PLEASE CITE AS FOLLOWS: Cioffi-Revilla, Claudio, Jong, Kenneth De, & Bassett, Jeffrey. 2012. Evolutionary Computation and Agent-based Modeling: Biologically-inspired Approaches for Understanding Complex Social Systems. Computational and Mathematical Organizational Theory, Vol. 18, no. 3, pp. 356–373. DOI 10.1007/s10588-012-9129-7.

Evolutionary Computation and Agent-based Modeling: Biologically-inspired Approaches for Understanding Complex Social Systems

Claudio Cioffi-Revilla, Kenneth De Jong, and Jeffrey K. Bassett*
Center for Social Complexity, Krasnow Institute for Advanced Study
George Mason University, Fairfax, Virginia 22030 USA
E-mail: ccioffi@gmu.edu

June 18, 2012

Abstract

Computational social science in general, and social agent-based modeling (ABM) simulation in particular, are challenged by modeling and analyzing complex adaptive social systems with emergent properties that are hard to understand in terms of components, even when the organization of component agents is know. Evolutionary computation (EC) is a mature field that provides a bio-inspired approach and a suite of techniques that are applicable to and provide new insights on complex adaptive social systems. This paper demonstrates a combined EC-ABM approach illustrated through the RebeLand model of a simple but complete polity system. Results highlight tax rates and frequency of public issue that stress society as significant features in phase transitions between stable and unstable governance regimes. These initial results suggest further applications of EC to ABM in terms of multi-population models with heterogeneous agents, multi-objective optimization, dynamic environments, and evolving executable objects for modeling social change.

*Funding for this study was provided by the Center for Social Complexity at George Mason University and by grant no. N00014-08-1-0378 from the Office of Naval Research. The authors are solely responsible for any errors in this study. Special thanks to Mark Rouleau for initial development of the MASON RebeLand model with the first author, to Sean Luke for support with the MASON system, and to members of the Mason-HRAF Joint Project on Eastern Africa for comments and discussions.

1 Introduction

Evolutionary Computation (EC) and agent-based modeling (ABM) are two important methodologies that hold significant promise for solving challenging problems in computer science and social science. This paper focuses on the biologically-inspired approach of EC, combined with the complex adaptive systems approach of ABM, to explore the potential of a hybrid framework that combines the two. The hybrid framework is intended to solve a class of social science problems that is relatively common: Understanding emergent dynamics in social systems composed of much simpler, heterogenous, autonomous agents interacting within a broader and changing ecology.

In the first section we describe the class of problems addressed by our bio-inspired EC-ABM approach and related earlier literature. In the second section we describe the EC-ABM hybrid framework, with emphasis on variation and selection mechanisms, and discuss its interpretation within computational social science. The third section illustrates the application of EC-ABM for understanding the structure and dynamics of a relatively simple polity system—the RebeLand model in MASON. In the fourth section we discuss our results, in terms of new domain-specific insights derived from the application of EC-ABM and other areas of possible application. The fifth section presents a summary.

1.1 Motivation

Two important methodologies, evolutionary computation (EC) and agent-based modeling (ABM), have independently reached a stage of maturity that is reflected in their wide-ranging areas of application. EC is based on biological ideas drawn from Darwinian evolutionary theory—i.e., adaptation through reproductive variation and natural selection (De Jong, 2006, 2009). These computational models of evolution have been successfully applied to a wide range of problems in engineering and science. ABM is a computational modeling technique that continues to be successfully used for bottom-up construction of models of complex systems (Gilbert, 2008; Terano et al., 2007; Takadama, 2010).

We are exploring the potential of a hybrid modeling framework that combines EC and ABM in an elegant and powerful way so as to capture both the "simplexity" of agent-based modeling (Kluger, 2008) and the robustness and adaptability of EC methods. This bio-inspired framework allows us to model the dynamic adaptations of a complex system as it evolves over time, and to discover the significance of critical factors of social stability and change. One or more of such factors may turn out to be critical, but none of them is hard-wired as a determinant of stability. In this paper we describe our hybrid modeling framework and illustrate its usefulness in understanding complex social dynamics by applying it to a relatively simple ABM of a polity, called RebeLand (Cioffi and Rouleau, 2010).

A polity is a complex adaptive social system consisting of a population, a system of government, and a policy process where optimization problems are common and complex ("wicked problems," as these are called in policy analysis and public administration; Ritter and Webber, 1973). The overall homeostatic purpose of a polity is to formulate and implement governmental policies for managing public issues. The RebeLand model represents a minimally realistic polity where some forms of severe stress can give rise to governmental instability problems, which some times lead to rebellious insurgency and state failure. Highly consequential governance problems arise in connection to optimization, such as deciding on tax rates, resource capacity reserves, policy planning, timing, implementation, and monitoring of policy effects. While contemporary social science provides a conceptual and theoretically grounded understanding of how a polity system operates, a major enduring challenge has been to develop and analyze formal and empirically grounded models of sufficient realism. A major specific challenge is exploration of the operating conditions that separate sustainable and stable polities from fragile or unstable and failing polities. EC offers a potentially valuable approach to analyzing these issues in conjunction with an appropriate ABM of a polity.

1.2 Prior literature

The combination of EC and ABM in social science research is still largely unexplored, especially compared to its widespread use in other scientific fields. Chattoe (1998), and Gilbert and Troitzsch (2005:ch. 10, pp. 230–255), provide overviews of EC and some early applications in computational social science and social simulation research. Examples of current areas include game-theoretic applications in social science, especially iterated PD-games (Axelrod 1987; Lomborg 1996); theoretical kinship effects in populations of neural networks (Parisi et al., 1995); budgetary evolution based on genetic programming (Chattoe and Gilbert, 1997); and the evolution of cultural algorithms that combine Darwinian and Lamarckian properties (Reynolds, 2008; Reynolds et al., 2002, 2008). A recent edited volume by Rennard (2006) contains numerous applications in economics, but applications of EC to other social sciences (anthropology, political science, sociology, psychology) have been relatively fewer. In particular, the challenge of realistically modeling and systematically analyzing a sufficiently complete polity has been an enduring challenge.

Our hybrid framework builds on earlier applications, specifically in the analysis of spatial ABMs, where human groups and organizations are geographically situated in a dynamic natural environment. Interestingly, while most earlier applications of EC have been non-spatial (e.g., organizational or game-based ABMs), evolutionary cultural algorithms have been developed to support socio-spatial ABMs (e.g., Reynolds' research on archaic states and origins of socio-political complexity in ancient Mesoamerica). However, whereas the cultural algorithm is part of a polity simulation system, the evolutional algorithm (EA) demonstrated in this paper is an analytical procedure.

2 An EC-ABM Hybrid Framework

EC and ABM are each based on a set of ideas (concepts, principles, modeling implementations, toolkits) from separate but increasingly related scientific communities. In this section we describe the general bio-inspired ideas behind EC, some specific areas of application that hold promise for social ABM, and the procedures followed for the illustrative results reported in the next section.

2.1 Bio-inspired Evolutionary Computation (EC)

The term Evolutionary Computation refers to a field and a collection of techniques (evolutionary algorithms) that use biologically-inspired models of Darwinian evolution to solve difficult problems in engineering and science. In these fields one typically has a model of a complex artifact or natural process that one is iteratively refining and improving over time. The ability to semi-automate this process of refinement and improvement is an important step for models of any size and complexity, and particularly true for models of complex social systems. The notion of iterative improvement over time maps nicely into Darwinian notions of adaptation through variation and selection.

This, of course, assumes that there exists a reasonable model of the object or phenomena that can be used as the starting point for the iterative refinement process. One of the virtues of the agent-based modeling approach is that the complexity emerges as a result of the dynamic interaction of entities during a run of the model, rather than requiring that complexity be explicitly captured in the model structure. This feature has been shown to be particularly valuable in modeling complex social systems (Axelrod, 1997; Epstein, 2006; Gilbert, 2008; Cioffi, 2010). So, the focus in this paper is to show how the combination of these two methodologies (evolutionary computation and agent-based modeling) can be exploited to provide a powerful framework for modeling complex adaptive social systems.

At a high level of abstraction this hybridization works as follows. An evolutionary algorithm maintains a population of "individuals" that represent potential model variations. Each model variation is run in order to assess the quality of that particular model with respect to a given modeling goal. Once each individual in the current population has been assigned a "fitness" based on the quality of the model it represents, a selection mechanism chooses the more fit individuals to produce "offspring" that represent variations on the parents to be evaluated as part of the next generation. Less fit individuals (model variations) die off, so the average fitness of the population increases over time.

In order to apply these notions to iterative model refinement we need to specify what parts of the model are changing and what the selection pressure is. Answers to these questions are described in the following subsections.

2.1.1 Model variation

Model variation includes parametric, structural, and behavioral variation. The first, and most prevalent, method of specifying model variation is via model parameters whose values are left unspecified beyond a range of "legal values." Manipulating parameter values to achieve modeling goals maps nicely onto a Mendelian notion of characterizing the important properties of an object in terms of a linear string of genes (parameters) that are inherited from parent to offspring with a modest dose of variation. Evolutionary algorithms have been shown to be particularly good at solving parameter optimization problems in which the parameters interact in highly non-linear ways (Hansen, et.al. 2003). This makes them particularly well-suited for tuning the parameters of complex, adaptive social systems.

Automating model refinement can also occur at a more fundamental level by allowing variation to occur with respect to structural elements of models. For example, variations in signal pathways in electronic circuits, network layout in critical infrastructure, or communication channels in social systems, all have a significant impact on overall model behavior. Evolutionary algorithms can and are used for improving model performance based on this kind of structural variation.

Perhaps the most fundamental level of agent-based modeling is the characterization of the behavior of agents that make up a system. Simple changes in these behaviors can have profound impact on the emergent behavior of a model as it plays out over time. Often, these agent behaviors are specified as a set of rules (Holland, 1986), as a decision tree, as a neural network (Nolfi and Floreano, 2000), etc. Evolutionary algorithms are also used for improving model performance based on this kind of behavioral variation.

Decisions regarding the most appropriate kind of variation to use vary from model to model and from application to application. What is clear is that, as we move from parameter variation, to structural variation, to behavioral variation, we are significantly increasing the complexity of the search space being explored by an evolutionary algorithm; with corresponding increases in both the risk and the reward associated with a successful outcome.

2.1.2 Model selection

Without an appropriate notion of fitness-biased selection, an evolutionary algorithm would be reduced to randomly searching a space of model variations. In order to achieve the desired model refinement goals, an evolutionary algorithm must receive quality-of-model feedback that it can use internally for selection purposes. Model refinement goals can and do differ. Three frequently used goals are briefly described here.

The first, and most obvious type of goal is to improve the fit of a model to data. The quality of a model is based on its ability to match both training (seen) data and testing (unseen) data. Evolutionary algorithms have been shown to be particularly good at solving model fitting problems in which the parameters interact in highly non-linear ways. As noted earlier, this makes them particularly well-suited for tuning the parameters of complex social systems.

Another goal of model refinement is to explore model variations with respect to important model behaviors. For example, we might want to know what model variations lead to instability in airplane design or governmental regimes. In this case selection is based on measures of instability, with higher levels of instability interpreted as "more fit" from a Darwinian perspective. Evolutionary algorithms have been shown to be quite efficient and effective at locating regions in search space that correspond to high fitness (Schultz, et.al. 1993).

A final example involves the exploration of model variations in search of new and unexpected behaviors and/or relationships. This requires that the modeler quantify properties such as "novelty" and/or "interestingness" in ways that provide useful feedback to an evolutionary algorithm. Successful examples of this are found in a variety of application areas, including data mining and technology innovation.

2.1.3 Evolutionary Algorithm (EA) variation

The simplest evolutionary algorithms (EAs) involve a single population evolving over time in response to externally generated fitness feedback. These simple EAs have been extended in a variety of ways to capture various application-specific properties. In this subsection we briefly describe a few that are particularly appropriate for modeling adaptive social systems.

The first obvious extension is that the fitness landscape in complex social systems is seldom a statically defined entity. Rather, it is dynamically changing over time as a model plays out. This is handled in EAs in several ways. First, like ecosystems, the best weapon against change is diversity. A variety of diversity-preserving mechanisms are part of a standard EA toolkit including niching, speciation, and co-evolution.

An important aspect of social system models is the notion of local interaction-neighborhoods. These notions of spatial proximity are handled in EAs by both fine-grained "cellular diffusion" models (Sarma, 1998) and coarse-grained "island" models (Cantù-Paz, 2001).

Finally, many applications involve exploring the tension between multiple conflicting goals (e.g., minimizing a tax rate while maximizing governance capacity). Some of the best multi-objective optimization algorithms today are evolutionary algorithms (Deb, 2001).

2.2 An Example ABM: The RebeLand Model

In order to illustrate our EC-ABM methodology we chose an ABM model that was both simple to understand, but exhibited a surprising amount of subtle complexity. MASON RebeLand is an ABM of a basic polity with the following mid-range fidelity features between a toy model and a high-resolution empirical model:

- Population of heterogenous agents, including households, government agents, insurgents, and security forces;
- Natural environment with explicit albeit simple geographical features, including mountains, valleys, and coastal regions, as well as weather;
- Natural resources, such as gold, diamonds, oil, or gas, linked by a basic transportation network connecting resource locations with population centers;
- Cities and other urban centers, including a capital and smaller settlements;
- Governance structure comprised of national and provincial systems for decison-making and public administration; and

- Public issues that affect the general population, such as inflation, drought, and other endogenous or exogenous stresses.

The main computational loop in RebeLand is intended to simulate recurring cycles of governance in a basic political system. Public issues arise with some frequency, intensity, and duration, causing stress on society. Government formulates and implements policies that provide public goods and mitigate or solve issues, using resources obtained from society (usually through taxation). Members of society in a stable polity earn sufficient resources to maintain well-being, support government, and be satisfied. However, when issues generate sufficiently high stress, individual members of society can rebel and join an insurgency. Depending on governmental capacity and the complex interaction of societal dynamics, an insurgency can be controlled by government or lead to instability and, in extreme cases, even state failure (Cioffi and Rouleau, 2010).

As a polity, RebeLand is based on the basic standard model of a political system (Almond et al., 2006; Cioffi, 2005). As an ABM, the purpose of RebelLand is to investigate the complex dynamics of political stability and state failure potential, given a variety of public issues that produce stress on society, and government that responds with policies based on available capacity. RebeLand is part of a broader project on computational modeling and comparative analysis of complex socio-natural systems in developing regions (Eastern Africa, specifically), as detailed elsewhere (Cioffi and Rouleau, 2010).

2.3 An Illustration of our EC-ABM Methodology

As indicated in the previous section the RebeLand model captures the basic principles of a polity rather than a model a particular political system. As a consequence, fitting the model to existing data was not a primary issue. Rather, we were interested in exploring variations of the model associated with governmental stability, instability (i.e., rebel insurgencies), as well as models that captured "tipping points" between the two. This space of model variation was defined by a set of six parameters:

1. *Government corruption*
2. *City tax rate*
3. *State tax rate*
4. *Issue onset interval*, measured by the number of time steps between the occurrence of a new public policy issue
5. *Issue onset duration*
6. *Terrorist success rate*

In order to explore the government stability issues discussed above, the model was implemented to collect the following data over the period of a simulation run:

1. *Government support*, or level of support that the population expresses for government

2. *Rebel support*, or level of support that the population expresses for insurgents

3. *Government durability*, or average maximal government duration in power, measured as the total number of time steps that the government remains in power.

4. *Government volatility*, or average number of power turnover events, measured as the number of times power changed hands between government and rebels.

With these dependent variables defined we explored the dependent variables (i.e., the model parameter space) in two ways. First, we performed a standard, but coarse, "parameter sweep" analysis to get a rough sense of how government stability was affected by the model parameters, giving us the opportunity to fix any obvious bugs in the model as well as do some fine tuning of the parameter value ranges.

Having completed this "sanity check" of the ABM model (technically, a verification test), we then focused on the EC component. This involved making a number of design decisions that are described here along with some insight into why they were chosen.

The goal of these experiments was to locate "interesting" parts of the parameter space. The key to achieving this with an EA is to define fitness functions that capture our notions of what we find interesting. We chose the following three functions:

$$f_1 = 1 + govSupport - rebSupport \quad (1)$$
$$f_2 = 1 + rebSupport - govSupport \quad (2)$$
$$f_3 = 1 - |govSupport - rebSupport| \quad (3)$$

where *govSupport* and *rebSupport* denote the average support for the government and the rebels, respectively, over a single simulation run. Function f_1 will return high values for environments that are the most accommodating for a stable government, while f_2 indicates environments that are the most hostile. Function f_3 was designed to identify environments that are somewhere between the two extremes, as a phase transition where support for the government and the rebels is about equal. The idea here was to identify polity features that represent a tipping-point between stability and instability, as a bifurcation set.

Within the EA we chose to represent model variations in standard parameter optimization form, namely, as a linear genome consisting of six real-valued real genes representing the model parameters. Most of the model parameters were bounded in the range [0,1], with the exception of *Issue Onset Interval* and *Maximum Issue Duration*. We chose to represent these normalized in the range [0,1] also so that the mutation operator could use the same mutation rate on each gene. The values of these two genes are then multiplied by 50 and 200 respectively when the genome is decoded.

The initial parameter sweep studies showed that the three fitness landscapes were not overly complicated. For example, they did not seem to contain a large number of sub-optimal peaks. Therefore we felt that a large EA population was not necessary. On the other hand the fitness evaluations were somewhat noisy in the sense that simulations run with the same parameter settings, but different random number seeds, produce different fitness values. This is a well-known effect when using stochastic simulations for fitness evaluations. As a result we thought an overlapping population scheme was in order. This would allow good parents to survive into the next generation, but would require that they continue to receive high fitness evaluations in future generations. We used a parent population size of $\mu = 5$, and an offspring population size of $\lambda = 10$. Once the children have been generated, these two populations are combined and they then compete via truncation survival selection to see which individuals survive into the next parent population. This is a well-known EA configuration referred to in the EC literature as a $(\mu + \lambda)$ Evolution Strategy.

Children are generated using a series of reproduction operators. The first is gaussian mutation, which adds a perturbation to each gene drawn from a gaussian distribution. Sensitivity studies showed a standard deviate of $\sigma = 0.03$ was appropriate for our problem. We also used a two-point crossover operator that mixes genetic material from two parents using two randomly chosen cut-points on their genomes. This is done to promote high variation in the population, especially early in the run.

With these EA design decisions in place, we are now in position to explore the model's parameter space using the adaptive search capabilities of EAs to find regions in parameter space that produce "interesting" model behavior as defined by the given fitness functions.

3 Some Illustrative Results

A complete description of the experimental results produced by this methodology is beyond the scope of this paper. Rather, in this section we illustrate the hybrid EC-ABM approach by reporting selective results obtained from (1) the preparatory parameter sweep of the MASON RebeLand ABM, and (2) the evolutionary algorithm, as described in the previous section. Although RebeLand is a relatively simple polity model, especially by comparison with other models in comparative political science, its resulting behavior shows both intuitive and counter-intuitive features.

3.1 Parameter sweep results

Because the results of the simulation were somewhat noisy, eight separate parameter sweeps were performed. In each sweep, the RebeLand ABM randomly generated a single physical environment, and used it over all combinations of the dependent variables. The plots below show the average surfaces of those eight sweeps. Accordingly, each point on the plot represents the average of the sweeps for (x, y) combinations variable settings.

Given the large number of results produced by the parameter sweeps and resulting plots, we highlight only the following main results grouped by dependent variables.

Government support. Figure 1 shows the parameter surface for *government support* (z-axis) as a function of six pair-wise combinations of *city tax rate*, *corruption*, *issue onset interval*, and *state tax rate*. The six plots in Figure 1 are reported for illustrative purposes. The plots show that support for government is most affected by *city tax rate*. Other parameters, such as *corruption* and *issue onset interval*, also have some effect. We also swept for other parameter combinations, such as *issue onset duration* and *terrorist success rate*, among others.

Rebel support. Support for rebels was also mostly affected by *city tax rate* and *issue onset interval*, in both cases as positive decreasing functions. In this case corruption had an even weaker effect. However, in this set of results the confidence surfaces show greater variation.

Government durability. The average duration of government in power was mostly affected by *city tax rate*, while *issue onset interval* had a weaker effect. Confidence surfaces also showed significant variation across plots.

Government turnover. The average number of times that the government changed hands was low (0 or 1) and mostly unaffected by any of the variables, including even *city tax rate*. On occasion, the confidence surface associated with each average surface also showed significant variation across plots, especially in the appearance of occasional sharp spikes, indicative of high variance across simulation runs for some specific parameter combinations.

Overall, *city tax* had the greatest effect on the state of the polity, whereas *corruption* and the federal *state tax rate* showed a weaker effect. In some cases, power can change hands many times (high volatility), although average turnover was low. Figure 2 shows one run that is a good example of this, for *city tax rate* $= 0.05$ and *issue onset interval* $= 1$.

Finally, we also found that independent variables such as *issued duration* and *terrorist success rate* did not have significant effects on government support or durability. We omit the surface plots for these results in the interest of space.

3.2 Evolutionary algorithm results

Ten runs of the EA were performed for each fitness function, resulting in thirty runs overall. Each run was allowed to run for 50 generations, and at the end of the run, the best individual from the last generation was examined. The independent variable settings (gene values) of those individuals were noted.

Figure 3 shows illustrative results from the EA analysis, focusing on the performance of the fitness functions for pairwise combinations of the following parameters: *issue onset duration*, *issue onset interval*, *city tax rate*, *state tax rate*, *terrorist success rate*, and *corruption rate*. The four plots show the same 30 runs, but plotted using different pairs of dependent variables in order to get a view on what is essentially a six-dimensional space.

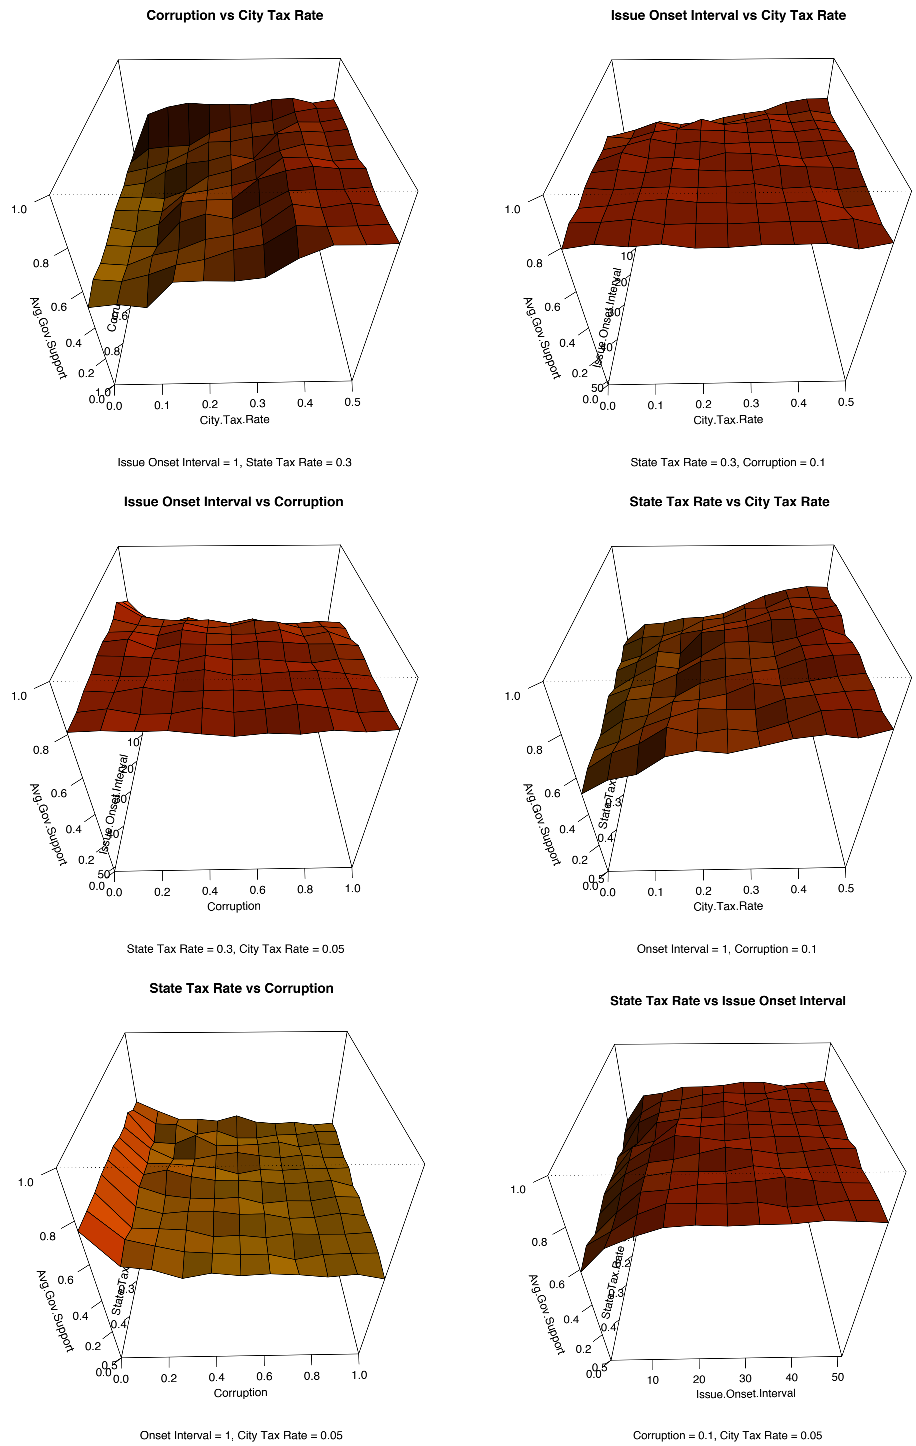

Figure 1: Citizen support for government as a function of government corruption, city tax rate, issue onset interval, and state tax rate. *Source*: Prepared by the authors.

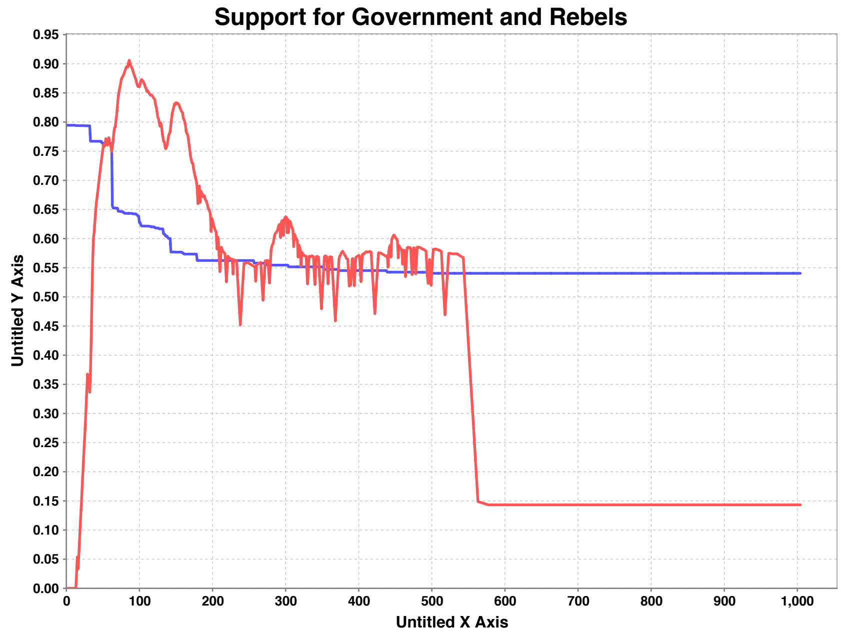

Figure 2: A plot of support for government (blue) and rebels (red) during one sample run of of RebeLand. *Source*: Prepared by the author.

Looking at these results we can get sense for which variables have the greatest impact on the simulation results. The purple dots are associated with fitness function f_1, and therefore show the conditions that are likely to allow successful governance. The red triangles represent function f_2, and indicate the conditions that will be the most difficult to govern under. Finally, the green squares indicate the runs that used function f_3, and approximate a mid-point ("tipping point" or phase transition) between f_1 and f_2.

The upper-left plot (in the parameter sub-space of *issue onset duration* and *issue onset interval*) shows a clear clustering of fitness functions, with the following results:

1. low values (< 15) of *issue onset interval* and mostly below-median values of *issue onset duration* are associated with f_1 (governmental stability);
2. mainly high values (> 25) of *issue onset interval* and mostly median values of *issue onset duration* are associated with f_2 (governmental instability); and
3. high values of (> 30) of *issue onset interval* and any value of *issue onset duration* are associated with f_3 (tipping points or phase transition).

Our results show also that dependence of political stability on the other variables is less pronounced, but the following patterns were observed for the plots shown in Figure 3:

- Mid-range values of *terrorist success rate* were associated with f_3 (green squares) or tipping points at the phase transition between stable and unstable regimes.

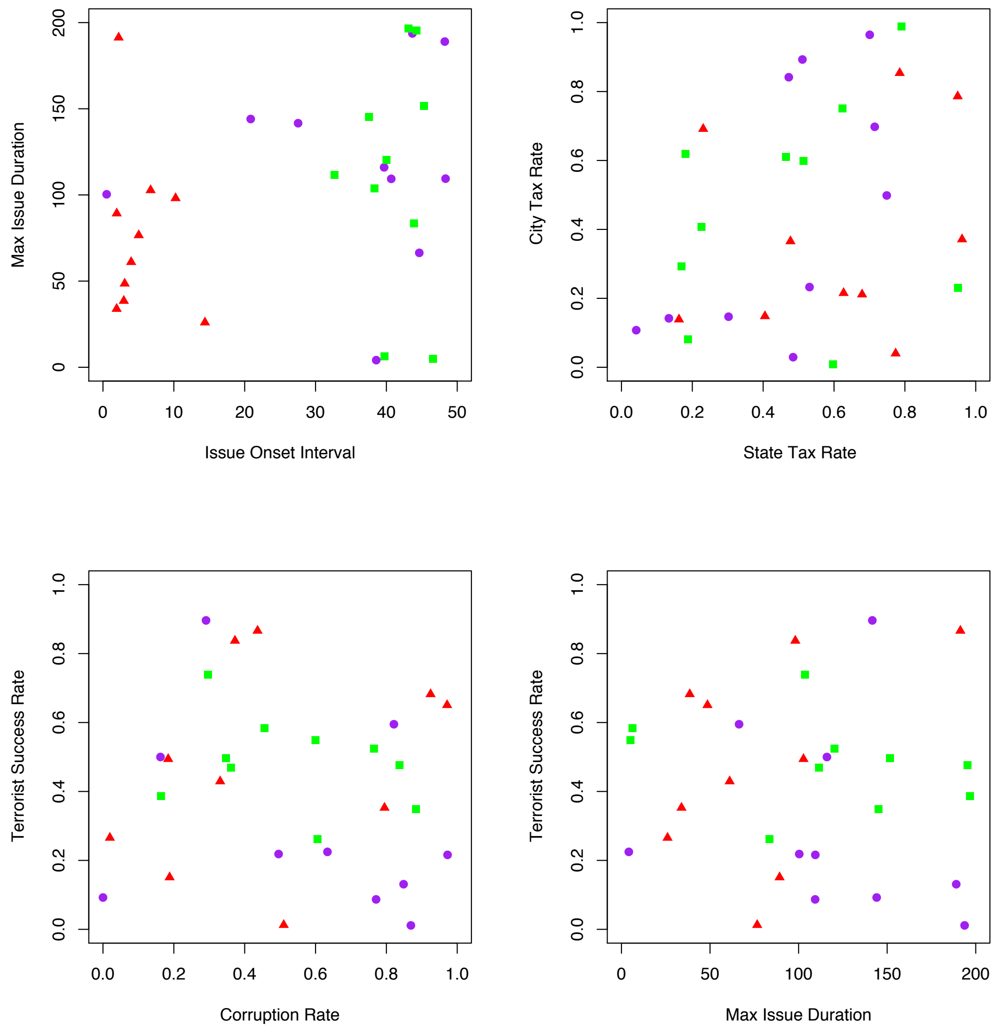

Figure 3: Each point represents the best individual in the last generation of a single run. The different fitness functions are represented as: f_1 (purple dots for political stability), f_2 (red triangles for political instability), and f_3 (green squares for tipping points or phase transitions between stable and unstable regimes). *Source*: Prepared by the authors.

- In stable regimes the two tax rates (city and state) tend to be positively correlated.
- Unstable regimes are associated with low values of issue duration.

Other fitness function plots similar to those in Figure 3 show various patterns consistent with these and are not shown here.

4 Discussion

In this section we discuss our main results from the parameter space exploration and the EA evaluations, followed by brief discussion of broader implications on EC-ABM hybrid applications mentioned earlier.

4.1 Main findings

Understanding the parameter space of any model (simple or complex) is essential for explaining the emergence of aggregate phenomena, understanding dynamical properties, and other important features—to say nothing of formulating viable policies. This is particularly true for complex or relatively complex social systems, or even for relatively simple systems such as RebeLand. Our initial analysis of the parameter space and fitness evaluations has explored only a small albeit meaningful sub-space, given the relative complexity of the model. Unfortunately, even simple polities such as RebeLand are not as simple as they may appear.

Of all the independent variables (parameters) we analyzed, *city tax rate* and *issue onset interval* were found to be the most critical in determining various aspects of political stability, including measures such as the population's support for government, as opposed to support for rebels, and temporal dimensions of political stability such as government longevity and turnover. Further analysis of these result is necessary, because intuition and empirical observation would also suggest other variables as being important. The pronounced effect of *city tax rate* could also be explained by the direct and negative effect it has on citizen's satisfaction, which in turn affects their approval for government.

Issue onset interval followed *city tax rate* as the second most important parameter affecting governmental stability. Interestingly, the effect of other variables, such as federal taxes on cities, corruption in the public administration, success rate of rebels, and the duration of public issues had far less or in some cases negligible effect on government stability.

Overall these findings seem to have high face validity, because a polity stressed by a high frequency of issues and citizen dissatisfaction due to high taxes could not remain stable for long; but low issue frequency could be associated with a stable regime, even when policy-making capacity is low.

4.2 Broader implications

The four EC applications mentioned earlier in the context of our hybrid framework also suggest several broader implications beyond the EC analysis presented in this paper.

4.2.1 Multi-population models

Social ABM is often about diverse or "heterogeneous" populations of agents. In RebeLand such agents are citizens of the general population, rebel insurgents, and government agents that, in turn, include policy-makers and members of various security forces (police, military). These four sub-populations of agents co-evolve, because the size of each group and the state of individuals depends, in part, on the agents in the other groups. For instance, the number of government security agents tends to grow when rebel numbers grow. When rebels grow, the number of non-contentious agents decreases.

A natural application of EC in this area would seem to be the use of co-evolutionary EAs to evolve (discover!) better dynamics that more closely resemble real-world social systems. The use of EC in such a "context of discovery" (paraphrasing C.G. Hempel) has not yet been attempted in social science, but would follow naturally from the application of EA. In addition, EAs could be applied to solve difficult institutional design problems, assuming some degree of near-decomposability (Simon, 1962) for a given collection of subproblems (De Jong, 2009: 55). For example, it would be valuable to discover how federal or national institutions could be better designed or reformed (i.e., re-engineered), given a set of known features for local governance institutions (Ostrom, 2009). This is an open, challenging issue in political science and related areas of social science (e.g., collective action theory).

4.2.2 Multi-objective optimization

The requirement of satisfying a set of objectives is fundamental and ubiquitous in even the simplest of polities. In polity models such as RebeLand the problem is compounded by multiple levels of public administration (local and national), multiple sources of public policy issues (natural, social, artifactual), and other complicating factors. Trying to improve the design of a polity with respect to one objective (e.g., decrease government turnover), let alone trying to optimize it, can cause problems with respect to other objectives (mitigating inflation, controlling the rate of growth in rebellion, etc.). Searching for Pareto-optimal points in the state-space of a polity model such as RebeLand is a computationally challenging problem, because even the parameter space of simple systems remains unexplored (Cioffi, 2005). This is arguably a point that may not be fully appreciate in the domain of contemporary political theory.

4.2.3 Dynamic environments

Traffic control problems, stock market trading, and similar domains are not the only EC applications where dynamic environments are common. Problems of governance are commonly affected by changing environments where the nature, variety, and timing of emerging public issues are constantly changing. For example, whereas the Weibull distribution provides a reasonable description for modeling the time between disruptive social events (riots, coups, civil wars), variation in the shape parameter has significant implications for the underlying dynamics. Other dynamic environments include demographic change, political campaigns, and international relations.

Although dynamic environments are challenging, suitable variations of standard EAs have done quite well in this area (De Jong, 2009: 55). Accordingly, computational social science could build on such experience, rather than "reinvent the wheel," and contribute new ideas that may come from social applications (e.g., Reynolds' "cultural algorithm"). Interestingly, coupled socio-techno-natural systems—notorious for their complex, dynamic properties—are an area of increasing interest in computational social science (Liu et al., 2007; Ostrom, 2009).

4.2.4 Evolving executable objects

In computer science the term "executable object" means an entity or "object" defined by static and dynamic features that can carry out instructions. A social instance of an "evolving" executable object is any social entity that can generate new beliefs, rules, values, roles, institutions, organizations, or other agents. Again, "suitably modified" EAs (De Jong, 2009: 55) have been used to evolve "objects whose execution behavior is quite complex."

The application of evolving executable objects to modeling social systems and processes is another unexplored frontier, and one that is quite consistent with the paradigm of "generative social science" (e.g., Epstein, 2009). Social theory is replete with hypotheses, as well as some empirically-validated theories, concerning the origin and evolution of social entities as executable social objects.

5 Summary

Computational social science in general, and social agent-based modeling (ABM) simulation in particular, are challenged by having to model and analyze complex adaptive social systems with emergent properties. Such systems and their associated processes are difficult to understand in terms of components, even when the organization of most or all of the component agents is know. In recent years evolutionary computation (EC) has become a mature field that provides a bio-inspired approach and a suite of computational techniques that are applicable to and provide new insights on complex adaptive social systems. EC has a rich record of applications in several domains that resemble that of social systems, including biological systems and information networks.

This paper demonstrated a combined EC-ABM approach to the challenge of understanding complex adaptive systems in social science. We illustrated our framework through the RebeLand agent-based model of a simple but rather complete polity system. Our results highlighted several variables as being more important than others, such as tax rates and frequency of public issue that stress society. These variables represent significant features in phase transitions between stable and unstable governance regimes.

These initial results obtained through evolutionary algorithmic analysis also suggested further applications of EC to ABM. Multi-population models with heterogeneous agents, multi-objective optimization, dynamic environments, and evolving executable objects for modeling social change provide additional EC approaches for future research.

The potential contribution of EC to understanding social systems could turn out to be greater than that of game theory, because of its scalable capacity to analyze small and large systems, including national and international systems where the largest numbers of agents are affected by human and social dynamics. While 2-person game theory is viable and fairly complete, scaling up to N-person game theory with the same degree of completeness in explanatory power has proven to be infeasible in spite of earlier hopes (Rapoport, 1970). Moreover, EC stands independent of game theory or even decision theory.

References

Abelson, Robert. 1958. Symbolic psycho-logic. *Behavioral Science* 3:1–13.

Abelson, Robert P. 1973. The Structure of Belief Systems. In *Computer Models of Thought and Language*, edited by R. C. Shanck and K. M. Colby. San Francisco: W. H. Freeman.

Almond, Gabriel A., Jr. G. Bingham Powell, Russell J. Dalton, and Kaare Strom. 2006. *Comparative Politics Today: A World View*. New York: Pearson Longman.

Axelrod, Robert. 1997. Advancing the art of simulation in the social sciences. *Complexity* 3 (2):193–99.

Cantù-Paz, Erick. 2001. Migration policies, selection pressure and parallel evolutionary algorithms. *Journal of Heuristics* 7 (4):311–334.

Chattoe, Edmund. 1998. Just How (Un)realistic are Evolutionary Algorithms as Representations of Social Processes? *Journal of Artificial Societies and Social Simulations* 1 (3). Available online.

Chattoe, Edmund, and Nigel Gilbert. 1997. A Simulation of Adaptation Mechanisms in Budgetary Decision Making. In *Simulating Social Phenomena*, edited by R. Conte, R. Hegelsmann and P. Terna. Berlin: Springer-Verlag.

Cioffi-Revilla, Claudio. 2008. Simplicity and Reality in Computational Modeling of Politics. *Computational and Mathematical Organization Theory* 15 (1):26–46.

Cioffi-Revilla, Claudio. 2010. On the Methodology of Complex Social Simulations. *Journal of Artificial Societies and Social Simulations* 13 (1):7. Available online.

Cioffi-Revilla, Claudio, and Mark Rouleau. 2010. MASON RebeLand: An Agent-Based Model of Politics, Environment, and Insurgency. *International Studies Review* 12 (1):31–46.

De Jong, Kenneth. 2006. *Evolutionary Computation*. Cambridge, MA: MIT Press.

De Jong, Kenneth. 2009. Evolutionary computation. *Wiley Interdsiciplinary Revies (WIREs) Computational Statistics* 1 (1):52–56.

Deb, Kalyanmoy. 2001. *Multi-Objective Optimization Using Evolutionary Algorithms*. New York: John Wiley and Sons.

Epstein, Joshua. 2006. *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton, NJ: Princeton University Press.

Gilbert, Nigel. 2008. *Agent-Based Models*. Thousand Oaks, CA: Sage Publishers.

Gilbert, Nigel, and Klaus Troitzsch. 2005. *Simulation for the Social Scientist*. Second edition ed. Buckingham and Philadelphia: Open University Press.

Hansen, Nikolaus, Sibylle D. Müller and Petros Koumoutsakos. 2003. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation* 11 (1):1–18.

Holland, John. 1986. Escaping brittleness: The possibilities of general purpose learning algorithms applied to parallel rule-based systems. In R. Michalski, J. Carbonell and T. Mitchell (Eds.), *Machine Learning II*. San Mateo CA: Morgan Kaufmann Press. pp 593–623.

Kluger, Jeffrey. 2008. *Simplexity: Why Simple Things Become Complex (And How Complex Things Can Be Made Simple)*. New York: Hyperion.

Liu, Jianguo, Thomas Dietz, Stephen R. Carpenter, Marina Alberti, Carl Folke, Emilio Moran, Alice N. Pell, Peter Deadman, Timothy Kratz, Jane Lubchenco, Elinor Ostrom, Zhiyun Ouyang, William Provencher, Charles L. Redman, Stephen H. Schneider, and William W.A Taylor. 2007. Complexity of Coupled Human and Natural Systems. *Science* 317 (5844):1513–1516.

Lomborg, Bjorn. 1996. Nucleus and Shield: The Evolution of Social Structure in the Iterated Prisoner's Dilemma. *American Sociological Review* 61 (2):278–307.

Nolfi, Stefano and Dario Floreano. 2000. *Evolutionary Robotics: The Biology, Intelligence, and Technology.* Cambridge, MA. MIT Press.

Ostrom, Elinor. 2009. A General Framework for Analyzing Sustainability of Socio-Ecological Systems. *Science* 325:419–422.

Parisi, Domenico, F. Cecconi, and A. Cerini. 1995. Kin-directed Altruism and Attachment Behavior in an Evolving Population of Neural Networks. In *Artificial Societies: The Computer Simulation of Social Life*, edited by N. Gilbert and R. Conte. London: UCL Press. Pp. 238–251.

Rapoport, Anatol. 1970. *N-Person Game Theory: Concepts and Applications.* Ann Arbor, Michigan: University of Michigan Press.

Rennard, Jean-Philippe, ed. 2006. *Handbook of Research on Nature Inspired Computing for Economics and Management.* Idea Group Inc.: Hershey, PA, USA, 2006.

Reynolds, Robert G. 2008. Computing with the Social Fabric: The Evolution of Social Intelligence within a Cultural Framework. *IEEE Computational Intelligence* 3 (1):18–30.

Reynolds, Robert G., A. Lazar, and S. Kim. 2002. The Agent-Based Simulation of the Evolution of Archaic States. In *Proceedings of the Agent 2002 Conference on Social Agents: Ecology, Exchange, and Evolution*, edited by C. Macal and D. Sallach. Chicago: University of Chicago and Argonne National Laboratory.

Reynolds, Robert G., Mostafa Ali, and Thaer Jayyoussi. 2008. Mining the Social Fabric of Archaic Urban Centers wth Cultural Algorithms. *Computer* 41 (1):64–72.

Ritter, Horst W. J., and Melvin M. Webber. 1973. Dilemmas in a general theory of planning. *Policy Sciences* 4:161–167.

Sarma, Jayshree. 1998. An Analysis of Decentralized and Spatially Distributed Genetic Algorithms. Ph. D. Thesis, George Mason University.

Schultz, Alan C., John J. Grefenstette, and Kenneth A. De Jong. 1993. Test and evaluation by genetic algorithms. *Intelligent Systems* 8 (4):9–14.

Takadama, Keiki, Claudio Cioffi-Revilla, and Guillaume Deffaunt, Eds. 2010. *Simulating Interacting Agents and Social Phenomena: The Second World Congress in Social Simulation*. Tokyo, New York, and Heidelberg: Springer Verlag.

Terano, Takao, and David Sallach, Eds. 2007. *Advancing Social Simulation: The First World Congress in Social Simulation*. Tokyo, New York, and Heidelberg: Springer Verlag.