

Improved Resource Exploitation by Collective Intelligence

Boschetti, F.

CSIRO Marine and Atmospheric Research, E-Mail: Fabio.Boschetti@csiro.au

Keywords: Collective Intelligence; Complex Adaptive Systems; Minority Game; Resource Management.

EXTENDED ABSTRACT

It is largely understood that the contribution which scientific research can provide to the management of limited renewable resources includes not only the study of the dynamics of the resource itself, but also the understanding of how humans interact with the resource and ultimately manage its exploitation. For long the realm of qualitative social science, the last decade or so has seen the pioneering of quantitative study of certain aspects of human behaviour and, despite still at its infancy, interesting applications start to appear (Little et al, 2004, Dreyfus-Leon 1999, Varis 1998).

One instance of such effort, applied to the competition for limited resources, has been pioneered by Arthur (1994) and later generalized in what is now called Minority Game (MG, see Challet and Zhang 1998). This tool has allowed scientists to highlight unexpected behaviours displayed by the community as a whole.

The fundamental reason for the complexity of the MG (and of the competition for limited renewable resource in general) is its self-referential and self-defeating nature (Batten, 2005). Imagine a group of individuals who routinely has to choose where to access a limited resource. The amount of the resource an individual can obtain depends on how many other individuals choose to access the resource at the same location (since the resource is limited and needs to be shared). How many individuals choose a specific location depends in turn by the expectation that the location will be more or less exploited. Such expectation will guide the individuals' choice and consequently its final level of exploitation. Consequently, *the expectation actually determines the outcome*: this is the self-referential aspect. Also, the more individuals expect a location to be profitable, the more will access it and the less profitable the location will result. *The expectation actually determines the opposite outcome*: this is the self-defeating aspect.

The self-referential and self-defeating nature of the MG has lead some authors to suggest that optimal

exploitation of a limited resource is not possible under purely competitive pressure (Hardin, 1968). The purpose of this work is to show that, at least under specific circumstances, this may not be necessarily so. The tool we employ is called Collective Intelligence (COIN, Wolpert and Tumer 1999). It has been proposed to address a number of engineering problems (mostly in computer science, see Tumer and Wolpert 2000, Wolpert et al, 2004) and has already shown good result on a generalized form of the MG (Wolpert et al 2000, Wolpert and Tumer 2001).

Our contribution is to simplify this method with the view of a possible use by actual human communities. We show that a COIN could potentially be employed by real people with no need of computer aid, by simply performing elementary calculation with pen and paper. We test its potential on a fishery exploitation problem and compare the results to different algorithms traditionally used in the MG literature, in which individuals act either randomly or rationally and can evolve their behaviour. The important result is that the use of COIN leads to improved resource exploitation not only for the overall community but also for each individual (on average), that is, no personal sacrifice is required for the good of the community. This may be a crucial factor in opening the road to a possible use of COIN in real applications. The results could naturally be extended to the exploitation of resources other than fisheries.

1. COMPETITION FOR LIMITED RESOURCES: PROBLEM SETTING

In this section we describe our test problem. We model a fishing fleet composed of competitive vessels, targeting a limited resource of a single fish species.

In our virtual fishery we employ N fishing vessels $n=1..N$ (agents in the general CAS literature). At each iteration of the game, each individual vessel needs to choose where to fish among Z available zones (only one zone can be chosen by a vessel at each iteration). In each zone a certain amount of

fish $Fish_z, z = 1..Z$ is available. This amount can change from zone to zone but, unless otherwise stated, is fixed in time. This means that we are not modelling the population dynamics of the fish species and we assume a constant amount of fish is present in the zones at each iteration, independently of how much has been caught in the past.

The amount of fish available at a particular zone is shared equally among all the vessels which chose to operate in that specific zone. Consequently, the catch of each vessel depends on the action of all other vessels. Also, each vessel has a maximum fishing capacity, that is, it is not able to catch more than a predetermined amount (the same for each vessel). Thus the catch for each vessel is given by:

$$Catch_n = \text{Min}(Fish_{zone_n} / Crowd_{zone_n}, MaxCatch) \quad (1)$$

where $Catch_n$ is the amount of fish caught by vessel n , $zone_n$ is the fishing zone chosen by vessel n , $Fish_{zone_n}$ is the amount of fish available in $zone_n$, $Crowd_{zone_n}$ is the number of vessels which chose to fish in $zone_n$ and $MaxCatch$ is the fishing capacity of each vessel.

Naturally, the total catch of the fleet is given by the sum of each individual catch,

$$TotalCatch = \sum_{n=1, N} Catch_n \quad (2). \text{ Notice that:}$$

$$\sum_{n=1, N} Catch_n \leq \text{Min}(\sum_{z=1, Z} Fish_z, N * MaxCatch) \quad (3)$$

that is, the fleet can catch the maximum allowed amount of fish only provided the vessels spread their fishing effort strategically among the different available zones. In this work we explore whether the vessels are able to adopt a strategy which allows the fleet to maximize $TotalCatch$.

2. THE MINORITY GAME (MG)

In the last few years a considerable body of work has been published on the study of the Minority Game (Zhang 1998), mostly within the physics of condensed matter community (see also <http://www.unifr.ch/econophysics/minority/> for an exhaustive collection of papers on the subject).

The crucial component of the MG is the process used by the vessels to determine which zone to fish. In our work, we modelled this by a method proposed Chau and Chow (2003). At the beginning of the simulation, a set of strategies is assigned randomly to each vessel. At each iteration,

different criteria are used by the vessels in order to select which strategy to employ among the ones available. The different criteria determine different algorithms which play the MG. In our work we tested 3 implementations:

- 1) 'Simple MG' (MG_{simple}): during the simulation each vessel keeps track of how many times in the past each strategies guessed correctly the next most profitable zone (independently on whether the zone was actually chosen by the vessel). At each iteration, the strategy with current highest score is employed to determine where to fish next. This implementation is the closest to the 'standard MG' discussed in the literature (Challet and Zhang, 1998).
- 2) 'Random MG' (MG_{random}): at each iteration each vessel picks a zone totally randomly. This implementation is useful as 'null' test case to evaluate the performance of other MGs.
- 3) 'Evolutionary MG' (MG_{evol}): every M iterations, each vessel's strategy set undergoes an optimization step via a Genetic Algorithm (GA) (see Boschetti et al 1996 for a description of the specific GA used in this work).

For the sake of clarity and conciseness, in the rest of the paper we will use the shorthand MG(s) to refer to the three algorithms described above, not to the Minority Game itself.

3. OPTIMISING THE WORK OF A COMMUNITY: THE COLLECTIVE INTELLIGENCE

In this paper we ask what strategy individual vessels should employ in order to achieve the best global exploitation of a resource (and possibly best individual return too). This can be cast as an optimization problem. Here the word optimization has a broader meaning that just maximizing an economic return. Rather we refer to the search, within the space of all possible vessels' strategies for set(s) which results in a specific global outcome (in our case best global catches).

One of the main challenges in numerical optimization is how to design a suitable 'cost function', that is a measure of how good a certain outcome is, compared to the outcome we wish to achieve. In our problem, one option for a 'cost function' could be to maximize the catch of each individual vessel. This option is represented by the MGs described above.

Another natural choice for a cost function could be to assign to each *individual* vessel a share of the *global* catch (that is, the better the global catch, the more each vessel is rewarded). In game theory, this is called a ‘team game’. This solution would work well with small fleets. When the number of vessels increases, ‘team game’ performances tend to worsen quickly (Wolpert and Tumer, 2001). The reason for this lies, once again, in the self-referential nature of the problem. When many vessels are modelled, it is hard for an individual vessel to determine how much its own action has affected the global catch. Vessel *n1* may have taken a very unprofitable action, but the actions of all other vessels may have compensated for it by still producing a good global catch. Vessel *n1* may thus ‘think’ its action was profitable and may adopt it again in the future, thereby preventing the fleet to improve its performance. In Wolpert et al. (2004) this problem is called lack of ‘intelligence’, and is referred to the inability of the vessels to obtain useful information (‘intelligence’) about the problem.

COIN addresses both problems (‘greed’ and lack of ‘intelligence’) and, remarkably, the solution it proposes is very simple. Here we illustrate briefly the method, while we refer the reader to Wolpert and Tumer (1999) for a detailed description of the mathematics, together with the theorems and their proof which give a solid base to the theory.

COIN overcomes the ‘team game’ problem by attempting to estimate the contribution of each individual vessel to the global catch. This can be achieved by calculating what the catch of the fleet would be if a specific vessel *n1* (say) did not exist. In the rest of the paper we will use the superscript ‘-*n*’ to refer to values calculated for the fleet as if vessel *n* did not exist. For example $TotalCatch^{-n1}$ is the (hypothetical) catch of the entire fleet in the absence of vessel *n1*. Notice that, in general, this is different from calculating a quantity for the entire fleet minus the same quantity for vessel *n1*: $TotalCatch^{-n1} \neq TotalCatch - Catch_{n1}$. Rather, the effect of vessel *n1* to the total catch should be calculated as:

$$E_{n1} = TotalCatch - TotalCatch^{-n1} = \sum_n Catch_n - \sum_{n \neq n1} Catch_n^{-n1} = Crowd_{zone_{n1}} * \text{Min}(Fish_{zone_{n1}} / Crowd_{zone_{n1}}, MaxCatch) - (Crowd_{zone_{n1}} - 1) * \text{Min}(Fish_{zone_{n1}} / (Crowd_{zone_{n1}} - 1), MaxCatch) \quad (4)$$

In Wolpert and Tumer (1999) it is proved that any difference function of the form of Equation 4 is

aligned, that is, by increasing E_{n1} we can not decrease *TotalCatch*. It results that if we could optimize Equation 4 for each individual vessel then we would automatically optimize the catch of the entire fleet.

In practice, this approach would be extremely complicated to implement, if not impossible. Because of the self-referential nature of the problem, it is very hard to calculate what the behaviour of the fleet would be if vessel *n1* did not exist. This is because the behaviour of each vessel in the fleet is affected by the presence/action of *n1*. Basically, we would have to account for the dynamics of the game and rerun the entire simulation from the starting point

It turns out that a remarkable (computationally feasible) simplification is possible which allows for the general idea to be useful. In Wolpert and Tumer (1999), it is shown that the effect of a set of agents can be approximated by simply replacing their action with a fixed one. In COIN jargon this operation is called ‘clamping’ and different options for such clamping operation are discussed (Wolpert and Tumer, 2001). This means that we can disregard the dynamics of the problem and still implement the idea with meaningful results. The simplest option for this approach (and the one which make more sense in our application) is to use 0 as ‘clamping’ value. In our example (Equation 4), this corresponds to vessel *n1* not going fishing and calculating the resulting $TotalCatch^{-n1}$. Notice that this calculation is fictional, because in our model (as in most realistic applications) vessel *n1* can *not* choose not to go fishing. In the COIN literature this approximation is given the fanciful name of Wonderful Life Utility (WLU), in the name of an old famous movie. Since the WLU is aligned and avoids the ‘team game’ problem, both self-referentiality and self-defeatingness are circumvented.

In order to appreciate the WLU it is important to understand the difference between E_{n1} and $Catch_{n1}$. If vessel *n1* chooses a zone which is under exploited, that is a zone where vessels can fish at their maximum capacity (this is clearly a profitable action to take), E_{n1} in Equation 4 will be large, since vessel *n1* will be able to catch its share without affecting the other vessels’ catch. In this case $E_{n1} \geq Catch_{n1}$. If vessel *n1* choose a zone which is over exploited, that is a zone where vessels can not fish to their maximum capacity (this is clearly an action to avoid), $E_{n1} \leq Catch_{n1}$ since all vessels need to share the limited amount of fish available and no extra catch is possible.

To solve Equation 4, a vessel n_l needs to access the catches in *the zone the vessel has fished in*, unlike in the MGs, in which it has to access the record of the global catches. However, unlike in the MGs, vessel n_l also needs to access the number of vessels which shared his fishing zone ($Crowd_{zone_{n_l}}$). Consequently, a vessel needs more detailed information, but needs to share it only with neighbouring vessels, not with the entire fleet.

The optimization of the global catch is obtained while *each individual vessel tries to maximize its own return*, since each vessel looks for underexploited zones. The vessels are still acting *selfishly*, but they are acting much less *competitively* than in the MGs. Also, unlike in the ‘team game’, no obvious sacrifice needs to be asked to each vessel in the name of a better community goal. This promises best outcomes for the fleet coinciding with best results for each individual vessel. Our experiments in section ‘Numerical Results’ below will verify whether this promise holds true.

4. COIN IMPLEMENTATION

In our COIN implementation, a strategy is represented as a set of weights, one for each fishing zone $W_z, z=1..Z$. W_z gives a measure of the expectation of obtaining a good catch at zone z . At each iteration, the vessel performs a random pick among the Z zones, with the probability of selecting zone z proportional to W_z . Basically, the more a vessel expects a specific zone to provide a good catch, the more likely the zone will be chosen. This parameterization allows for a balance between *exploration* and *exploitation*. In the statistical physics parlance, this corresponds to a Boltzmann sampling with no temperature decay.

At each iteration t , each vessel evaluates the catches it obtained in the past T iterations and it uses this information to set W_z (that is, to predict which will be the most profitable fishing zone at time $t+1$). In order to account for non stationarity in the predictions, the past catches at times $t_p = t-T+1..t$ are discounted proportionally to $t_d = t-t_p$, that is, the longer ago the catch was taken the less it influences the prediction. In all our test we used $T=20$. This implementation can be seen as a simplification of the ones proposed in Wolpert and Tumer (1999) and Wolpert and Tumer (2001).

Finally, in order to better evaluate the performance of the COIN, a ‘team game’ implementation is also tested, in which the ‘cost function’ for each vessel

is simply a share of the global catch ($Catch_n = \frac{TotalCatch}{N}$), that is, vessels are fully cooperative. In the rest of the paper this algorithm will be referred to as TG.




Figure 1. Mean Global Catch for the different algorithms in the test with $N=50$, $Z=4$, $MaxCatch=5$. The curves are averaged over 10 runs

5. NUMERICAL RESULTS

In this section we compare the performances of the different algorithms we introduced above. In the first test we model a fleet of 50 vessels ($N=50$) and 4 available fishing zones ($Z=4$). The resources available in each zone differ: zone 1, 3 and 4 can provide 50 units of fish ($Fish_{z \neq 2} = 50$), while zone 2 can offer 100 ($Fish_2 = 100$). Each vessel can catch at most 5 units ($MaxCatch=5$). Following Equation 3, we can estimate the maximum possible catch for the fleet:

$$TotalCatch \leq \min(\sum_{k=1,K} Fish_k, N * MaxCatch) = 250 \text{ unit.}$$

A summary of the results can be seen in Table 1. For each algorithm, we show the mean catch for the vessel which caught the largest amount of fish (column 1), the mean catch for the vessel which caught the smallest amount (column 2), the total mean catch of the entire fleet over the entire simulation (column 3), and the mean catch of the entire fleet at the last iteration step (column 4). In order to account for stochastic fluctuations inherent in the algorithms, all results are averaged over 10 different runs.

We can notice the following:

- 1) COIN provides the best catch for the entire fleet as well as for each individual vessel: the worst performing vessel obtains a better catch than the best performing vessel with any other algorithms. This suggests that no individual sacrifice is necessary in order to achieve best

global catches when using the COIN approach.

- 2) The second best catches are provided by the Team Game. The TG approach of putting the fleet's interest above the individuals results in circumventing the self-defeating nature of the MGs, and thus outperforms them. However, the inability of the TG vessels to discriminate the effect of their own action from the ones of all other vessels results in performing worse than the COIN.
- 3) The performances of the three MGs are quite similar and considerably worse than on the COIN and TG.

Table 1. Summary of the comparison of the different algorithms in the test with $N=50$, $Z=4$, $MaxCatch=5$, $Fish_{z \neq 2} = 50$ $Fish_2 = 100$. In this and following table Max IC= maximum Individual catch, Min IC= minimum Individual catch, GMC=Global Mean Catch, FGC=Final Global Catch.

	Max IC	Min IC	GMC	FGC
COIN	4.89	4.80	242.40	250.00
TG	4.77	4.53	232.33	237.50
MG_{simple}	4.23	4.07	207.63	204.50
MG_{evol}	4.23	4.06	207.08	210.00
MG_{random}	4.23	4.07	207.58	207.00

It is also worthwhile to analyze the time series of the Global Mean Catch (GMC) produced by the different algorithms (Figure 1):

- 1) The MGs' curves oscillate around a fixed exploitation baseline. This is characteristic of the Minority Game behaviour (Arthur, 1994). No global adaptation can be seen as a function of time, not even for the algorithm employing evolution (MG_{evol}).
- 2) The TG quickly reaches good performance and slowly improves for the rest of the run, with oscillations of much smaller amplitude than the three MGs. This shows a certain level of adaptation to the modelled conditions. Despite the oscillation, the TG's GMC is always well above the MGs' ones.
- 3) The COIN's performance is the worst at the very beginning, but within 10 time steps it overtakes all other algorithms. It then keeps on improving, converging towards the maximum allowed catch (250) with oscillation of decreasing sizes.




Figure 2. Mean Global Catch for the different algorithms in the test with $N=50$, $Z=4$, $MaxCatch=5$. At the beginning of the run we have and $Fish_{z \neq 2} = 50$ and $Fish_2 = 100$. After 200 steps the fish stock distribution changes to $Fish_{z \neq 3} = 50$ and $Fish_3 = 100$. The curves are averaged over 10 runs.

In the second test we examine how the algorithms adapt to change in resource distribution. We use the same fishery scenario as in the first test ($N=50$, $Z=4$, $MaxCatch=5$). At the beginning of the run we set $Fish_{z \neq 2} = 50$ $Fish_2 = 100$. However, after 200 time steps, we modify $Fish_2 = 50$ and $Fish_3 = 100$, while we leave zones 1 and 4 unaltered. Basically the larger stock moves from zone 2 to zone 3. The time series of the Global Mean Catch can be seen in Figure 2. Obviously, the behaviour of the curves in the first 200 steps is similar to the one described in Figure 1. Soon after time step 200 the catches of both COIN and TG decrease suddenly and sharply. Naturally, the COIN and TG fleets had 'trained' themselves to exploit more heavily zone 2 and its drop in stock affects negatively their catches. Both COIN and TG are able to adapt to the new resource distribution and adjust their fleet allocation to exploit it. However, the COIN fleet is much faster in retraining to the new scenario and its GMC recovers faster than the TM's one. This explains the final better performance displayed in Table 2.

6. IMPLEMENTING COIN WITH PEN AND PAPER

Conceptually, the core of the COIN approach lies in Equation 4, which implements the idea of evaluating the effect that vessel nI has on the fleet by comparing the current catch to that of a hypothetical fleet without nI . Computationally, the crucial step lies in mimicking the behaviour of the hypothetical fleet without vessel nI by using the Wonderful Life Utility with clamping factor 0. The fact that this (very rough) approximation works, immensely simplifies the approach and makes the conceptual idea behind COIN actually

implementable. In this work, we simplified the COIN algorithm further by imposing some (albeit minor) modifications. In particular:

- 1) unlike in Wolpert et al. (2000), the weights W_z are not *adjusted*, but rather *reset*, at each iteration according to previous catches;
- 2) in previous applications of COIN to the Minority Game (Wolpert et al., 2000, Wolpert and Tumer, 2001) the return R of each agent for attending a certain zone is given by $R = y * \exp(-y/c)$, where y is $Fish_z$ in our problem. In this equation, c plays the part of a pre-determined 'ideal' number of agents attending each zone. The presence of the parameter c helps to guarantee a balanced allocation of the vessels among the zones. In order to simplify the calculation, we did not use this option;
- 3) finally, in Wolpert et al. (2000) and Wolpert and Tumer (2001), in order to account for non stationarity in the predictions, the past catches are discounted by an exponential function while we used a simpler, linear one.

Table 2. Summary of the comparison of the different algorithms in the test with $N=50$, $Z=4$, $MaxCatch=5$. At the beginning of the run we have $Fish_{z \neq 2} = 50$ and $Fish_2 = 100$. After 250 steps the fish stock distribution changes to $Fish_{z \neq 3} = 50$ $Fish_3 = 100$. All values are averaged over 10 runs.

	Max IC	Min IC	GMC	FGC
<i>COIN</i>	4.87	4.61	237.74	250.00
<i>TG</i>	4.80	4.25	225.74	237.00
<i>MG_{simple}</i>	4.26	4.09	208.49	210.50
<i>MG_{evol}</i>	4.23	4.06	207.04	210.50
<i>MG_{random}</i>	4.22	4.07	207.46	204.50

The net result of these modifications is that all is needed in order to implement a COIN strategy are some basic bookkeeping and a handful of elementary operations (+, -, *, /). While this makes little difference when COIN is run on a computer, it may make a difference if real human agents want to test/use the procedure. Basically, the entire COIN approach could be performed with pen and paper by agents with only primary school training. This may not be relevant to modern fishing fleet in developed countries (which are geared with sophisticated equipment), but it could broaden the application to other resource management problems in the developing world. In particular,

here is the 'pen and paper' pseudo-algorithm which vessel nI 's crew needs to perform:

- 1) keep a record of the day catch, and of how many other vessels fished in the same zone.
- 2) Calculate the WFU function E_{nI} and store it.
- 3) Retrieve the values E_{nI} from the last T fishing days/periods and call them $E_{nI}^t, t=1..T$ where t is day/period, with $t=1$ being the most recent day/period and $t=T$ the least recent.

- 4) Calculate the weights W_z as

$$W_z = \sum_{t=1..T} \delta_z^t E_{nI}^t \frac{T-t+1}{T}, \quad (6)$$

where $\delta_z^t = 1$ if the vessel fished in zone z at time t in the past, and $\delta_z^t = 0$ otherwise.

- 5) Normalise the cumulative sum of the weights

$$W_z, CumW_z = \sum_{i=1}^z W_i \text{ and } CumW_z = \frac{CumW_z}{CumW_z};$$

- 6) Pick a random number r between $[0,1]$.
- 7) Find the smallest z for which $CumW_z \geq r$.
- 8) Next day/period, fish in zone z .

7. DISCUSSION, LIMITATIONS, AND DIRECTION FOR FURTHER WORK

The results we presented are very encouraging. Nevertheless, before some serious expectations can be placed on this technique more detailed modelling needs to be done. Here we present a (short) high priority list of items for further work:

- 1) the resource dynamics needs to be included into the simulations. How the fish population responds to the COIN exploitation and how the COIN approach can adapt to the population dynamics itself is a crucial factor which needs to be carefully explored.
- 2) Human behaviour also needs to be modelled more accurately. For example, we should explore what would happen if one or more vessels in the COIN behave greedily (by following standard MG strategies), or communication between nearby vessels was disrupted (either purposely or accidentally).
- 3) In current implementation we are not modelling the availability of public information or 'a priori' expectations of fish population behaviour. It is worthwhile to study how such information could alter the effect of COIN.
- 4) The potential receptivity of real communities to a COIN approach needs to be assessed. To

our knowledge, no experience is reported in the literature on this subject.

We conclude by discussing by far the most important concern: in this work we assumed that the resource is fully renewable. In many resource management scenarios, maximizing the exploitation of a limited resource is often *not* what we want, especially in environments under considerable ecological pressure. In a certain sense, this issue is not directly related to whether COIN is a valuable approach, unless a less effective fishing strategy is seen as a management tool for stocks control. What is important is that COIN is an optimization technique. This, as mentioned before, does not necessarily imply maximizing exploitation, rather *finding a best strategy for a goal of our choice*. From this perspective, the challenge is to formulate an ecologically responsible and economically valuable goal and see whether COIN can help achieving it. Technically, this means finding a WLU 'cost function' which would result, *at the same time*, in a sustainable behaviour for the entire fleet and an individual 'selfish' goal which is worthwhile for each vessel to pursue. The results we presented suggest that such combined aims are not necessarily contradictory, rather they may be *aligned* and COIN can offer a valuable approach to finding such alignment. This is surely the single most important direction for future work we aim to address.

8. CONCLUSIONS

We propose a simplified version of the Collective Intelligence which can be easily employed by a community of human agents in order to plan the exploitation of a limited resource. We compared COIN against other game theoretical approaches on a number of virtual fishery scenarios, characterised by varying fleet size, fish stock distribution and fishing capacity. In all tests the COIN not only guaranteed optimal global catch but also maximised the catch of each individual vessel. Achieving this in a competitive environment may be a key factor in this method's acceptance by real communities. In the view of actual implementations, we described a pseudo algorithm, which allows the COIN to be carried out by 'pen and paper', with minimum bookkeeping and only elementary calculations.

9. REFERENCES

W.B. Arthur, (1994). Inductive behaviour and bounded rationality. *The American Economic Review* 84, 406-411.

Batten, D., (2005), Are some human ecosystems self-defeating?, *Environmental Modelling and Software* (forthcoming)

Boschetti, F., M. Dentith, and R. List, (1996), Inversion of seismic refraction data using Genetic Algorithms, *Geophysics*, 1715-1727

Challet, D. and Y-C, Zhang, (1998), On the minority game: analytical and numerical studies. *Physica A* 256, 514

Challet, D., M. Marsili, and R. Zecchina, (2000), Phase Transition in a Toy market, *Proceeding, Int. J. of Th. and Appl. Fin.* 3-3

Chow, F. K. , and H. F. Chau, (2003), Multiple Choice Minority Game, *Physica A*, Vol. 319, pp. 601-615.

Dreyfus-Leo'n, M.J., (1999), Individual-based modelling of fisherman search behaviour with neural networks and reinforcement learning, *Ecol. Modelling* 120, 287-297.

Hardin, G., (1968) The tragedy of the commons. *Science* 162, 1243-8.

Little, L.R., S. Kuikka, A.E. Punt, F. Pantus, C.R. Davies, B.D. Mapstone, (2004), Information flow among fishing vessels modelled using a Bayesian network, *Environmental Modelling and Software*, 19, 27-34.

Tumer, K., and D. Wolpert, (2000) Collective intelligence and Braess' paradox. In Proceedings of the Seventeenth National Conference on Artificial Intelligence, pages 104-109, Austin, TX, 2000.

Varis, O., (1998), A belief network approach to optimization and parameter estimation: application to resource and environmental management., *Artif. Intell.* 101, 131-163.

Wolpert, D. and K. Tumer, (2001), Optimal Payoff Functions for Members of Collectives, *Advances in Complex Systems*, 4(2/3):265-279.

Wolpert, D., K. Wheeler, and K. Tumer, (2000), Collective Intelligence for Control of Distributed Dynamical Systems. *Europhys. Lett.*, 49 (6), p. 708.

Wolpert, D. and K. Tumer, (1999), An Introduction To Collective Intelligence, Tech Report: NASA-ARC-IC-99-63, http://ic.arc.nasa.gov/ic/people/kagan/coin_pubs.html

Wolpert, D., K. Tumer, and D. Bandari, (2004), Improving Search Algorithms by Using Intelligent Coordinates, *Physical Review E*, 69, 2004.

Zhang, Y., (1998). Modeling Market Mechanism with Evolutionary Games, *Europhys. News*, 29, 51.