Man on Earth
– discovering viable ecological survival strategies

Bruce Edmonds

Centre for Policy Modelling,
Manchester Metropolitan University

bruce@edmonds.name

ABSTRACT
Many previous societies have killed themselves off and, in the process, devastated their environments. Perhaps the most famous of these is that of "Easter Island". This suggests a grand challenge for the agent community: that of discovering what kinds of rationality and/or coordination mechanisms would allow humans and the greatest possible variety of other species to coexist. In particular, solving this challenge consists of designing and releasing a society of plausible agents into a simulated ecology and assessing: (a) whether the agents survive and (b) if they do survive, what impact they have upon the diversity of other species in the simulation. The simulated ecology needs to implement a suitably dynamic, complex and reactive environment for the test to be meaningful. Agents, as any other entity have to eat other entities to survive, but if they destroy the species they depend upon they are likely to die off themselves. Up to now there has been a lack of simulations that combine a complex model of the ecology with a multi-agent model of society. A suitable dynamic ecological model and simple tests with agents are described to illustrate this challenge.

Categories and Subject Descriptors
G.3 [**Simulation and Modeling**]. J.4 [**Social and Behavioral Sciences**]: – *sociology*. K.4 [**Computers and Society**]. I.2.11 [**Distributed Artificial Intelligence**] – *multiagent systems*. J.3 [**Life and Medical Sciences**] – *biology and genetics*.

General Terms
Design, Experimentation, Human Factors.

Keywords
Agent-Based Modelling, Ecology, Social Intelligence, Ecological Intelligence, Resilience, Grand Challenge.

1. THE DANGERS OF ECOLOGICAL DAMAGE AND SOCIETAL COLLAPSE
The evidence is overwhelming that, many times, humans have destroyed the ecologies they inhabited to their own and other species detriment – sometimes causing whole settlements or civilizations to disappear. Examples include: the inhabitants of Easter Island who built its famous stone statues in a race for status and killed all trees on the island in the process [12], or the Mayan civilization where a combination of increasing climatic aridity, demands of agriculture and societal conflicts lead to an abandonment of their impressive step temples in the jungle [14]. However you look at it, humans have a profound effect upon the ecosystems they come into contact with, even to the extent that (as some have argued) we are in the middle of the sixth great extinction event – the *Holocene* [8].

However, *how* humans will effect a particular ecosystem is not always clear – sometimes it seems that a balance between humans and the rest of the ecosystem is established, but at other times, the arrival of humans can only be described as catastrophic [7]. The "Social Intelligence Hypothesis" [11] suggests that the main adaptive advantage that our brains give us is our ability to socially organize. From this view our brains provide us with social intelligence first (for example abilities to: recognize other individuals, to develop a personal identity relative to a group, to be able to communicate, to be sensitive to status, to imitate, to train our offspring for a long time, and to adsorb a whole culture when experienced over a long time); any "general" intelligence we have as individuals is a by-product of these social abilities.

Due to these social abilities, groups of humans can inhabit a variety of ecological niches. They do this by adapting to a niche in terms of developing a body of knowledge, including words, ideas, techniques, social norms, systems of value and ways of organizing, that enables the group to survive there [13]. Once established, this body of knowledge can be passed down to new members of the group so that the group can retain its ability to survive in that niche over time. Broadly, this set of knowledge can be associated with the culture of the group. Thus, the abilities of groups of humans can change far more rapidly than that of most animals that have to rely on genetic evolution. Humans are thus at a distinct advantage in terms of any adaptive "arms race" with other species. Their social intelligence has equipped them to survive in a hostile and unpredictable world, ensuring their own immediate survival as their priority (as with other species).

However, they do not necessarily plan for the long-term and can cause such a degree of environmental damage to the niches they inhabit that they endanger the survival of their own group [7]. In this way, the arrival of humans within a system of ecosystems can have a profound impact – not merely changing the extent of extinction but also the whole way that the dynamics of that ecosystem works. The abilities of humans are over-tuned towards immediate survival, with the contrary result that, in the longer-term, they grab resources to themselves in a way that can jeopardize their own group survival.

Now that humans, using their technology, can inhabit almost any ecological niche on earth, any ecological disaster that we cause might well not be limited to a particular niche but may affect us globally. This challenge is thus to use agent-based simulation techniques to help address this problem, and thus contributing to the survival of our and other species on this planet. This challenge can be seen as an amplification of that implicitly posed in [2] or else a contribution to the wider challenge posed in [7].

2. THE PARTICULAR CHALLENGE FOR THE MULTI-AGENT COMMUNITY

There have been many agent-based simulations addressing the interaction of man with the environment, going back (at least) to 1994 [1] (see [3] for a review). Individual-based ecological models go back even further (see [9] for a review). However to fully address this challenge we need to have a multi-agent model concerning human decision making and social interaction *combined with* an individual-based model of an ecology that more fully reflects the dynamism and complexity of real ecologies. Up to now, models of humans interacting with their environment have had either a relatively simple model of human interaction or a simple model of the ecosystem they are embedded in (such as a systems dynamic model). As it said in [6] in 2012:

"…The more serious shortcomings of existing modelling techniques, however, are of a structural nature: the failure to adequately capture nonlinear feedbacks within resource and environmental systems and between human societies and these systems." (p. 523)

In other words, to fully address this challenge we need to start to understand how the complexity of human cognition, the complexity of human society and the complexity of dynamic ecologies might interact. Otherwise, we might miss some of the complications that might affect our and other species survival.

To ensure that the environment in which the agents representing humans and their society is sufficiently challenging, we require a model of this environment that satisfies the following criteria.

a) The environment needs to include space, so that there can be a differentiation in terms of niches and allow for some spatial migration between parts of the landscape
b) The environment needs to include niches with different characteristics, for example deserts (which can not sustain life) and natural barriers (which impede migration)
c) Complex food webs of species need to be able to develop within each niche either extracting resources from the environment or other individuals (predation)
d) New species need to be able to evolve in response to the pressure of the environment, other species and humans
e) Agents representing humans need to be embedded within these niches, needing to use/eat other species to enable their own survival

Once such a test bed is established, the challenge would be implemented in several phases:

1. *Bed in the ecology*. Run the ecological model for a while to allow a rich and dynamic ecology to evolve.
2. '*Freeze'* the ecology and save this state as the starting point for different evaluative runs.
3. *Inject the human agents*. Then place a small society of agents with given cognitive and social abilities into the ecological simulation.
4. *Assess the result*. After a suitable period of simulation time assess the outcome.

The assessment of the final state of the simulation could be done in a variety of ways, including:

- Measuring the diversity of the ecology, for example the average genetic difference between individuals, as in [5] (excluding humans).
- The species-number distribution – how many species are there with a population of at least 2^n, where n varies (the "Species Abundance Distribution" of [10]).
- The number of trophic layers that have survived for a period of time since the injection of human agents, shown by the distribution of trophic layers.
- The health of the society of agents, in terms of the number of surviving humans and its variability over time.

Measures such as these can be brought together to assess the sustainability/health of the socio-ecological system as a whole.

Thus this challenge can be encapsulated as follows:
To design plausible cognitive and social abilities that, when implemented in agents and assessed in the above way, reliably result in a sustainable and healthy socio-ecological system.

3. TOWARDS A CHALLENGE TEST-BED

To illustrate this challenge I describe a simulation test-bed that meets the stated criteria, has been tested with simple agents representing humans, and assessed in some of the above ways.

3.1 The Basic Set-up

In this, entities, plants, herbivores and predators, are represented as individual objects. They inhabit one of a number of patches arranged in a 2D grid that makes up the world. Each patch is well mixed so that interactions within that patch are random, but there is a probability that each individual can migrate to one of the four neighbouring patches each tick. Each patch and individual has a binary bit-string that represents its characteristics. There is a basic energy economy; so that energy is 'rained' down into the world (each tick), divided equally between patches, and which ultimately drives the whole ecology. These bit-strings and a fixed random interaction matrix, described below, determine whether an individual can extract energy from a patch or predate upon another. The bit-string of any individual is passed to any progeny but there is a probability that one of the significant bits of their characteristic is flipped at birth.

3.2 Species Abilities and the Energy Extraction/Predation they Allow

Key to this understanding this simulation is how it is determined whether individuals can extract energy from a patch or predate upon another. This method is adapted from that in [4]. A random interaction matrix with the dimensions of the length of individuals' bit-strings is generated at the start of a simulation. It is filled with normally distributed random floating-point numbers (mean 0, SD 1/3). This interaction matrix determines which entity can eat another entity as follows: (1) the non-zero bits of the predator select the columns of the matrix, the non-zero bits of the potential prey select the rows; (2) the intersection of the selected rows and columns determine a set of numbers, (3) these are summed; (4) if the sum is greater than zero the predator can eat the prey, in which case the prey dies and the predator gains a percentage of its energy value (the rest is lost). This calculation is illustrated in Fig. 1.

Essentially the same process is used to determine which entities can extract energy directly from the environment, except that the part of the prey is taken by the patch with its bit string (padded with zeros to reach the appropriate length). In this case only those with scores greater than zero get any of the patch's energy. The patch's energy is divided between all qualifying individuals in proportion to their score against the patch. This scheme has the consequence that no individuals can extract energy from a patch with a bit-string of all zeros (a 'dessert').

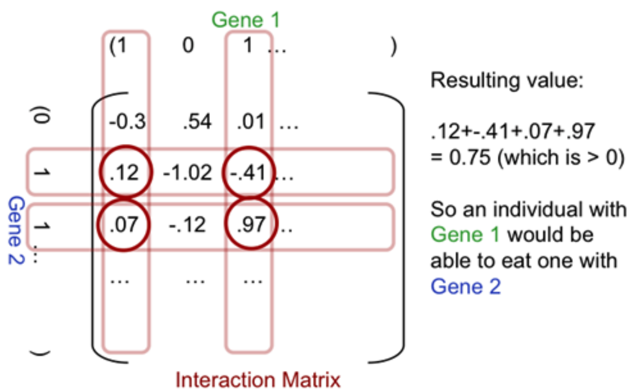

Fig. 1. The use of the interaction matrix to determine predation as well as energy extraction from a patch to give its relative fitness

This interaction scheme allows complex food webs to be evolved, for example via a genetic "arms-race" between predator species and prey species, since it allows for adaption with respect to another specific species. It also allows for competitive adaption to particular kinds of patches. In other words fitness is not an absolute number but relative to the environment and the other existing species, if it extracts energy from this, or another species. [4] showed that this kind of scheme can be used to evolve complex ecologies with plausible characteristics including food webs with similar network characteristics to those of observed food-webs (however this was for a single patch).

3.3 Simulation Execution

At the start of the simulation, the random interaction matrix is generated. Each patch is allocated a random bit-string with the given number of bits, padded out with zeros to make it the same length as individuals' bit-strings. The "environmental complexity" is the number of significant characteristics that patches can have – the number of bits in their bit-string. Bit strings of length 2 allow for 4 types of patch, of length 3 8 types etc.

The simulation starts with no individuals. Each tick:

- *Energy Distribution*. A fixed amount of energy is added to the model, equally divided between all the patches.
- *Death*. A life tax is subtracted from all individuals, if their total energy is less than zero it is removed.
- *Initial seeding*. (In the initial phase), until a viable population is established, a single random individual is introduced with a given probability each tick.
- *Energy extraction from patch*. The energy stored in a patch is divided among the individuals on that patch that have a positive score when its bit-string is evaluated against the patch's bit-string (in the above manner) in proportion to its relative fitness, at the simulation's efficiency rate.
- *Predation*. In a random order, each individual is randomly paired with a given number of others on the patch. If it has a positive dominance score against the other, the other is removed and the individual gains a fixed proportion of its energy, given by the "efficiency" parameter.
- *Maximum Store*. Individuals can only retain so much energy, so any above the maximum level set is discarded.
- *Birth*. If an entity has a level of energy > the "reproduce-level", it gives birth to a new entity with the same bit-string as itself, with a probability of mutation. The new entity has an energy of 1, subtracted from the energy of the parent.
- *Migration*. With a probability determined by the "migration" parameter, the individual is moved to one of the neighbouring 4 patches.

3.4 The 'Human' Agents

Broadly speaking, the agents representing humans should be processed in manner similar to any other individual with only a few differences. The most important difference is in the acquisition and passing on of techniques among their own group. Thus their "bit string" that determines their ability to predate upon (or resist being predated upon) is not determined genetically but can be learned socially by imitation from parents and/or peers. Whether an agent predates upon another individual and when it moves to a neighbouring patch could be part of what is determined by the agent's decision processes. One might well set the required minimum energy that humans need to give birth as much higher than for other individuals and allow them to store more energy. They might have a complex social structure with food passing between themselves according to its rules (e.g. an internal economy). They may have a tribal structure that allows each individual to recognise others from their own tribe and those who are outsiders, which may affect their behaviour. Many other extensions are possible to reflect other human attributes, e.g. warfare between groups, deliberate planting of crops or hoarding.

3.5 Some Illustrative Results

Fig. 2 shows the a graph of the number of species in a typical initial stage of the model, showing the development of plants, then herbivores and finally predators, providing a suitably complex and dynamic environment, with a range of trophic levels, ready for the injection of agents representing humans.

To give a simple flavour of some of the possible results, some

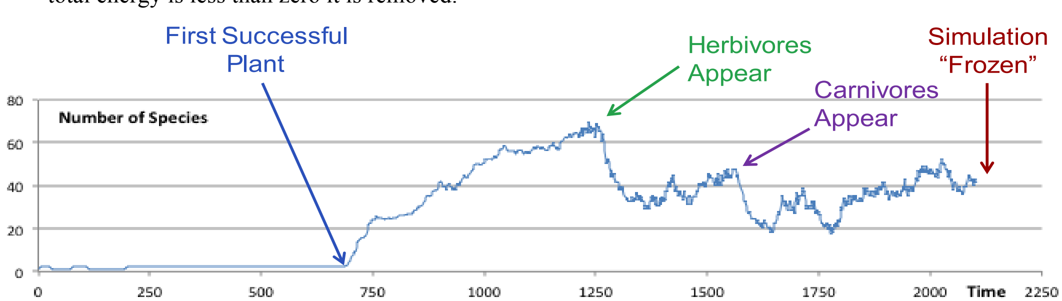

Fig. 2. A typical run of the model during the "Bedding In" phase, making the state of the simulation suitable for the injection of agents representing humans

very simple agents were injected at the point indicated in Fig. 2. Then the simulation was run for a further 1000 simulation ticks with different migration rates (the probability any entity or agent would move to a nearby patch in the 2D grid). 25 otherwise independent runs were performed both with and without 'human' agents added, and the final mean ecological diversity measured.

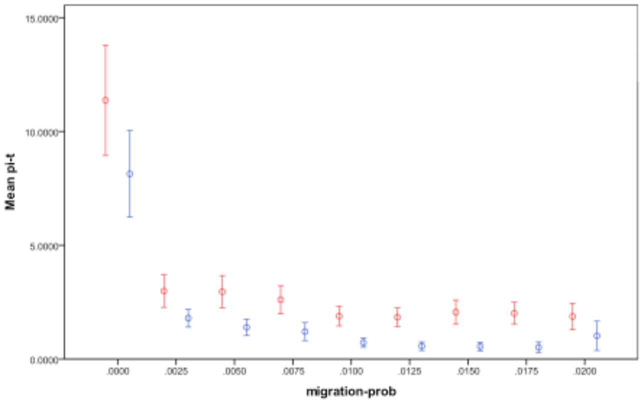

Fig. 3. Mean diversity for different migration rates with (blue) and without (red) human agents (error bars indicate a 95% confidence interval).

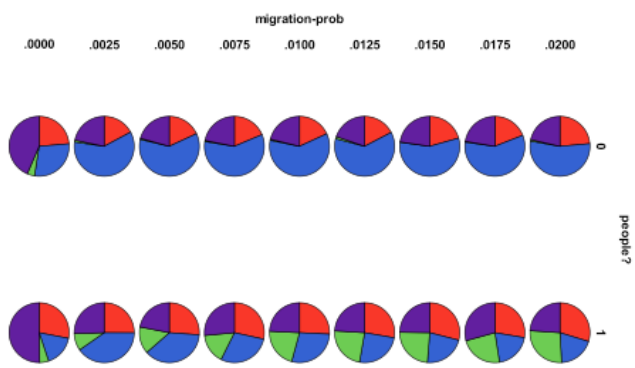

Fig. 4. Proportions of final ecological states at final tick over independent runs for different migration rates, with and without agents, red=plants only, blue=with higher trophic levels, purple=monoculture, green=non viable.

As we see from Fig. 3 and Fig. 4 above, the agents have a consistent and negative impact upon the ecologies they invade, but, generally, a higher negative impact at higher levels of migration, which tends to make the ecologies more uniform. However, as Fig. 4 indicates, this is far from a uniform effect, reducing the diversity a bit in each run. Rather it indicates an increasing proportion of ecological catastrophes (the green proportion in Fig. 4) that occur in many cases.

4. CONCLUSION

Multi-agent simulation could apply its expertise in terms of specifying and exploring the cognitive/social abilities of agents with respect to such a test bed, and start to tease out the complex and often counter-intuitive effects of such abilities. Knowledge about this could play a real part in helping us understand our own, fragile and complex, relationship with the ecologies we inhabit and exploit. It is time to show that agent-technology can deliver tangible benefits to our environment and our chances of survival.

5. ACKNOWLEDGMENTS

This research was partially supported by the Engineering and Physical Sciences Research Council, grant number EP/H02171X/1. Many thanks to those at the Manchester Complexity Seminar, especially Alan McKane, and those at iEMSs conference in 2012 for discussions on this subject.

6. REFERENCES

[1] Bousquet, F., Cambier, C. et P. Morand. 1994. Distributed artificial intelligence and object-oriented modelling of a fishery. Mathematical Computer Modelling, 2018, 97-107.

[2] Bousquet, F., Barreteau, O., Le Page, C., Mullon, C. & Weber, J. 1999. In Blasco F. & Weill A. (eds.) Advances in Environmental and Ecological Modelling, Elsevier, Paris, 113-122.

[3] Bousquet, F., & Le Page, C. 2004. Multi-agent simulations and ecosystem management: a review. Ecological Modelling, 176(3-4), 313-332.

[4] Caldarelli, G., Higgs, P.G. and McKane, A., 1998. Modelling Coevolution in Multispecies Communities. Journal of Theoretical Biology, 193, 345-358.

[5] de Aguiar, M.A.M., Baranger, M., Baptestini, E.M., Kaufman, L. and Bar-Yam, Y., 2009. Global patterns of speciation and diversity, Nature, 460, 384-387

[6] Deffuant, Guillaume et al., 2012. Data and models for exploring sustainability of human well-being in global environmental change. European Physical Journal Special Topics, 214(1), 519-545.

[7] Diamond, J., 2004. Collapse: How Societies Choose to Fail or Succeed, Viking.

[8] Eldridge, N. 2000. Life in the Balance: Humanity and the Biodiversity Crisis. Princeton University Press.

[9] Grimm, V. (1999). Ten years of individual-based modelling in ecology: what have we learned and what could we learn in the future? Ecological Modelling, 115:129-148.

[10] Hubble, S. P. 2001. The Unified Neutral Theory of Biodiversity and Biogeography. Monographs in Population biology, 32. Princeton University Press.

[11] Kummer, H., Daston, L., Gigerenzer, G., & Silk, J. 1997. The social intelligence hypothesis. In P. Weingart, S. D. Mitchell, P. J. Richerson, & S. Maasen (Eds.), Human by nature. Between biology and the social sciences. Mahwah, NJ: Erlbaum, 157-179.

[12] Loret, J. & Tancredi, J. (eds.). 2003. Easter Island: Scientific Exploration into the World's Environmental Problems in Microcosm. Kluwer/Plenum.

[13] Reader, J. 1990. Man on Earth. Penguin Books.

[14] Turner, B.L. & Sabloff, J. A. 2012, Classic Period collapse of the Central Maya Lowlands: Insights about human–environment relationships for sustainability. PNAS 109(35), 13908-13914.