

MUSIC FROM FRACTAL NOISE

Michael Bulmer
University of Queensland
<mrbl@maths.uq.edu.au>

There are many interesting connections between music and mathematics, though these are rarely used when teaching maths in schools. In this paper we look at one example which involves using musical motivations to introduce some mathematical ideas. The aim is to develop a random method for making music which produces pleasant results.

INTRODUCTION

Folklore has it that music and mathematics are somehow related as human skills. However when it comes to teaching mathematics we rarely use examples from music, preferring more “practical” applications from physics, engineering, or finance. This is perhaps a shame since it gives little credit to the artistic side of mathematics.

In this paper we will look at a short teaching activity that introduces a variety of ideas, such as autocorrelation, through the goal of creating pleasant random music. While having an obvious conclusion, this is very much an open-ended task, reflecting the desire to encourage an ongoing artistic appreciation of mathematics.

The materials required for this activity are some dice and an instrument for listening to the generated music. Alternatively, a computer can be used to simulate the dice rolls or to play the music. Some experience with statistical measures will be useful for the mathematical aspects. The activity also works well after a discussion of the more traditional kinds of fractals.

RANDOM MUSIC

This activity should start with some discussion and brainstorming about how music could be created randomly and the kind of properties that such music should have. There can also be some discussion about how maths might help in this task. Below we describe three methods for creating music. Each is based on some kind of *noise*, a random process in time. Students may well come up with other methods that go beyond these.

White Noise

One of the easiest methods is to generate notes one at a time using dice. Suppose we have 6 dice, so when we roll them together and add up the results we get a number between 6 and 36. For each of these possibilities, assign it a note of some music scale. For example, 6 could be the C below middle C, 7 could be the following D, 8 the E, and so on up to 36. This is a standard major scale, but you can also use other scales or modes, as described by Kandell (1984). You can write the result on normal music paper, or just make a plot of the raw numbers.

If you listen to this music it will sound pretty bad, almost like static that has been slowed down. This noise is termed *white* because of this. Figure 1 shows a plot of 256 notes generated by this method.




Figure 1. Time plot of white noise

After you have generated a sequence of pitches, you can also generate a sequence of durations for your notes using a similar method.

Brown Noise

An obvious problem with white noise is that there is no connection between successive notes. To overcome this failing, students may suggest methods that result in various kinds of *brown noise*. This noise gets its name because it corresponds to the random walks of physical Brownian motion. The standard example is of a drunk who staggers randomly back and forth, sometimes moving a bit in one direction and sometimes moving a bit in the other direction.

For example, we might start our music at middle C. To make each new note we roll a die. If the die comes up with a 1 then we go down two notes from where we are; if it comes up 2 we go down one note; for 3 we go up one note; for 4 we go up two notes; and for 5 and 6 we stay where we are. Figure 2 gives an example of noise generated by this method.




Figure 2. Time plot of brown noise

Again, you can use a similar process to produce lengths of notes. Brown music is less painful to the ear but is still rather boring.

Pink Noise

White noise and brown noise can be seen as two extremes for random music. In white noise there is no association between successive notes while for brown noise there is a very strong association. White noise is dull because it is too unpredictable but brown

noise is also dull because it is too predictable. Interestingly, neither is predictable in the long run.

Traditional music, on the other hand, seems to achieve a balance between these extremes. A composer might sit down with a vision for a whole piece of music, devise finer structure for smaller sections, and then write the notes for each section. This gives the pattern over time a *long-range dependence* while still involving short-term randomness. This construction is reminiscent of the construction in the plane of the Koch snowflake, as described in Mandelbrot (1982). The first four steps in making the snowflake are shown in Figure 3. The result is an object which is *self-similar*, possessing similar structures as you look closer and closer at it. Such objects, whether they are in space, like the snowflake, or in time, like musical notes, are called *fractals*.




Figure 3. Creating a Koch snowflake

Neither white nor brown noise has the long-range dependence or self-similarity required by “nice” music. We will create a type of noise that lies in between the extremes of white and brown, sometimes known as *pink* noise (or $1/f$ noise, as described by Voss and Clarke (1978)).

To achieve true pink noise is actually very difficult, as described by Mandelbrot (1971). However, Gardner (1978) describes a simple method, invented by Richard Voss, which approximates pink noise and which is easy enough for students to both carry out and understand. Using n dice, this method will generate 2^n notes. We will illustrate it here with 3 dice, labelled A, B, and C. Make a list, as in Table 1, of the numbers from 0 up to $2^n - 1$ with their binary representations. Start by rolling all three dice and adding up the results to give the first note (note 0). To generate each subsequent note, look at the binary digits that change from row to row in the table. For example, when moving from note 0 to note 1, the C digit changes while the A and B digits stay the same. Follow this by rolling the C die again while leaving the A and B dice as they were. Add up the three results to give note 1. To get note 2, roll both B and C but leave A alone. Continue doing this until all 8 notes have been generated.

Note	A	B	C
0	0	0	0
1	0	0	1
2	0	1	0
3	0	1	1
4	1	0	0
5	1	0	1
6	1	1	0
7	1	1	1

It is clear that this method will give a series of notes that exhibit long-range dependence. The higher digits change less frequently and so the corresponding dice provide a long-term stability in the sequence of notes. Compare Figure 4, showing an example of pink noise created with 8 dice, with the pictures in Figures 1 and 2.




Figure 4. Time plot of pink noise

DESCRIBING NOISE

So far we have motivated and described the three colours of noise in general terms. Of course we can also use mathematics to help explore these ideas more concretely. In particular, the notion of musical notes being “related” over time is captured by the definition of *autocorrelation*. More advanced students can also use a variance calculation to try and determine the fractal dimension of our pink noise.

Autocorrelation

This activity requires an understanding of the standard correlation coefficient between two variables. This is an easy idea to introduce in isolation and can be motivated by looking at real data sets. For example, you could get the class to measure the lengths of their feet and their heights, display the data, and then look at the calculated correlation.

Autocorrelation uses the same calculation, measuring the correlation coefficient r_k between values in the noise sequence and the values k time points ahead. (Write this down as a usual data set to convey the idea.) This number k is called the *lag*. For example, the lag would have no effect for white noise.

Since the two data sets in autocorrelation are actually the same, the standard correlation formula can be simplified to the following, as described by Chatfield (1996):

$$r_k = \frac{\sum_{t=1}^{N-k} (x_t - \bar{x})(x_{t+k} - \bar{x})}{\sum_{t=1}^N (x_t - \bar{x})^2}$$

To visual the autocorrelation structure of noise, students can make a correlogram, a plot of autocorrelation against lag. This realistically requires a computer; even if students cannot make the plots themselves there is still much room for discussing the plots.

The correlograms for white and brown noises follow patterns that are easy to guess. Students can be encouraged to sketch the pattern they would expect to see beforehand. The white noise correlogram in Figure 5 captures the fact that there really is no association between values in the sequence, giving autocorrelations close to 0. Figure 6 shows the correlogram for brown noise. The “random walk” nature of the noise means that values close together will be highly correlated. As time increases values wander away from each other and correlation declines, ultimately tending to 0.




Figure 5. Correlogram for white noise




Figure 6. Correlogram for brown noise

The picture for pink noise in Figure 7 shows what we would like, moderate correlation over the short term which do not disappear to 0 over the long term.




Figure 7. Correlogram for pink noise

The patterns in these correlograms are more easily visible when longer runs of values are generated. (A run of length 1024 works well, requiring 10 dice for the pink noise.)

Variance Plots

The correlograms described above give an intuitive feel for the correlation structure of the different types of noise, requiring only a minimal background in statistical ideas when presented in conjunction with a discussion of correlation. However, the activity can be extended if students know the important rule for the effect of sample size on the variability of the average :

$$\sigma_{\bar{x}}^2 \propto \frac{1}{n}$$

Unlike correlation, it is probably unwise to use this activity to introduce this activity since it turns out not to always hold! To see this, think of the run of values as a series of samples of size m , for each of which we can calculate the mean. That is, for the sequence x_1, x_2, \dots , calculate the averages $(x_1 + \dots + x_m)/m$, $(x_{m+1} + \dots + x_{2m})/m$, \dots , and then calculate the sample variance of these numbers, $\text{Var}(m)$. (This is certainly something that is better suited to a computer!) Repeat this for a range of m values, say 2 to 30, and then plot a graph of $\log(\text{Var}(m))$ against $\log(m)$.

If the standard rule held we would expect to see a line of points with slope equal to -1 . Figure 8 shows this plot for the white noise, where the least-squares line through the points is $1.29 - 1.07x$, a slope very close to -1 . This is not surprising since in white noise there is no association between adjacent values and so the consecutive samples really are independent.




Figure 8. Variance plot for white noise

The picture for brown noise, shown in Figure 9, is quite different. Here the variability of the sample average seems to be independent of the sample size! This happens because the variability of the values in a sample increases as the sample size increases, since the random walk can cover more ground, and this perfectly cancels with the decrease in variability from having a larger sample.




Figure 9. Variance plot for brown noise

The Koch snowflake is made up of a line (an object of dimension 1) which seems to fill up the plane (dimension 2). Thus it is said to have a *fractional dimension*, somewhere between 1 and 2 (in fact it is 1.26). The analogous fractal dimension of a noise is given by the Hurst parameter, H . If β is the slope of the line in the variance plot, then H is equal to $1 + \beta/2$. White noise has $H = 0.5$ while brown noise has $H = 1.0$. As is to be expected, the parameter is somewhere in between for pink noise. Figure 10 shows the variance plot for pink noise, giving the least-squares line $1.40 - 0.29x$. Thus the pink noise has a fractal dimension of $H = 0.86$.




Figure 10. Variance plot for pink noise

EXTENSIONS

One of the great advantages in involving an obviously creative area such as music in mathematics teaching is that it leads immediately to many extensions. Students are naturally keen to come up with improvements on the basic algorithm which make the resulting music more pleasant to the ear. These improvements could be mathematically motivated, such as trying to change the autocorrelation structure. They could also be

musically oriented, such as changing the scale used to produce modal music or introducing a second instrument.

There are also extensions away from the original musical task, looking at other time series to develop an understanding of the kind of patterns that might emerge. Students could be set a project of making some measurement over time and then discussing the behaviour that they observe. Stock prices typically give brown noise while anything to do with independent observations will give white noise. Mandelbrot and Wallis (1968) originally found pink noise when looking at the flooding patterns of rivers. Other patterns observed may motivate a general discussion of time series analysis, including seasonal variation and trend.

The study and use of pink noise and $1/f$ phenomena is currently of broad interest. For example, traffic on the Internet exhibits the long-range dependence of pink noise that current models do not properly capture; Jeong et. al (1999) give an overview of the role of such self-similar noise in teletraffic research. There is an extensive bibliography of other $1/f$ phenomena, in such areas as astronomy, ecology, economics, electronics, and DNA sequences, on the web at <http://linkage.rockefeller.edu/wli/1fnoise>.

Technical Notes

The fractal noise used in this paper was generated by simulating the structured dice rolls in Mathematica. This could have been done easily in almost any programming language. To listen to the resulting music the numbers were converted into a standard MIDI file using the Perl package MIDI-Perl by Sean Burke. This was then imported into the QuickTime Player on the Macintosh and played. The instrument sounds available in QuickTime are lovely, but again there are many other simpler ways to generate notes on a computer.

REFERENCES

- Chatfield, C. (1996), *The Analysis of Time Series: An Introduction* (4th ed.). London: Chapman & Hall.
- Gardner, M. (1978), White and brown music, fractal curves and one-over- f fluctuations. *Scientific American*, April, 16–31.
- Jeong, H.D.J., McNickle, D., & Pawlikowski, K. (1999), Fast Self-Similar Teletraffic Generation Based on FGN and Wavelets. *Proceedings of IEEE International Conference on Networks (ICON'99)*, Brisbane, Australia.
- Kandell, J. (1984), Computer Music Worth Listening To. *inCider*, February, 68–72.
- Mandelbrot, B. (1971), A Fast Fractional Gaussian Noise Generator. *Water Resources Research*, 7, 543–553.
- Mandelbrot, B. (1982), *The Fractal Geometry of Nature*. New York: W.H. Freeman and Company.
- Mandelbrot, B., & Wallis, J.R. (1968), Noah, Joseph and operational hydrology. *Water Resources Research*, 4, 909–918.
- Voss, R.F., & Clarke, J. (1978), “ $1/f$ noise” in music: Music from $1/f$ noise. *Journal of the Acoustical Society of America*, 63 (1), 258–263.