

An ontology for experimentation with a social simulation of land use change

J. Gary Polhill Macaulay Institute, Craigiebuckler, Aberdeen. AB14 0PR. United Kingdom. g.polhill@macaulay.ac.uk

Edoardo Pignotti, Department of Computing Science, University of Aberdeen, Aberdeen. AB23 3UE United Kingdom.
epignott@csd.abdn.ac.uk

Nicholas M. Gotts Macaulay Institute, Craigiebuckler, Aberdeen. AB14 0PR. United Kingdom.
n.gotts@macaulay.ac.uk

Pete Edwards, Department of Computing Science, University of Aberdeen, Aberdeen. AB23 3UE United Kingdom.
pedwards@csd.abdn.ac.uk

Alun Preece, Department of Computing Science, University of Aberdeen, Aberdeen. AB23 3UE United Kingdom.
apreece@csd.abdn.ac.uk

Abstract

In the FEARLUS (Framework for the Evaluation and Assessment of Regional Land Use Scenarios) project, we have developed a methodology for experimenting with our agent-based model of land use change. In a collaboration with the University of Aberdeen (one of the ESRC's e-Social Science Initiative Pilot Demonstrator Projects), we have developed an ontology for this methodology, and enabled access to it via the Semantic Grid (Pignotti et al., 2005). Here, we present the ontology, showing how it enhances the rigour of reported outputs from agent-based models with stochastic elements, and briefly describe how the ontology and its realisation on the Grid contributes further to developing good science (Moss, 2004; Moss & Edmonds, 2004) in the area of Agent-Based Social Simulation.

Introduction

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., in press).

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 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.

Ontologies¹ (Fensel, 2003) are used to capture the meaning of metadata terms and their interrelationships. 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 Schema² (RDFS)—a vocabulary for describing properties and classes of RDF³ resources, with a semantics for generalisation hierarchies; and OWL⁴ (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.

The Semantic Grid (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, provides access to computing,

¹ The term 'ontology' has been borrowed from Philosophy, but in Computer Science it means something different. An ontology is defined by Gruber (1993) as 'a formal explicit specification of a shared conceptualisation'.

² <http://www.w3.org/TR/rdf-schema/>

³ RDF stands for Resource Description Framework. See <http://www.w3.org/RDF/>

⁴ <http://www.w3c.org/2004/OWL/>

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.

Key to good science is 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 detailing 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 promotes 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.⁵ 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 without the need for reprogramming.

The experimentation ontology

The experimentation ontology is depicted in figure 1. The ontology contains high-level classes to support scientific debate, and should be appropriate for any application in agent-based social simulation. These high-level classes reflect the publication of scientific information, and the agreement and disagreement between scientific information and hypotheses, allowing the experiments to be fitted within the wider context of on-going discussions in the community.

The user specifies the location of the model to be used, and sets up the FEARLUS experiment by collecting

⁵ Note that the ontology described here is that of the experimental methodology rather than the FEARLUS model itself.

together the parameter settings to be used, and creating an instance of each of the FearlusExperiment and FearlusHypothesis classes. Though not shown in figure 1, we have so far used three types of hypothesis in experimental work with FEARLUS. These pertain to comparing various different algorithms used by the agents for making decisions in FEARLUS, in environments with various spatio-temporal patterns of change:

- Type 1: Agents with decision algorithm X do better than agents with decision algorithm Y in an environment with spatio-temporal pattern of change A.
- Type 2: Agents with decision algorithm X do better against agents with decision algorithm Y in an environment with spatio-temporal pattern of change A than in one with pattern of change B.
- Type 3: In an environment with spatio-temporal pattern of change A, agents with decision algorithm X do better against agents with decision algorithm Z than do agents with decision algorithm Y.

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 pattern of change of the environment in the case of type 2, and the use of decision algorithm X or Y in the case of type 3).

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) to obtain the result of the experiment.

Without the ontological support provided by the Semantic Grid, however, 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. 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.

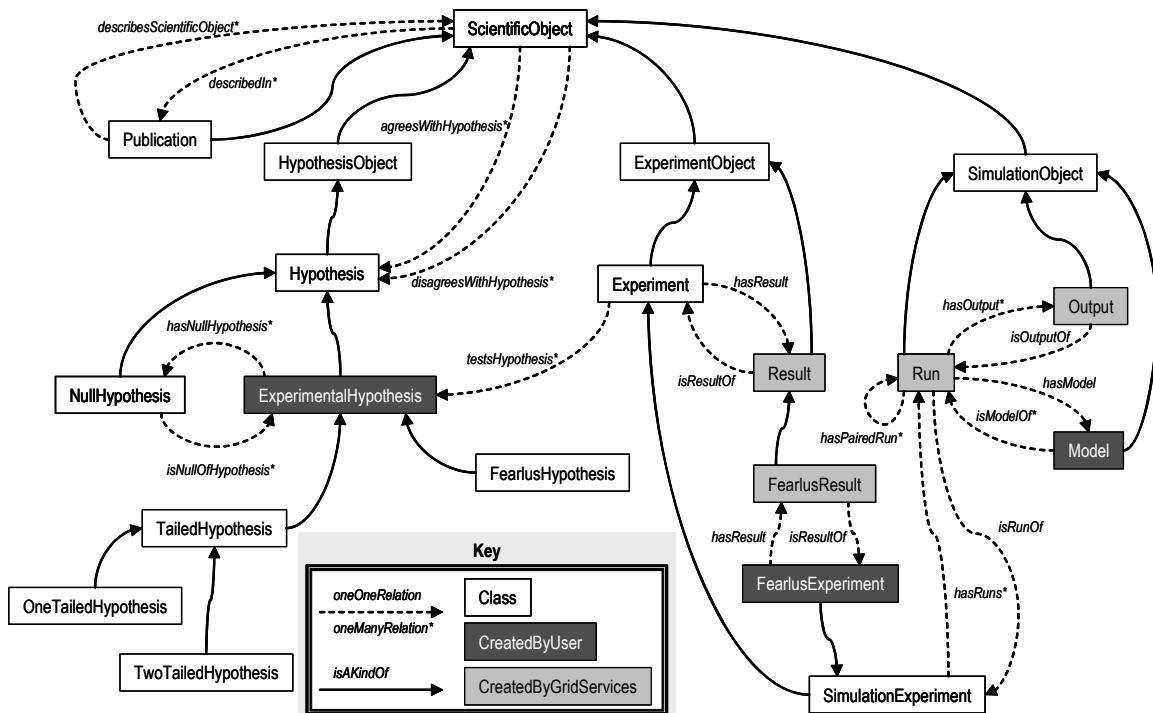


Figure 1. The FEARLUS-G experimentation ontology.

Demonstration

In this section we show screenshots from the web-based user interface, demonstrating the access to the three levels of repeatability that the Semantic Grid services provide.

In figure 2, 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.

In figure 3, 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. 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.

In figure 4, 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 4, an investigator could select different strategies to use in the land use selection algorithm contexts shown.

By supporting the investigator’s work in figure 4 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 1). Should the investigator make their work public, the original researcher will be able to pick up on these developments and respond to them.


The screenshot shows a web interface for an experiment titled 'InnovatorsVSImitators0'. The interface includes a header with a logo and a user icon. Below the title, there are several fields: Name, Description, Comment, Tests Hypothesis (with a long URL), Uses Model, Number of Runs (5), Grid Resource Status (Experiment completed), hasRuns (a list of five run IDs), Significance Level (0.005), Reporting Year (year200Report), hasEnvironment, hasSubpopulations, and an Experiment Summary table. At the bottom, there are three buttons: 'Publish Instance', 'Remove Instance', and 'Add to Project...'.


Run	Mo1Ho4R80oY0000oI.sp	Mo1Ho8R80oY0000oI.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


Figure 2. Browsing a completed experiment.


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: 

Subpopulation for the experiment: 

Number of Runs: 

Single Subject Hypothesis: true false 

Reporting Year: 

Uses Model: 



Comment: 

Figure 3. Creating a new experiment to run.

Subpopulation

M01H04R800Y00000I.ss
 [About this instance...](#)

Name:

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
```




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

Figure 4. Editing a parameter file.

Conclusion

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.)

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.

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.

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).

Bibliography

- Aamodt, A., & 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., & Delworth, T. 2000, "Quantifying the uncertainty in

- forecasts of anthropogenic climate change". *Nature* 407, pp. 617-620.
- Berners-Lee, T., Hendler, J., & 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
- Fensel, D. 2003, *Ontologies: A Silver Bullet for Knowledge Management and Electronic Commerce*. New York: Springer-Verlag.
- Foster, I., Kesselman, C., & 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., Polhill, J. G., & 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., & 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.essa.eu.org/ESSA2004/files/index.htm>.
- Lee, S., Cho, M. K., Jung, J. W., Kim, J. H., & Lee, W. 2004, "Exploring protein fold space by secondary structure prediction using data distribution method on Grid platform". *Bioinformatics* 20 (18), pp. 3500-3507
- McClelland, J. L., Rumelhart, D. E., & 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_moss_goodscience_fall04_final.pdf
- Moss, S., & 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., & Gotts, N. 2005, "Semantic support for computational land-use modelling." *Cluster Computing and Grid 2005*, Cardiff, UK.
- Polhill, J. G., Gotts, N. M., & 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., & 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., & Gotts, N. M. in press, "What every agent based modeller should know about floating point arithmetic". *Environmental Modelling and Software*.
- Rao, A. S., & Georgeff, M. P. 1991, "Modeling rational agents within a BDI-architecture." In Allen, J., Fikes, R., & Sandewall, E. (eds.) *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*. San Mateo, CA: Morgan Kaufmann Publishers.
- Roure, D. D., Jennings, N., & 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*.
- Yang, X. J., Chang, Z. M., Zhou, H. F., Qu, X. L., & 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.

Author biographies

Gary Polhill is a research scientist in the Landscape Change Programme at the Macaulay Institute.

Edoardo Pignotti is a research assistant at the Department of Computing Science, University of Aberdeen.

Nick Gotts is a research scientist in the Landscape Change Programme at the Macaulay Institute.

Pete Edwards is a Senior Lecturer and currently Head of Computing Science at the Department of Computing Science, University of Aberdeen.

Alun Preece is a Senior Lecturer at the Department of Computing Science, University of Aberdeen.