EXTENDED ABSTRACT

THEODORE DAVID CARMICHAEL. Complex Adaptive Systems and the Threshold Effect: Towards a General Tool for Studying Dynamic Phenomena Across Diverse Domains. (Under the direction of DR. MIRSAD HADZIKADIC)

1. Introduction.

Most interesting phenomena in natural and social systems include transitions and oscillations among their various phases. A new phase begins when the system reaches a threshold that marks a qualitative change in system characteristics. These threshold effects are found all around us. In economics, this could be movement from a bull market to a bear market; in sociology, it could be the spread of political dissent, culminating in rebellion; in biology, the immune response to infection or disease as the body moves from sickness to health. Complex Adaptive Systems (CAS) has proven to be a powerful framework for exploring these and other related phenomena. Our hypothesis is that by modeling differing complex systems we can use the known causes and mechanisms in one domain to gain insight into the controlling properties of similar effects in another domain. To that end, we have created a general CAS model; one that is flexible enough so that it can be individually tailored and mapped to phenomena in various domains, yet retains sufficient commonality across applications to facilitate a deeper, cross-disciplinary understanding of these phenomena. In this work, we focus on the threshold effect. We show that the general model successfully replicates key features of a CAS. And we demonstrate its general applicability by adapting the model to three domains: cancer cells and the immune response; political dissent in a polity; and a marine ecosystem.

2. Motivation.

While there have been many implementations of CAS models used in various domains, these models are usually created in isolation, such that an economics model is used only to study market effects and dynamics, or a population model is used to describe and predict the causes and limits of population growth. There have also been examples of ideas and concepts inherent in CAS being transferred from one field to another. However, there has not been a general CAS model developed that is intended to replicate phenomena simultaneously in more than one field or area of inquiry. As Neil Johnson writes:

> "In particular, the connections between such systems have not been properly explored – particularly between systems taken from different disciplines such as biology and sociology. Indeed it is fascinating to see if any insight gained from having partially understood one system, say from biology, can help us in a completely different discipline, say economics."

Therefore, the main research goal of this work is to pursue such generality; to look closely at some of these connections across domains by working towards a general CAS model, one that can serve as a common language, even for fields that are far apart. We do this in the context of threshold effects, a phenomena common to many domains. By specifically identifying general properties in a common framework this CAS tool aims to: 1) facilitate the transfer of knowledge from one domain to another; and 2) stimulate a deeper understanding of the properties in one system by using the general model for mapping like properties from another system.

This is, by necessity, a challenging and long-term goal. In its current state, CAS modeling is more art than science, and it is difficult to capture all the salient features of even a single complex system. However, there is a tremendous amount of potential in moving towards a "language of CAS" that is widely applicable. What we show here is not intended to be a final solution, but rather a significant first step towards the development of this language.

3. Background Information.

CAS and ABM (Agent-based Modeling) techniques are powerful due to their inherent flexibility. But most of these models are not developed with general applicability in mind. Rather, they tend to be designed in a “stove-pipe” manner, without much thought given to the transfer of knowledge from one domain to another. Some exceptions to this include Schelling’s classic model on segregation, which is useful for explaining this fundamental property in a wide variety of domains; and the Axelrod, et al., model on state-level alliances during WWII that was later applied to corporation-level alliances regarding two competing versions of the UNIX operating system.

The model described here is designed to be broader than either of these examples. It is a general model for entire systems, rather than a piece or aspect of these systems. And, unlike the political alliance model, our system utilizes thousands or tens of thousands of agents, rather than just a handful. Further, our general tool is designed to produce the emergent properties of these systems as they appear over time, not just the end-state configurations.

We achieve these advances by focusing on the threshold effect: a phenomena found across systems and across disciplines – and by ensuring our model exhibits all the most salient features of CAS. We define a threshold effect as a change in sign or abrupt change in magnitude (either enduring or a spike) in the first or second derivative of a system variable. We characterize three distinct threshold processes: 1) the ratchet mechanism, 2) cumulative causation, and 3) contagion.

Although definitions of CAS tend to vary somewhat, there is largely agreement on what properties these systems have: a large number of self-similar agents that, through simple rules and simple interactions, produce self-organization and emergent, global-level features, and exhibit non-linear dynamics in their overall behavior.

4. Previous Work.

Two research projects we collaborated were used to help inform the design of the general CAS model: a computer simulation laboratory for social theories, in the context of an insurgency in Afghanistan; and a detailed model of soft-tissue cancer. These ABM simulations share two common traits: they both were designed with a single domain in mind, and they both utilize as much domain information as possible. While successful, the large number of agent attributes and system characteristics make it hard to deconstruct these models and determine which variables and which interactions have the greatest impact. Thus, our general model is both designed to be widely applicable and to use only the fewest number of attributes and relationships required to simulate and explain the desired outputs, thus following Occam's Razor. Further, our model is designed to be generative in nature; i.e., the global properties of the system are not imposed in a top-down manner, but rather with bottom-up emergence of features based on the agent interactions. This not only allow the model to replicate a key feature of CAS, but also gives both greater flexibility in the model's use and improves the model explanatory power.

5. Iterative Design of the General Model.

The first iteration of our model is not applied to any specific domain, but rather uses only generic agents: A-agents which are immobile and laid out in a grid pattern, and thus represent the simulation environment; and B-agents, which freely move around the simulation and interact with the A-agents. While this iteration does not exhibit all the properties of CAS, it does make a significant contribution. We show that self-organizing behavior among the B-agents – in this case, clustering behavior – does not depend on the B-agents interacting directly with each other. Instead, these agents interact *indirectly*, by the feedback mechanisms from B-agent to A-agents, and back to B-agent. Yet they still manage to self-organize. We contrast this with similar clustering behavior of agents found in the literature Figure 1. This new result could have interesting implications not only for ecological modeling and simulations, but also for application in a general sense to any "flocking" model. For example, in a cancer/immune system study, it may be assumed that clusters of immune agents are interacting with each other simply because they are clustered together. However, our results indicate that that is not necessarily the case, which could help us understand the physical processes involved in cellular activity.

The second iteration model is an improvement over the first in that it has the capacity for long-term stability and behavioral persistence. The previous version would always (by design) result in one of two end-states [Figure 2]. Thus, the ability to form and maintain a basin of attraction makes this model a more powerful CAS simulation, and allows for greater flexibility in research application and experimentation.

In this model, as with the first, the A-agents move in steps between the 0-state and the 1-state. Once they reach the 1-state they can affect neighboring agents towards the 1-state. B-agents still move, and still affect A-agents towards the 0-state, much like sheep moving around a field and eating grass. The refinement of this model is that these agents are limited in lifetime and eventually die out. However, they also reproduce as a positive function of how many A-agents are affected, thus preserving a resilient system. If the number of B-agents dwindle, then the A-agents "grow" more; if the B-agents are too many, then their "food" source is restricted, which reduces the number of new B-agents being introduced into the system.

6. Application to Soft-tissue Cancer.

In this mapping, the A-agents represent a soft-tissue cancerous growth, while the B-agents represent immune cells that attack the cancer, moving it towards the 0-state (also referred to as the "healthy" state). Figure 3 and Figure 4 shows that our model has successfully reproduced the behavior of other *in silico* cancer and immune-system experimental results found in the literature. The flexibility of the model is also highlighted here, in that a wide range of growth rates and response dynamics are simulable, even though the general model remains quite simple in terms of the agent-types and interactions. Intriguingly, this model also exhibits long-term transitions from one steady state to another [Figure 5]. These transitions are neither periodic nor – thus far – predictable.

7. Application to Political Dissent in a Polity.

Using the same general tool, we then applied this model to political dissent. The civilians are represented by the A-agents and the government forces, which act to quell the dissent, are represented by the B-agents. During the process of mapping the general model, two changes were made, both of which enhance the model's flexibility while still preserving its general applicability. The first is that the neighborhoods – e.g., each A-agent and its surrounding A-agents – are no longer restricted to simply the adjacent agents. Each defined neighborhood can be both a larger radius and include more neighbors. In this way, a wide variety of agent networks can be represented, and not only those that are restricted by geography. The second change is that the concept of "resources" was introduced into the system. For example, the available resources of the central government may impose a limit as to how many government agents can be created and maintained by the system. Conversely, a minimum number of agents can also be maintained regardless of the amount of dissent. This concept could indicate a police force or standing army.

While these refinements were produced to more realistically match political dissent, they also increase the model's generality, even in the previously discussed domain of soft-tissue cancer. Cancer does not necessarily spread to only nearby or adjacent cells. Furthermore, the immune system not only maintains a baseline level even in the absence of sickness, but it also inherently limited by the body's resources. In this case, the resource is "energy" rather than "money," but the concept is the same.

Multiple experiments with the dissent model application show that the contagion effect is qualitatively similar to the known properties of the spread of dissent in a population, as found in the literature. Further, the correlations between the number of dissenters and the government response are shown to be similar to real-world political dissent [Figure 6].

Another important similarity between the cancer application and the dissent application is that the transmission of effects between A-agents is additive in nature; unlike some other contagion models, both the political dissent and the spread of cancer require multiple contacts to spread, not just one. This commonality increases the possibility that a greater understanding of the dynamics in one of these domains may be applicable to the other.

8. Marine Ecosystem Application.

Both the soft-tissue model and the political dissent model exhibit some of the classic characteristics of the predator-prey dynamic, such that the cancer (or the dissent) is the prey, and the immune cells (or the government agents) act as the predator. Thus, the application of this model to a true ecosystem is a fairly straightforward process. We used the example of a marine ecosystem here, although other predator-prey systems would be equally appropriate.

In this mapping we again extended the model, this time to account for three, and later four, types of agents. The A-agents are analogous to plankton; the B-agents are fish that feed on the plankton; and the new 'predator' agents feed on the fish. Additional experiments were conducted with an "egg" agent, such that the fish would reproduce in a more realistic fashion. The reproduction rate, however it is simulated, remains a positive function of how many A-agents are affected (or, in this case, eaten). This is a standard assumption of the Lotka-Volterra equations on predator-prey dynamics.

The main results from this series of experiments are derived by relaxing another assumption from these classic equations: that of unlimited food available to the prey population. Thus, we adjusted the amount of plankton growth in the simulation and quantified the effects on the other two agent population. What we found was surprising, and not previously found in the literature. As expected, the prey population size increases along with the increases in food available, and thus the predator population also increases (reflecting the fact that there are more fish available to eat). However, the fish population did not remain elevated for long. Instead, *all* the gains from the extra resources at the lowest tropic level accumulate to the predators; the fish population quickly reverts to the previous level, such that their gains are only temporary [Figure 7]. This is possible because the fish, though they are eating more, and reproducing more, are also dying much faster. Thus, their average age decreases even as their overall population levels remain the same [Figure 8].

9. Significance and Future Work.

As Johnson, Epstein, and others have pointed out, we have not previously seen a model that is specific enough to capture the dynamic properties of a complex system, yet still general enough to be applied to multiple domains. Thus, what we have accomplished here is to increase the range of complexity available to very simple, general models, while simultaneously increasing the generality of very domain-specific models, in order to meet somewhere in the middle. It is not, as we have stated, intended to be a final solution; rather, it is an important first step.

Our general CAS model has successfully reproduced a particular class of phenomena, defined herein as threshold phenomena, found in disparate domains. Furthermore, the few key attributes, as listed in Table 1, do not represent all the variables, functions, and interactions that our model represents and that the three domain mappings – cancer, political dissent, and marine ecosystems – all have in common. Thus, there are certainly many more potential systems that can be simulated with this general tool, as the extent of its range and flexibility has not yet been fully explored.

The ultimate goal for this research is to build more and better bridges across the disciplines. In discussing reports on interdisciplinary creativity, van Raan states that "Eminent scientists strongly emphasize the crucial role of instruments for the progress of science, particularly the 'bridging' role between disciplines, by transferring instruments from one discipline to another."

The general tool described here is not simply intended to allow researchers from multiple fields to work on a single issue (although it may certainly be useful to that end). Rather, as we illustrated with the multiple systems that can be simulated, it is a way to bridge the gap between different complex systems themselves. There are many particular problems that require a multi-disciplinary approach to solve, but these are often done by splitting a multi-faceted problem into parts: the biologists tackle the biology parts, the economists study the economic parts, and so on, bringing the individual contributions back together for the final solution.

Our model instead acts as a common language: a detailed method for describing and discussing threshold phenomena. By emphasizing a common language rather than a common problem, our model allows research in one area to inform research in other, distant areas; further, it can lend insight into domain-specific issues simply by providing a methodology for reframing a research question in terms of our general model specifications.

Work is progressing on this front in many areas. The ecosystem mapping – where resource increases at the lowest tropic level accrue to the highest tropic level – may have application in economic systems as a sort of "trickle up" economic theory. Research is also underway to apply these lessons to concept acquisition in student learning, in the context of deductive proof construction, as well as the study of terrorism and human rights. We have also pending proposals to apply CAS techniques and this general tool to data mining in a hospital environment, and to local political engagement in a community.

These collaborative efforts, perhaps more than anything else, exemplify the potential of a robust, general CAS model as a common language for bridging the gap between disciplines, and leveraging insights found in one domain to inform research in many others.

10. Figures and Tables.

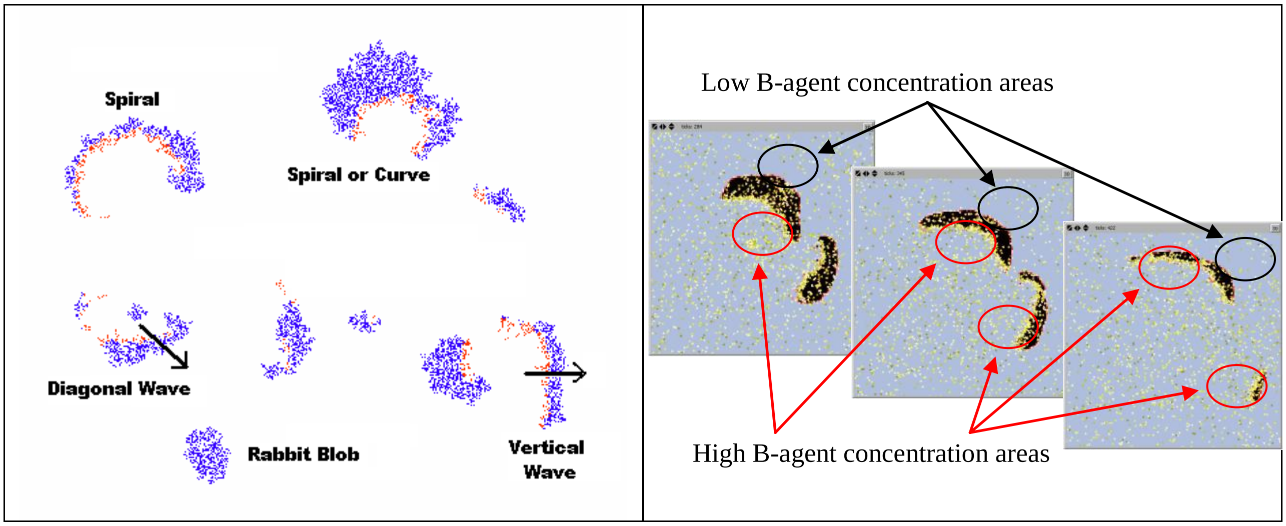

Figure 1 - On the left, a typical agent clustering behavior, compared to (right) clusters found in the general CAS model.

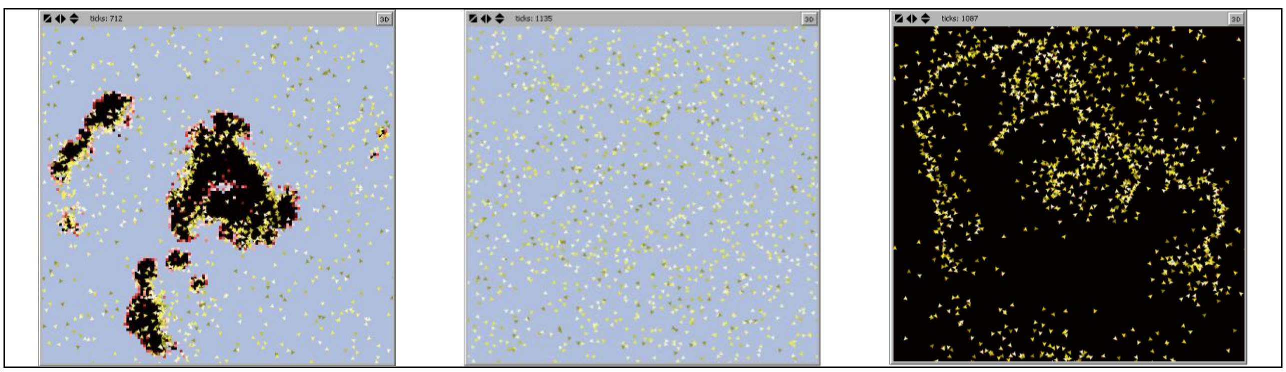

Figure 2 - First iteration model during a simulation run (left); and two possible end-states (middle, right).

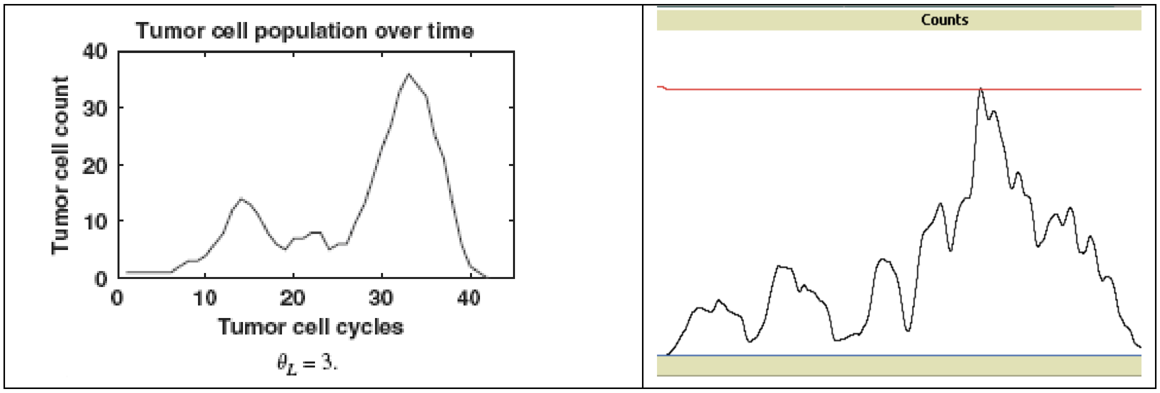

Figure 3 - Unsuccessful tumor growth from the literature (left), compared to unsuccessful tumor growth in the general CAS model (right).

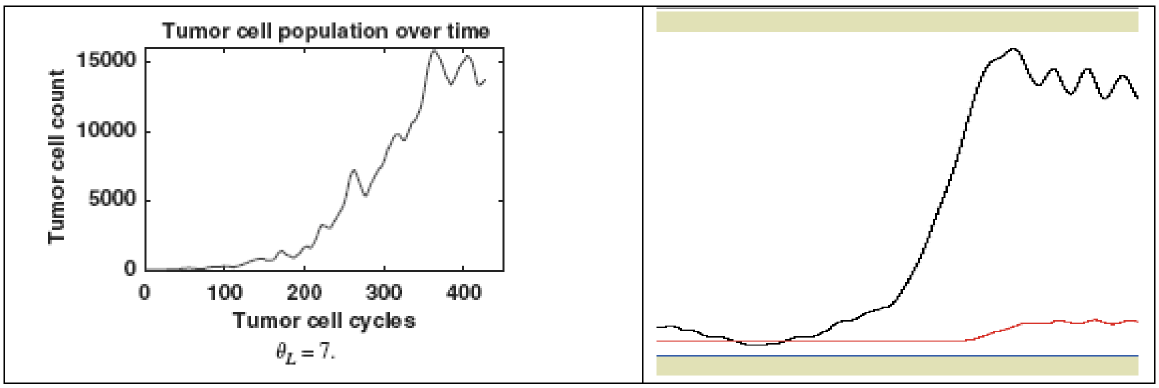

Figure 4 - Successful tumor growth from the literature (left), as compared to successful tumor growth in the general CAS model (right).

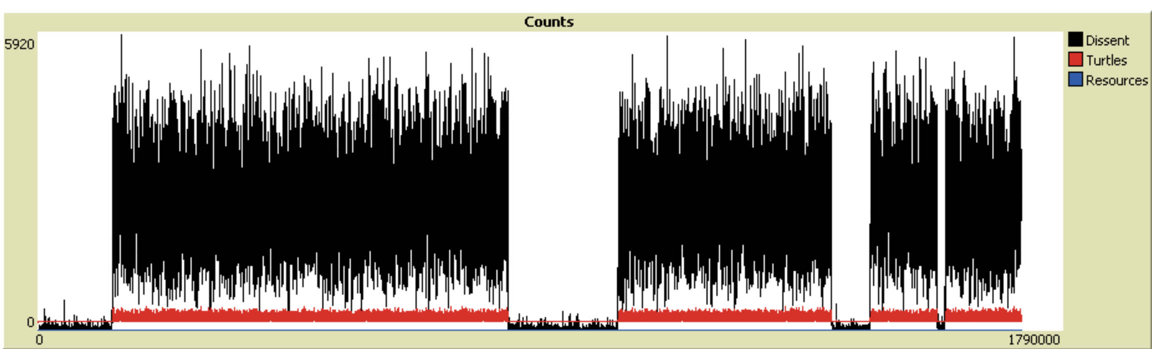

Figure 5 - Extended model run displaying stochastic transitions between two steady states.

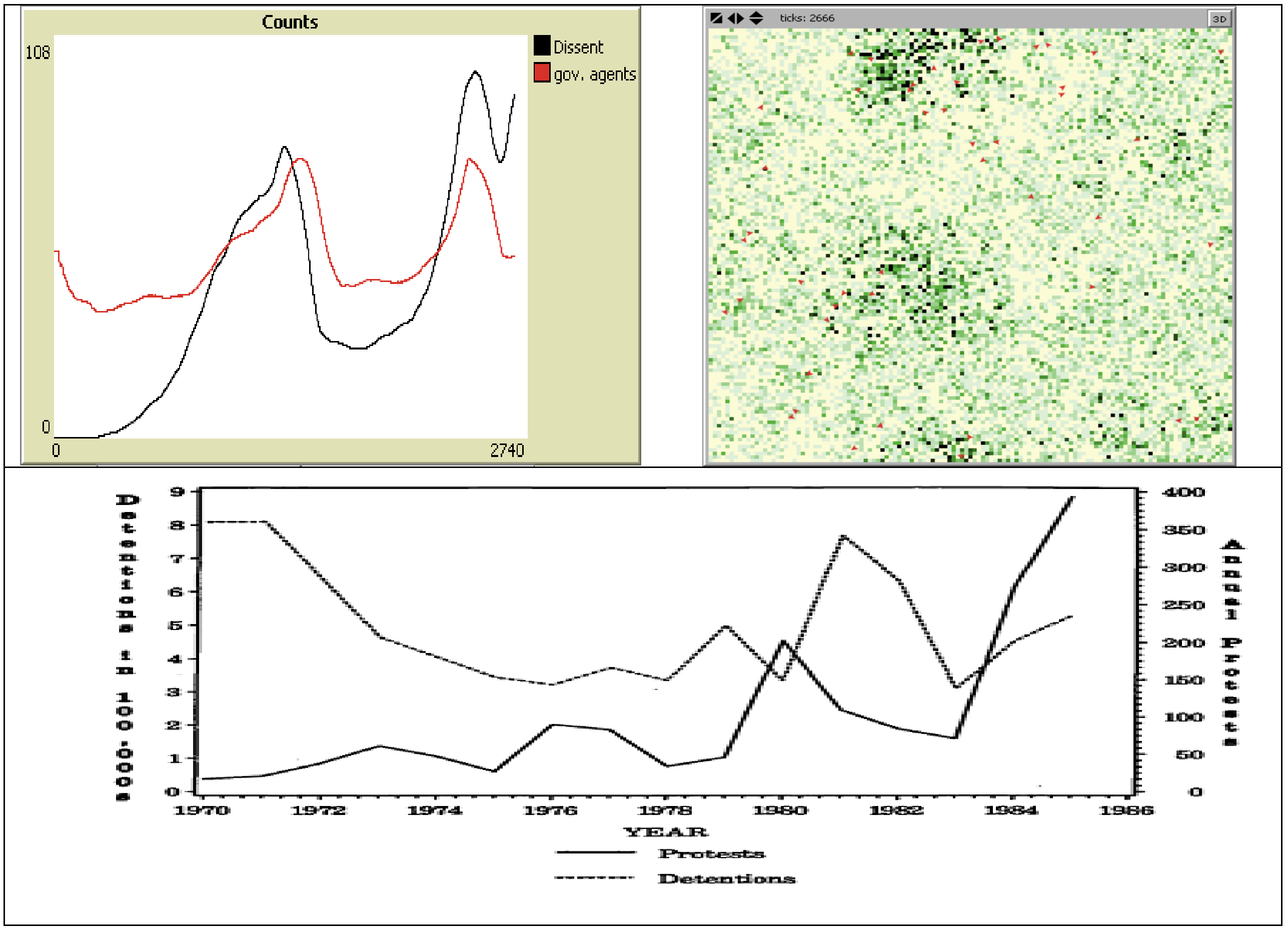

Figure 6 - CAS model mapping for political dissenters and government agents (top left); graphic representation of dissenters and government agents (top right); detentions and protest in South Africa, 1970-1986 (bottom).

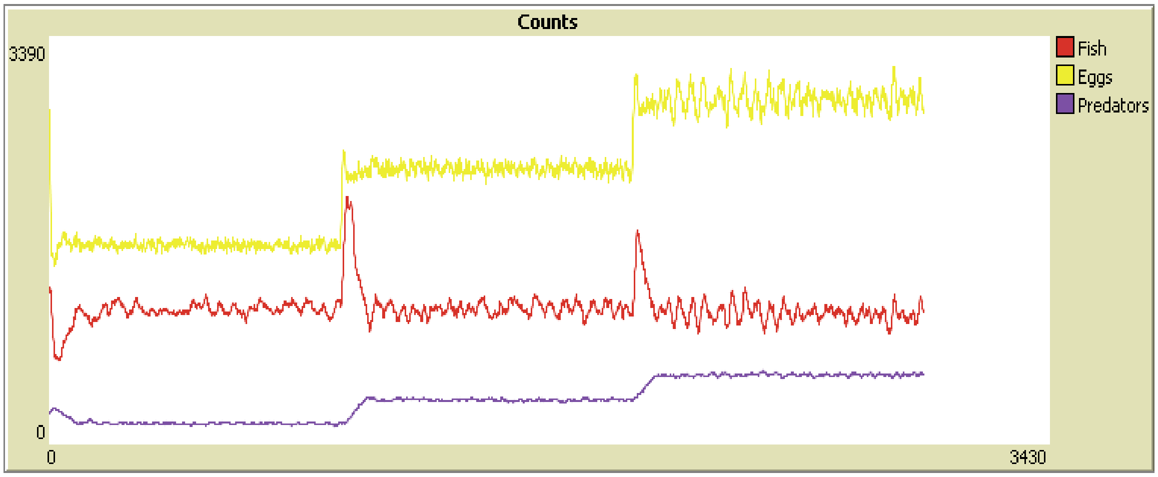

Figure 7 - 3000 simulation time-steps showing population counts for fish (red), fish-eggs (yellow) and predators (purple), at 20%, 30% and 40% food levels (1000 steps per level).

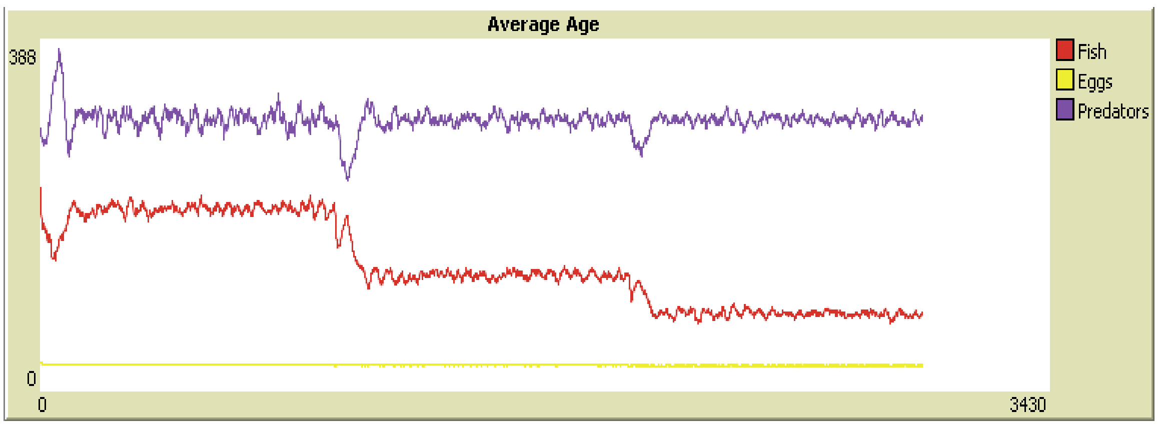

Figure 8 - 3000 simulation time-steps showing mean population age at 20%, 30% and 40% food levels (1000 steps per level).

Table 1 - Key parameters and attributes for each model.

	1st iteration	2nd iteration	cancer	dissent	marine
number of agent-types	2	2	2	2	3-4
self-regulating?	no	yes	yes	yes	yes
A-agent neighbors	adjacent only	adjustable	adjacent	nearby, random	n/a
B-agent reproduction basis	none	based on positive function of successes	based on positive function of successes	based on positive function of successes/ limited by resources	based on consumption levels
B-agent lifetime	infinite	limited	limited	limited, possibly resource related	limited or predation
min. # of B-agents?	initialized	adjustable	yes, adjustable	yes, adjustable	no