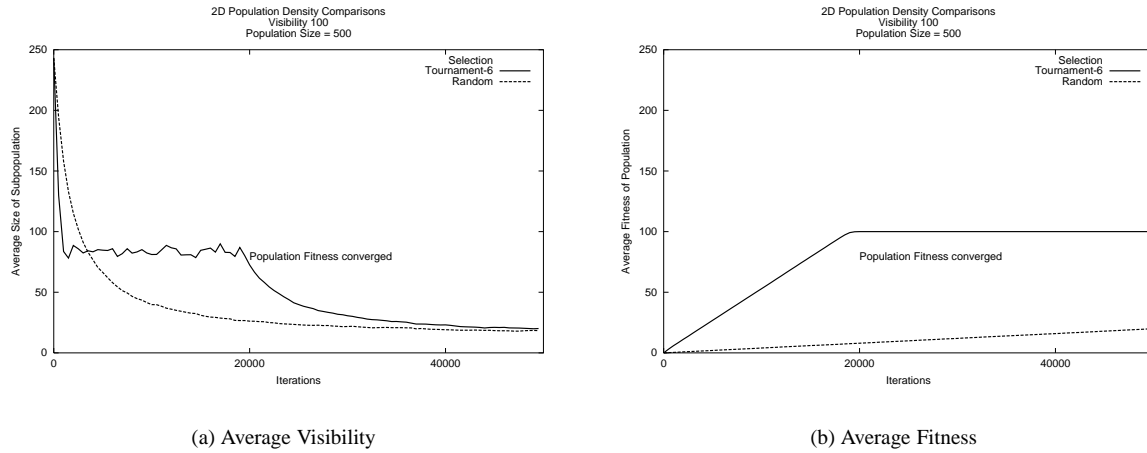


Figure 6: Behaviour using gradual fitness pressure.

3.6 Two-Parent Selection

EC systems often use 2 parents to create a new individual, where the mixing of parent representations is generally referred to as crossover. For our experiments a simple version of crossover will be used, called 1-point crossover (Yao 1999). This crossover, when applied to two bit strings, selects a random point to cut the bit strings, and swaps the two cut strings to produce a new individual. In addition to crossover, a random mutation is also usually applied that maintains some diversity in the population. For our experiments a simple random mutation of bit values ($0 \rightarrow 1$ and $1 \rightarrow 0$) with a low probability will be applied. Both crossover and mutation are generally applied with some probability. The following algorithm describes the two-parent selection mechanism with spatial selection.

1. Create an initial *randomly distributed* population $P = \{p_1, p_2, \dots, p_N\}$ of N individuals over a 2-dimensional euclidean plane, bounded by $[x_{min}, y_{min}] \rightarrow [x_{max}, y_{max}]$.
Define V_{p_i} as the set of individuals $\in P$ that are visible from p_i .
2. **While not finished Do**
 - (a) M-Tournament(best) select an individual $p_i \in P$.
 - (b) M-Tournament(worst) select an individual $p_j \in P$.
 - (c) M-Tournament(best) select an individual $p_l \in V_{p_i}$.
 - (d) Create a new individual p_k using crossover and mutation, based on p_i and p_l , with position randomly assigned within the circle of radius v from p_i .
 - (e) $P = P \cup \{p_k\} - \{p_j\}$

To demonstrate the behaviour of the previous algorithm, a simple problem can be used to illustrate how as selection pressure varies the dynamics of the spatial distribution of the population changes. Assume that the goal of the EC system is to discover a particular 100-bit gene pattern - for simplicity let the maximum fitness be when all bits are 1's. The initial population will be randomly assigned bit strings with each bit set to 0 or 1 with equal probability. The crossover probability is set to 1.0, so that the 2 parents, p_i and p_l , are always combined. A low mutation probability is also given so that the diversity of the population is maintained. The parameter settings are summarised in Table 2, and a typical run using these parameters is shown in Figure 7. A number of properties of the system can be observed from Figure 7. The average sub-population size increases rapidly at several points in the evolution. This occurs each time there is a new population member that has a fitness advantage over the average population. This creates a larger sub-population to allow the evolution to proceed more rapidly. Note that after approximately 80000 iterations the population reaches the global optima, and therefore the sub-population distribution begins to behaviour as if no selection pressure exists.

4 DISCUSSION

The previous descriptions of a spatial sub-population model for evolution have shown a number of attractive properties:

Parameter	Value
Population Size	500
Visibility	100
Initial Spatial Distribution	$[-100,-100] \rightarrow [100,100]$
Base Selection	6-Tournament Selection
Replacement	Steady-State
Iterations	100000
Crossover	1-point
Crossover Prob.	1.0
Mutation Prob.	0.01
Mutation Bit Prob.	0.01

Table 2: Settings for the 100-bit problem.

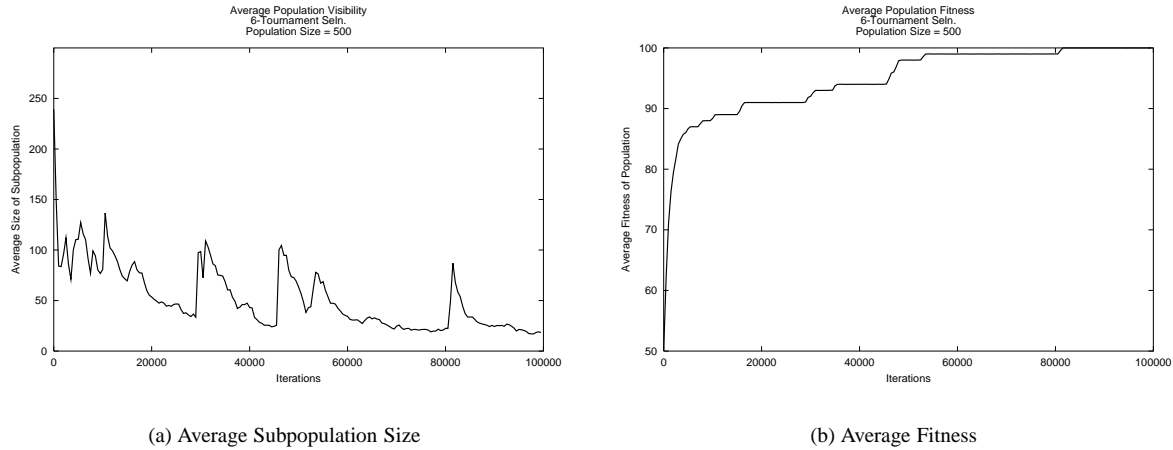


Figure 7: Behaviour using 100-bit fitness.

1. Average sub-population size is independent of initial spatial scale;
2. The same type of behaviour is observed for a variety of spatial dimensions;
3. Population distribution is self-similar at all scales;
4. Average sub-population size is independent of v after a certain number of iterations; and
5. The sub-population size automatically configures itself when a fitness pressure is applied to the population.

The implications of this work, especially given the promising performance using spatial selection (Dick & Whigham 2002), imply that the approach is both interesting and worthy of extension. The fact that the system does not rely on a specific spatial size or visibility to structure the subpopulation means the approach is parameter free, and is of relatively low cost to implement (complexity scales linearly with the number of individuals in the population). Since the described spatial system allows the automatic configuration of subpopulations it is also a simple approach to the issue of population distribution, and how this could be structured based on individual fitness during evolution.

5 FUTURE WORK AND CONCLUSION

There are a number of directions that this work can take in the future:

1. A mathematical description of the random spatial model, to explain why the average long run behaviour is independent of spatial extent and v .
2. Extending the standard schema theorem for Genetic Algorithms (Vose 1991) to incorporate spatial structure, with the aim of explaining why performance improvements have been observed.
3. Implement a distributed model of this approach, to demonstrate how the spatial structure can be taken into account.

4. A study of subpopulation dynamics for biological systems needs to be undertaken to determine whether similar patterns to the described spatial model can be observed in physical systems.

In conclusion, this paper has introduced the basic properties of a spatial visibility model for constructing dynamic subpopulations for evolutionary computation. The behaviour of the system is independent of initial conditions, and adjusts the subpopulation size in response to fitness pressures. This work extends concepts in distributed evolution and may also be a useful tool for exploring properties in biological systems.

ACKNOWLEDGEMENTS

Colin Aldridge must be thanked for raising the issue of random models and multiple dimensions.

References

- Adami, C. (1998). *Introduction to Artificial Life*. Springer-Verlag. New York.
- Cliff, D. & Miller, G. (1996). "Co-evolution of Pursuit and Evasion: Simulation Methods and Results" In P. Maes, M. Mataric, J. Meyer, J. Pollack & S. Wilson (eds), *From animals to animats 4*. MIT Press. Cambridge, MA pp. 506–515.
- Dick, G. & Whigham, P. (2002). "Population Density and Spatially Constrained Selection in Evolutionary Computation" *Technical Report Information Science Discussion Paper 2002/08*. Dept. of Information Science, University of Otago.
URL: <http://divcom.otago.ac.nz/infosci/publctms/complete/papers/dp2002-08.pdf.gz>
- Forrest, S. & Jones, T. (1994). "Modeling complex adaptive systems with Echo" In R. Stonier & X. Yu (eds), *Complex Systems: Mechanisms of Adaptation*. IOS Press. Amsterdam pp. 3–21.
- Grefenstette, J. (1981). "Parallel adaptive algorithms for function optimisation" *Technical Report TR CS-81-19*. Vanderbilt University, Nashville, TN.
- Haynes, T. & Sen, S. (1996). "Evolving Behaviour Strategies in Predators and Prey" In G. Weis & S. Sen (eds), *Adaption and Learning in Multiagent Systems*. Lecture Notes in Artificial Intelligence Springer Verlag. Berlin.
- Hraber, P. & Milne, B. (1997). "Community Assembly in a model Ecosystem" *Ecological Modelling*. **103**: 267–285.
- Hull, D. (1980). "Individuality and Selection" *Annual Review of Ecology and Systematics*. **11**: 311–332.
- Ray, T. (1992). "An approach to the synthesis of life" In C. Langton, C. Taylor, J. Farmer & S. Rasmussen (eds), *Artificial Life II*. Addison Wesley Reading, MA pp. 371–408.
- Robertson, G. (1987). "Parallel Implementation of Genetic Algorithms in a Classifier System" In J. Grefenstette (ed.), *Proceedings of the 2nd International Conference on Genetic Algorithms*. Cambridge, MA. pp. 140–147.
- Rosin, C. & Belew, R. (1997). "New Methods for Competitive Coevolution" *Evolutionary Computation*. **5**(1): 1–29.
- Schwehm, M. (1992). "Implementation of genetic algorithms on various interconnection networks" In M. Valero, E. Onate, M. Jane, J. Larriba & B. Suarez (eds), *Parallel Computing and Transputer Applications*. IOS Press. pp. 195–203.
- Vose, M. (1991). "Generalizing the Notion of Schema in Genetic Algorithms" *Artificial Intelligence*. **50**: 385–396.
- Yao, X. (1999). "Introduction to Evolutionary Computation" In X. Yao (ed.), *Evolutionary Computation: Theory and Applications*. World Scientific. Singapore pp. 1–26.