

Greedy Algorithms and the Making Change Problem

Abstract

This paper discusses the development of a model which facilitates the understanding of the 'Making Change Problem,' an algorithm which aims to select a quantity of change using as few coins as possible. The paper introduces the Empirical Modelling approach to generating software for educational use and discusses its benefits due to its extensive scope for redefinition, customisation and extension.

1 Introduction

This paper discusses the educational issues of a model of the Making Change Problem which requires you to pay a given amount using as few coins as possible. In general people can easily derive the correct solution to this problem, however they do not know the logic they are using in selecting the appropriate coins. This model exploits the users understanding of the action of giving change in order to introduce the abstract mathematical concepts of greedy algorithms. The model allows the user to make a meaningful connection between the mathematical logic and their experiences of these actions.

This paper begins by discussing the theories of learning, before discussing why Empirical Modelling is an appropriate method of constructing educational models. Current educational technologies are often felt to be limited in scope for development and adaptation to personal needs. In contrast models developed the EM way can be relatively easily adapted to the users needs.

1 Empirical Modelling and Learning

1.1 Theory of learning

Learning involves making connections between experiences of artefacts, and the development of skills in their manipulation. It begins as a private experience of interactions that reveal features and relationships between objects. Gradually we gain know-

ledge of important features and dependencies between artefacts. We learn to understand our role as an agent through interaction with the artefacts and observing how our actions cause change. (Beynon 1997)

Kolb's cycle (Fig 1), suggests that through active experimentation and reflection on our personal, private, experience we are able to develop new concepts and abstract ideas. These new ideas in turn stimulate further experiments with the concrete artefacts that typically lead to new views of our experiences. The greedy model discussed in this paper exploits our personal understanding of real world artefacts by allowing users to experiment with a model which references these artefacts in order to develop an abstract conceptualisation.

Fig 1 : Kolb's experiential learning cycle

1.2 Benefits of Empirical methods

This section of the paper discusses the benefits of empirical modelling as an educational tool with particular regard to its empirical basis and the extensive possibilities for redefinition, customisation and extension.

private experience / empirical / concrete
 interaction with artefacts: identification of persistent features and contexts
 practical knowledge: correlations between artefacts, acquisition of skills
 identification of dependencies and postulation of independent agency
 identification of generic patterns of interaction and stimulus-response mechanisms
 non-verbal communication through interaction in a common environment
 directly situated uses of language
 identification of common experience and objective knowledge
 symbolic representations and formal languages: public conventions for interpretation
public knowledge / theoretical / formal

Box 1 : An empiricist Perspective on Learning

1.2.1 Similarity to the real world

Many people criticize the use of computers in education because the virtual nature of computer interaction leads children to a less engaging learning experience (Cronick 2003, Roe 2003). It cannot be denied, however, that the interactive nature of computer based models can offer advantages over real-world artefacts. Experimentation can often be more easily performed on a computer as there are no real world limitations. For example users are more free to experiment through exploration, trial and error, as the outcome of an incorrect experiment cannot cause a negative outcome.

As the name suggests Empirical Modelling (EM) is based on experimentation and observation rather than abstract theory. It can be thought of as an approach to developing models which we interact with in a similar fashion to artefacts in the real world. This is a great advantage in the area of educational technology where it can be argued that learning through experimental interaction, as opposed to learning from written materials, can be more effective.

Constructivist learning theory maintains that knowledge is not received from outside, but that we construct knowledge in our head. To varying degrees constructivists believe that our understanding is individual and changing, based on our private experiences (Allesi 1991). EM supports a similar view and is in agreement with William James's idea that knowledge is rooted in personal experience. Box 1 above shows how learning experiences can be classified on the scale between private and public knowledge. (Beynon 1997, Roe 2003)

1.2.2 Customisation

The demand for adaptable educational software arises from the different needs and abilities of those using the software. In this respect Empirical Modelling is a very useful technique for the development

of educational models. In traditional educational software, a learner interacts with the software but does not modify it. In contrast EM models allow the user to customise and extend the system as appropriate to the teacher and student's needs.

When developing models for use in education there are three different perspectives to consider; the teacher's, student's and developer's (Beynon.2003, Roe 2003). Each naturally has different concerns and as a result the development of a model which all parties deem successful is a significant challenge. For example computer programmers are usually concerned with matters such as requirements, usability, and low level implementation details, their focus is on providing polished software products (Beynon 2001). In contrast the teacher is concerned with the learning outcomes for their students that the software facilitates. Teachers have little regard for the specifics of the implementation. Indeed they may prefer software that may lack the sophistication of a commercial product but allows a higher degree of adaptability (Beynon 2001). Section 3.2.3 discusses how the greedy model was developed with the ability to adapt the model easily.

1.3 Empirical Modelling process

The empirical modelling process is an incremental and iterative one. It is not appropriate to view our models as a collection of conventional programs with fixed-input output relations, it is more like a laboratory in which different kinds of models can be conceived and developed (Roe 2003). When developing a new model the modeller first analyses the actions and perceptions of the user. They then begin by constructing a basic model expressing these ideas as a combination of observables, dependencies and agents. Interaction with this model teaches the modeller more about their understanding of the situation and leads to further development. For example they may realise their original

perceptions were wrong, or they may extend the definitions to add further functionality. Such incremental development may not be possible when developing models in other ways where it is necessary to specify the whole system for the model to work, or redefinition is difficult.

An excellent example of the Empirical Modelling concept is the development of a spreadsheet. The user creates a spreadsheet in order to fulfil some function, for example to make inferences about some real world artefacts. The sheet is developed based on their perception of these observables and evolves as the user interacts with the spreadsheet and uses it to discover more about the situation. Changes to the spreadsheet are not viewed as specification of an entirely new sheet, rather they are perceived as an expression of a change in our views. Construction of a spreadsheet is simultaneous and intertwined with its use; changes are made dynamically as insights occur. In this respect each model is an arbitrary point in a continuum of models where the association between one model and its neighbours is established on the basis of how we understand the model to work. (Roe 2003)

2 Modelling Study

This section discusses the development of the Greedy model which accompanies this paper and the issues that became apparent during the modelling process.

2.1 Aim of the model

The aim of the model was to teach the users the mathematical concepts of the 'Making Change Problem' greedy algorithm. In the main this model assumes the user is already well aware of the real world artefacts, and aims to exploit these private experiences and formalise them by teaching a related abstract mathematical principle. The learner can only progress to using symbolic representations meaningfully when they have a degree of experience gained through interaction with the real domain. Of course the model could be used to teach children how to make change in the first place, but this paper concentrates on the educational benefit to those who already know how to make change.

2.2 The Greedy Model

As introduced above the development of a model follows an evolutionary process, starting with a basic model with limited interactivity and culminating

in a sophisticated interactive model with a derived set of learning outcomes.

At the end of each major change the model was reviewed by two users; Sean (age 10) and Hannah (age 7). Further developments were made based on observations of them using the model and their feedback.

2.2.1 Initial analysis

Experiments performed during the early stages of modelling an artefact are tentative and exploratory; they are a reflection of our provisional construal. The simple aim of the first study was to create a model which would list a set of values and use the greedy algorithm to create a given total. In this respect during the early stages of development the model was a complete abstraction of the observables as set of values. When the user presses a button the model simply selects the best 'coin' until it reaches the total

The algorithm used was :

$$\text{Next coin} = \text{Highest valued coin} \leq (\text{target} - \text{total coins already selected})$$

Figure 2: The Greedy Algorithm

These early studies introduced dependencies between the total the user was trying to collect, the sum of coins already selected, and the total left to collect. Although the user was not able to select the coin themselves in the early stages, the model still had educational value as the users were able to predict which value the model would select, if correct their understanding would be reinforced, if incorrect they would be able to try again to understand why they were wrong.

User tests at the end of Study 1 showed that although Hannah had grasped the mathematical concept of the algorithm her mental arithmetic was not reliable enough to trust that she calculated the total amount left to collect correctly. This meant that she did not always select the most appropriate coin and her understanding was not reinforced. As the aim of the model was not necessarily the ability to teach mental arithmetic, but rather the understanding of the principles, it was felt that certain users could benefit from a label showing how much there change still needed to be chosen, and this was included in the second study. However, where the user can be relied upon to perform the mental arithmetic correctly the teacher may wish for them to do it themselves. A later study recognised

this by making this label dependent on a boolean value. When this value was true the model showed the amount that still needed to be collected, but when false did not. The teacher would be able to customise the model by changing the value in the TkEden window depending on the desired difficulty and learning outcome.

Study 3 embellished the earlier studies by introducing a set of targets. When the user had calculated the change for the first target they were able to press a button and move on to another. These targets were read in from a file which ensured that the teacher was easily able to change the targets without searching for the appropriate code. In reference to the real world where you are able to see the coins you have already selected this study also included a label showing the values of the coins that had already been selected.



Figure 3 : Interface of Study 4

This study also included a reset button which reset all values. This gave the user more freedom to experiment and make mistakes as they could always try again if the attempt was unsuccessful.

Until this point the coins have been selected automatically and the user has only been able to predict which coin will be picked. Study 4 made significant improvements in the interface by allowing the user to select the coins they feel are appropriate. This was achieved by making each value into a button, selecting the button added the corresponding value to the current change.

2.2.2 Improving real world references

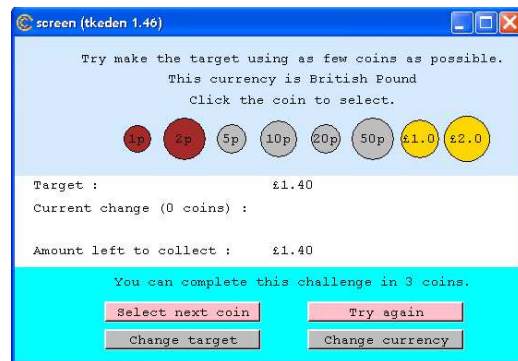


Figure 4 : Interface of Study 7

As previously mentioned the values used in the model have so far been abstract with no reference to other observable features of the coins, for example currency, size and shape. It was decided that the user's experiential knowledge of the making change problem could be better exploited by making the model more similar to the real world artefacts. Stage 5 improved the user interface by creating a set of coin symbols.

One drawback of the user having an in-depth understanding of the problem is that they may not consider it closely enough to understand the new concepts the model is introducing. This problem was addressed by enabling the user to change the currency. By including different denominations the model is still exploiting the users mathematical of making change, however it now requires them to stop and think about the mathematics in greater detail. When study 5 was tested Hannah found the model more interesting and commented on the different denominations of Dollars. Sean noticeably found the calculations in Dollars more difficult and was required to think about his actions in greater depth.

Another great benefit of a variety of currencies was the ability to introduce pre-decimal British Currency. When using the greedy algorithm on 'old money' one does not always derive the solution with the minimum number of coins. Introducing this currency made the model more interesting as Sean and Hannah were unaware of the 'old money' and seeing the algorithm fail increased their understanding of the algorithm's limitations.

In an earlier study Hannah observed the values of the model and questioned the existence of the "100p coin?" Study 7 addressed this by introducing a function to format each currency with the correct denomination. This function was stored in the currency definitions so was redefined whenever the currency was changed.

By this stage there was a clear link between the users experiences of giving change and the actions of the model.

2.2.3 Derived learning outcomes

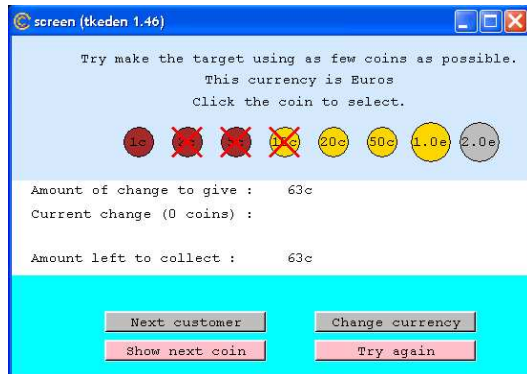


Figure 5 : Interface of Study 9

The next stage of development limited the number of coins available which allows the teacher to contrive situations where the greedy algorithm fails. Seeing the greedy algorithm fail helps to show its limitations and improve understanding. For example if you use the greedy algorithm to create 63 when you cannot use 10, 5 or 2 uses 13 coins ($50+13*1$), however it is possible to only use 6 coins ($3*20+3*3$) if you do not apply the greedy algorithm.

2.2.4 Study 4a - Change the algorithm

The abstract nature of the early studies facilitates further development as they are not limited with irrelevant details. Study 4a emphasised this by revisiting the Study 4 and changing the model to aid in the understanding of the Knapsack problem. This problem aims to select a set of items such that their value is maximised but the total weight of the items is less than a specified limit. This was achieved by redefining the 'currency' to be a set of pairs of values and weights. It involved limited redefinition of the interface.

2.3 Worksheet

After the final stage of development a worksheet was created to lead users through a series of contrived examples to help them learn features of the algorithm. The worksheet introduced the scenario that the user is giving change in a shop and clarified

the aim of the model to use as few coins as possible.¹

Values to pay and coin limitations were contrived to allow the user to learn interesting things about the behaviour of the algorithm. Also by limiting the coin supply of certain denominations or changing the denominations themselves (e.g. to pre-decimal currency, or American currency) users can investigate the impact of the particular denominations in making the change.

Conclusion

The paper highlighted the benefits of the empirical modelling approach to developing tools for educational use though the discussion of the implementation of the 'Greedy' model.

It has been shown that the incremental development process lends itself well to the implementation of such models due to the extensive scope for customisation and extension. In particular the teacher has the ability adapt the model to satisfy the particular needs of their students. Overall the system was developed to provide an existing framework for teachers to use but a relatively simple interface is provided for changing the currencies, targets and availability of coins in order to allow the teacher to derive problems that are appropriate to the curriculum. The final stages of development were concerned with allowing a teacher with limited computer programming experience to be able to easily customise the model. Throughout development any code that can be customised by the teacher, e.g. targets and currencies, has been kept in a separate file which makes it easy to recognise what is modifiable. The system allows flexible adaptation, extension and re-use by users who aren't computer specialists.

The ability to customise the model has also been shown in the development of Study 4a which changed the aim from learning the mathematical principle of the 'Making Change Problem' to the principle of the 'Knapsack problem'.

Acknowledgements

¹ The worksheet and the description of the contrived problems can be found with the source code of the model.

I'd like to thank Sean and Hannah for their patience in the repeated testing of the models discussed in this paper.

References

S. Alessi, S Trollip (1991) *Multimedia for Learning*, Allun and Bacon, 3rd ed.

W.M. Beynon (1997) *Empirical modelling for educational technology*. Available at <http://www.dcs.warwick.ac.uk/research/modelling/papers/reports.html>

W.M. Beynon (2001) *Liberating the computer arts*. Available at <http://www.dcs.warwick.ac.uk/research/modelling/papers/reports.html>

W.M.Beynon.(2003) *Radical Empiricism, Empirical Modelling and the nature of knowing*. Available at <http://www.dcs.warwick.ac.uk/research/modelling/papers/reports.html>

A. Cronick (2003) *Learning through Empirical Modelling*. Available at <http://empublib.dcs.warwick.ac.uk/projects/fraction/sCronick2003/>

C. Roe (2003) *Computers for Learning: An Empirical Modelling perspective*, PhD thesis, Department of Computer Science, University of Warwick, UK.