Open-Endedness: Definitions and Shortcuts

Wolfgang Banzhaf[1], Guillaume Beslon[2], René Doursat[3] and Susan Stepney[4]

[1] Department of Computer Science, Memorial University of Newfoundland, St. John's, NL, A1B 3X5, Canada
[2] Université de Lyon, INSA-Lyon, INRIA Beagle, CNRS LIRIS UMR5205, Villeurbanne, France
[3] Informatics Research Centre, School of Computing, Manchester Metropolitan University, M1 5GD, UK
[4] Department of Computer Science, and York Centre for Complex Systems Analysis, University of York, YO10 5DD, UK

Introduction

The open-endedness of a system is often defined as *the continual production of novelty*. Here we report on recent work by Banzhaf et al. (2016) that pins down this concept more rigorously. We define several types of novelty that a system may exhibit, and classify these as *variation*, *innovation*, and *emergence*. We also define an architecture suitable for building simulations of open-ended novelty-generating systems.

A meta-model for open-ended systems

Our definition of open-endedness (OE) is relative to the model and meta-model of the system under investigation. The concepts of "model" and "meta-model" are used for building both scientific and engineering models of systems: *Models* provide an abstract language for the relevant concepts. A scientific (descriptive) model is used when observing and experimenting on a (natural or artificial) system; An engineering (prescriptive) model (often a computational model) is used to design and implement an artificial system. In the case of a computational model, the model defines the concepts to be implemented in code. *Meta-models* provide the analogous abstract language to define models, comprising the concepts that can be used to build the model. For example, in an object-oriented system, the meta-model would contain the concepts of 'class', 'object', 'method', 'association', and so on.

There are many possible models and meta-models for capturing and analysing the behaviour of systems. These models and meta-models may be *implicit* or *explicit*. Here we assume that we will implement and analyse agent-based (or entity-based) systems. A suitable meta-model for such an analysis, which allows us to capture intuitions about major transitions, contains the following concepts:

- **Entity**: an identifiable integrated whole within the model: a "thing" with structure (organisation) and behaviour (activity, processing). An entity may be an **atomic entity**, with no (modelled) internal structure, or a **system entity**, composed of internal components.

- **Environment**: part of the domain of interest not modelled as explicit entities, for example: space, fields, flows.

- **System**: a local environment plus a collection of interacting entities, forming some identifiable whole. A system may be an **aggregate system**, comprising a collection of entities in a local environment but not considered to form an entity in its own right; or a **system entity**, modelled as an entity at a higher level than its component entities.

- **Interaction**: entities interact with each other and with their environment, potentially forming systems

- **Level**: level-0 entities are atomic entities; level-$N > 0$ entities are system entities that contain lower-level entities, including at least one level-$(N-1)$ entity.

This meta-model can be used to provide the concepts used in building specific entity-based models of system.

Definitions

We define types of novelty and OE *with respect to the system's current model and meta-model. Novelty* in an observed system is classified as:

0. *Variation: novelty within the model.* Variation changes an instance of the model, such as a change to the values of a variable that exists in the model, without changing the model itself. Variation explores a pre-defined (modelled) state space, producing new values of existing variables.

1. *Innovation: novelty that changes the model.* Innovation changes the model: for example, adds a new type or relationship (that conforms to the meta-model), or eliminating an existing one. Innovation changes the combinatorics and the size/structure of the state space, thereby growing/shrinking the possibilities of variation.

2. *Emergence: novelty that changes the meta-model.* Emergence changes the meta-model: a change that adds a new meta-type or relationship, or possibly eliminates an existing one. A change to our meta-model that adds a new level is a *major transition*. Another change to our meta-model

might include the addition of the concept of *process* as a first-class thing.

An *open-ended event* is an event that results in innovation or emergence. An *open-ended system* is a system with the ability to continually produce open-ended events.

Implementation shortcuts

To study a system such as the ones captured by our meta-model, we need to be able to simulate multi-level systems. We may need to introduce new components in our simulation to simplify or accelerate the dynamics of the higher levels. We call these new elements **shortcuts**, since they directly implement some properties of the higher level that, in an ideal simulation, would emerge from the generative level-0 entities. Shortcuts are hard-coded design optimisations manually introduced into the simulation.

Shortcuts provide optimisations by explicitly constraining structures and behaviours at their level, rather than requiring these constraints to emerge from the system's behaviour. Hence, specific shortcuts will enable or constrain certain classes of open-endedness.

When designing the simulation of an open-ended system, one crucial design step is then the identification and implementation of relevant shortcuts. These will be research-dependent: they will depend on the specific question one wants to answer with the simulation. Examples of shortcuts used (usually implicitly) in simulations are:

- **Individuality**: entities at level $N > 0$ are hard-coded rather than emergent.

- **Replication**: the explicit implementation of operators that replicate the entities at level N external to the entities.

- **Fitness**: replacement of the differential reproductive success emerging from the difference between entities by an explicit computation of reproductive success according to some target task.

Open-ended simulations

We have defined OE in terms of innovation and emergence: changes to models and meta-models. In classical software simulations, a model is designed (say, in UML) conforming to some meta-model (e.g. of object orientation, or agent-based systems), and then implemented. The running code conforms to the model. That model does not change as the system runs. Therefore, can such simulations ever exhibit innovation or emergence? Can they be open-ended?

The short answer is that they exhibit innovation and emergence. The implementation model is an engineering model. However, the running system can be analysed with a scientific model, and this model *can* change. Consider Conway's Game of Life. The engineering model consists of dead/alive cells, which interact with their neighbours. This model is

(typically) fixed. However, the model used to analyse the running system can include concepts such as blocks, and gliders, and other higher-level components. These higher levels are *emergent*, according to our definition.

The longer answer, however, is that they probably cannot exhibit OE, i.e. the continual production of open-ended events. The reason is that the code, conforming to the fixed model, cannot directly exploit the (scientifically modelled) innovations and emergents. It may well be the case that changes to the scientific (observational) model need to be fed back into changes to the engineering model. Steps along the way include:

1. The emergent novelty is recognised outside the simulation, by analysing the scientific model.

2. The emergent novelty is anticipated and recognised, using pre-coded recognisers present in the code (hence in the engineering model), available to report it once it appears.

3. The type of emergent novelty is anticipated and captured, using pre-coded rules available to recognise it, and make it a component of the simulation once it appears.

4. The emergent novelty is somehow emergently recognised by the simulator, and new code is generated by the simulator to capture the recognition.

5. The emergent novelty is somehow emergently captured by the simulator: once recognised, new shortcut is generated by the simulator to capture the specific emergence.

How to realise the later steps remains an open research question. Until then, simulations can use hard-coded shortcuts to overcome some of these issues.

Conclusion

We have defined open-endedness in terms of novelty, and identified three classes of novelty in terms of models and meta-models. We have sketched a meta-model for entity-based simulations, and a shortcut architecture for efficient implementation. We have outlined research requirements for fully open-ended simulations that can incorporate model changes and meta-model changes automatically.

References

Banzhaf, W., Baumgaertner, B., Beslon, G., Doursat, R., Foster, J. A., McMullin, B., de Melo, V. V., Miconi, T., Spector, L., Stepney, S., and White, R. (2016). Defining and simulating open-ended novelty: Requirements, guidelines, and challenges. *Theory in Biosciences*. DOI: 10.1007/s12064-016-0229-7, 54pp.