Learning about and through Empirical Modelling

Russell Boyatt, Antony Harfield and Meurig Beynon
Department of Computer Science, University of Warwick, UK
{rboyatt, ant, wmb}@dcs.warwick.ac.uk

Abstract

Empirical Modelling is a body of principles and tools that has been developed for the construction of interactive environments. Our previous research has indicated respects in which Empirical Modelling is intimately linked with learning activity of many different varieties. In this paper, we recount informal evidence in support of this claim that can be drawn from the assessment exercise attached to the "Introduction to Empirical Modelling" module offered in Computer Science at the University of Warwick. This assessment takes the form of an open-ended modelling and paper-writing exercise. Such an exercise is shown to be effective for learning about Empirical Modelling. It also promotes self-motivated exploration in unknown domains that is one of the key skills for life-long learning. The extent to which students not only learnt about Empirical Modelling, but also about the domain which they chose to model was unexpected. This leads us to suggest that Empirical Modelling could be effective in facilitating learning in other domains.

1. Introduction

"Introduction to Empirical Modelling"' is a module that has been run in the last 4 years for final year undergraduates on the 4-year MEng Computer Science course at the University of Warwick. Empirical Modelling (EM) is an approach to creating interactive environments developed extensively at the University of Warwick. EM places a greater emphasis on the human-centered construction of models that embody the elements of observation, dependency and agency that are encountered in every-day experience [1]. The process of constructing models (or artefacts) leads the modeller to a personal understanding of situations in the spirit of constructionism [4]. In 2004-5, a new form of assessment for the module was introduced. This involved the informal publication of an online journal [2] – the first Warwick Electronic Bulletin on EM (WEB-EM-1) – to which the students were required to submit papers and associated EM artifacts.

This paper has three principal components. We first briefly present the format of the module, the student submissions, the marking, and the results. We then identify some features of effective learning that are illustrated in the submissions. Finally, we discuss the significance of EM in relation to the kind of learning exhibited in WEB-EM-1, and review issues arising that suggest possible directions for future research and module development.

2. The Module Assessment

The module ran for 10 weeks, with 2 hours of lectures and 1 hour of computer laboratory sessions per week (with extra laboratory time available). Students were introduced to the concepts of observation, dependency, and agency that they were expected to use when analysing their problem domain. The laboratory gave students experience of the principal EM tool, the tkeden interpreter [6], with its associated family of built-in notations for framing dependencies between scalars, strings, lists, geometric entities and screen display elements. The module also introduced the LSD notation for accounting for inter-agent communication, and more advanced utilities such as can be used to support agent-oriented parsing and to depict networks of dependencies.

For the module assessment, we issued a Call for Papers requiring two submissions. Students first submitted a paper title and abstract. These submissions were reviewed and feedback was given. Students subsequently submitted their full paper and accompanying model.

The coursework had two objectives. The first was to assess the students' understanding of Empirical Modelling through written and modelling exercises based on a common theme of the students' own choice. The second was to equip the students with basic research skills that would be useful in further education. In the Call for Papers, we requested that students submit original and high quality papers relating to Empirical Modelling and its applications supported by a relevant documented modelling study.

2.1 Background

In the academic years 2002-03 and 2003-04, the coursework assessment required students to build a model using the EM tools. The students were all given the same task – in 2002-03 to implement a board game, and in 2003-04 to make a model of a heating system. Although many good submissions reflected the hard work of the students we felt the scope was not wide enough for capable students. On the evidence of their submissions, many students were keen to put effort into developing their submissions beyond the original specification even though their knowledge of board games and the workings of heating systems was limited. On that basis, it seemed natural to give students the opportunity to apply the EM tools and principles covered in the module to a subject of their choosing. Since many fourth year students are likely to proceed to research we also wanted to promote research skills that would assist their future studies.

2.2 Submissions

The selection of the submissions to which we shall refer in this paper is listed below. A complete list of submissions is available from the first Warwick Electronic Bulletin on Empirical Modelling [2].

- Tournament. A notation and model for the organisation of knockout tournaments.
- IceCube. An exploration of IceCube, a technology developed by Microsoft that deals with reconciling divergent replicas of some shared system state.
- Grid computing. A simulation to allow exploration of the efficiency of a computing grid.
- Bridges. A model exploring basic engineering principles behind bridge building.
- Non-decimal bases. A learning artefact to aid the understanding of non-decimal bases.
- Greedy algorithms. A learning artefact to demonstrate a greedy algorithm for the 'making change' problem with different currencies.
- Wumpus. A model that illustrates the game of "Hunt the Wumpus" first introduced in AI research. A screenshot of this model can be seen in Figure 1.
- Poker. A model studying the communication of information in a game of poker.
- Frisbee. A model exploring the interaction in a game of frisbee.
- Human Biology. A dependency-based simulation to illustrate how the lungs function.

2.4 Marking

Out of the 25 abstracts initially submitted and approved, all but six led on to final submissions to WEB-EM-1. Our analysis is based on the final submissions. Each submission comprised of a model and a paper; in the assessment process, these were marked together. The marks served as a good discriminant of skill and understanding in EM, lying in the range 45-80%, with an overall average of 63%.

3. Analysis

This section describes aspects of learning that were highlighted by the assessment:

- Learning can occur and skills can be developed without a preconceived objective
- Learning is stimulated by personal interest.
- Learning is reinforced when practice and principles are combined.
- Learning is aided by exploration.

In the following section we shall consider respects in which Empirical Modelling is well-suited to supporting learning that exhibits these characteristics.

3.1 Learning can occur and skills can be developed without a preconceived objective

As in previous years, the coursework helped to develop practical skills with the EM tools. However, the potential for emphasising different aspects of Empirical Modelling was apparent with this new style of assessment. Some students stuck to the basic tools whilst others made use of other, often more technical, tools and notations. The Frisbee model made use of only of the basic notations for data manipulation and line drawing, introduced at the beginning of the course, and the student was clearly proficient in building models with these notations. Other models, such as Tournament, involved the development of special-purpose notations which exercised a different skillset associated with agent-oriented parsing. Another student modelled the game of poker from different viewpoints using the distributed EM tool. Others emphasised the incremental aspect of model-building in their model. For example, the Making Change model used incremental development to show how a learning artefact might be adapted to situations that arise and evolve as the understanding of the learner develops. Each student developed the same basic skills but some also demonstrated extra skills. The fact that students could choose what skills to develop within certain broad constraints contributed to the diversity and richness of the submissions.

The journal-style of assessment demands a different skill set from the typical computer science coursework. Coursework is usually a specific task which the student should tackle in a preconceived way and hence often results in similar submissions. In our assessment, the students were given a set of tools and asked to develop their own theme within a general framework of possible applications of EM. This required the student to be self-motivated and to think for themselves about how they should approach the coursework. As can be seen from section 2.3, a wide range of topics and interests was represented in the submissions.

We found that students were able to direct their own learning without being given a specific coursework objective. The student who submitted the Wumpus model initially set out to reconstruct the original Wumpus game using the EM tools provided. This proved successful but furthermore the interactive, open-ended nature of the environment allowed the student to model different scenarios they had not originally considered, e.g. by changing the rules of the game and/or manipulating the information presented to the player. In his submitted paper, he explained how through this interaction with his model he had begun to appreciate how his ability to win the game depended upon the rules of the artificial Wumpus world and how, outside such a constrained environment, pure reasoning was not always sufficient.

3.2 Learning is stimulated by personal interest

In his account of constructionism, Papert [4] stresses the importance of personal interest as a motivating factor for active learning in a constructionist idiom. In this spirit, the students were encouraged to think about issues of which they had particularly rich experience or were particularly interested in learning about. One student explored applications of greedy algorithms by carrying out empirical research into how her younger siblings learnt about giving the correct change. Coursework often forces students to study situations with which they are unfamiliar or topics that do not interest them. By choosing their topic, students were able to draw and build upon a wide range of prior knowledge, interests and experiences. EM actively encourages this type of learning [5]. Because students

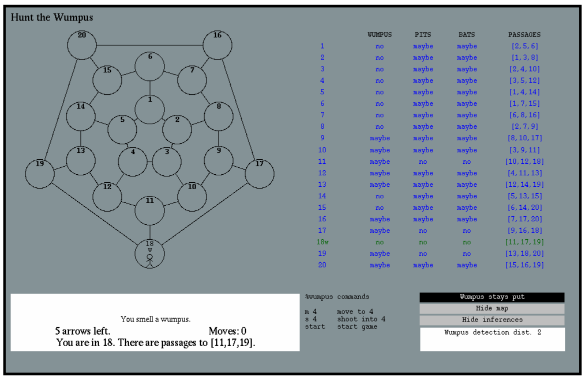

Figure 1. The Wumpus model

worked on topics of their own choice, the focus of their effort could be on EM principles and tools and not on an arbitrary topic prescribed for them.

All of the submissions showed evidence of an interest in domains other than Empirical Modelling. These domains ranged from personal hobby interests to aspects of the computer science curriculum. The Poker and the Frisbee models were inspired by recreational interests. One student made use of his Grid Computing model to complement his coursework for another computer science module. Another developed some research by Microsoft into the IceCube framework. From the depth and quality of his submission, it is apparent that this student spent as much time learning about IceCube as they did about Empirical Modelling. This contributed significantly to the quality and ingenuity of his final model; it also demonstrated how the model-building could stimulate learning in other domains. Yet other students chose to model phenomena in other academic subject areas. One submission relating to human biology was a model of the lungs that incorporated a primitive simulation to expose the effects of damage to organs or of cigarette smoking. The simple but effective use of dependency in this model highlighted the extent to which naive medical knowledge of bodily functions is knowledge of basic inter-relationships between physical conditions and parameters. This underpinned the educational purpose behind the model – just one of many references to education in the written submissions.

3.3 Learning is reinforced when practice and principles are combined

The practical element of a subject can often become divorced from reflection on principles. Although model-building is a useful tool for developing basic EM skills, it should be guided by higher-level motivations and interpretative activities. In previous years, the written component of coursework had been primarily oriented towards the technical documentation of models. Introducing the paper-writing exercise into the written component of the coursework helped to promote a broader awareness of the thinking behind EM and its implications.

Several models successfully illustrated deeper concepts of direct relevance to EM. In the Wumpus model, for instance, the environment can be configured so as to expose the limitations of logic outside a context of stable expectations and reliable knowledge. In his extension to the traditional AI game, the student was able to expose problematic aspects of a purely logicist outlook on intelligence in a manner that had not been preconceived. The Grid Computing model was a good example of a student using a model to convey concepts. In this case the model served to illustrate the basic principles of grid computing by generating animations of the kind of diagram that would typically be found in an introductory textbook.

3.4 Learning is aided by exploration

The quality of the submissions was such that most of the students were able to grasp the use of the EM tools and, in some cases, exploit their more advanced aspects. This is one reason why exploration into other domains occurred so naturally in the coursework. Once the tools had become familiar, the student no longer had to focus on the technicalities of modelling, but could make use of the tools to communicate or develop their domain understanding. As is to be expected, the better the student's EM skills, the more they were able to explore their problem domain.

To demonstrate this, we have categorized the submissions by the extent to which they explored their problem domain: little/no exploration, controlled exploration, and free exploration. Submissions that showed little/no exploration were generally based on the style of implementation that is quite familiar in computer science. A typical example is a model of a game that concentrates more on the implementation than exploring the observations, rules and interfaces that shape the interactions within the game. In the 'controlled exploration' category, the submissions often related to a problem domain with which the student was familiar, possibly drawn from the academic field. An example is the Human Biology model that enabled the user to explore the effect of smoking on the oxygen intake via the lungs. Models in the 'free exploration' category typically resulted from a student's engagement in unfamiliar problem domains. For example, the Bridges model originated in a basic study of the strength of bridges and ended up modelling complex issues in suspension bridges. Applying a statistical T-test at a 99% confidence level shows that students who engaged in exploration achieved higher marks than the students who did not show signs of exploration in their coursework.

These observations suggest that students who had a good grasp of EM tools were able to engage fully with the problem domain and produce coursework of a higher standard. This is what you would expect as we were evaluating Empirical Modelling ability rather than expertise in the problem domain.

4. The significance of EM for learning

The merits of EM as a vehicle for learning have been discussed and illustrated in detail elsewhere [5]. In particular, EM has been viewed as providing more effective support for constructionism than conventional approaches to computer-based modelling [3]. Crucial to this claim is the scope that EM offers to engage with activities that relate to the most primitive aspects of learning. EM artefacts can embody tacit pre-articulate knowledge that is expressed only through the modeller's personal interaction and interpretation.

The characteristics of the learning exhibited in the WEB-EM submissions accord well with the characteristics of EM. Where conventional programmers are encouraged to assemble a secure base of knowledge prior to writing the first line of code, EM practitioners are encouraged to initiate their exploration of the application domain and their construction of a computer-based model simultaneously. The fundamental reason for this distinction in outlook has to do with the perception of knowledge that underlies thinking about conventional programming and EM. The conventional programmer targets sophisticated knowledge that is sufficient to provide a robust logical framework ('knowing *that* certain relationships hold') and complementary precise recipes for action ('knowing *how* to achieve specified goals'). By contrast, EM is primarily concerned with a much more primitive conception of knowing (cf. [1]) – with conjunctions between experiences as personally encountered by the modeller. The qualities of EM artefacts stem from this grounding in experienceable connections that pervades the context within which all 'knowing that' and 'knowing how' is subsequently rooted.

The principal technical contribution of EM to moulding one experience so that it 'knows' another is to be found in the notions of observable, dependency and agency. The diversity of the domains represented in WEB-EM-1 is further evidence of the pervasive relevance of these primitive notions, which are viewed as conceptually prior to the identification of formal objects and structures. The integrity, functionality and interpretation of an EM artefact is framed by the meaningful interactions that the modeller can carry out with it, as guided and ostensibly constrained by past experience of interaction. In this respect, EM artefacts are ontologically quite unlike computational objects and structures, even though in practical settings they may resemble them closely. The character of EM artefacts is consistent with the features of the learning activity described in the previous section: not necessarily being associated with a specific learning objective; connecting closely with personal experience; synthesising empirical and theoretical elements; drawing upon extensive exploratory activity. It is also relevant to whether students chose to pursue the "Introduction to EM" module to its completion, since it epitomises an orientation towards knowledge that is congenial to some but alien to others. Unsurprisingly, there is only a loose correlation between good overall performance in computer science and aptitude for EM.

5. Concluding remarks

As our module was an introduction to Empirical Modelling, we intended and expected the students to learn about the basic concepts of EM. Students demonstrated their proficiency with EM tools and techniques through the models they were able to construct. What is more surprising is the extent to which some students learnt about the topic area chosen for their modelling exercise in carrying out the assessment. Learning evidently occurred in both domains – in EM and in their area of interest. We believe that this will be the case in other domains beyond those covered by student submissions to date.

Furthermore, we believe that our assessment exercise promotes those meta-skills relating to self-motivated, self-directed exploration of a problem domain that are most needed by life-long learners. It is characteristic of life-long learners that they are not necessarily following a formal path of education. They are much more likely to have personal goals and individual learning objectives. Although more rigorous research beyond this preliminary appraisal is needed, the above review indicates that EM tools have the potential to address the needs of life-long learners.

EM and our assessment strategy seem to work well together, and to offer prospects for learning in other disciplines. Whether the style of assessment would be so effective with traditional programming tools is unclear, since it is hard to interpret embodying observation, dependency and agency in models in that context.

When all students built models from the same specification, assessment was more straightforward because it was easier to identify the relative merits of each piece of work. Our new style of assessment has meant that there is a greater focus on the use of EM principles and how they have been applied to the specific topic area chosen by the student. This presents challenges for our marking procedures as it requires additional time to familiarize with each individual model and its topic. However, we have found that the students who focused on unusual subject areas for their model building tended to build models that aided rather than hindered our own understanding of that area.

Further empirical data is needed to support wider generalisations about EM as facilitating learning in other domains. Early indications from the latest WEB-EM coursework submissions (2005-2006) show a similar trend to results discussed in this paper.

Future plans for developing the WEB-EM concept include incorporating conference presentations into the assessment. Subject to resources and interest, there may also be scope for making module material available online and inviting submissions from external students.

Acknowledgements

We would like to thank all of the students taking the 'Introduction to Empirical Modelling' module whose assessment submissions have informed this paper.

References

[1] W. M. Beynon. Radical Empiricism, Empirical Modelling and the nature of knowing. 13(3):615-646, December 2005.

[2] W. M. Beynon (ed.). The First Warwick Electronic Bulletin on Empirical Modelling (WEB-EM-1), 2005. http://www.dcs.warwick.ac.uk/~wmb/web-em-1/.

[3] W. M. Beynon and A. Harfield. Empirical Modelling in Support of Constructionist Learning: a Case Study from Relational Database Theory. In Proc 5th IEEE International Conference on Advanced Learning Technologies (ICALT'05), 396-398, 2005.

[4] S. Papert and I. Harel. Situating Constructionism. In I. Harel and S. Papert, editors, Constructionism, 1-12. Ablex, Norwood, NJ, 1991.

[5] C. Roe. Computers for learning: An Empirical Modelling perspective. PhD thesis, University of Warwick,November 2003.

[6] A. Ward. Interaction with meaningful state: implementing dependency on digital computers. PhD thesis, University of Warwick, May 2004.