

Evolutionary Economic Agents

Fergus Nolan, Jarek Wilkiewicz, Dipankar Dasgupta, Stan Franklin

The University of Memphis
758 North McLean Boulevard
Memphis TN 38107 USA.

fergus@fnolan.com, jwilkiewicz@acm.org, dasgupta@msci.memphis.edu, stan.franklin@memphis.edu

Abstract

An empirical work is described which compares the optimization levels produced by a group of economic agents versus those of a similar group of economic agents which additionally employ a genetic algorithm (GA) to attain a higher level of optimization. The problem domain is multimodal. It incorporates multiple hard and soft constraints and dynamical behaviors. It also has areas of infeasibility and non-linear behaviors. The simulated model environment provides several types of sensors, actuators and opportunities for inter-agent resource mediation. Evidence is offered to support the theory that multiple weak methods operating in concert, on a shared problem, can produce better results than the individual weak methods acting alone. The problem area is resistant to the use of strong methods.

Introduction

Many real-world domains require the ability to deal with environments and problems which include multimodal features, non-linear functions, environmental interactions, feasibility constraints both hard and soft, and often, dynamical elements.

An example of one of these problems is natural language understanding. An interpretation of human natural language understanding proposed by cognitive science researchers (Bates and Goodman, 1997, Tomasello, 1998) could suggest that the human natural language process possesses many or all of these real-world features.

Such systems are referred to as complex systems, following nomenclature established by the researchers at the Santa Fe institute (Holland and Mimaugh 1996, Kauffman 1993).

The initial idea for this study is from Holland and Mimaugh, where they suggest that Genetic Algorithms could be combined with an economic agent (Wellman 1996, Mullen and Wellman 1996) to solve complex systems problems. However, the Echo agent discussed in Holland and Mimaugh's work lacks essential elements of the Wellman economic agent model, and Holland never followed up on his idea. This paper is an attempt to do some work in this area.

The essential enhancement obtained by adding the economic agent capability to the hybrid agent/GA architecture

exemplified in Holland's Echo is the market-based ordering behavior of the economic agent model. A market is a mechanism used by agents to mediate the distribution of constrained resources. In complex systems terms, a market allows Pareto optimization in the presence of hard constraints (quality requirements applied by a consumer prior to acceptance of a good or service) and soft constraints (a preference of a consumer to accept the less expensive of two goods which both meet hard constraints).

Using Luc Steels (1990) example of agents operating locally using weak methods, and adding the enhancement of a market for resource conflict mediation, yields the idea of economic agents. As a group, they can achieve a stable state of equilibrium which, according to General Equilibrium Theory, has the properties of Pareto Optimality and Pareto Quality. This implies that a stable solution meets all hard and soft constraints, and that the fitness or welfare of one agent cannot be improved without worsening the condition of at least one other agent. The emergent effect of the economic agent community is to provide a satisficing solution to a complex optimization problem. Economic agents are weak methods because the summation of their fitness or welfare values is not guaranteed to be optimal.

Genetic Algorithms are best at optimization of general functions using an exploratory approach. They become more complex and lose generality when faced with constrained problems or multi-modal problems. They are also weak methods, but they can produce satisficing optimization.

Holland's suggestion of combining GA with Economic Agent processing has several interesting potential features, when combined with Franklin (1995, 1997). Agents are autonomous, they have goals and drives, they can sense their environment and change it. Economic agents additionally obey hard and soft constraints and allocate resources efficiently. A group of economic agents can obtain a globally feasible solution to a multi-modal problem with a good degree of global optimization. Evolutionary economic agents can breed better agents that can improve the overall global level of optimization. This is group-level learning.

The question to be addressed by this study is: Does the addition of an evolutionary process to a community of economic agents improve the overall level of optimization?

We note that the existing mathematical models used for GAs cannot be relied on, as the assumption of a stable envi-

ronment is violated. In addition, we cannot rely on General Equilibrium Theory for this system, as equilibrium might never be achieved in a non-static environment. The complex systems mathematical analysis (Kauffman, 1993) addresses mainly inter-species genetic interactions, such as herbivore kills by predators, and does not easily look extendable to economic analysis. Issues of computability are likely to arise. As a practical matter, experimentation offers a promising opportunity to explore these techniques.

Related Work

There is a large literature on hybrid genetic algorithms. Hybrid schemes are available with neural networks (Genetica, NeuroForecaster, Moriarty and Miikkullainen, 1998), cellular automata (Dolan 1997-1999, Agre and Chapman 1990, Holland 1996), subsumption or simulated subsumption, case-based reasoning, hillclimbing (Lobo and Goldberg 1996), and fuzzy systems (Wilson 1994) as the lower layer.

The complex system nature of economies is discussed persuasively (Anderson, Arrow and Pines, 1988, Arthur, Durlauf and Lane, 1997, Kollman and Miller, 1996 and Krugman, 1996). These works focus on modeling real economies, which provides useful techniques and modeling prototypes. It should be noted that this empirical study uses an economic simulation mainly as a metaphor and that modeling real economies is not our objective.

Memetic Algorithms use a population-based approach with local heuristic search in optimization problems, and is described with examples at www.densis.fee.unicamp.br/~moscato/memetic_home.html, the Memetic Algorithms Home Page (accessed on 12/06/1998).

Pure genetic algorithms exist which deal with multi-modal problems and seek to achieve Pareto optimality within the GA (Horn, Nafpliotis and Goldberg 1994). This type of problem typically involves multiple constraints. The solution space in these problems therefore contains infeasible areas. The GA may generate infeasible points in the optimization process. The technique is successful in allocating the population over multiple modes and in sharing fitness over the modes, but can be difficult.

Parallel GAs are an alternative technique (Cantu-Paz 1998, Cantu-Paz 1997) with application to multi-modal problems. Motivated by a need to distribute processing over multiple nodes in a parallel computer, this technique maintains parallel populations and introduces interplay between the populations using new genetic operators like migration.

There is no evidence that the combination of Economic Agents and GA has been tried before.

Simulation Objective

After the example of Echo (Holland and Mimnaugh 1996), and with respectful gestures to Animat (Wilson 1994), a simple simulation of a complex system was chosen as a test case for

the theory.

A community of agents is proposed which inhabit an environment, which they can sense to some extent and which they can change by their actions. They produce commodities and trade them in an effort to survive, to consume and to create wealth. The wealth of each agent is used as a fitness value for the GA.

Each agent has a set of 10 chromosomes. The GA acts on these chromosomes to simulate breeding between agents and the selection of agents for replacement (or death). The fitness value being optimized by the GA is the wealth measure produced by the economy.

The chromosomes influence the decisions of the agents, as does the market prices and other information from the environment, and fuzzy logic adds an exploratory, stochastic factor.

The GA rates for selection, crossover and mutation are controlled by parameter. These parameters can be set to zero to prevent GA processing.

The objective is to measure the wealth summations produced by several runs of the model with GA active and compare with similar measures produced by a community of economic agents without evolutionary capability.

Model Description

This section describes the environment, the products, the economic agents (producer, consumer and auctioneer), the GA, and the evaluation cycle of the model. The metaphor of a simplified primitive human economy is used, as this is a complex system whose workings are familiar to most.

Environment

Figure 1 is an edited screen shot of the simulated agent environment. It shows an overall wealth indicator as a bar on the left, and a 10 X 10 grid representing a 100 square mile area. There is a town near the center shown as a black square. It is the location of the annual market and the auctioneer is based there.

The four shaded areas represent farmland and the two farmer groups live there. The farmland can produce food, clothing and beer when worked by a farmer. The remaining 95 white areas are forest, which has two populations of animals, herbivores and predators. The forest localities object maintains the two animal populations in a simple cycle which is deterministic except for the actions of the hunters.

Figure 1 also shows the four groups of economic agents, as small pie charts. F1 and F2 are the farmers and H1 and H2 are the hunters.

The Products in the Economy

Wellman's economic agent model defines the resources in the system as products or goods.

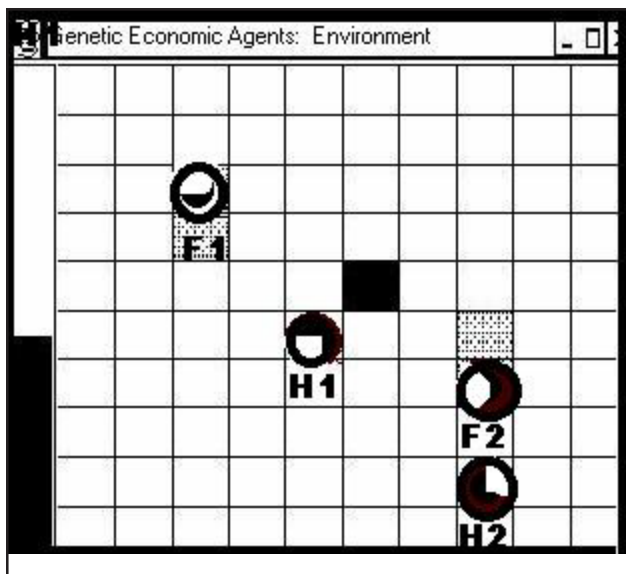


Figure 1. The simulated agent environment

The simulated products in the economy are food, clothing and beer, and these are generic commodities. They are accounted in arbitrary units, where one unit of clothing is the equivalent of a large outer garment, a unit of food supports an agent for six weeks, and the unit of beer is a keg.

Farmers can produce all products, a total of about 10 to 20 units per year. The amount produced per farmer is a function of genetic factors, the farmer's production decisions, environmental factors, group size, crop specialization and a large measure of fuzziness. The farmer decides what to grow at the beginning of the season based on the latest available prices which he evaluates as market-follower or contra, depending on his genetic disposition. He can raise or lower the percentage of a given crop 10% in any year.

Hunters try to fuzzily sense their environment for presence of other groups and animal population levels. They sense all squares contiguous to their current position, with an error level based on their genetic ability for hunting and sensing, and decide fuzzily where to move. They try to hunt where their activities will affect animal population sizes beneficially, in accordance with the predator/prey population cycle status at that locality. They hunt all year as a group, then divide the food and clothing produced by a tournament process in which a pecking order is established based on random factors and genetic aggression and hunting ability.

Each agent is endowed initially with an arbitrary number of gold pieces, and it is the universal medium of exchange, one of the items exchanged in each transaction. Unlike the other commodities, negative gold levels may be carried by agents. This is done because a fixed gold standard appears to suppress economic activity and we are not interested right now in simulating monetary effects. This can and does mean that individual agents can end up with big negative balances, and even inherit them from an agent replaced in a genetic

operation. These individuals are at increased risk of selection for death each round, but the issue seems not to effect the model. To ameliorate the problem, we added a solvency drive to the agents. The behavior appears to favor the non-GA simulations slightly, as the GA tends to thresh in the presence of large negative wealth balances.

The Economic Agents

The producers are the only true agents in the system. They are autonomous, they sense their environment and global price data, they change their environment, including their peers, by taking resources, killing animals and helping set market prices. They have drives (eat, be warm, drink beer, breed). They have goals (survive, accumulate wealth). They decide what to do, based on their sensor data, on their genetic dispositions, on their goals and drives and in a fuzzy manner. They participate in a market for all goods. They offer their goods to the market and bid on goods available.

A limitation of the model is that the population, number and size of groups are fixed during a run. There are two groups of farmers and two of hunters, of equal size. The overall size of population is set by parameter, and fixed during a simulation. Populations of 80 were used in this study.

The Consumer

Wellman specifies the consumer as an agent. In this model, there is one consumer per producer, and it is implemented as a simple, deterministic method in the Producer agent bean. Each year it accesses the producer's storehouse and consumes some amount of goods. It then awards a fitness score to the agent, which is a scaled measure of wealth incorporating some constraints.

The wealth score of 1 denotes the basic survival level. It is currently set at 8 units of food and 2 of clothing. The consumer sets a score of less than 1 if the basic survival level is not met, based on the proportion of the survival quota the agent has. This reduces the agent's chances of survival.

If the (hard) survival constraint is met, the consumer then assesses an additional amount of beer and clothing to consume. If the agent has this quantity, it is consumed and the agent's score increases according to the value of goods and gold he owns, expressed as a multiple of the survival quota.

The fitness score therefore incorporates wealth, the hard survival constraint and the additional wealth-accumulation and consumption objectives, which are soft constraints.

The Auctioneer

An economic agent in the Wellman scheme, the auctioneer is a simple, deterministic mechanism in this model. It collects bids and offers, compares them, sets the price which will allow maximum sales, then performs the transactions implied by the offers less or equal to the price, and the bids greater than or equal to the price.

There is a simple, iterative process of auction rounds with an arbitrary deadline. The auctioneer also supplies agents with current global commodity prices on demand.

The Genetic Algorithm

We use a simple steady state GA, which is invoked once per economic agent year. There is a parameter to control the percentage of the population selected for breeding. Agents are selected one at a time until this quota is met, with probability equal to their wealth. The selected agents are paired randomly and crossed over at a random bit in the gene. Mutation is then performed on each bit with a probability determined by parameter. Selection rates of 0.25 and mutation of 0.001 were used in the GA version of the test runs. During testing, higher levels of crossover were very disruptive.

A group is selected for replacement by using a weighted probability equal to the highest wealth minus the agent's wealth. This results in the wealthiest agent always surviving until the next year, while the least successful agents are more likely to be selected for the replacement group.

The offspring is assessed for farmer and hunter genetic aptitudes and matched with a replacement in a corresponding occupational group. The offspring then replaces the culled individual, inheriting his (probably meager) wealth, and receives one third each from the parents accumulated goods and gold. The splitting of parental wealth in this way is arbitrary and is intended to reflect the economic costs of child rearing while keeping the goods in the system constant. We found during testing that improved genes produced late in the run had little effect on the overall results unless they were given some initial wealth, and this allocation results in the offspring having the mean wealth of the parents.

This keeps the population constant, which is the result of a technical limitation, not an objective of the model.

There are 10 genes, each of which occupies 2 bits. The allele values, from 00 to 11, are Low, Medium, High and eXtra. Each gene controls an attribute and contributes to economic agent's performance at various sensory, productive and decision-making tasks. The genes helpful to farmers are: Defense, Focus, Physical Strength, Work Ethic. Hunters are helped by Aggression, Swiftiness, Tool Use and Wanderlust. Both groups are helped by Cunning and Explore. The chromosome has ten genes occupying a 20-bit string.

The Yearly Processing Cycle.

The number of processing cycles in a model run is determined by parameter, typically 30. Each cycle corresponds to a calendar year in the model. See figure 2 for an outline of the processing.

In the production part of the cycle, different processing is provided for farmers and hunters. Farmers decide what to grow based on what was grown last year, current produce

prices and various technical points. The growth is based on farmer genetics, farmers preferences, group factors, crop specialization, and a major random component. Hunting genes are a negative to a farmer as he wants to go hunting during the growing season.

All hunting activities and decisions are done as a group. There are arbitrarily 140 hunting days per year. At the beginning of the season the hunters decide what to hunt based on market prices and preferences based on their genetics. Each day of hunting, the hunter group collectively estimates the predator/prey populations in all adjacent squares, with accuracy based on genetic and random factors. They then decide which adjoining locality to move to, based on an estimation of the potential bag in adjoining localities, the product preference they decided at the beginning of the season and genetic factors including wanderlust. The hunters then perform hunt-

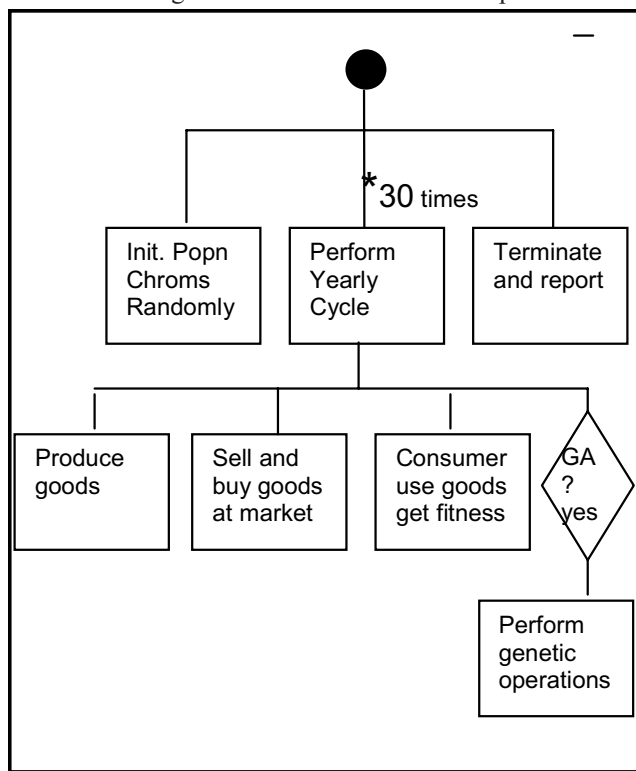


Figure 2: Outline of Agent Processing Cycle.

ing. The day's bag is determined by genetic factors relating to hunting ability, randomness and the numbers of predator and prey in the location. The long-term well-being of the hunters depends on their innate hunting ability and the accuracy at which they estimate the predator/prey positions on their characteristic population cycle. If the hunters choose well, they tend to act as a dynamical element in the animal population cycles, reducing the fluctuations and optimizing the potential bag in the future. If they exacerbate the natural predator/prey cycles, they can crash all animal populations in a location, thus making famine more likely in the future.

One wonders at our forebears expertise in these complex life-and-death decisions, and wishes that more of this expertise had survived the scientific era.

At year end, the bag is divided among the hunter group. Each gets a survival ration, if possible, and the rest is fought over in a random tournament, weighted by individuals hunting aptitudes and aggression gene level.

Implementation Details

The implementation uses Java on Intel equipment, Windows 98 and Windows NT 4.0, Borland Jbuilder 2 and a package entitled GA Playground (Dolan 1997-1999).

Additional beans were written to implement the Agent processing and the Agent Environment GUI, and these were combined with minor changes in the GA Playground source, as kindly provided by the author. See (Dolan 1997-1999) for details of additional freely-distributed class libraries used by the GA Playground.

The simplest possible technical approach was taken to the agent processing, implying that concurrency issues, reusability and performance opportunities were avoided.

Basic reporting tools, including a GUI, ASCII diagnostics and a report file in spreadsheet format are provided.

Empirical Results

Two series of tests were run and statistics gathered for comparison. Each series was 30 runs of the model. Each run simulated 30 years in the agent environment. The test population was 80. Endowment was 100 gold pieces. The chromosomes for all agents were generated randomly and an assessment of hunting versus farming capability was used to assign individuals to the groups.

The only difference between the two series was that no GA was employed in the second series. This was done by setting the parameters for selection, crossover and mutation to zero. For the with-GA series, selection for crossover was 0.25 and mutation 0.001. In the with-GA tests, each year, from the population of 80, 20 become parents, 10 die, 10 offspring are born and 70 individual economic agents survive into the next year. This is a steady-state GA.

Figures 3 and 4 show the results of both test series. The minimum value for charting purposes was forced to 2 to allow for the logarithmic axis in the charts, as there was a small number of negative wealth values caused by major price fluctuations and 'recession'.

The graphs show a clear superiority in the version with GA, although this requires fairly close study. The maxima and end points for the GA run dominate the upper reaches, for example, with 8 runs versus 5 for the no-GA series ending up in the 100k-1M band. The average maximum wealth over the 30 runs for the no-GA scenario was 732,992 versus 1,174,727 for the with-GA run, an improvement of 60%. The variation between individual runs is large, and it is unlikely

that the current data establishes a statistically significant result. The inherent instability of the model precludes reliability in results from one run to the next, and it is clear that the model has enough random factors in it. It is also apparent that the type of dynamic factor that Wellman noted in the operation of the speculator/arbitrageur agents is badly needed.

The model does not converge in the 30 annual cycles used in the experiment. We would expect more consistent results if that occurred. However, the chaotic and emergent nature of the model would not guarantee that a convergent run would not contain large deviances also.

Several epiphenomena were noted: cycles of recession (or famine), inflation, individual agent insolvency and animal population crashes, as well as the emergence of tycoons. Many of these are predicted by Wellman or by general economic theory. The emergence of super-performing individuals was interesting. These individuals typically gained wealth of 100 times the population average. We would be interested

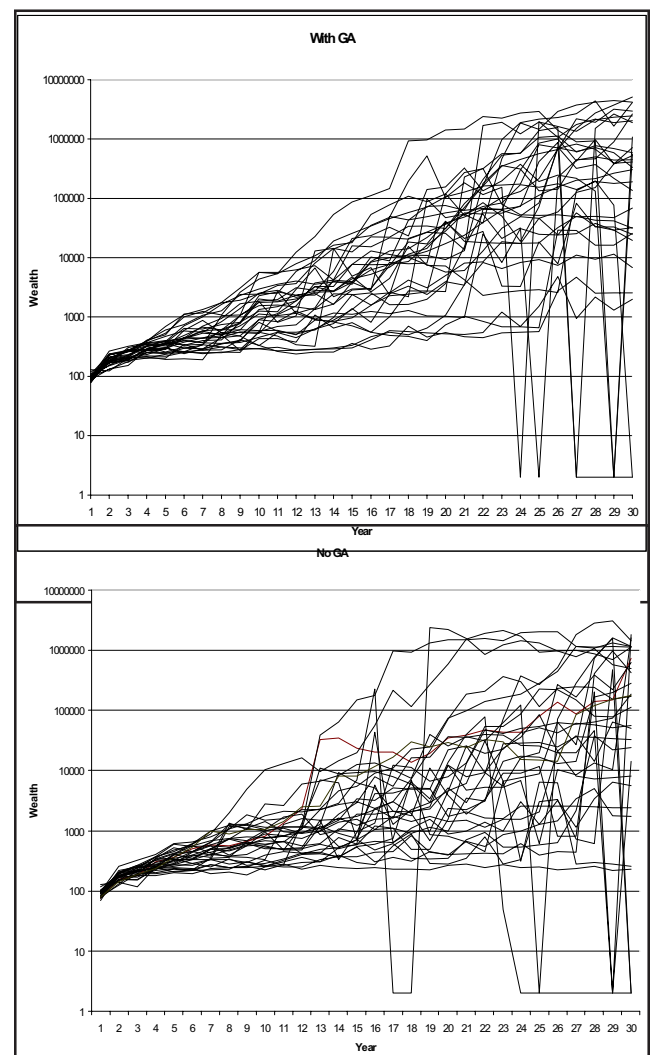


Figure 3, model run with GA, Figure 4 without GA.

in formulating future problems in such a way that they would represent outstanding solutions to the problem or search at hand. The typical execution of the model for 30 cycles took one and a half minutes on a 300MHz Pentium II.

Future Work

Although the initial results, given the crudeness of the model, are encouraging, this work raises more issues than it answers. We need better design techniques for emergent systems, a fuller economic agent implementation, a real-world environment to work with, and an interesting problem to solve.

One of the most critical issues during development was the lack of a comprehensive, easy-to-use set of monitoring and visualization tools. Time spent developing these at the outset of an emergent system development will be saved many times over later in the project.

Acknowledgments

The authors acknowledge the kindness and expertise of Ariel Dolan (Dolan 1997-1999), who shared the unpublished Java source code for his software, GA Playground (Dolan 1997-1999), and who helped with guidance and support. We recommend his example and his web page.

References

- Agre, Philip E. and Chapman, David. "What Are Plans For?". *Designing Autonomous Agents*, Maes, Pattie, ed. Amsterdam, Elsevier Science Publishers, 1990. 17-34.
- Anderson, Philip W.; Arrow, Kenneth J. and Pines, David (ed). *The Economy as an Evolving Complex System*. Proceedings Volume XXVII, Santa Fe Institute. Reading, Mass: Perseus Books 1997.
- Arthur, W. Brian; Durlauf, Steven N. and Lane, David A. (ed). *The Economy as an Evolving Complex System II*. Proceedings of the Evolutionary Paths of the Global Economy Workshop, 1987. Reading, Mass: Perseus Books 1988.
- Baars, Bernard J. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge Press, (reprinted) 1990.
- Baars, Bernard J. *In the Theatre of Consciousness*. Oxford: Oxford Press, 1997.
- Bates, Elizabeth and Goodman, Judith. 1997. On the inseparability of grammar and the lexicon: evidence from acquisition, aphasia and real-time processing. In *proceedings, Language and Cognitive Processes*, 12(5/6), 507-584.
- Cantu-Paz, Erick. *Designing Scalable Multi-Population Parallel Genetic Algorithms*. IlliGAL Report No. 98009. May 1998. <http://www-illigal.ge.uiuc.edu>
- Cantu-Paz, Erick. *A Survey of Parallel Genetic Algorithms*. University of Illinois, 1997. <http://www-illigal.ge.uiuc.edu>
- Dolan, Ariel. *GA Playground*. <http://www.aridolan.com>. Java code and HTML document, 1997-1999.
- Franklin, Stan. 1995. *Artificial Minds*. Cambridge Mass. MIT Press.
- Franklin, Stan. Graesser, Art. 1997. Is it an Agent, or just a Program? : A Taxonomy for Autonomous Agents. In *proceedings of Third International Workshop on Agent Theories, Architectures and Languages*. Intelligent Agents III. Berlin. Springer-Verlag. 21-35.
- Genetica. commercial GA software by NIBS Inc. at <http://sunflower.singnet.com.sg/~midaz/Nfga611.htm>, viewed 11/1/1998.
- Holland, John H., Mimmnaugh, Heather. 1996 *Hidden Order : How Adaptation Builds Complexity*. Reading MA. Perseus Press.
- Horn, Jeffrey; Nafpliotis, Nicholas and Goldberg, David E. A Niche Pareto Genetic Algorithm for Multiobjective Optimization. In *Proceedings IEEE-ICEC 1994*.
- Kauffman, Stuart. *The Origins of Order: Self Organization and Selection in Evolution*. Oxford: Oxford University Press, 1993.
- Kollman, Ken and Miller, John H. *Computational Political Economy*. Santa Fe: Santa Fe Institute, 1996.
- Krugman, Paul. *The Self-Organizing Economy*. Malden: Blackwell Publishers, 1996.
- Lobo, Fernando G., and Goldberg, David E. Decision making in a hybrid genetic algorithm. IlliGAL Report No. 96009, September 1996. <http://www-illigal.ge.uiuc.edu>
- Moriarty, David E. and Miikkulainen, Risto. Forming Neural Networks through Efficient and Adaptive Coevolution. *Evolutionary Computation*, 5(4), 1998.
- Mullen, Tracy and Wellman, Michael P. Some Issues in the Design of Market-Oriented Agents. In M. Wooldridge, J. Mueller, and M. Tambe (eds.), *Intelligent Agents II: Agent Theories, Architectures, and Languages*, Berlin Springer-Verlag, 1996.
- NeuroForecaster. software by NIBS Inc. Website at <http://sunflower.singnet.com.sg/~midaz/Nfga611.htm> viewed 11/1/98.
- Steels, Luc, *Exploiting Analogical Representations, Designing Autonomous Agents*, (ed) Maes, Pattie. Amsterdam, Elsevier Science Publishers, 1990. 71:88.
- Tomasello, Michael (ed). 1998. *The New Psychology of Language: Cognitive and Functional Approaches to Language Structure*. Mahwah, NJ. Lawrence Erlbaum Associates.
- Wellman, Michael P, *Market-Oriented Programming: Some Early Lessons*. In *Market-Based Control: A Paradigm for distributed Resource Allocation*, (ed) Scott H. Clearwater. Singapore: World Scientific, 1996. 74-95.
- Wilson, Stewart W. *ZCS: A Zeroth Level Classifier System*. 1994.