

[J. Gary Polhill, Edoardo Pignotti, Nicholas M. Gotts, Pete Edwards and Alun Preece \(2007\)](#)

A Semantic Grid Service for Experimentation with an Agent-Based Model of Land-Use Change

Journal of Artificial Societies and Social Simulation vol. 10, no. 2, 2
<<http://jasss.soc.surrey.ac.uk/10/2/2.html>>

For information about citing this article, click [here](#)

Received: 01-Sep-2006 Accepted: 20-Nov-2006 Published: 31-Mar-2007



Abstract

Agent-based models, perhaps more than other models, feature large numbers of parameters and potentially generate vast quantities of results data. This paper shows through the FEARLUS-G project (an ESRC e-Social Science Initiative Pilot Demonstrator Project) how deploying an agent-based model on the Semantic Grid facilitates international collaboration on investigations using such a model, and contributes to establishing rigorous working practices with agent-based models as part of good science in social simulation. The experimental workflow is described explicitly using an ontology, and a Semantic Grid service with a web interface implements the workflow. Users are able to compare their parameter settings and results, and relate their work with the model to wider scientific debate.

Keywords:

Agent-Based Social Simulation, Experiments, Ontologies, Replication, Semantic Grid

Introduction

1.1

Collaborations between large groups of scientists, who may be dispersed internationally, are increasingly seen as essential to enhance the scientific process. While research has always involved collaboration between individual scientists, there is now even greater necessity for tools to support sharing of knowledge, resources, results and observations. This paper describes the FEARLUS-G project, involving social scientists at the Macaulay Institute in Aberdeen investigating land-use change and computer scientists at the University of Aberdeen. Central to this collaboration is FEARLUS ([Polhill et al. 2001](#); [Gotts et al. 2003](#)), an agent-based model of land-use change which contains objects that represent human decision-makers in the real-world (land managers) and takes into account attributes such as yield from land parcels. Parameters to the modelling environment allow a variety of land-use strategies and their outcomes to be explored.

1.2

The reported outputs of agent-based social simulations are often given on the basis of only a single run of the model, or too few runs ([Izquierdo & Galan 2004](#)). This can be unsatisfactory for two reasons: Firstly, there is the issue of sensitivity of the model to the particular parameter settings that happen to have been used for that single run. Secondly, the run may not itself be typical of the output one might get for those specific parameter settings. The former case is a particular issue when parameter settings are either derived from observed data with non-zero error bounds, or are otherwise non-observable parameter settings, such as coefficients, or values of parameters that describe an abstract concept. The latter is an issue when there are any particularities associated with the run, typically a particular random seed, though there may potentially be other particularities due to, say, the order in which data is presented to the model, or platform-specific issues such as a non-IEEE 754 compliant floating-point environment ([Polhill et al. 2005](#); [Polhill et al. 2006](#)).

1.3

The FEARLUS experimental methodology, described by Polhill et al. ([2001](#)) and Gotts et al. ([2003](#)), is designed to address some of these issues by facilitating multiple runs of the model in batch mode using different random seeds, and testing hypotheses about the outputs of these runs using statistical tests. In this way, we are able to confirm to a measurable level of statistical significance that the model typically behaves in the hypothesised manner. The experiments in these papers have involved tens if not hundreds of thousands of runs of FEARLUS models, which have been managed through the use of Perl scripts (which can only use one machine at a time), requiring access to local machines to run them. This poses a challenge to collaborating with colleagues from outside the local institution.

1.4

Ontologies (not to be confused with 'Ontology' in philosophy) are defined by Gruber ([1993](#)) as "formal explicit specification[s] of a shared conceptualisation", and can be used to capture the meaning of metadata terms and their interrelationships ([Fensel 2003](#)). The main benefit of using ontologies is that they facilitate access to heterogeneous and distributed information sources by defining a machine-processable semantics for those information sources. Important technologies include RDF Schema^[1] (RDFS) — a vocabulary for describing properties and classes of RDF^[2] resources, with a semantics for generalisation hierarchies; and OWL^[3] (Web Ontology Language) — which adds more vocabulary for describing properties and classes, e.g. relations between classes, cardinality, etc. One application of ontologies is in describing metadata about simulation models, including agent-based social simulations: i.e. the classes of object or agent that appear in the model and how they interrelate ([Polhill & Gotts 2006](#); [Gotts & Polhill 2006](#)), though that is not the application discussed here.

1.5

The Semantic Grid ([de Roure et al. 2001](#)) annotates the Grid ([Foster et al. 2001](#)) with metadata describing the resources it makes available, just as the Semantic Web ([Berners-Lee et al. 2001](#)) does to the Web. The Grid, in contrast to the Web, uses a set of standards and protocols to provide access to heterogeneous resources such as computing, application, data, storage or network resources. It is largely used for complex computational problems, such as exploring protein fold space ([Lee et al. 2004](#)), running climate change simulations ([Allen 1999](#); [Allen et al. 2000](#)), and processing particle physics ([Cortese et al. 2004](#)) and remote sensing ([Yang et al. 2004](#)) data. The Semantic Grid, through providing ontological support to the Grid, creates a searchable, reusable resource that is understandable by a wider community.

1.6

Beyond collaborations, good scientific practice requires that results and experiments are repeatable and verifiable, as attested by the proposed use of the Forum in JASSS for replication ([Gilbert 2004](#)). Typically, however, the kind of detail involved in showing results data, parameter settings used, and descriptions of the model and its source code means such information, essential to repeatability and verifiability, is not suitable material for publication in a scientific journal. In the case of the FEARLUS experimental methodology, the ontology and its availability on the Semantic Grid would promote rigour through facilitating (the first) three of four key levels of repeatability that are relevant to promoting good science in agent-based modelling:

1. Provision of access to the results data from the model runs on which authors base their conclusions. This allows checking that the authors' claims (hypotheses) based on these results are justifiable. Such claims may be based on many thousands of runs of the model, creating results data that simply cannot be practically reproduced in published scientific material.
2. Provision of the capability to re-run the experiments that generated the results using the same parameter settings and software, to check that the results themselves are typical of the model's behaviour and not based on carefully selected runs that support the authors' claims. The contribution of the Grid is key to this provision. Whilst many researchers supply the source code or executable versions of their model for others to download, platform compatibility issues may prevent replication of results, as may a lack of access to parameter settings. The former could apply even in the case of models written in Java where, for example, a model required faster CPUs or more RAM than might be available to a typical researcher, or where appropriate Java Runtime Environments and versions of Java library were not available on the researcher's platform for some reason. The latter is again an issue where many different parameter settings are used, and the information cannot practically be reprinted in journals.
3. The capability to re-run the experiments using different parameter settings. This facility enables researchers to check that the authors' results are not based on carefully selected parameter settings. This is typically the domain of sensitivity analysis.
4. Though not provided by the work described here, the ability of a researcher to modify the source code or reimplement the model from scratch allows them to check for what might be called 'algorithmic sensitivity' of the model: e.g. would different data structures used to implement the concepts in the model result in different behaviour? Other examples are changing scheduling arrangements, or using different algorithms to implement the agents' cognitive behaviour (e.g. Neural Networks ([McClelland et al. 1986](#)), Beliefs-Desires-Intentions ([Rao & Georgeff 1991](#)), or Case Based Reasoning ([Aamodt & Plaza 1994](#))). Releasing the source code of published models on the web is a necessary step in enabling this requirement, but if modifications were also Grid-enabled, there is scope for allowing them to be used from the same Grid-supported experimental ontology as the original model, which would greatly facilitate repeating results using the same methodology and parameter settings. Reimplementations would be facilitated by adequate descriptions of the models, something not always possible in the context of a journal, as the level of detail required could be too great for anything other than an overview. Ontologies also have a potential contribution to make here as stated earlier through providing a formal description of the model. In the context of appropriately programmed modular frameworks for simulation models, Semantic Grid/Web services could be used to bolt together options for modelling a system, allowing exploration of implementation variants for subcomponents of a model without having to reprogram the entire system. As it stands, FEARLUS is a modelling framework providing some scope for exploring implementation variants through parameter settings, and hence FEARLUS-G does provide the beginnings of this level of repeatability.

1.7

These levels of repeatability are also pertinent to collaborating with simulation models. Clearly, when collaborating, a remotely-accessible platform that facilitates sharing of results and parameter files, and automatically distributes simulation runs across the available machines, has a significant contribution to make to the ease with which the shared study is conducted.



The FEARLUS-G ontology and application

2.1

Though not germane to the contribution this paper is making, a brief summary of the FEARLUS modelling framework will nevertheless assist in understanding the experimental methodology used to study it. For the sake of consistency with other descriptions of FEARLUS models in the literature, terms referring to objects in the model are distinguished from their real-world counterparts by an initial capital letter, and are italicised on first use. Figure 1 shows the annual cycle of a FEARLUS model, in which two or more *Subpopulations of Land Managers* (the agents in FEARLUS) with different *Land Use Selection Algorithms* are compared in an *Environment*. The Environment consists of *Land Parcels* with spatially varying but temporally unchanging *Biophysical Characteristics* (intended to represent such things as soil type, moisture, and gradient), and *External Conditions* (intended to represent the climatic or economic factors influencing the success of a land use), which change over time, but are spatially homogenous. The parameters determining the spatial pattern of the Biophysical Characteristics, and the rate of change of the External Conditions, are termed the *Spatio-Temporal Heterogeneity Type* (STHT) of the Environment. In each annual cycle, Land Managers use their Land Use Selection Algorithm to choose a *Land Use* for each Land Parcel they own independently, and then the *Economic Return* is calculated from that *Year's External Conditions* and accrued to Land Managers in their *Account*. Finally, Land Managers with a negative Account must sell their Land Parcels, the potential buyers of which include new Land Managers not currently in the model. Since the number of Land Parcels owned by its member Land Managers reflects the success of a Subpopulation in choosing Land Uses that are a good match for the local Biophysical Characteristics and changing External Conditions, it is this property that is used to compare the Subpopulations after a predetermined number of Years in a simulation run.




Figure 1. The FEARLUS model

2.2

The experimentation ontology is depicted in Figure 2. Instances generated from the classes in the ontology are of three types: those which need to be created by the user, e.g. the definition of an experiment; those instances created by the Grid service, e.g. the set of simulation run instances associated with an experiment and the experiment result instance summarizing the results; the third group are instances created to support scientific argumentation, sharing and reuse. For example, a hypothesis can be associated with related publications and experiments.




Figure 2. The FEARLUS-G experimentation ontology. (© 2005 IEEE; Reproduced from Pignotti et al. (2005a) with permission.)

2.3

The key class in our generic representation is the `Hypothesis`, which we consider to be a scientific concept that has not yet been fully verified. A hypothesis is "a tentative explanation that accounts for a set of facts and can be tested by further investigation" ([American Heritage Dictionaries 2004](http://www.ahdictionary.com/entry/article.html)). This implies that a scientific community works to support a hypothesis by contributing publications, experiments and other related hypotheses that agree or disagree with the "target" hypothesis. In our vision, the classes

HypothesisObject, ExperimentObject, SimulationObject and Publication are subclasses of a generic ScientificObject class.

2.4

We define two properties `agreesWithHypothesis` and `disagreesWithHypothesis` which link any given scientific object to `Hypothesis` instances. These properties capture the idea that Science is a discourse that does not necessarily present a consistent domain of knowledge, allowing us to support scientific argumentation in the FEARLUS context by linking specific evidence (`FearlusExperiment`, `Run`, `FearlusResult`) to hypotheses. The generic scientific object representation can, of course, be adapted for different scientific domains by extending it with more specific subclasses and properties. The FEARLUS-specific classes, meanwhile, expose the experimental workflow, making its use transparent to other users.

2.5

The Protégé^[4] tool was used during the ontology generation process. Protégé is an open-source tool developed at Stanford Medical Informatics and provides a highly customizable ontology editor that enables conceptual modelling with Semantic Web languages such as OWL and RDFS. The architecture of Protégé is cleanly separated into a "model" part and a "view" part. The model is the internal representation of ontologies and knowledge bases. The view component provides a user interface to display and manipulate the underlying model. The user interface contains a number of useful components: class editor, property editor, instance editor, ontology graph. An important component of Protégé is the OWL plug-in that allows the ontology (model) to be represented with OWL syntax. Moreover the OWL plug-in provides access to the Dublin Core^[5] OWL ontology, which has been used to provide basic annotation of `SimulationObject` instances. For example `dc:creator` defines the creator of an hypothesis, publication or experiment, `dc:contributor` the contributor(s). This ability of OWL to import ontologies is one of its key advantages, enabling ontologists to build on each others' work through referencing it over the web.

2.6

The ontology described above was integrated into the collection of Grid services that make up FEARLUS-G; together, the full range of technologies allow sharing of computational and data resources across different hosts. Figure 3 presents an overview of the FEARLUS-G architecture which is built on top of Globus Toolkit 3^[6] (Foster et al. 1998). Grid services in FEARLUS-G are used for uploading Scientific Objects (see Figure 2) to the Grid application (Upload Service) and executing simulation runs and experiments (FEARLUSG Service, FGSimulation Instance). One key component in FEARLUS-G is the Scientific Object Repository: an open source reusable semantic data service based on Jena2 (McBride 2001) to provide semantic data storage, query and retrieval functions. This service uses ELDAS (Baxter et al. 2003) to manage data repositories for RDF models; ELDAS was developed by the UK National e-Science Centre eDIKT project. Jena2 is a Java framework for writing Semantic Web applications. It provides programmable access to RDF and OWL sources, ontologies, documents, ontology reasoning and RDF query capabilities. The advantage of enabling Jena2 as a Grid service is that semantic resources can be distributed and used in a dynamic environment. Searching a large collection of RDF resources can be a computationally intensive task; the Grid offers the potential for distributed processing of such queries.




Figure 3. The FEARLUS-G service architecture

2.7

Figure 4 shows a typical FEARLUS experiment cycle in the context of the co-laboratory created by deploying FEARLUS and the experimental ontology on the Semantic Grid. The user specifies the location of the *Model* to be used, and sets up the FEARLUS experiment by collecting together the parameter settings and creating an instance of each of the *FearlusExperiment* and *FearlusHypothesis* classes. Though not shown in Figure 2, we have so far used three types of hypothesis in experimental work with FEARLUS. These pertain to comparing various different Land Use Selection Algorithms used by the Land Managers for making decisions in FEARLUS, in Environments with various Spatio-Temporal Heterogeneity Types (STHTs):

- Type 1: Managers with Algorithm X do better than Managers with Algorithm Y in an Environment with STHT A. This hypothesis enables straightforward comparisons between two Subpopulations of Land Managers on the basis of the Land Use Selection Algorithm each employs. For example, do Land Managers with random experimentation of new Land Use types in their Land Use Selection Algorithm outperform purely imitative Land Managers in situations of a moderate rate of change to the External Conditions?
- Type 2: Managers with Algorithm X do better against Managers with Algorithm Y in an Environment with STHT A than in one with STHT B. This hypothesis enables the Environmental context of the relative performance of two Subpopulations to be compared. For example, do randomly experimenting Land Managers do better against imitating Land Managers in a situation of rapid change to the External Conditions than they do when the External Conditions change slowly?
- Type 3: In an Environment with STHT A, Managers with Algorithm X do better against Managers with Algorithm Z than do Managers with Algorithm Y. This hypothesis enables the social context of the relative performance of two Subpopulations to be compared. For example, do imitating Land Managers do better against randomly experimenting or specialist Land Managers in situations of high spatial heterogeneity but slow change to the External Conditions?

2.8

Types 2 and 3 hypotheses use a series of paired-replicate runs — each pair of runs keeping as many aspects of the simulation as similar as possible, whilst varying the aspect of interest (the Spatio-Temporal Heterogeneity Type of the Environment in the case of Type 2, and the use of Land Use Selection Algorithm X or Y in the case of Type 3).




Figure 4. The FEARLUS-G workflow and (in blue) the facilities provided by the Semantic Grid. (© 2005 IEEE; Reproduced from Pignotti et al. (2005a) with permission.)

2.9

Knowing which Type of hypothesis is being used enables the runs to be generated automatically. The fact that several runs are used (according to the required power of the statistical test) means that this process is well-suited for the Grid to manage the computation, as each run can be conducted in parallel. Once the runs are completed, Grid services can also be used to collate the output from all the runs, and conduct the appropriate statistical test (a Binomial test for Type 1, and a Sign test for Types 2 and 3 (Seigel & Castellán 1988)) to obtain the result of the experiment.



Demonstration

3.1

The FEARLUS-G semantic grid service was made available for a short period on Aberdeen University's terabyte machine at the Department of Computer Science.

3.2

As well as testing by the computational land-use modellers at the Macaulay Institute, we conducted an evaluation of FEARLUS-G with the help of a group of 'assessors' from the international scientific communities interested in integrated agent-based socio-environmental simulation. Assessors were asked to use the FEARLUS-G interface to create and apply metadata to Grid-based land use simulation experiments; in addition, they were asked to explore and comment upon the mechanism for linking resources together to form an argument structure. We also tested the wider usability of the service by asking assessors to conduct land-use experiments and collaborate with other members of the evaluation group using the FEARLUS-G tools. The evaluation provided a number of insights into the utility of Semantic Grid technologies and methodologies in the social simulation context (Edwards et al. 2005): several assessors proposed extensions to the metadata framework to support other models and methodologies; the 'scientific argument structure' concept was understood and accepted.

3.3

In this section, we provide screenshots from the web-based user interface to demonstrate the collaboration and repeatability facilities the software provides.

3.4

In Figure 5, a researcher has completed an experiment to test a hypothesis, for which there is insufficient evidence. By clicking on the 'Publish Instance' link, they can make the experiment public, and available to other researchers. Each of the links (underlined) in the experiment form can be clicked on to explore the associated resource, including reports from individual runs, and parameter settings. The ontological support means that other researchers can know that the experiment, hypothesis, runs, parameters and results all belong together, and are associated with a particular publication. Without this support these would just be files in a database on the web, with nothing to necessarily indicate what they pertain to.

Experiment

InnovatorsVSImitators0

Name: InnovatorsVSImitatorso

Description:

Comment:

Tests Hypothesis: [InnovatorsVSImitators-Mo1Ho4R8ooYoooooI.sp-better-than-Mo1Ho8R8ooYoooooI.sp-against-Mo1H12R8ooYoooooI.sp-in-v-Po-E16u-BET8-H8RvH12R.model](#)

Uses Model: [modelo-6-6unix](#)

Number of Runs: 5

Grid Resource Status: Experiment completed

hasRuns: [InnovatorsVSImitatorso1120125856839](#)
[InnovatorsVSImitatorso1120125859680](#)
[InnovatorsVSImitatorso1120125902650](#)
[InnovatorsVSImitatorso1120125905560](#)
[InnovatorsVSImitatorso1120125928210](#)

Significance Level: 0.005

Reporting Year: year200Report

hasEnvironment: [Po-E16u-BET8-H8RvH12R.model](#)

hasSubpopulations: [Mo1Ho4R8ooYoooooI.sp](#)
[Mo1Ho8R8ooYoooooI.sp](#)

Experiment Summary:

Run	Mo1Ho4R8ooYoooooI.sp	Mo1Ho8R8ooYoooooI.sp
1	114	0
2	1	224
3	21	204
4	21	204
5	23	202

Number of trails in experiment : 5
Number of success of Subpopulation 1 : 1
Number of success of Subpopulation 2 : 4
Probability of success : 0.5
One-tailed p-value : 0.187500


[Publish Instance](#) [Remove Instance](#) [Add to Project...](#)


Figure 5. Browsing a completed experiment


3.5


In Figure 6, another researcher creates an experiment to rerun the published experiment. This time, ten, rather than five runs have been selected to increase the power of the test, which does not have a significant result. Typical FEARLUS experiments involve 30 or more runs. Note also that the model version is available from a drop-down menu. A partial contribution to the fourth level of repeatability is introduced here — though it does not allow the code of the model to be edited, different versions of the model could have implementation variants embedded that would allow the effect of these to be checked.


New Type 2 Experiment Set

Click on the  symbol if you need help with filling in a field.


Name: 


Description: 


Significance Level:
(For example: 0.005) 


Environment for the experiment:
Select the Environment(s) for the current experiment
 


Subpopulation for the experiment:
Select the Subpopulation(s) for the current experiment

Number of Runs: 

Single Subject Hypothesis: true false 

Reporting Year: 

Uses Model: 



Comment: 

Figure 6. Creating a new experiment to run

3.6

In Figure 7, a parameter file is being edited, allowing the investigator to use the new parameter file as part of an experiment to test whether some published results are robust to that change. Note that a further partial contribution to the fourth level of repeatability is also introduced here: though access to changing the code is not provided, a modelling framework with a modular design (such as FEARLUS) allows modelling components such as elements in decision algorithms to be swapped in and out by specifying them in the parameter files. For example, in Figure 7, an investigator could select different strategies to use in the land use selection algorithm contexts shown.

Subpopulation



[About this instance...](#)

Name: M01H04R800Y00000I.ss

Description:

Comment:

File:

```

NumberOfStrategyClasses: 3
Class  AboveThresholdProbability
      BelowThresholdNonImitativeProbability
      BelowThresholdImitativeProbability
      InitialProbability
HabitStrategy  1.0    0.0    0.0    0.0
RandomStrategy 0.0    1.0    0.0    1.0
NoStrategy     0.0    0.0    1.0    0.0
  
```

Figure 7. Editing a parameter file

3.7

By supporting the investigator's work with an underlying ontology, it will be possible to add it to scientific debate about the original hypothesis through appropriate use of the `agreesWithHypothesis` and `disagreesWithHypothesis` properties (see Figure 2). Should the investigator make their work public, the original researcher will be able to pick up on these developments and respond to them. In the meantime, if the work is being done as part of a collaborative effort, co-workers can use the properties to record the relationships between their respective findings.



Discussion

4.1

Through our experiences to date we have learned valuable lessons about the provision of Semantic Grid tools to the social simulation community; these are now discussed:

1. The demonstrator Grid service can provide scientists interested in modelling land-use phenomena with a means to run much larger-scale experiments than previously possible on standalone PCs, and also gives them a Web-based environment in which to collaborate and share simulation results. The Grid is clearly a valuable tool for managing the computation involved in running large-scale experiments. In fact the Globus environment is ideal for simulation modelling because the cost of a simulation run is less than the cost of the communication necessary to run a distributed experiment and aggregate the results. However, the Globus environment is not well-suited for the fine grained exploratory work that the existing FEARLUS desktop application is able to achieve when used in GUI mode. Scientists still want to use the patterns of workflow related to the desktop application. For this reason, we have extended the existing desktop application to be able to upload models to the Grid service, creating a link between the exploratory tool and the Grid experiment platform.
2. We have defined an initial collection of ontology elements that describe the scientific objects necessary to enable collaboration between members of a community of land-use scientists. Our ontology represents generic scientific concepts such as hypotheses and experiments, as well as more domain-specific concepts tailored to the use of FEARLUS. We have demonstrated that social scientists can run FEARLUS experiments on the Grid and collect the results; a private interface (My Workspace) allows them to manage their activities (experiments, investigations, etc.) while their work is in progress. Scientific resources can also be made public, allowing others to investigate them further.

Not all categories of users require access to the same meta-data. For example, we have exposed FEARLUS model parameters to land-use scientists familiar with the existing FEARLUS application in such a way that the complexity of the Grid services is hidden to them. For this reason we provide, together with the desktop application, a web application (My Workspace) which allows users to access and manipulate low-level simulation model parameters needed to run a FEARLUS experiment. However, we also need to provide high-level data to promote accessibility to the wider scientific community. For this reason we also provide access to our public interface where resources contain high-level generic properties (e.g. `describedIn`, `agreesWithHypothesis`, `hasDescription`).

3. Throughout this project the computational land-use modellers at the Macaulay Institute have been encouraged to lead the process of engineering the FEARLUS ontology. This important activity forced them to think about what conducting an experiment actually meant. Prior to this project commencing, Perl scripts were used to conduct experiments with FEARLUS and the experimental design was thus hidden. One positive side-effect of the ontology building exercise is that the experimental workflow is now exposed, making its use transparent to other users.
- 4.
5. If access to FEARLUS-G is to be provided outside the host institution, then security is a particular priority, though it is still a concern in a closed environment. Results and experiment data, both during research and after publication when they are available for scrutiny, should not be corruptible by naive or malicious use of the system. The Globus toolkit provides standard facilities for handling security. Globus Toolkit 3 uses identity certificates for both users and Grid services to enable secure authentication and communication over an open network. [\[7\]](#)

4.2

Without the ontological support provided by the Semantic Grid, the results and output would just be files, with nothing to say even that these files necessarily pertain to results or output of an agent-based simulation model, let alone to inform anyone who is interested that they have been created during the course of an experiment to test a hypothesis that is described in a publication. Annotating the results and output files, along with the models, hypotheses, experiments and publications with the ontological metadata allows the work to be inspected, re-used, and modified by other researchers, and hence also with collaborators. With appropriately recorded time-stamp data, it can also be used to create an audit trail, which can be used as part of conformance to quality standards such as ISO 9001, something that is increasingly required by funding bodies.



Conclusion

5.1

Reporting the output of agent-based simulation models with stochastic elements or with large numbers of potentially legitimate parameter settings should generally be done on the basis of statistical tests that show that such output is typical of the behaviour of the model. (An exception might be when the researcher is trying to demonstrate that a particular phenomenon is possible in a model.)

5.2

This pilot project has provided Semantic Grid support to experimentation with an agent-based model of land use change. We have demonstrated how this provision enables three levels of repeatability in social simulations: inspecting the original results on which a researcher's conclusions were based; rerunning experiments to check the same results are attained; and to make modifications to parameters to check the sensitivity or particularity of the results. Partial contributions to a fourth level of repeatability, making changes to the code, are provided through enabling different model versions to be used, and through a modular design approach to the modelling framework. The facilities provided improve the rigour of agent-based social simulation work, through complementing material published in journals, in which detailed descriptions are not always practicable due to the high volume of outputs generated by the models, and the complexities and subtleties of their implementation. These same features useful for disseminating simulation experiments to the wider community are also essential to smooth collaboration among colleagues. Through deploying the service on the Semantic Grid, colleagues can collaborate internationally, with all of the security concerns being dealt with by the Globus toolkit.

5.3

The Semantic Grid is clearly useful in providing support for investigations using social simulations. The 'Semantic' part provides an ontologically supported context for all of the information associated with a piece of work, from hypotheses through experiments to

publications. The 'Grid' part provides access to computing resources to run the simulations, without requiring users to download, compile and install software and parameter files on their own machines, with all the compatibility issues that this raises.

5.4

Future work will focus on using ontologies to describe the models themselves, as well as the experiment workflow to study them, and on deploying other social simulations. In the near future, we also hope to make the pilot grid service available on a longer-term basis. A link will be made available from the FEARLUS-G project website at <http://www.csd.abdn.ac.uk/research/fearg/>.



Acknowledgements

This work was funded by grants from the Scottish Executive Environment and Rural Affairs Department, and the Economic and Social Research Council (Award Reference: RES-149-25-0011). This paper contains material from papers presented at ESSA 2005 (Polhill et al. 2005) and ICeSS 2005 (Edwards et al. 2005; Pignotti et al. 2005b).



Notes

¹ <http://www.w3.org/TR/rdf-schema/>

² RDF stands for Resource Description Framework. See <http://www.w3.org/RDF/>

³ <http://www.w3.org/2004/OWL/>

⁴ <http://protege.stanford.edu/>

⁵ <http://dublincore.org/>

⁶ The next deployed version of the system is expected to use Globus Toolkit 4.

⁷ <http://www-unix.globus.org/toolkit/docs/3.2/security.html>



References

AAMODT A and Plaza E (1994) Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications* 7 (1), pp. 39-59.

ALLEN M (1999) Do it yourself climate prediction. *Nature* 401, p. 642.

ALLEN M, Stott P, Mitchell J, Schnur R and Delworth T (2000) Quantifying the uncertainty in forecasts of anthropogenic climate change. *Nature* 407, pp. 617-620.

AMERICAN HERITAGE DICTIONARIES (2004) (Eds.) *The American Heritage Stedman's Medical Dictionary 2nd edition*. Boston, MA: Houghton Mifflin.

BAXTER R, Ecklund D, Fleming A, Gray A, Hild B, Rutherford S and Virdee D (2003) Designing for broadly available grid data access services. In: *UK e-Science All Hands Meeting* (CD-ROM).

BERNERS-LEE T, Hendler J and Lassila O (2001) The semantic web. *Scientific American* 284, pp. 28-37.

CORTESE P, (and several hundred others!) (2004) ALICE: Physics performance report, volume I. *Journal of Physics G: Nuclear and Particle Physics* 30 (11), pp. 1517-1763

EDWARDS P, Preece A, Pignotti E, Polhill G and Gotts N. (2005) Lessons learnt from the deployment of a social simulation tool to the semantic grid. *Proceedings of the First International Conference on eSocial Science, NCESS/ESRC, Manchester*.

FENSEL D (2003) *Ontologies: A Silver Bullet for Knowledge Management and Electronic Commerce*. New York: Springer-Verlag.

FOSTER I and Kesselman C (1998) "Globus: A toolkit-based grid architecture." In: *The Grid: Blueprint for a Future Computing Infrastructure*. Morgan-Kaufmann. pp. 259-278

FOSTER I, Kesselman C and Tuecke S (2001) The anatomy of the Grid: Enabling scalable virtual organizations. *International Journal of High Performance Computing Applications* 15 (3), pp. 200-222.

GILBERT N (2004) Plenary presentation to European Social Simulation Association, 2nd Annual Conference, Valladolid, Spain.

GOTTS N M and Polhill J G (2006) Narrative scenarios, mediating formalisms, and the agent-based simulation of land use change. *EPOS 2006 Epistemological Perspectives on Simulation, University of Brescia, Italy, October 5-6, 2006*. <http://www.eco.unibs.it/epos/Papers/GottsPolhillEPOS2006Abstract.pdf>

GOTTS N M, Polhill J G, and Law A N R (2003) Aspiration levels in a land-use simulation. *Cybernetics and Systems* 34 (8), pp. 663-683.

GRUBER T R (1993) A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5, pp. 199-220.

IZQUIERDO L R and Galan J M (2004) Appearances can be deceiving: lessons learned reimplementing Axelrod's 'evolutionary approach to norms'. *European Social Simulation Association, 2nd Annual Conference, Valladolid, Spain*. http://www.uni-koblenz.de/~essa/ESSA2004/files/papers/Izquierdo_Galan_ESSA04.pdf.

LEE S, Cho M K, Jung J W, Kim J H and Lee W (2004) Exploring protein fold space by secondary structure prediction using data distribution method on Grid platform. *Bioinformatics* 20 (18), pp. 3500-3507

MCBRIDE B (2001) Jena: Implementing the RDF model and syntax specification. In Decker S, Fensel D, Sheth A and Staab S (eds.) *Proceedings of the 2nd International Workshop on the Semantic Web (SemWeb2001), Hong Kong, China, May 1, 2001*. CEUR-WS Vol-40. <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-40/mcbride.pdf>

MCCLELLAND J L, Rumelhart D E and The PDP Research Group (1986) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 2: Psychological and Biological Models*. Cambridge, MA: MIT Press.

MOSS S (2004) Search for good science: A personal memoir. *New School Economic Review* 1 (1), pp. 1–10. http://www.newschool.edu/gf/nser/articles/0101_mosss_goodsience_fall04_final.pdf

MOSS S and Edmonds B (2004) Towards good social science. *CPM Report 04-135, Centre for Policy Modelling, Manchester Metropolitan University, UK*. <http://cfpm.org/cpmrep135.html>

PIGNOTTI E, Edwards P, Preece A, Polhill G and Gotts N (2005a) Semantic support for computational land-use modelling. *Cluster Computing and Grid 2005, Cardiff, UK*.

PIGNOTTI E, Edwards P, Preece A, Polhill G and Gotts N (2005b) Providing ontology support for social simulation. *Proceedings of the First International Conference on eSocial Science, NCeSS/ESRC, Manchester*.

POLHILL J G and Gotts N M (2006) A new approach to modelling frameworks. *Proceedings of the First World Congress on Social Simulation, August 21-25 2006, Kyoto University, Kyoto, Japan*. Vol. 1, pp. 119–126.

POLHILL J G, Gotts N M and Law A N R (2001) Imitative versus nonimitative strategies in a land use simulation. *Cybernetics and Systems* 32 (1–2), pp. 285–307.

POLHILL J G, Izquierdo L R and Gotts N M (2005) The ghost in the model (and other effects of floating point arithmetic). *Journal of Artificial Societies and Social Simulation* 8 (5) <http://jasss.soc.surrey.ac.uk/8/1/5.html>

POLHILL J G, Izquierdo L R and Gotts N M (2006) What every agent based modeller should know about floating point arithmetic. *Environmental Modelling and Software* 21, pp. 283–309.

POLHILL J G, Pignotti E, Gotts N M, Edwards P and Preece A. (2005) "An ontology for experimentation with a social simulation of land use change". In Troitzsch K (Ed.) *Representing Social Reality: Pre-Proceedings of the Third Conference of the European Social Simulation Association, Koblenz, , 5-9 September 2005*. Germany: Verlag Dietmar Fölbach, pp. 324–330

RAO A S and Georgeff M P (1991) "Modeling rational agents within a BDI-architecture." In Allen J, Fikes R and Sandewall E (Eds.) *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*. San Mateo, CA: Morgan Kaufmann Publishers.

DE ROURE D, Jennings N and Shadbolt N (2001) Research agenda for the semantic grid: A future e-science infrastructure. *Technical Report, UK e-Science Series UKeS-2002-02, National e-Science Centre, Edinburgh, UK*.

SEIGEL S and Castellan N J (1988) *Nonparametric statistics for the behavioural sciences: International edition. Statistics Series, 2nd ed.* Singapore: McGraw-Hill.

YANG X J, Chang Z M, Zhou H F, Qu X L and Li C J (2004) "Services for parallel remote-sensing image processing based on computational grid". *Proceedings of the Grid and Cooperative Computing Workshops 2004, Lecture Notes in Computer Science 3252*, pp. 689–696.

[Return to Contents of this issue](#)

© [Copyright Journal of Artificial Societies and Social Simulation, \[2007\]](#)

