

Virtual worlds for teaching statistics

Dr Michael Bulmer
Department of Mathematics
The University of Queensland
Queensland 4072
Australia

m.bulmer@uq.edu.au

Introduction

Data analysis and statistical reasoning are an integral component of modern scientific and medical research. However, students typically take a statistics course in the first or second year of their programs, at which stage it is not at all clear from their other courses why statistics might be useful to them. This can make teaching such a statistics course a challenge. It is not helped by the fact that most students have seen statistics before as part of their mathematics at secondary school, and take from this an obsession with mathematical detail that is not required at the tertiary level (Higgins 1999).

The ideal approach to get students aware of the holistic nature of statistical investigation is to have them carry out studies from beginning to end. They should design experiments or surveys, perform them, explore and summarise the results, and then make appropriate statistical conclusions. This approach is feasible for in-class work in small classes or for project work in larger classes. Mackisack (1994) gives an overview of the other benefits of experimental work, including the appreciation of the practical issues involved in carrying out experiments and collecting data. However, it is difficult to supervise, guide, and discuss these activities in large classes, primarily because of resource and time limitations.

To address this issue, a collection of 'virtual worlds' has been developed to efficiently give students a scientific context for their statistic work, with some control over the experiment itself. The focus here is on the life sciences but the principles are equally useful in other applications of statistics, as well as in applications of general mathematical modeling.

Virtual rats

Most introductory statistics textbooks are rich with real data sets, allowing students to relate the results of their explorations and analyses back to real scenarios. This is certainly desirable. Cobb and Moore (1997) suggest that 'statistics requires a different *kind* of thinking [to mathematics], because data are not just numbers, they are numbers with a context'. However, it is still a somewhat passive experience because the students have to take the context for granted. They have not been involved in obtaining the data and so lack ownership of the setting. The first aim of using virtual worlds is to engage the students in thinking about the design of the experiment and the origin of the variability in the data, while not allowing this to be so time consuming that the rest of their learning suffers.

As an example of a virtual experiment, consider a setting described by Moore and McCabe (1998) of an experiment for exploring the effects of calcium and magnesium on blood pressure. Three levels of each mineral in the diet of rats were considered, giving nine possible treatments to try. A standard textbook exercise would give the resulting data, or simply the summary statistics from the nine treatment groups, and then use questions to have the students visualise the data or test for main effects and interactions using ANOVA.

A simple interface for a virtual version of this experiment is given in Figure 1. Here the user can specify the calcium and magnesium levels (each Low, Medium, or High) for a particular rat and then click the measure button to find out its blood pressure after the treatment. Clicking again returns the blood pressure for another rat. An appreciation for the effect of the treatments and of the variability of results

can be obtained easily in this simulation; carrying out this actual experiment would be impractical in a statistics course, especially with a large first-level class. One student commented ‘it was an interesting way of collecting the data rather than simply being given data to work with’. Another noted that ‘it was a good way to investigate the effects on rats without using actual rats’, suggesting that they really identified with the virtual setting, despite the simplicity of the interface.

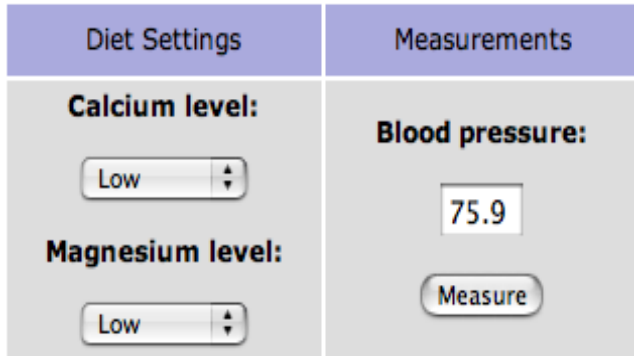


Figure 1. Interface to virtual rats

The virtual rats are implemented using *JavaScript*, requiring only a few lines of code to specify a matrix of means and standard deviations for each combination and then randomly generating a blood pressure when requested. *JavaScript* was chosen over alternatives, such as Java applets or scripts running on the web server, because it was easy to develop and maintain. A new virtual experiment can be made readily by copying the web page and changing the appropriate text, numbers, and the underlying statistical model for the response. All textbook data sets used in this course will soon have a virtual equivalent on the course web page. There is then a limitless range of data for students to practice with, all under their experimental control.

A second reason for adopting *JavaScript* is that it is easily visible to any interested student through their web browser. Students can look deeper at how the blood pressures are being generated and by so doing gain an insight into the nature of statistical models.

Virtual plants

The virtual rats have proven successful in getting students to think about statistical issues beyond the mechanical techniques that are typically emphasised. However, they leave out the measurement process. The way the measurements are carried out, with possible errors and biases, can have profound effects on the statistical analyses. It is not clear, for example, whether the variability that the students observe in their virtual blood pressures is coming from differences in the rats or from errors in the measurements.

Simulating the measuring of blood pressure on rats is a difficult activity to capture in a concrete manner. An alternative setting was created using virtual plants, generated using L-systems (Prusinkiewicz and Lindenmayer 1990). An L-system was developed for a field

of 12 plants using the *L-Studio* software, based on the flowering *Lychnis* plant that came as an example with the software. To allow a similar approach to the virtual rats, a randomised growth parameter was introduced to govern the overall growth of each plant. The 12 plants were treated as two plots of 6 plants, and with 20 growth levels for each plot this gave 400 possible growth scenarios.

Simple movies were initially produced from screen dumps of the plant growth. These had a parallel projection and allowed height to be easily measured. However, the screen graphics were not clear enough to see high detail, such as was needed for counting the number of leaves. This was partly useful, keeping the measurement process difficult, but it did not reflect reality. A new set of virtual plants was developed using ray-traced graphics to give much clearer images of the plants. One disadvantage of this is that it uses a perspective projection, making measurements of height more difficult. This was offset by the inclusion of shadows that were more perpendicular to the student.

The plants have been used in a similar way to the rats. The interface, shown in Figure 2, has two factors that can be controlled. Four levels of nitrogen fertilizer can be specified (None, Low, Moderate, and High) along with three levels of irrigation (None, Some, and Lots).

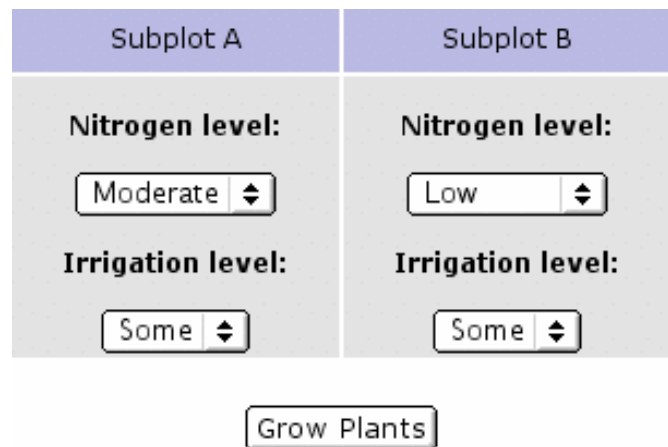


Figure 2. Plant growth experiment

However, rather than being able to obtain a series of measurements by clicking a button, the user instead gets an animation that shows the growth of the field. Since the ray tracing was time-consuming, plant growth animations were delivered using pre-generated *QuickTime* movies. Treatment combinations were then mapped onto one of the possible 20 growth levels in each plot to determine which movie to display. Further details about the creation of the plant growth movies are given in Bulmer (2003).

Figure 3 shows one of the *QuickTime* movies, based on the settings in Figure 2; all plants received some water, but the plants at the rear had moderate fertilizer while the plants at the front had low fertilizer. A student noted that this was ‘better than [the rats practical]—could actually see the results in picture form which gave a better idea of what happened.’

[Click Here](#) to view video animation



Figure 3. Example animation of virtual plants

In dealing with the measurement task, students needed to start by thinking about why they might be measuring plant growth. They could measure the heights of plants if they wanted to see which treatment gave taller plants, or they could count branches, leaves, or flowers if they wanted to see which gave higher yields.

Measurements of height can be made from the screen using a plastic ruler. It is quite a satisfying experience to walk into a computer laboratory and see a room full of students with rulers up to the screens. They are physically engaging with the setting, rather than passively taking a set of mysterious numbers from a textbook exercise, or even from a real study.

Measuring height with a ruler is difficult because it is hard to know where the top of a plant is; they branch outwards after an initial vertical growth. Counting flowers or leaves is difficult because it is hard to know you've seen them all, just as in real life. The measurement process should be difficult and students should have to think about what simplifications or estimates they are making. One student complained that 'it would have been more helpful and easier to use if some kind of variable was presented with the movie as a result [rather than making student measure it off the screen]' but most saw the purpose of the exercise in that it 'made it more interesting than copying information out of a data set'. Bulmer (2002) gives additional discussion of the evaluation of the virtual plants.

As with the JavaScript for the rats, using L-systems as the basis of the virtual plants gave the opportunity to broaden the understanding of curious students.

The University of Queensland hosts the Centre for Plant Architecture Informatics, a group that uses such models for a range of applications, and students could look at their web pages to find out more about this area of contemporary mathematics.

Virtual humans

While the virtual rats and plants have been highly successful in their aims, there are many students in the course whose interest lies in human biology or medicine. To cater for these students, a 'Virtual Human' project is currently in development. This aims to give similar settings to the plants for statistical investigations, but is also in collaboration with the School of Human Movement Studies that will use this media for non-statistical work in laboratories and lectures.

Of course creating a suite of virtual humans is a more challenging task than growing virtual plants. It is simple to make a field of flowers that look pretty and exhibit different growth responses without having to have the model perfectly match the real world. The L-Systems are ideal for this task, giving a structural model that can be controlled through simple scripting. Creating a 2D or 3D model of a human that is convincing to a user is more difficult since everybody knows about humans. An initial plan was to overcome this by stepping back and simply generating stick figures that allowed measurement of the necessary attributes, similar in essence to a basic L-System. However, through an Apple University Development Fund grant the resources became available to create more sophisticated models and ray-traced graphics of humans.



Figure 4. Rafat and George, two virtual humans

Figure 4 shows two virtual males produced as part of the project. These come from models created in *Poser*, a software package specifically designed for making 3D models of humans and then rendering them as images or animations. Most importantly, *Poser* is scriptable. A collection of 6250 rendered human images has now been generated, based on parameters such as height, shoulder width, hip width, and arm girth.

The parameter used for creating these virtual humans are a subset of the measurements made by Human Movement Studies students about themselves in one their first-year laboratory sessions on anthropometry. An initial use of these images has been to create a 'virtual class', where each student has a virtual counterpart based on their own personal body measurements. The lecturer for that course can then use this virtual class for illustrations, allowing students to identify directly with the content of lectures.

Discussion

The key ingredient missing for the virtual humans is a functional aspect that would be able to model basic processes such as heart rate, blood pressure, and so on. As a result their use is currently limited to the same kinds of measurement tasks as plants. Simple anthropometric studies, such as looking for relationships between body dimensions, are straightforward. Experimental scenarios are more difficult to motivate. An initial statistical laboratory involves a comparison between children who received a new growth hormone and a control group, though this is rather contrived. More sophisticated virtual studies relating to population genetics are being developed. However, to capture interesting medical experiments these images would need to be augmented by other visualisations of function.

There is a trade-off here though. The original aim of the virtual worlds was to efficiently allow students to have some control over experimental design and data collection. Adding further features to the interface may have the negative effect of focusing the students too much on the virtual reality at the expense of thinking about the statistical issues. In this way the virtual worlds are not intended to be simulation models. Despite all the effort that has gone into

preparing them, the intention is that they would only be used for 5–10 minutes at the start of a 1 hour statistics practical. It is certainly desirable to have graphics that are attractive and engaging, but once engaged the students should move on to the learning activity at hand.

Acknowledgements

I would like to acknowledge the assistance of Jim Hanan (Centre for Plant Architecture Informatics) and Craig Engstrom (School of Human Movement Studies) in developing the virtual models, and the Apple University Development Fund for financial assistance with the virtual human project.

References

- Bulmer, M. (2002) Teaching and Learning Mathematics with Virtual Worlds. *Proceedings of the 2nd International Conference on the Teaching of Mathematics (at the undergraduate level)*. [Online] Available: <http://www.math.uoc.gr/~ictm2/Proceedings/pap42.pdf>
- Bulmer, M. (2003) Growing virtual plants for teaching and learning statistics. *Proceedings of the Apple University Consortium Academic and Developers Conference, Adelaide*, 3-1 – 3-7.
- Cobb, G.W. and Moore, D.S. (1997) Mathematics, statistics, and teaching. *American Mathematical Monthly*, **104** (9), 801-823.
- Higgins, J.J. (1999) Nonmathematical statistics: a new direction for the undergraduate discipline. *The American Statistician*, **53** (1), 1-6.
- Mackisack, M. (1994) What is the use of experiments conducted by statistics students? *Journal of Statistics Education*, **2** (1) [Online] Available: <http://www.amstat.org/publications/jse/>
- Moore, D.S. and McCabe, G.P. (1998) *Introduction to the Practice of Statistics* (3rd ed.). New York: W.H. Freeman
- Prusinkiewicz, P. and Lindenmayer, A. (1990) *The Algorithmic Beauty of Plants*, New York: Springer-Verlag.