Sonification of depictions of numerical data

Karen Vines Chris Hughes Hilary Holmes Vic Pearson
Claire Kotecki Laura Alexander Chetz Colwell Kaela Parks*

December 8, 2016

Executive Summary

Sonifications are non-verbal representation of plots or graphs. This report details the results of an eSTEeM-funded project to investigate the efficacy of sonifications when presented to participants in study-like activities. We are grateful to eSTEeM for their support and funding throughout the project.

Contents

1 Introduction

The depiction of numerical data using graphs and charts play a vital part in many STEM modules. As Tufte says in a key text about the design of plots and charts "at their best graphics are instruments for reasoning about quantitative measurement" [29]. In this report we will focus on static images, such as graphs and plots in printed materials. Dynamic images in which the user can change features of the diagram or graph are not considered. In order to meet the OU's mission of being open to [all]

---

*Kaela Parks is Director of Disability Services at Portland Community College, Portland OR, USA; every other author listed here works at The Open University, UK.

people, such plots and graphs need to be accessible to *all* students, as some students may otherwise be disadvantaged in their study.

The Equality Act 2010 [8] requires universities to avoid discrimination against people with protected characteristics, including disability, and to do so by making 'reasonable adjustments'. The Equality and Human Rights Commission offers guidance for Higher Education providers [9]. The Act created the Public Sector Equality Duty [10], which requires universities to promote equality of opportunity by removing disadvantage and meet the needs of protected groups. In the context of The Open University, this means that the authors of module materials should ensure that plots and charts (or alternate versions of them) are accessible to all students with visual impairments, including those students with no vision at all.

Individuals who are blind have often had limited access to mathematics and science [6], especially in distance learning courses [21]; in part this is because of the highly visual nature of the representations of numerical relationships. Methods commonly used to accommodate learners who are blind or have low vision include: use of sighted assistants who can describe graphics verbally; provision of text-based descriptions of graphics that can be read with text-to-speech applications (for example JAWS [22], Dolphin [23]); accessed as Braille, either as hard copy or via refreshable display, or through provision of tactile graphics for visual representations.

Desirable features of an accessible graph include the following [28]:

- Perceptual precision: the representation allows the user to interpret the plot with an appropriate amount of detail.
- First-hand access: the representation allows the user to directly interpret the data and is not reliant on subject interpretation by others (bias).
- Works on affordable, mainstream hardware.
- Born-accessible: the creator of the plot would not have to put extra effort into creating the accessible version.

The first two of these features aim to provide the user of an accessible graph with an *equally effective* alternative to a 'standard' graph; the representation should give the user a "true" picture of the data depicted. The user should be able to form their own opinion about the data without it being mediated through another's (possibly subjective) view. The remaining two criteria are about the practicality of the solution: it does not require the user to have access to specialist equipment and that, as far as possible, it does not require the creator of the original plot to do anything 'special' to create the accessible alternate version.

Arguably the first two features, with their emphasis on effectiveness, are far more important than the other two features. An *effective* accessible graph that is expensive to produce and use might be worth producing, whereas a cheap accessible graph which leaves the user with *no information* about the original plot or chart certainly won't be worth it. Born-accessibility is important as it enables authors to produce an accessible version of a plot or graph easily and quickly.

Another important desirable feature, not given by [28] is as follows.

- Quick: the representation allows the user to quickly extract the information required.

When module materials are designed, the time required for an average sighted student to digest plots and figures is factored into the estimation of the student workload [25]. If the accessible alternative slows down the rate at which the information in a plot or graph can be assimilated, this could have a significant impact on the workload for students reliant on them. It could also distort the flow of the material presented, thereby impeding learning.

Figure descriptions

Within the OU, figure descriptions are the standard way to provide accessible alternate versions of plots and graphs [5]. A figure description is a written description of the graph or plot; such descriptions are intended to convey the essential information contained in the graph whilst being concise. By converting a plot or graph into a plain text representation, it only requires users to have access to a screen reader; it is less clear, however, that figure descriptions have the other three desirable features of an accessible graph.

Figure descriptions are not currently born-accessible as they have to be written as a separate, manual task by the author, after the graph or plot has been created. There has been research into the automatic writing of figure descriptions [31], but this does not appear to have come into mainstream use even within the OU. Certainly figure descriptions do not provide first-hand access to the data –

instead access is mediated through the author of the description. The exercise is often subjective, as 'important' features may well be ambiguous; for example whether a line connecting data sufficiently conforms to a mathematical shape and if so which one, or whether particular data points appear to be sufficiently different to the rest so as to be noteworthy. Furthermore, the author may have to be trained in order to write effective figure descriptions, which clearly adds time to the production of the figure.

The need to provide users of a figure description with a concise, yet rich enough, description can be a difficult, if not impossible, task. Perhaps more importantly (from a teaching point of view), the requirements of a useful figure description can conflict with learning outcomes. For example, consider the following learning outcome: "recognise patterns in a residual plot". When a residual plot only contains a few points it is possible to detail where all these points lie on the plot, and hence allow the reader of the figure description to form their own opinion of what the pattern might be (or whether there is no pattern). The task of picking out patterns in data can be transformed from one using a plot to one using a table of data. With more points on a residual plot the description then has to include detail about the pattern they form; but how can a student demonstrate, let alone learn, this skill if someone else's opinion about the pattern is always given to them? Furthermore, long figure descriptions can be very time consuming for the student to use, and may not satisfy the 'quick' criteria.

Tactile graphs Tactile graphics can be created with some models of Braille embosser, or with swell paper that produces raised-line drawings by exposing the capsuled paper to heat, or through lower tech methods such as collage. There are standards for the creation of tactile graphics as well as opportunities to purchase already-created sets of instructional materials.

To be effective, tactile graphics cannot simply be given to a person who is blind with the expectation that meaning-making will commence; rather, there tends to be a need to verify existing knowledge and describe the tactile in a way that leverages existing understanding of the information being depicted [26]. At that point the tactile graphics can increase understanding. Producing tactile graphics and making them available for the person who is meant to access them takes time; this means, in the case of a distance learning student, the tactile graphics potentially need to be created at one location, then sent by mail to another location, with interactive communication throughout.

Considering tactile graphics against the desirable features of an accessible graph, assuming that the user has the ability to sense through touch, they allow both first-hand access and perceptual precision. Assuming that the *institution* is responsible for the financial investment of the embosser, then from the *user*'s perspective, tactile graphs work on mainstream hardware. The work involved in distributing the tactile graphs to the user would typically be shared between the institution and the author of the graph.

Can tactile versions be accessed quickly? There is typically a calibration or orientation process undertaken by the user (e.g, to find the axis lines). There must also be a verification step to ensure that the user is analysing the appropriate graph: there may well be scenarios in which a user is given multiple tactile graphs, and they first must identify which one to use.

Sonifications Figure descriptions and tactile graphs are not the only means of attempting to produce an accessible version of a plot or graph. An alternative non-verbal representation may also be provided by sonification.

Sonifying data dates back to at least the 1980s [32, 3]; for example, Bly suggested using sound to represent data to allow the listener to try to distinguish differences between datasets. The principle underlying this approach is to map aspects of sound, such as pitch, duration and timbre to different variables. Although the motivation given by Bly was to simultaneously represent more dimensions of the data than can generally be done using most visual representations, the application of this approach to accessibility has been noted by others (see, for example, [24, 7]).

Bly and Yeung were using sonification in a context where printed statistical graphs do not work very well. Sonifications of more traditional plots and graphs have also been proposed; for example, scatter plots [13, 7]; box plots [14, 7]; histograms [15]; pie charts [16]; plots of mathematical functions [19].

The basic principle underlying sonification in this context is to map the various dimensions of the data (for example different variables) to different aspects of sound (for example the pitch and the time when the note is sounded. The different aspects of sound that can be used include:

- when the note is sounded;

- the pitch of the note;
- the loudness of note;
- the timbre of the note; for example, whether it sounds like a plain note, or like an instrument such as a piano or a violin.

(See, for example, [30, 17].)

It is important that the chosen option enables the listener to distinguish different values; to help with perception it is possible to map more than one aspect of sound to a dimension of the data.

1.1 Aims and scope of this study

Sonifications of plots and graphs can meet at least two of the five desirable feature for accessible graphs. The ubiquitous inclusion of sound cards and the requisite software on PCs, tablets and smart phones, means that sonifications can be played on a wide array of media. The sonifications, by translating dimensions of data directly into sound, attempt to provide the same first-hand access to the data that a printed (possibly tactile) plot or graph does. Furthermore this direct translation of data into sound means that, in theory at least, the creation of a sonification could piggyback on the creation of a plot or graph – not adding an additional, manual, task for the author. This just leaves the question of perceptual precision: in the context of providing accessible alternate formats for plots and graphs in OU teaching and assessment material this means whether the sonifications enable the *listener* to achieve the same insights as *viewers* of the corresponding plot and graph can. It is this question that our study set out to answer.

The aim of this project was to explore the potential of sonification in OU STEM modules; we aimed to do so by producing sonifications of depictions of numerical data and examining the extent to which these sonifications were 'effective'; explicitly, with reference to the five desirable criteria of an accessible graph, we are particularly interested in attempting to evaluate users' ability in using sonified versions of graphs: to gain perceptual precision of the graph being portrayed; to have first-hand access to the information in the graph; and to be able to access the version quickly. We focused upon the types of depictions that most-commonly occur in module materials: scatter plots and line charts.

2 Our study

This study was designed to investigate the effectiveness of sonifications (with reference to perceptual precision and first-hand access) in the more general context of student study — in particular level 1 study in mathematics or science, areas that make heavy use of plots and line graphs in their teaching.

Materials from four OU modules were examined for suitable plots:

- MU123 'Discovering mathematics'
- MST124 'Essential Mathematics 1'
- M140 'Introducing statistics'
- S104 'Exploring science'

Of particular interest were plots where it was felt that written descriptions of the plots struggle to provide an adequate alternative for VI students. From these modules a total of six examples were identified as having teaching points based round plots which did not require too much background knowledge.

1. 'Ferris Wheel': exploring the features of a mathematical function. An example taken from MU123.
2. 'Radioactive decay': exploring the impact of radioactive decay on the numbers of atoms. An example taken from S104.
3. 'Radial speed of a star': using observations to investigate a physical phenomenon. An example taken from S104.
4. 'Correlation': using plots to learn about a statistical concept. An example taken from M140.
5. 'PISA': summary of a scatter plot. An example taken from M140.
6. 'Fit of lines': using plots to decide whether a line fits some data. An example taken from M140.

The general structure of each example was designed to mimic activities that include graphs or plots as an integral part; each example started by giving some context in which the plot or plots arose. This required adaptation from the original module materials as we did not expect participants to have completed study of the module up to that point. As a practical note, there was a break (about 20 minutes) for our participants between examples 3 and 4.

2.1 Participants

As the aim of the study is partly to gain information about the usability of the sonifications, it was decided to study a small number of participants in depth. A total of 12 participants were included: 5 sighted OU students currently studying at least one of MU123, MST124, M140 and S104; 5 adults, who were either blind or severely visual impaired, with an interest in mathematics and/or science (but not necessarily OU students); 2 further blind or visually impaired people were selected by Kaela.

The decision to select sighted OU students studying these particular modules, was to capture the efficacy of sonifications for OU students early on in their study journey and hence still in the process of honing their study skills. In particular, it means that they are at the stage where it would be reasonable for students to start learning the skills to utilise sonifications. It also means that the students are currently studying high-population OU modules, which are the modules on which sonifications, if successful, would be first deployed.

The VI adults were chosen to investigate whether there appears to be any difference in the effectiveness in this key group – the group whom it is hoped that sonifications would most benefit. Ideally these participants would have also been current OU students, however it was not clear that sufficient numbers of OU students with severe visual impairments could have been recruited from the current base of OU students studying mathematics and/or science.

In the interest of widening the 'impact' of our study beyond the OU context, we collaborated with Kaela Parks from Portland Community College (USA), who replicated the study with two visually impaired participants; Kaela is the Director of Disability Services at PCC.

We are keen to anonymise our participants, and will detail them in what follows as S1, VI2, . . . , VI12; a brief summary of each is given below. The participants S1, S3, S5, S7, S9 are *current*, sighted OU students; VI2, VI4, VI6, VI8, VI10 are visually impaired or blind, and not necessarily students at the OU; VI11 and VI12 are known to Kaela at Portland Community College, and are visually impaired or blind.

- S1 Started MU123 in October 2015; at time of testing (February 2016) S1 has been through 3 books, with just 1 left; has just started MST124 (16B); does *not* enjoy listening to music, finds it distracting, does not play any instruments (although was encouraged to do so).
- VI2 Sat A-level pure maths & stats 42 years ago; chemistry and biology A-level; went to university to study pharmacy, when sight started to deteriorate; maths was always their favourite subject at school. Has not done much maths since then, but helped daughter with school maths. Does like music, does not play any instruments.
- S3 A-levels in maths & chemistry, got to the end but wasn't really enjoying it that much; joined the RAF as an engineer, electrical avionics; 6 month technical college course with RAF; lot of physics, lot of maths. OU-modules: engineering of the future (T174), discovering mathematics (MU123). Does 'get benefit' from music, but doesn't go out of way to listen to it. Does not play musical instruments.
- VI4 Recently did a maths course at a college in Milton Keynes; feels reasonably proficient; most studying was before they lost their sight. Has radio (Radio 2) on all the time.
- S5 Has biology degree (1973); has taught science for 40 years, part of it included teaching A-level Psychology which included statistics; taught maths 1 year of career, now studying Economics and Mathematical sciences at the OU. Enjoys listening to music a lot; goes to concerts, always has radio on, sings in choir, plays piano.
- VI6 An OU student in Maths and IT, studying MU123 in 15J. Enjoys music, tends not to have it on the whole time; when listening to radio, it tends to be Radio 4. Does not play musical instruments.
- S7 GCSEs 15 years ago, studying S104 in 15J. Enjoys listening to music to a large extent, but generally not while studying.

VI8 Has studied a basic computing module, TU100. Enjoys listening to music but 'can live without it'; does not play musical instruments.

S9 Currently doing three OU maths/science modules; has GCSEs in maths and science; did AS levels then left school. Enjoys listening to music 'quite a lot', sometimes while studying so as to reduce surrounding distractions; used to play flute.

VI10 Did maths and science to GCSE level, did psychology as A-level and then at degree level; subsequently has done masters degree (subject unknown). Knowledge of Excel, data analysis. Quite likes listening to music; used to play tenor horn (grade 8), read music using Braille.

VI11 High-school mathematics courses, and a masters' level multivariate statistics course. Enjoys music 'passionately', studied it at university.

VI12 Has taken statistics at undergraduate level. Loves music, listens to it every day, 'helps with my stress level'.

2.2 Survey methods

In the study a facilitator (either Karen, Chetz or Kaela) went through each of the six examples (detailed on page 4) with each participant. At the beginning of each session general background about the participant was gathered; this information included the participant's study history with the OU (for the sighted participants) or experience in studying maths and/or science (for the VI participants). Participants were also asked the extent to which they enjoyed listening to music; at the end of the session the participants were asked general questions about the usability of the sonifications.

The participants were asked to listen to the sonification(s); by not allowing the participants to see the plot initially, each participant, whether sighted or visually impaired, started without a given visualisation. Having heard the sonification the participants were asked about the plot; the questioning probed the participants' perception of how they visualised the plot or line graph, together with the conclusions that they were able to draw.

Each example followed the same general pattern:

1. Introduction/scene-setting
2. Listening to the sonification. (Here participants were free to listen to these as many times as they wished.) At this stage participants usually asked to draw or otherwise describe what they thought the plot looked like based on the sonification.
3. Reading the figure description. Participants were then asked to what extent their impression of the plot had changed based on the extra information contained within the figure description.
4. The 'big reveal', that is, looking at the plot. For the VI participants, tactile versions of the plot were made available instead of the visual version.

Finer details of the transcript used for each example are given in appendix A on page 23.

The same ordering of the examples was used for all the participants as it was felt that the later examples demanded more skill in the interpretation of sonifications than the earlier. The earlier (easier) examples were necessary to help participants calibrate themselves both to the sonification, and to the nature of the activities. It was clear to us that some (if not *most*) of our participants exhibited signs of nervousness, and the first example helped to put them more at ease.

In Examples 4 and 6 the general pattern was complicated as the learning objective necessitated the use of multiple plots and therefore multiple sonifications. What follows is a brief description of each example; complete transcripts are given in section 4 on page 23.

2.2.1 Example 1: Graph of a sine function (MU123)

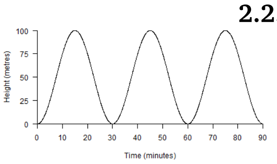

FIGURE 1

The first example was taken from MU123 *Discovering Mathematics*. This module is aimed to suit unconfident mathematics learners and is the first module in mathematics many students study in their degree program. In MU123, the scenario of modelling height of an individual on a ferris wheel over time is used to get students to think about the properties of a key function in mathematics — the sine function. These properties include the following (see fig. 1 and the larger version in fig. 17 on page 24):

- The function is periodic: the same pattern is repeated over time. In the case of the ferris wheel, one repetition (or cycle) corresponds to one complete rotation of the wheel.
- There are a series of peaks and troughs: the values of the peaks and troughs remain constant.
- There are smooth transitions between the peaks and troughs; the way in which the function increases is the mirror image of the way in which it decreases.

2.2.2 Example 2: Radioactive Decay (S104)

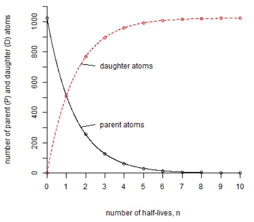

FIGURE 2

This example was taken from S104 *Discovering Science*, a module that is aimed at students embarking on a science degree; as such it aims to provide an introduction to a range of different science subjects. This particular example focuses on the decay of radioactive atoms (the 'parent' atoms) to form new atoms (the 'daughter' atoms). In a time period every parent atom has the same probability of decaying as any other parent atom. The rate at which the parent atom decays depends on which element it is from, and for any element it is measured by the half-life; the definition of the half-life is the time for half the atoms to decay.

The aim of the graph (see fig. 2 and the larger version in fig. 18 on page 26) is to demonstrate the consequences of this decay on the number of parent atoms and the number of daughter atoms. For the purposes of the teaching an artificial situation, concerning 1024 atoms of a element with a half-life of 30 days is considered. Important features of the plot include the following:

- The number of parent atoms declines rapidly at first and then flattens out.
- The number of parent atoms is close to 0 after 7 half-lives.
- The number of daughter atoms increases rapidly at first and then flattens out.
- Towards the end the number of daughter atoms equals the number of parent atoms at the start.
- The number of daughter atoms equals the number of parent atoms after just one half-life.

Although the figure plotted the number of parent atoms and the number of daughter atoms on the same plot, for the purposes of the sonification these were split into two separate sonifications so as to enable listeners to clearly hear pattern. The timing of the two sonifications was the same, as was the range of notes used to depict values on the vertical axis. Furthermore, in keeping with the data given in S104, the sonification only sounded the values at every half-life.

2.2.3 Example 3: Radial speed of a star (S104)

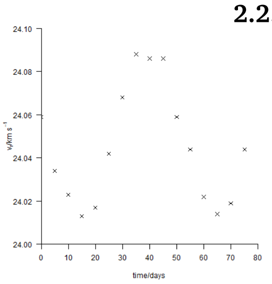

FIGURE 3

This example was taken from material at the *end* of S104, and focuses on a situation where collected data are used to infer facts about physical phenomenon. In this case observations about the radial speed of a star are used to deduce whether there is a planet orbiting it. In the unit students are expected to plot the data (see fig. 3 and the larger version in fig. 19 on page 27) and then make the following deductions.

- The radial speed appears to be periodic. (In fact looking like a sine function.)
- There appears to be some variation about this pattern.

2.2.4 Example 4: Correlation (M140)

This example was taken from M140 *Introducing Statistics*, in which students are given their first introduction to University-level statistics; the focus of the module is statistical concepts and general issues surrounding statistical modelling. Example 4 surrounds the introduction of the concept of *correlation*; the teaching involves *multiple* scatter plots so as to demonstrate the difference between datasets that have different *correlation coefficients*. *One* of the graphs used in this example is shown in fig. 4; the complete set of graphs for this example can be seen in figs. 20 to 22 on page 28 and on page 29.

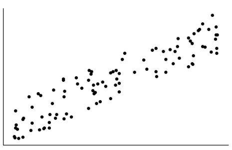

FIGURE 4

The first three plots (see fig. 20 on page 28) give examples of correlation coefficients: −1 the minimum possible, 0, and +1 the maximum possible. The plots are intended to demonstrate that −1 corresponds to an exact negative association between two variables, +1 corresponds to an exact positive association and 0 corresponds to no association.

The assessment of correlation from a sonification represents a departure from the strategies required for the sonifications used in Examples 1, 2 and 3, in which the interpretation relies upon following a tune. In this example, the focus is upon the extent to which a tune stands out from noise; in the case of the correlations −1 and +1 the tune is clear, with no noise present. (The sign of the correlation is dependent upon whether the notes go up or down.) However for the correlation of 0 there is no tune, only noise.

Plots D–F (fig. 21 on page 29) involved three different values of positive correlation: +0.1, +0.5 and +0.9. These examples are intended to provide participants with a feel of what different associations sound like.

Finally one graph (graph G, see fig. 22 on page 29) was presented for students to estimate the correlation displayed in the plot. In M140, such an activity is used to allow students to assess whether they have understood the general principle; estimating correlations from scatter plots is known to be a difficult task, so a ballpark estimate is all that is looked for.

2.2.5 Example 5: Interpreting a scatter plot (M140)

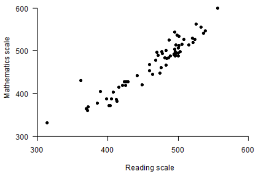

FIGURE 5

This example was also taken from M140, in which the focus is to *interpret* the information from a scatter plot, and examine the relationship between two variables. In M140 students are introduced to the following check-list:

- Is the relationship positive, negative or neither?
- Is the relationship linear or non-linear?
- Is the relationship strong or weak?
- Are there any outliers?

The particular graph used in this example (see fig. 5 and the larger version in fig. 23 on page 31) is one to which students are expected to be able to compare this checklist. In this case the expected answer is that the relationship appears to be positive, linear and relatively strong. There is one outlier: the point corresponding to a reading scores of roughly 375 and a mathematics score of roughly 425; furthermore in the plot there are two extreme points, one corresponds to a point with particularly low scores for reading and mathematics and one corresponding to a point with particular high scores for reading and mathematics. This is because this point does not fit with the general trend set by the other points. At the point in M140 at which this occurs students are not expected to assess to the degree of strength of the relationship.

2.2.6 Example 6: Assessing the fit of lines (M140)

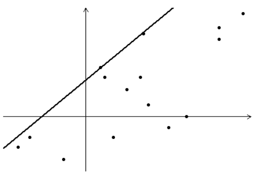

FIGURE 6

The focus for of the final example was fitting lines to data, another topic from M140; in particular, assessing whether a straight line drawn on a scatter plot represents a reasonable summary of the relationship between two variables. We used four graphs for this example; *one* of the graphs used in this example is shown in fig. 6; the complete set of graphs for this example can be seen in fig. 24 on page 32.

In M140 the teaching activity involved just a single scatter plot with four straight lines added. The students are then expected to say which lines, if any, fit the data well and which lines do not. Additionally, for those lines that students decide do not fit well, they are then expected to justify their choice.

Early experimentation with sonifying this plot indicated that trying the depict all four lines and the data points on the same sonification was likely to be too difficult to interpret. This was not surprising given that distinguishing and mentally separating the information from three or more streams of sound is known to be extremely difficult [12]. Instead the representation of the plot was split into five sonifications.

The first sonification represented only the points in the scatter plot, not any of the lines, so as to allow the listener to concentrate on the pattern of the points without being influenced by any suggested relationship. Each of the other sonifications gave a representation of one of the lines as well as the points; the representation of two graph elements (points and lines) simultaneously was thought to make the sonification hard to interpret and was the reason why this example was left until last. The lines were represented by a continuous tone, and the points by discrete notes; although this helped listeners distinguish between the two, it is unclear that this is the *best* way to help listeners compare the line and points.

2.3 Implementation of sonification

The principles of sonification creation have been implemented by others; in some cases the sonification has been designed to work well with data arising in one particular context or even just one dataset. Other implementations, and the ones we consider here, have been designed to work in a more generic context.

MathTrax MathTrax [1] is a standalone sonification package produced by NASA, and is designed to sonify mathematical functions. As such, it is good at sonifying line graphs, including those specified as mathematical functions (as opposed a series of points that are joined together). It also includes features that are of particular relevance when conveying mathematical shapes and forms, such as highlighting when values fall below the horizontal axis. MathTrax also allows user to complete graphical tasks using sonification. For example, to determine at which point or points a line crosses the horizontal axis (a graphical means of solving an equation).

PlayitbyR PlayitbyR is a suite of functions designed to work within the statistical software environment R [27]; much of the code provides a link between R and CSound. As might be expected for code working within R, the emphasis in playitbyR was *statistical* rather than *mathematical* plots. Unfortunately playitbyR only appears to work in R version 2; it does not appear to have been updated to work with the current R version 3.

FLOE Some of the newer work involves projects that are part of the international effort to leverage html5 in ways that improve access by giving users control over flexible delivery options [11]. An example of a current project is a sonified pie chart wizard [4], which currently works best in the Chrome browser. This allows users to populate fields for categories and their counts to produce a pie chart that can be accessed visually, or as a combination of text to speech with audible tones to represent the percentage of each category .

Given the disadvantages of these other implementations, the implementation used in this project was one written by one of the project team members, and takes the form of a suite of functions designed to work in R. The underpinning reasons for the choice to use R were as follows.

- The statistical analysis of data is a highly visual process. Plots often play an integral role in the 'Analyse' phase of the modelling cycle; there is a need for sonification of plots not to be separated from the rest of data analysis and to form part of a statistical package.
- R is open-source and available on a range of different platforms, meaning that it is generally available to those who want or need to use it.
- There is evidence that R is more accessible by VI individuals than other widely-used statistical packages [18].
- The environment provided by R makes it relatively easy for users to increase functionality by writing their own code.

Technical At the heart of this code is the transformation of a horizontal line on a graph into a sound. Positions along the horizontal (x-axis) are translated into elapsed time of the sonification. This was simply done in a linear fashion. So for a sonification designed to be of duration T, taking x_l to be the lowest value depicted on the x-axis and taking x_h to be the highest value, a position x on the x-axis corresponds to the time $t(x)$ in the sonification where

$$t(x) = T\frac{(x - x_l)}{(x_h - x_l)}.$$

The height of the line is mapped to the pitch of the note, and hence the frequency of the note. As with the timing of the note, the frequency is determining via a linear transformation; the frequency f of a line at position y is given by

$$f = (f_1 - f_0)\frac{(y - y_l)}{(y_h - y_l)},$$

where f_1 and f_0 are the frequencies of the highest and lowest note to used respectively and y_h and y_l are the highest and lowest values displayed on the vertical axis.

All the other forms on a plot are built up using this basic form. For example points are just taken to be very short lines, and curved lines are discretised and approximated by short sections of horizontal lines. Provided this discretisation is done on a fine enough scale it is not audible, and the change in pitch sounds smooth.

Stitching together short segments of sound together creates an audible ‘click’ if the beginning of the new segment does not have the same value as the end of the previous segment. To overcome this, the point at which the new segment starts in the cycle is carefully matched to that where the previous segment ends.

The shape of the waveform dictates the timbre of the resulting sound. For the sonifications used in this study, only simple waveforms were used, principally just a standard sine wave. This leads to sonifications that sound like plain tones rather than more complex tones such as those generated by traditional musical instruments.

3 Findings

The six examples that we chose each highlighted different points of interest; we will go through each example in turn, focusing on the feedback from our participants at the various engagement points. The quotes and visualisations we have chosen to present here focus on participants’ perceptual precision, their first-hand access to the medium, and the understanding (or not) that they were able to gain; we have tried to provide a balanced account.

3.1 Example 1: Graph of a sine function (MU123)

We were very keen for our first example to serve multiple purposes, which included: put each participant at ease, both with the process, and with the facilitator; allow each participant to calibrate their hearing to the sonification. As such, we deliberately tried to make our first example as conceptually easy as possible; indeed, each one of our participants engaged with every part of this example as we had intended.

Example 1 was in the context of a person travelling on a ferris wheel; in particular, we were interested in modelling the height above the ground of a passenger on board a ferris wheel at time t.

The complete facilitator transcript for this example is given on page 23; each and every participant answered the ‘engagement’ questions in the *introduction/background* section correctly, which helped participants to gain confidence.

When first presented with the sonification only half of the participants said that it sounded like how they expected. Reasons for it not being so were generally that the sonification was shorter than they expected or that they expected separate notes rather than a continuous tone.

When asked to describe the sonification, we received responses such as:

> *Sounds like my little boy going down the slide . . . I can follow it along in my mind. I can visualise a wavy line . . . it goes up and down 3 times*

(S7)

and,

> *Rises from the base 0 line, rise, reach peak, come down . . . sweeping series of arcs . . . going between the 100 on the y-axis and 0*

(VI6)

Throughout our examples, we asked our participants to ‘sketch’ an interpretation of the graph that the sonification was supposed to represent. In the case of our sighted participants this meant sketching on paper using a pen or pencil; for our visually impaired or blind participants, we used Wikki Stix [20] for the UK-based participants, and a Draftsman Tactile Drawing Board [2] for the US-based participants. An example of our participants’ sketches is given in fig. 7.

Each participant was asked for feedback about the figure description; VI8 said:

> *my initial reaction is that feels like a good figure description. . . but I wouldn’t easily or very confidently be able to draw that as a result of hearing that*

(VI8)

When asked to describe which medium (sonification, description or visual graph/tactile version) is more helpful, our participants had the following comments:

> *I feel, for me personally, the [tactile] graph gives more information; if I’d had the description and just need an idea of the graph, the sound is much easier to interpret*

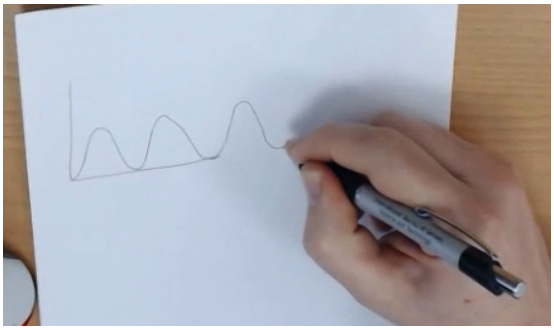

FIGURE 7: Visualisation of the sonification from S7.

(VI6)

> *I actually really like the sonification. Listening to the sound, I could actually see it in my mind what it was going to look like.*

(S7)

> *similar, really, because the beauty of a physical graph is that I can stop it at any point and interrogate the points... whereas this [sonification] I can't really interrogate either of these*

(VI8)

Finally, each participant was shown the graph, either in visual or tactile form; see fig. 8 (and fig. 17 on page 24).

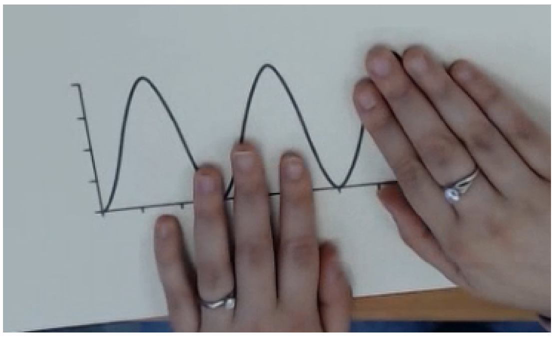

FIGURE 8: VI10 examining the tactile version of the graph for Example 1.

> *That's kind of what I expected. I'm not sure of a better description of it, so I think the sound actually made me see that better than description*

(S5)

> *Hey, look at that! [positive response, counts cycles]. That is exactly what I was expecting from the sound*

(S7)

> *I appreciate the other things [sonifications, description] are all saying that but to me they don't come across as clearly as this [tactile] does*

(VI8)

3.2 Example 2: Radioactive Decay (S104)

As in Example 1, an introduction to the topic was given (see complete facilitator transcript on page 24), with several early engagement questions to make try to ensure that the participants were following. The main difference as compared with Example 1 is that the sonification represented the graph of a *scatter plot*, and *not* a function that obeys a formula; this lead to a more 'discrete' sounding sonification, rather than changes in a continuous noise. In Example 2, each participant was given a total of 2 sonifications to which they had to listen and interpret: a 'daughter' isotope and a 'parent' isotope.

Daughter isotope Reactions to the sonification of the daughter isotope included:

> *I think it would be easier if I had more of a musical background; sounded like it was going down with time, like the pitch was going down. Not sure if using right terms. Also seemed to be getting quieter with time.*
>
> (VI2)

> *Can sort of grasp what it's trying to do; what I'm trying to visualise is, is it a decreasing gradient as such? Is it getting less - starting up here and coming down, but I'm not sure if it's coming down or if it spirals off. The notes were hard. They seem to plateau*
>
> (S3)

> *it changes... maybe steeper at first... I think it kind of levelled out a little bit and then it didn't change at all towards the end. [plays audio file again] so it's started to level out around the seventh data point. [plays audio file multiple times] so it really started to, obviously, begin to really slow down around the fourth data point. And by the seventh, it pretty much was flat, as far as I could tell... [participant then created sounds and talked about how the tones weren't a full octave apart]*
>
> (VI11)

A sample of the participants' visualisation is shown in fig. 9; note that S1 and VI2 illustrated that they heard *discrete* data points (figs. 9a and 9b), while S3 seemed to understand the *overall* shape of the curve, but has drawn their visualisation as a *continuous* curve (fig. 9c). Note that figs. 9b and 9d demonstrate the use of Wikki Stix.

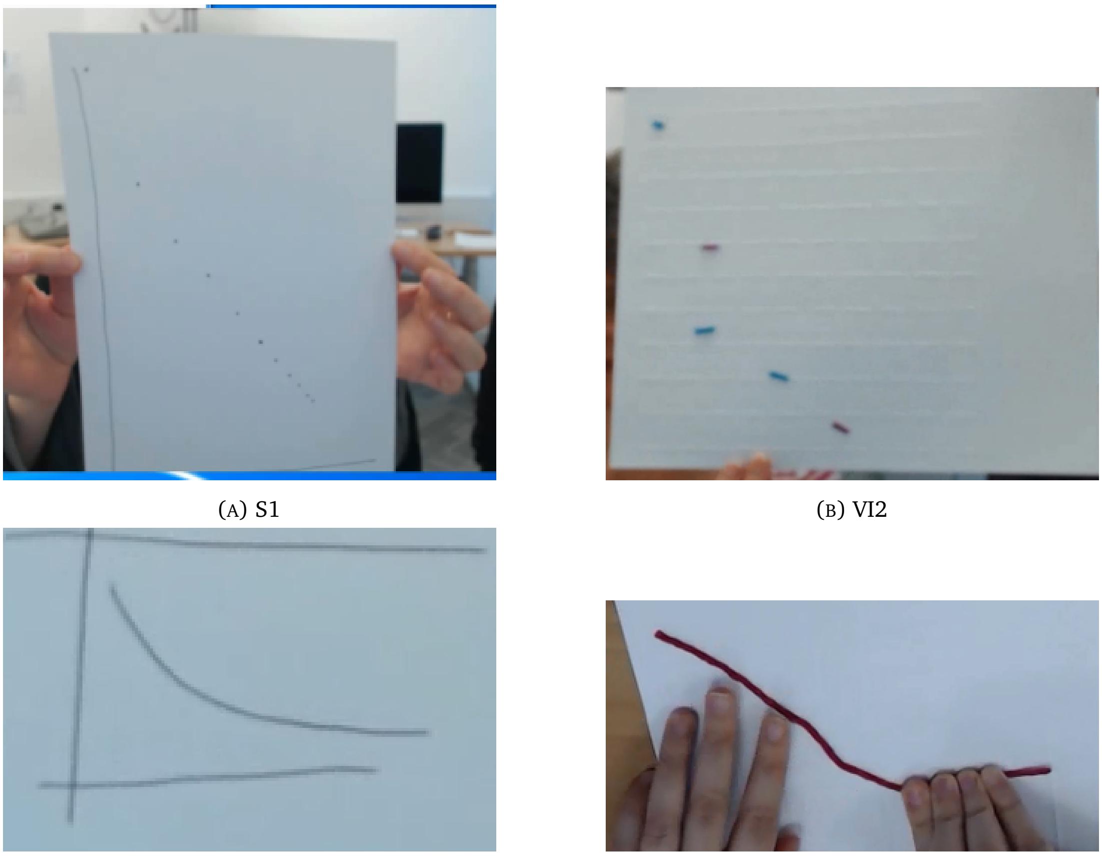

(A) S1 (B) VI2

(C) S3 (D) VI10

FIGURE 9: Participant's graphical interpretation of the sonification of the daughter isotope (Example 2).

Parent isotope Participants described the sonification of the parent isotope as, for example:

> *I sense an increase, but again I'm expecting it to be basically the opposite [Karen: based on your science knowledge?] no, purely going off how it sounds*
>
> (S3)

She starts off up here, then she goes up and levels off.

(VI4)

It's increasing until you get to a certain point, then it levels out.

(VI12)

Sample sketches of the way in which the participants visualised the sonification are shown in fig. 10.

(A) S7 (B) VI10

FIGURE 10: Participant's graphical interpretation of the sonification of the parent isotope (Example 2).

Having heard the figure description (detailed on page 25), our participants said

oh its got curves!

(S1)

For me this encapsulates the whole issue of figure descriptions, it assumes you can make an accurate model in your mind . . . from hearing. . . that description and then be able to manipulate it, and I think that is an unfair burden on blind students. . . you wouldn't ask a sighted student to do that

(VI8)

On being given the tactile or visual version of the graphic,

oh wow! It is pretty much what I was expecting - ish. I was expecting the parent atom to be a steep curve, but level sooner

(S3)

Somewhat confused by it [sonification], I suppose; This one [references tactile] gives me a lot clearer idea for understanding the material you read out

(VI6)

Sharper than what my curve was but I think the general principle of what I had in my head was pretty close. . . That one[daughter] looked very much like I thought it would

(VI10)

Yeah, this is what I was thinking. But I still like the tactile better. I guess I trust the tactile more. But I think the more I would use the sonification, I would get used to it. But because I'm used to tactile, that's why I think I'm more drawn toward the tactile. It's more tangible. Literally. That's one of the key piece with anything. You have to practice with it for you to trust it.

(VI12)

3.3 Example 3: Radial speed of a star (S104)

The complete transcript for this example is given on page 25; as in Example 2, the graph in this example represents a discrete data set (rather than a continuous curve as in Example 1).

Samples of the way in which our participants interpreted the sonification are shown in fig. 11.

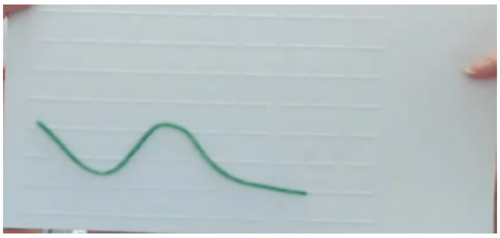

(A) VI2

(B) S3

FIGURE 11: Participant's graphical interpretation of the sonification of the radial speed of a star (Example 3).

When asked for a verbal description of the sonification, we received responses such as

> *It's not quite The Clash, is it?*
>
> (VI6)

> *it fluctuates... At times it looks like it's going faster than others... It didn't seem like it was doing the first same short pattern repeating... With the first example [ferris wheel]... it was a very clear repeat ... it [this example] didn't seem like it was the same short sequence repeated ... It didn't feel like it was the same pattern repeating.*
>
> (VI10)

Does the figure description help?

> *It does because from listening to that and reading that it does. I imagine if I knew more about what was going on it would help.*
>
> (S3)

> *Strengthens my initial argument: to throw someone [who is] blind ... only with figure descriptions ... is unacceptable.... It's what I call a 'sighted solution'.*
>
> (VI8)

3.4 Example 4: Correlation (M140)

The complete transcript for this example is given on page 26; as described previously on page 7, this example was made up of 7 graphs, which were delivered in the batches A–C, D–F, and finally G.

Plots A–C When asked *Which one sounds like the positive relationship, which one the negative relationship and which one where there is no relationship?* we received responses such as

> *Positive relationship is Plot B - it sounds happy [laughs]; I'm struggling a bit with this one. Plot C sounds like the worst possible case because it sounds miserable but I'm not sure if that is the 0 or the -1 as I'm confused between the distinction between those two. -1 was the negative and 0 was no relationship, so I would say that Plot C is where there is no relationship at all and plot A is where there is a negative relationship.*
>
> (VI2)

> *first one, I would say is positive relationship; the middle one, I would say, is the no relationship... I can picture the points [does visual] quite close together but no pattern to them... the negative one would be the last one*
>
> (S7)

> *Well, there's an upward trend. It seemed pretty clear in the first one and the opposite in the third one. And the second one sounded a little crazy... So the first one, I would say, is 1. And I guess the second one is zero, and the third one will be minus 1.*
>
> (VI11)

When asked to interpret the sonifications visually, S5 and VI6 did so as in fig. 12; additionally, VI6 said about Plot B *the middle one is going to be tricky; if we give the dog a wikki stick and wait a couple of hours, probably give you a good example of what I think of the middle*

(VI6)

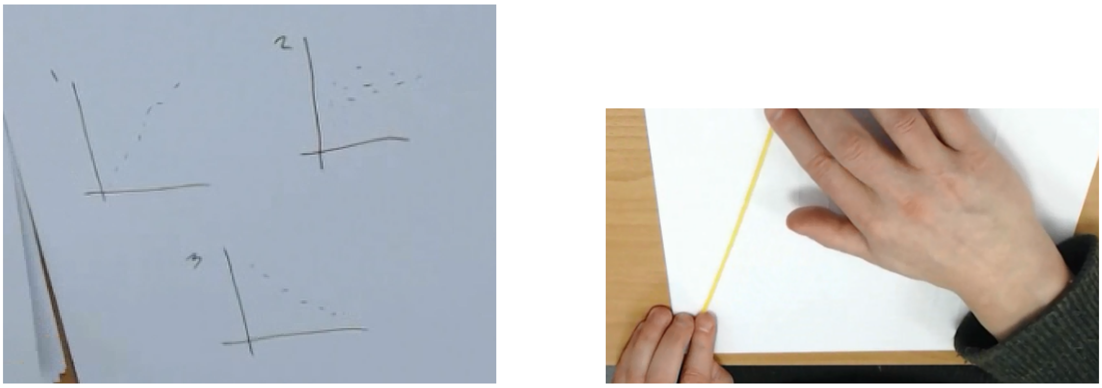

(A) S5 (B) VI6 (plot A)

FIGURE 12: Participants interpretations of plots A–C for Example 4.

When compared with the visual graph

Wow. I was not expecting that; [Karen: What surprised you most?] [Refers to Plot B:] Totally random. I can see (I know) why it's like that, but I really didn't get that from listening to the sonification. [listens again] Personally I struggled to interpret that; [refers to B] I expected it to be on a negative axis. B shocked me, to be honest. That was a surprise.

(S3)

I thought the way they were represented... It was quite intuitive... That's why I went on gut instinct with how the sound made me think and I think that worked quite well.

(VI10)

Plots D–F This group of sonifications was more challenging to the participants, as they were asked to distinguish between three positive values of the correlation coefficient; each was played the three sonifications, and asked to match each one to a correlation of either 0.1, 0.5 or 0.9.

All definitely positive as they're all going up. First one is strongest correlation, middle weakest because most disjointed sounds

(S5)

fairly sure that E is 0.9, not sure about D and F [listened to D again] D is 0.1, F is 0.5 and E is 0.9

(VI6)

So the first I think is the 0.9 because it seems to be more distinct... [listens to E and F again]... OK the last one is .1 and the other one, E, is 0.5

(VI12)

After hearing the figure description:

my translation of the scatter graphs from the sound is completely different from the description, but this is an area that didn't mean a lot to me at school, and the only other time has been since studying here

(VI6)

They'd all be a line going up, but for 0.9 all the scatter plots would be quite close, then with each you'd get increased distance, but still see they are moving in an upward direction.

(S9)

Comparing with the graphs:

Can feel general positive slope... Might be helpful to have best line drawn to understand degree of scatter.

(VI2)

> *Yeah, positive but scattered. That is what I was expecting but I suppose that's from knowledge of those [points]*

(S3)

> *It definitely looks how it sounds. It definitely feels how it sounds. You know, I don't even know if I would have been able to tell them apart by just a graphic. [interviewer asks if the sound helped] Yeah, I think the sound did more with this one.*

(VI12)

Plot G This was designed to challenge the participants: they were played a sonification, and then asked to guess the value of the correlation coefficient (anything between -1 and +1). As Example 4 was already quite long, not all participants were shown this example; in particular, VI4, VI6 and VI8 were not given this example.

The guesses for the final correlation (the actual correlation was 0.92) by the participants were as follows:

Around +0.7 or +0.8	3
Around +0.9	3
More than +0.9	1

This indicates even with their brief exposure to sonifications of correlations the participants who got this far were able to use start to gauge the strength of correlations.

> *pitch is getting louder towards the end which suggests a degree of correlation heading towards the 1. [Asks for repeat of sonification] Think it could be quite high, so could be 0.95*

(VI2)

> *I'd expect it to be like plot D?*

(S3)

> *I think it's definitely a positive correlation, . . . closer to a 1 than F was, definitely. . . I think it's closer than E was. . . I think it's probably about 0.7 or 0.8, I'd reckon, quite close*

(S9)

After hearing the figure description

> *you could tell there wasn't as big a scatter as with the [0.]5 and with the .1 the scatter was all over the place . . . Maybe that's why I was thinking it sounded more positive than the .9*

(VI10)

Viewing the graph

> *Yes, that's right [HAPPY]. Quite easy to do that one having just listened to those 3 so as long as you have reference ones, it would be fine.*

(S5)

> *plot G is probably even closer to 1 than plot D; . . . from the sound I did manage to hear the scattering, I did think that they were closer together. . . not off by much*

(S7)

3.5 Example 5: Interpreting a scatter plot (M140)

The complete transcript for this example is given on page 29. This example immediately followed Example 4, so our participants sometimes chose to carry over some of the terminology they had just been using.

On hearing the sonification:

> *OK, so we've got correlation . . . so maths on one axis and reading on the other, it won't matter which way round, will it. Felt like there was one odd country that was down there all on its own. Bit of a bunch up here. An odd little . . . one up there. Sounded fairly positive because it was going up.*

(S5)

Samples of how participants interpreted the sonification *graphically* are shown in fig. 13. Note, in particular, that VI10 and S5 (figs. 13a and 13c) appeared to visualise the scatter plot as discrete data points, whereas S3 (fig. 13b) seemed to 'fill in the gaps' by drawing a continuous curve. Note the different interpretations in the overall shape of the sonification; VI10 and S5 have drawn what seems to be a straight line, and S3 has drawn a curve.

(A) VI10

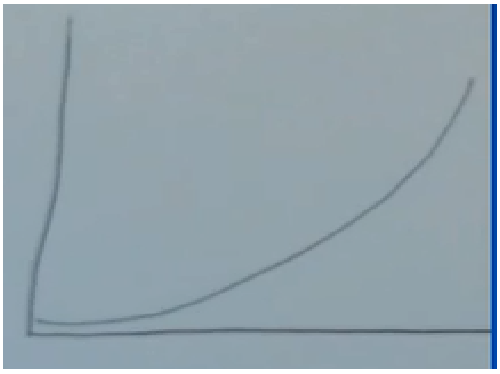

(B) S3

(C) S5

FIGURE 13: Participant's graphical interpretation of the sonification representing a scatter plot of Maths and English scores(Example 5).

Further thoughts on the sonification include

> *One on its own... then a group... then more like lines going up... then a big group... then smaller group, all equally spaced... then one on its own.... Small clustery bits, then a great big blob... with all countries going up... then one on its own.*

(VI10)

> *When you read Braille, your fingers go cold. Lots of points on a graph... you lose sensitivity so graph points will be lost*

(VI6)

> *Quite strong... the higher the maths score the higher the reading score, but there is still quite a bit of change within them.*

(S9)

> *It sounds like mathematics generally increases as reading scores increase. That's what I would get out of that. And it's not a perfect trend.*

(VI11)

> *Well, I notice that there's one country that's probably really, really low– the lowest. There's one that's close. And then there's a little bit of a cluster. Then there's a bigger cluster as you're*

> *going. All this is as you're going higher, percentages of math and reading getting closer together – clustered in the high range. And then there's one that's really at the top.*
>
> (VI12)

After the participants were allowed to view the graph,

> *It does help to see it, or feel it, in this case. I did hear it... mathematical outliers were more difficult to spot with the sound than the reading ones. Since the notes are all jumbled together, it was kind of difficult to pick out those, where the notes were relatively high for its position.*
>
> (VI11)

Surprising aspects of the sonification:

> *Don't like dots being fused together, might be easier for sighted person to distinguish. Didn't really get feeling for graph at all till I got the tactile one.*
>
> (VI8)

> *It's strange. It appears that maybe that's just a phenomenon related to the hearing, you know? Just how it sounds. It fools the ear because the data points are becoming more sparse. But that doesn't necessarily mean that there are outliers, because they're not tactilely separate. You can see, from touching, that they're very close to others. They're still lumped together with a few other scores*
>
> (VI11)

3.6 Example 6: Assessing the fit of lines (M140)

This example certainly appeared to cause the most difficulty for our participants. The protocol changed and evolved as we worked through our participants; initially, all of the graphs/tactile versions were not shown until the very end of the example, but it became clear that it was much more helpful to show the participants the 'points only' graphic so as to help them calibrate.

We believe that this example pushed the boundaries of the usefulness of sonifications; each one was 'dual track', as it contained the sonification of the points, together with the sonification of the line of best fit.

When asked, *What appears to be the trend? Do they* [the points] *go up or down?*, we received responses such as:

> *General trend as the graph moves along the x-axis, rises; ... early on there was a slight rise, and then a slight dip... rising again with a slight cluster, little gap, and then another cluster*
>
> (VI6)

Some participants did demonstrate that they were able to visualise the rough pattern of the points, see fig. 14.

(A) S7 (B) S9 (C) VI10

FIGURE 14: Participant's graphical interpretation of the sonification representing the points only (Example 6).

Line of best fit The crux of this example was to give the participants four different lines, and to see which one they thought best fitted the points (see fig. 24 on page 32 for details). Some participants seemed to grasp the idea quite well

> *OK I have a good idea of what I think is being represented. (listens to the others twice each) OK I think I've got a good grasp of those. A was easy to grasp. It's a positive, relatively steady gradient, quite positive; (listened to B again) B quite low, very neutral gradient, then it suddenly peaked and a real strong trend and went off the scale as such; (listened to C again) Again more positive start than B, and then again really peaked towards the end and went up quite steep; (listened to D again) Plot D found it very similar to C, pretty much the same, about the same as C*
>
> (S3)

while for others, it seemed to cause confusion:

> *Interesting. To me, you need to have your points clearly defined to start. Then, put your line through it. Its confusing things to have the line going through and you have the line going through but you haven't finished defining some of the points. It is the find the mean?*
>
> (VI4)

> *OK, those three long lines sound the same to me. They sound - I don't know, I can't distinguish any difference. OK, maybe I have to hear them side-by-side. [listens to them all again in reverse order.] Well that one ended faster. I still can't tell the difference between the last two. the last two do a better job.*
>
> (VI12)

Some participants attempts at visualising the line of best fit are shown in fig. 15.

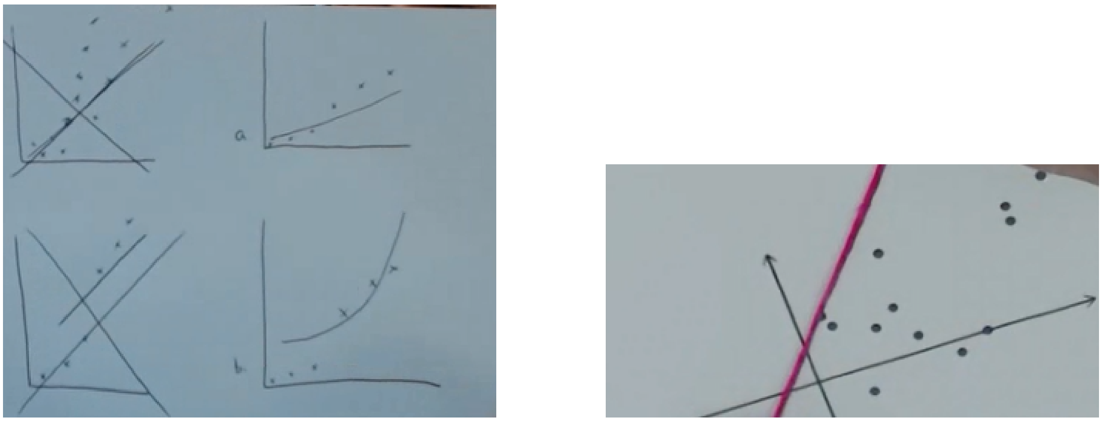

(A) S3 (B) VI10 (line A)

FIGURE 15: Participant's graphical interpretation of the line of best fit (Example 6).

After hearing the figure description, S3 said (with reference to fig. 15a)

> *I'd leave B as it is, obviously there's more points, but I'd say the line is roughly what I think. . . Wow. I must have gone totally off the rails here [referring to D]. D I totally thought was really steep but reading this I would change it and it would be [redraws] and the points are more even. . . . . I didn't get that from the sonification.*
>
> (S3)

And after seeing the graphs/tactile versions

> *So I didn't quite get that (implies A). I didn't really. It's difficult to know what is positive and what is negative. C matched quite well. B doesn't fit well. C and D fit quite well. A doesn't fit - it's way on the outside - it's not in the middle of the points.*
>
> (S1)

3.7 Overall impressions

At the end of each session, we had a discussion with the participant about their overall impressions of sonifications.

They were OK to listen to. . . sometimes it was quite fast; it did give an impression of an overall line; but it was over quite quickly. maybe if was a bit more spaced out then I could put more detail on the graph as to where the points were.

(S1)

wouldn't rate very highly personally. Tactiles are better. Tactile more tangible. Would prefer to rely on tactiles and a verbal description.

(VI2)

I thought at times they were difficult but you can tell at times what they were trying. . . .a picture in your mind

(S3)

What you're doing with the sound its good. If our students are plotting points and the sound is going in the right direction, or suddenly the sound is wrong, it can help [to confirm]

(VI4)

They were fine. Quite comfortable. Wouldn't want to do this for hours on end, but you don't do graphs for hours on end anyway. It does give you a good picture of where things are going.

(S5)

It's helpful to get an idea with some of them. . . . Would certainly use it if I had access to it for where I needed to get a general idea. . . . If I had to extract specific data from a graph, I wouldn't even bother looking at it. . . . If I had the graph, I'd probably feel on the graph. . . give me the data, and if there's a formula, I'll grab the formula

(VI6)

Tones aren't . . . annoying. Random ones do sound a bit crazy . . . but wouldn't particularly annoy me. . . it gives you a good indication of whether the data is close together or spaced out

(S7)

Figure description is just describing, and it is one person's interpretation. Tactile allows me to explore it, and easier to compare graphs. Both is best, almost as good as experience of sighted reader.

(VI8)

They were all OK, I wouldn't listen to them for fun, but they weren't annoying. . . . Would listen to them when I was on those questions. . . every time I needed a reminder of the graph. . . I think I would have it on a loop.

(S9)

For simpler data sonification could be fine without description or tactile. More complex would need description. If had to measure anything then have to have tactile. Depends on what is required from the diagram. A basic overall picture sonification could be very useful.

(VI10)

4 Impact

As mentioned in the introduction, the aim of this study was to see how effective sonifications can be as alternate accessible versions of plots and graphs in module materials. The results in this study show that the sonifications did enable most of the participants to get the gist of the plot; this was despite being initially unused to being presented with plot and graphs in this format. Greater experience with sonifications should only increase participants' ability to interpret plots and graphs given in this format.

The sonifications also enabled the participants to gain an impression of the plot or graphs quickly; each of the sonifications was only 6 seconds long. Although participants generally listened to sonifications more than once, using them did not add significantly to study time. Participants indicated that listening to sonifications was not an unpleasant experience, and expecting students to cope with multiple sonifications in a single study session does not appear to be an unreasonable ask. Based on this we feel that sonifications of plots and graphs should, where possible, be made available to students.

However, the results from this study also show that sonifications, even when presented alongside figure descriptions, are not sufficient to provide students with an adequate replacement of the original plot or graph. Instead what is needed is a *blended* approach: sonifications, figure descriptions and tactile diagrams. Each of these alternate representation provide a different, but complimentary, aspect of the original plot of graph (fig. 16).

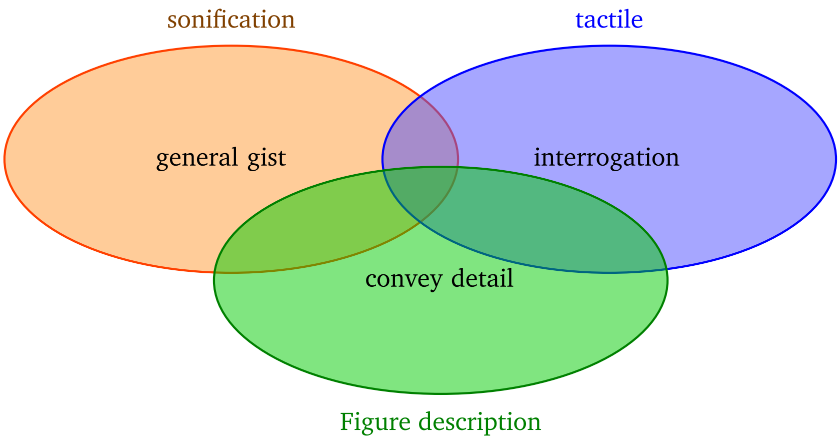

FIGURE 16: A blended approach

The sonification quickly provides the *general gist* of a plot, giving the general sense of the shapes that form the plot. Not only is this important for getting an overview of the relationships that are depicted in the graph, but it also allows the user of a tactile version to know where to look when interrogating the patterns in the plot.

The tactile diagram allows *interrogation* of the plot; for example reading of the position of points or lines against the scales, or against other points/lines. Perhaps more importantly, tactile diagrams give the user the freedom to explore the plot in the way that they want to, rather than having the information presented in a fixed order.

Figure descriptions *convey detail* such as axis labels and scales – information that is lost in sonifications. Furthermore although this information can be included on tactile diagrams, it presupposes that all users can read Braille, which may not always be the case. For example Edwards [7] cites a report that suggested that in 1991 the proportion of blind people who could read Braille might have been as low as 2%.

The concentration on a blended approach also allows the production of each element to be simplified. For sonifications, the need to try to include the numerical information given by the scales is avoided, as the production of tactile diagrams can focus on the best representation of the symbolic information. Finally, figure descriptions may become shorter and easier to write as they can concentrate on conveying any textual information in the plot, such as axis labelling.

Evaluation of method

Each of our participants went through each of the six examples in exactly the same order:

- they heard the sonification;
- they heard/read the figure description;
- they interrogated the tactile/visual version.

By the time that each participant reached the interrogation stage, they had already been exposed to two different interpretations of the graph. As such, it is impossible to discern just how influenced each participant was by each of the three different media. A future study may examine particpant's visualisation/understanding after, for example, allowing participants *only* to access one of three media types.

References

[1] Nation American Space Agency. *MathTrax*. URL: http://prime.jsc.nasa.gov/MathTrax/index.html (visited on 09/19/2016).

[2] American Printing House for the Blind Inc. *DRAFTSMAN Tactile Drawing Board*. URL: https://shop.aph.org/webapp/wcs/stores/servlet/Product_DRAFTSMAN%20Tactile%20Drawing%20Board_1-08857-00P_10001_11051 (visited on 09/19/2016).

[3] Sara Bly. "Presenting information in sound". In: *CHI '82: Proceedings of the 1982 conference on human factors in computing systems*. Gaithersburg, Maryland, USA, 1982, pp. 371–375.

[4] *Chart Authoring Demo*. URL: http://build.fluidproject.org/chartAuthoring/demos/ (visited on 09/19/2016).

[5] *Describing visual teaching materials*. URL: https://learn3.open.ac.uk/mod/oucontent/view.php?id=78466 (visited on 09/19/2016).

[6] U.S. Department of Education. *Advisory Commission on Accessible Instructional Materials in Postsecondary Education for Students with Disabilities*. 2011. URL: http://www2.ed.gov/about/bdscomm/list/aim/publications.html (visited on 09/19/2016).

[7] Alistair D N Edwards. "Auditory display in assistive technology". In: *The Sonification Handook*. Ed. by Thomas Hermann, Andy Hunt, and John G Neuhoff. Berlin, Germany: Logos Verlag, 2011. Chap. 17, pp. 431–453.

[8] Equality and Human Rights Commission. *Equality Act 2010*. URL: https://www.equalityhumanrights.com/en/equality-act/equality-act-2010 (visited on 09/19/2016).

[9] Equality and Human Rights Commission. *Higher education provider's guidance*. 2016. URL: https://www.equalityhumanrights.com/en/advice-and-guidance/higher-education-providers-guidance (visited on 09/19/2016).

[10] *Equality and Human Rights Commission: Public Sector Equality Duty*. 2016. URL: https://www.equalityhumanrights.com/en/advice-and-guidance/public-sector-equality-duty (visited on 09/19/2016).

[11] *Flexible Learning for Open Education*. 2016. URL: http://www.floeproject.org/ (visited on 09/19/2016).

[12] John H Flowers. "Thirteen years of reflection on auditory graphings: promises, pitfalls and potential new directions". In: *Proceedings of ICAD 05 – Eleventh meeting of the International conference on auditory display*. Limerick, Ireland, 2005.

[13] John H Flowers, Dion C Buhman, and Kimberly D. Turnage. "Cross-modal equivalence of visual and auditory scatterplots for exploring bivariate data samples". In: *Human Factors* 39 (1997), p. 340.

[14] John H Flowers and Terry A Hauer. "'Sound' Alternatives to visual graphics for exploratory data analysis". In: *Behavior Research Methods, Instruments & Computers* 25.2 (1993), pp. 242–249.

[15] John H Flowers and Terry A Hauer. "The ear's versus the eye's potential to assess characteristics of numeric data: Are we too visuocentric?" In: *Behavior Research Methods, Instruments & Computers* 24.2 (1992), pp. 258–264.

[16] Keith M Franklin and Jonathan C Roberts. *Pie Chart Sonification*. Tech. rep. Computing Laboratory, University of Kent, 2003.

[17] Steven P Frysinger. "Proceedings of ICAD 05 - Eleventh meeting of the International conference on auditory display". In: *First*. Limerick, Ireland, 2005.

[18] A Jonathan R Godfrey and M Theodor Loots. "Statistical software (R, SAS, SPSS and Minitab) for blind students and practitioners". In: *Journal of Statistical Software* 58 (2014).

[19] Florian Grond, Trixi Drossard, and Thomas Hermann. "SonicFunction: experiments with a function browser for the visually impaired". In: *Proceedings of the 16th International Conference on Auditory Display*. Washington DC, USA, 2010, pp. 15–21.

[20] Omnicolor Inc. *Wikki Stix*. URL: http://www.wikkistix.com/ (visited on 09/19/2016).

[21] *IT Accessibility Risk Statements and Evidence*. 2015. URL: https://library.educause.edu/Resources/2015/7/IT-Accessibility-Risk-Statements-and-Evidence (visited on 09/19/2016).

[22] *JAWS screen reader*. URL: http://www.freedomscientific.com/Products/Blindness/JAWS (visited on 09/19/2016).

[23] Dolphin Computer Access Ltd. *Dolphin screen reader*. URL: https://yourdolphin.com/ (visited on 09/19/2016).

[24] David Lunney and Robert C Morrison. "Auditory presentation of experimental data". In: *Proceedings of the society of photo-optical instrumentation engineers conference*. Vol. 1259. 1990, pp. 140–146.

[25] *Online Student Workload Tool*. URL: http://learning-design.open.ac.uk/ (visited on 10/25/2016).

[26] K. Parks, G. Dietrich, and L. Wadors. "3D Printing: how to create and use tactile learning objects". In: 2015. URL: https://www.ahead.org/conferences/2015/pre-conf.

[27] *PlayitbyR*. URL: https://www.r-project.org/ (visited on 09/19/2016).

[28] Ed Summers, Julianna Langston, Robert Allison, and Jennifer Cowley. *Using SAS/GRAPH to create visualizations that also support tactile and auditory interaction*. Tech. rep. Paper 279-2012. SAS Global Forum, 2012.

[29] Edward R Tufte. *The visual display of quantitative information*. Cheshire, Connecticut, USA: Graphics Press, 1983.
[30] Bruce N Walker and Grogory Kramer. "Mappings and metaphors in auditory displays: an experimental assessment". In: *ACM transactions on applied perception* 2 (2005), pp. 407–412.
[31] Michel Wermelinger. *iChart — interactive exploration of data charts. eSTEeM Final Report*. Tech. rep. The Open University, 2013.
[32] Edward S Yeung. "Pattern recognition by audio representation of multivariate analytical data". In: *Analytical Chemistry* 52 (1980), pp. 1120–1123.

A Appendix

This appendix provides the transcripts for each of the six examples that we studied with each of our 12 participants.

Example 1: Graph of a sine function

Key idea Modelling how far someone is above the ground. (And as it's a model, time is speeded up)

Introduction/ background In this example imagine that you're on a Ferris wheel that completes one revolution every 30 minutes; you board the Ferris wheel at ground level and reach a maximum height of 100 metres. You are interested in trying to model your height above the ground depending on the time; in other words, you'd like to be able to predict your height above the ground for any time. Before you listen to the sonification, think about the following.

engagement questions

- If the Ferris wheel completes a revolution every 30 minutes, how long would it take us to reach the maximum height? (15 minutes)
- How long would it take us to come back to ground level? (30 minutes)
- If we were to plot our height above the ground on the vertical axis of a graph, and time (from 0 to 100 minutes) on the horizontal axis, what kind of features would we expect to see on our graph? (Answers will vary widely! Be prepared to give guidance.)

Sonification is a means of drawing a graph using sound. For this particular sonification, the pitch of the note corresponds to height above ground. So high notes are when you are far above the ground and low notes are when you are close to the ground. The time when a note is sounded corresponds to how far a point is along the horizontal axis.

Let's listen to the sonification!

engagement questions

- Is the sonification as you'd expect?
- What does it sound like to you?
- What do you think the high points in the sonification represent? (Reaching maximum height on Ferris wheel)
- How about the low points? (Back to ground level)
- How many revolutions do you think are represented on the sonification? (Answers will vary, but point out that we hear more than one cycle.)
- What do you think would happen to the sonification if we were to complete one revolution in 10 minutes (as opposed to 30 minutes)? (The sonification should speed up)
- What do you think would happen to the sonification if the Ferris wheel was larger? (Perhaps the pitch would reach a higher level?)

Figure Description The figure description for this figure is as follows.

This is a graph of height against time. The horizontal axis is marked from 0 to 90 in intervals of 10 and is labelled 'Time (minutes)'. The y-axis is marked from 0 to 100 in intervals of 25 and is labelled 'Height (metres)'.The graph shows 3 oscillations of a sine-like curve, starting at the minimum point (0, 0). The minimum values of the curve are 0 and the maximum values are 100. The curve repeats itself every 30 units along the horizontal axis *engagement questions*

- Do you think that this fits well with the visualisation that you have from the sonification?
- Which is more helpful?

The graph! Participants were given either a visual representation of the relevant graph, or a tactile version; for the purposes of this report, the graph is shown in fig. 17.

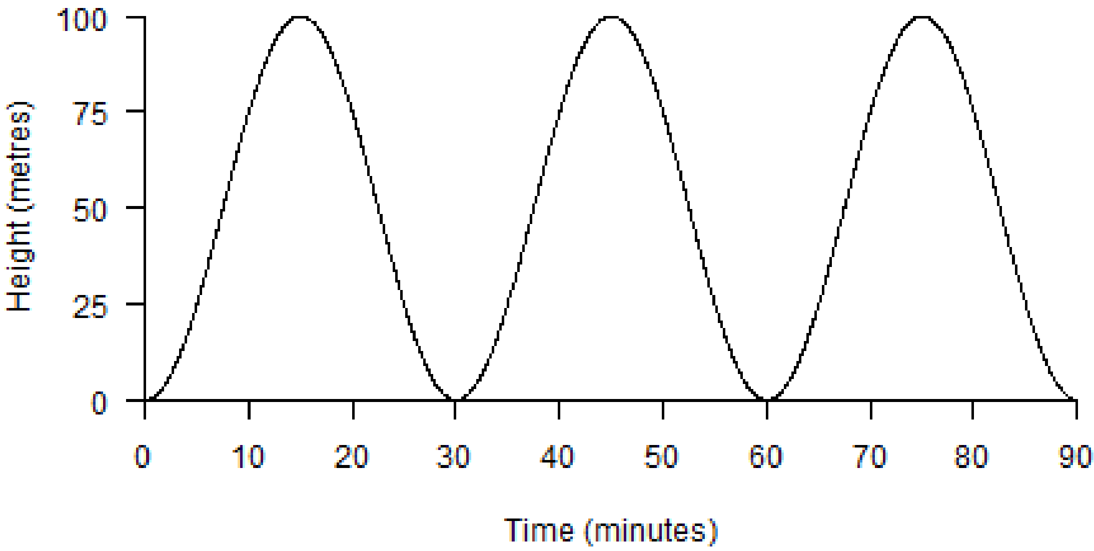

FIGURE 17: Graph for Example 1 (Graph of a sine function).

Example 2: Radioactive decay

Key idea What are the numbers of "parent" atoms over 11 equally spaced time points? (Assume start with 1024 of them.) Similarly what are the numbers of "daughter" atoms?

Introduction Many elements consist of several stable isotopes. Isotopes of an element have the same number of protons and electrons, but different numbers of neutrons. For example, hydrogen has two stable isotopes; hydrogen with an atomic nucleus of one proton and deuterium with a nucleus of one proton and one neutron. A few elements have isotopes that are inherently unstable. The more unstable an isotope is, the greater the chance that an atom of it will break down at any instant by radioactive decay, to form a stable isotope of another element. For now, you just need to know that the end product of such radioactivity is a stable isotope called a daughter isotope, which has formed by the decay of the original, radioactive parent isotope.

By decaying, the abundance of a radioactive parent isotope decreases with time, while its daughter isotope becomes more abundant. The time taken for half the radioactive nuclei in an isotope sample to decay is called the isotope's half-life. Half-lives range from fractions of a second for some rare, artificially produced isotopes, to over a hundred billion years for a few naturally occurring isotopes.

Whatever the half-life actually is, this phenomena is an example of exponential decay.

Suppose a mineral sample contains, in addition to the elements that make up its bulk, 1024 atoms of a radioactive isotope whose half-life is 3 days.

Let's listen to the sonifications! The first sonification represents the number of atoms from the parent isotope after every 3 days (1 half-life) for a period of 30 days. Play the sonification of the change in number of atoms from the parent isotope first of all (Example 2 - parent). Listen to the first sonification.

engagement questions

- Based on the sonification, try sketching what the plot look like.
- So would you say that the number of atoms of the parent isotope is increasing or decreasing or a mixture of the two?
- What about the rate at which the number of parent atoms changes? Is this constant or does it change? (And if so in what way does it change?)

The second sonification represents the number of atoms from the daughter isotope after every 3 days (1 half-life) for a period of 30 days. Now play the sonification of the change in number of atoms from the daughter isotope (Example 2 - daughter)

Listen to the second sonification

engagement questions

- Based on the sonification, try sketching what the plot look like.
- So would you say that the number of atoms of the daughter isotope is increasing or decreasing or a mixture of the two?
- What about the rate at which the number of daughter atoms changes? Is this constant or does it change? (And if so in what way does it change?)

Figure descriptions On the original plot, the number of parent atoms and the number of daughter atoms are plotted on the same graph. The corresponding figure description is as follows.

The horizontal axis is marked from 0 to 10 in intervals of 1. It is labelled number of half-lives, n. The vertical axis is marked from 0 to 1000 in intervals of 100. It is labelled number of parent (P) and daughter (D) atoms. Two curves each with 11 plotted points, one for each of $n = 0$ to $n = 10$. One curve, labelled daughter atoms starts at (0,0) and rises steeply initially, reaching half way up the vertical scale in $n = 1$ half life, then gradually reduces in angle and flattens out appearing horizontal by about $n = 7$. The other curve, labelled parent atoms does exactly the opposite, beginning at the point (0, 1024), falling steeply initially, to cross the daughter atoms line at (1, 512) and then gradually lessening in angle to approach the horizontal axis. It is essentially on the horizontal axis between after $n = 7$.

engagement questions

In the light of this description do you want to change either of your sketches. If so in what way?

Comparison with the graph Now look at the graph.

engagement questions

Based on the sonification, is the graph as you expected? Why or why not?

The graph! The graph is shown in fig. 18.

Example 3: Radial speed of a star

Key Idea Do measurements of “radial speed” vary over time? And if so, in what way?

Introduction Imagine that a distant planet is orbiting around a star, in the same way that the Earth orbits round our star, the Sun. If we want to look at planets outside our solar system, i.e. those orbiting round other stars, we need a way to distinguish the very dimly lit planet from the very bright star it is orbiting. At the distances involved, we cannot do this by direct imaging, we would only see the bright star. So, how can we determine if a distant star we’re studying has a planet orbiting it?

One way of doing so is to consider the wavelength of the light received by the observer. Effectively the combined system of the star and planet will rotate. Half the time the star will be moving towards you and half the time it will be moving away from you. This would lead to a change in the wavelength of the light — which can be used to calculate what is known as radial speed.

The following sonification details the radial speed of a star (in kilometres per second) against time (days). The pitch is dependent on the radial speed; the higher the radial the speed, the higher the pitch. It is OK for the participant to listen to the sonification as many times as desired (Example 3). And its OK to return to it during the discussion. However it would be nice to know how many times the student felt it was necessary to listen to it.

Let’s listen to the sonification!

engagement questions

- Based on the sonification, try sketching how radial speed changes over time. That is, do the notes go up? Or down or up and down? Or does there seem to be no pattern in the pitch over time? Or something else?
- How would you describe the relationship between the radial speed of the star and time?

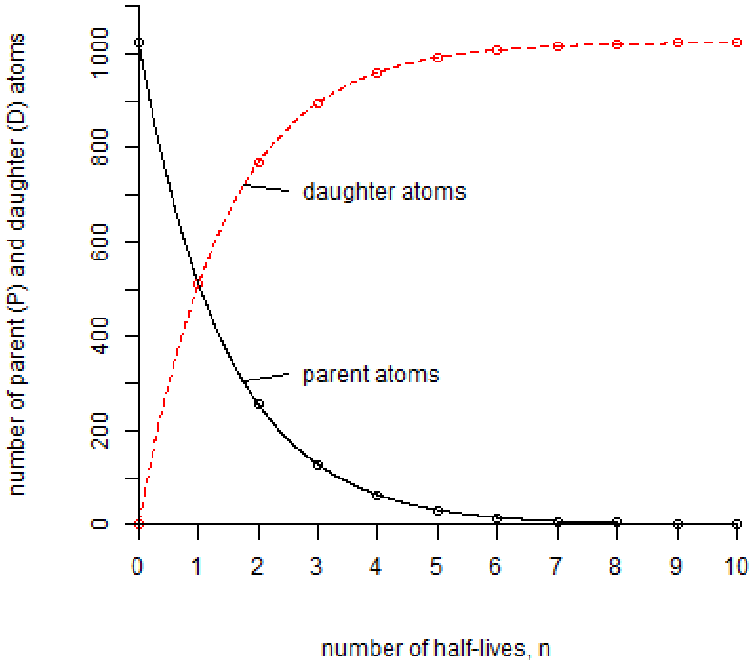

FIGURE 18: Graph for Example 2.

Figure Description The figure description for this graph is as follows.

A graph in which the horizontal axis is marked from 0 to 80 in intervals of 10. It is labelled time /days. The vertical axis is also marked from 24.00 to 24.10 in intervals of 0.02. It is labelled Vr/km s^{-1} .There are 15 plotted points. The lowest is at approximately (15, 24.01), the highest at (35, 24.09). The overall pattern is that the plotted points form a sine wave about the mid point on the vertical axis, i.e. 24.05 km s^{-1}. The wave goes down initially to the low point described previously around 15 days, then back up all the way to the high point described previously around 35 days, before heading back down. The plotted points are not exactly on what would be the best fit line, but are close to it.

> *engagement questions*
>
> In the light of this description do you want to change your sketch? If so in what way?

Comparison with the graph Now look at the graph.

> *engagement questions*
>
> - Based on the sonification, is the graph as you expected? Why or why not?
> - How could the sonification be improved?

The graph! The graph is shown in fig. 19.

Example 4: Correlation

Key Idea On a scatter plot how strong is the "signal" (pattern) compared to the "noise" (random scattering)? $+1/-1$ means its all signal, no noise. 0 means its all noise, no signal. All other values somewhere in between.

Introduction How strong is a relationship between two variables? That is, to what extent does knowing the value of one variable provide information about what the value might be for another variable? For example suppose you have a group of children aged 0 to 5 years. Looking at the child and measuring their height would allow you to make a good guess at their age. However this situation is different if you had a group of people aged 40 to 45. Then knowing someone's height is unlikely to help when

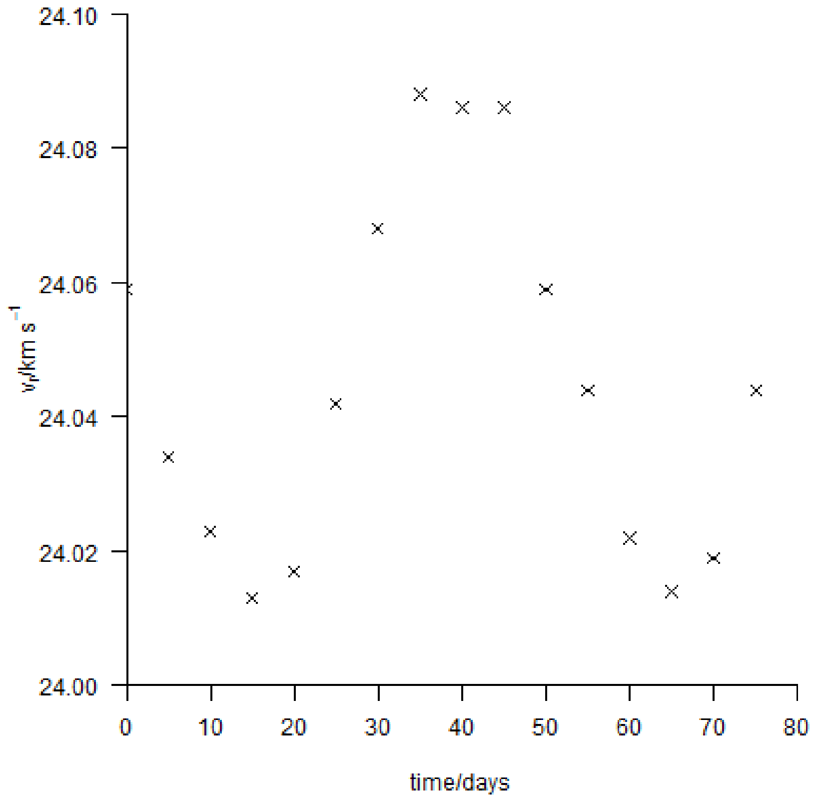

FIGURE 19: Graph for Example 3.

guessing their age. This is because for young children there is a strong relationship between height and age but for adults there is, at best, a weak relationship between height and age.

So how is the strength of relationships measured? One common measure is the correlation coefficient. The correlation coefficient takes values between −1 and 1.

Positive values of the correlation coefficient suggest a positive relationship. That is high values of one variables are linked with high values for the other variable. (And equally low value of one variable are linked with low value for the other variable.) Similarly negative values of the correlation coefficient suggest a negative relationship. That is high values of one variable are linked with low values of the other variable. The size of the correlation (that is how close it is to either +1 or −1) indicates how strong the relationship. A correlation of +1 or −1 suggests that the relationship is exact. That is knowing the value of one variable allows you to exactly deduce the value of the other variable. A correlation of 0 suggests that the relationship is so weak that we can say there is no relationship.

A good way to get a feel for this to think about what different correlations sound like.

Let's listen to some correlations!

First allow the student to listen to a sonifications of a scatter plot where there correlation coefficients are +1, −1 and 0 respectively. (These are files Example 4 - Plot A, Example 4 - Plot B and Example 4 - Plot C.) The following are sonifications of scatter plots corresponding to the correlations +1, −1 and 0.

engagement questions

- Which one sounds like the positive relationship, which one the negative relationship and which one where there is no relationship? Here Plot A is the positive relationship, Plot C is the negative relationship and Plot B is where there is no relationship.
- In what way does the sonification for the correlation of 0 differ from that those for the correlations of +1 and −1. The key is focus on how 'jumbled' the notes sound. No clear 'tune' emerges as the notes seem to go randomly up and down.
- Based on what you have heard, trying sketching each of the graphs that these sonifications correspond to.

Figure Description

The figure description for the graphs are as follows.

Each scatter plot shows the positive axes, but there are no scales and the axes are not labelled. On each plot are a number of points.

On plot A the points lie on a straight line with a positive slope.

On Plot B the points appear to be randomly scattered.

On plot C the points lie on a straight line with a negative slope.

> *engagement questions*
> In the light of this description do you want to change any of your sketches? If so in what way?

Comparison with the graphs

Now look at the graphs.

> *engagement questions*
> Based on the sonification, is the graph as you expected? Why or why not?

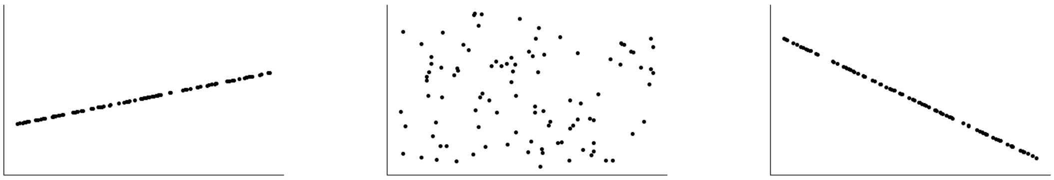

(A) (B) (C)

FIGURE 20: Graphs for Example 4.

Let's listen to some more correlations!

The following are sonifications of scatter plots corresponding to the correlations 0.1, 0.5 and 0.9. (These are in files Example 4 - Plot D, Example 4 - Plot E, Example 4 - Plot F.)

> *engagement questions*
> - Which one do you think is which? Here Plot D has correlation 0.9, Plot E has correlation 0.5 and Plot F has correlation 0.1.
> - Based on what you have heard, trying sketching each of the graphs that these sonifications correspond to.

Figure Description

The figure description for the graphs are as follows.

Each scatter plot shows the positive axes, but there are no scales and the axes are not labelled. On each plot are a number of points.

On plot D the points lie roughly either side of a straight line with a positive slope.

On Plot E the points lie roughly either side of a straight line with a positive slope, but there is a great deal of scatter.

On Plot F the points are scattered around with no discernible pattern, although there is some increase in the x value associated with an increase in the y value the points lie on a straight line with a positive slope.

> *engagement questions*
> In the light of this description do you want to change any of your sketches? If so in what way?

Comparison with the graphs

Now look at the graphs. (Allow the participant to look at Plots D, E and F.)

> *engagement questions*
> Based on the sonification, are the graphs as you expected? Why or why not?

Again at this point there is the option to just move straight on to the next example.

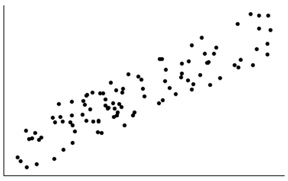

(D)

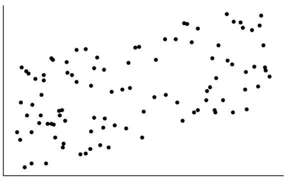

(E)

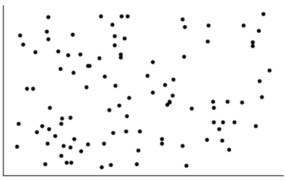

(F)

FIGURE 21: Graphs for Example 4.

Let's listen to one last correlation!

Finally listen to the sonification of one last sonification depicting a scatter plot. (This in file Example 4 - Plot G.)

engagement questions

- Based on what you have heard, trying sketching the graph that this sonification corresponds to.
- Guess a value for the correlation that this scatter plot depicts. (This is difficult, so a ballpark figure is fine.)

Figure Description

The figure description for the graph is as follows.

The scatter plot shows the positive axes, but there are no scales and the axes are not labelled. On the plot are a number of points. The points are quite closely grouped and rising.

Comparison with the graph

Now look at the graph. Allow the participant to look at plot G.

engagement questions

- Is the graph as you expected? Why or why not?
- What is your estimate of the correlation now?

As has already been mentioned, it turns out for the scatter plot G the correlation coefficient is 0.92, with anything between 0.8 and just under 1.0 being a good guess. But it something that students would be expected to struggle with (and experienced statisticians for that matter!).

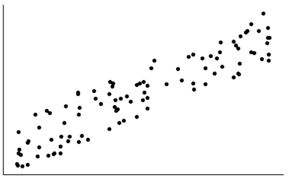

(G)

FIGURE 22: Final graph for Example 4.

Example 5: Interpreting a scatter plot

Key Idea

What is the general pattern of points on a scatter plot? And are any points differ to the rest?

Introduction

How do education systems in different countries compare? This is one question that The Programme for International Student Assessment (PISA) tries to answer. As part of research they administer tests in reading, mathematics and science to 15-year-olds in various countries. The tests are standardised so that attainment can be compared across countries.

The data we will consider in this example relates to testing that was done in 2009, relating to reading and mathematics. In total there are data from 65 countries and regions, including the UK.

In the following sonification each country/region corresponds to a note. The notes are ordered with respect to the reading score. So the higher the reading score for a country/region the later its note is sounded. The pitch of the note corresponds to the mathematics score. The higher the mathematics score is, the higher the pitch of the note.

Let's listen to the sonification!

It is OK for the student to listen to the sonification as many times as desired. (Its in file Example 5.) And its OK to return to it during the discussion. However it would be nice to know how many times the student felt it was necessary to listen to it.

engagement questions

- Try sketching the scatter plot.
- What pattern, if any, do you pick up in the sonification? That is do the notes go up? Or down or up and down? Or does there seem to be no pattern in the pitch over time? Or something else?
- What shape does the relationship between the reading and mathematics scores appear to have? Do the notes appear to go up steadily over time? (This suggests a linear relationship.) Or do the notes appear to go up more quickly over time? Or perhaps the notes go up less quickly over time? Or is it not possible to tell?
- How strong does the relationship between the reading and mathematics scores appear to be? That is, how clear is the 'signal' compared with the 'noise'.
- Do any points seem to be outliers? That is do any points seem not to fit in the pattern set by the rest of the points. If so roughly when do these outliers seem to occur, and in what way do they seem different?

Hopefully students will pick up that there seems to be one point near the beginning that seems to be a bit too high for a point occurring at that time. That is for one country the mathematics score was unusually good given its (quite poor) reading score. But I wouldn't be surprised if they miss this at this stage.

Figure Description

The figure description for this graph is as follows.

A scatter plot in which the horizontal axis is marked from 300 to 600 in intervals of 100. It is labelled Reading scale. The vertical axis is also marked from 300 to 600 in intervals of 100. It is labelled Mathematics scale. There are some extreme values. These are at approximately (310, 330), then (360, 430) and then (560, 600). The rest of the dots are fairly closely packed and show that an increase on the reading scale is generally associated with an increase on the mathematics scale.

engagement questions

In the light of this description do you want to change your sketch? If so in what way?

Comparison with the graph

Now look at the graph.

engagement questions

- Based on the sonification, is the graph as you expected? Why or why not?
- Are there any aspects that surprised you? If so what are they?

It wouldn't surprise me if its only at this stage that they notice the outlier. (The points that got a high mathematics score given its reading score.)

The graph!

The graph is shown in fig. 23.

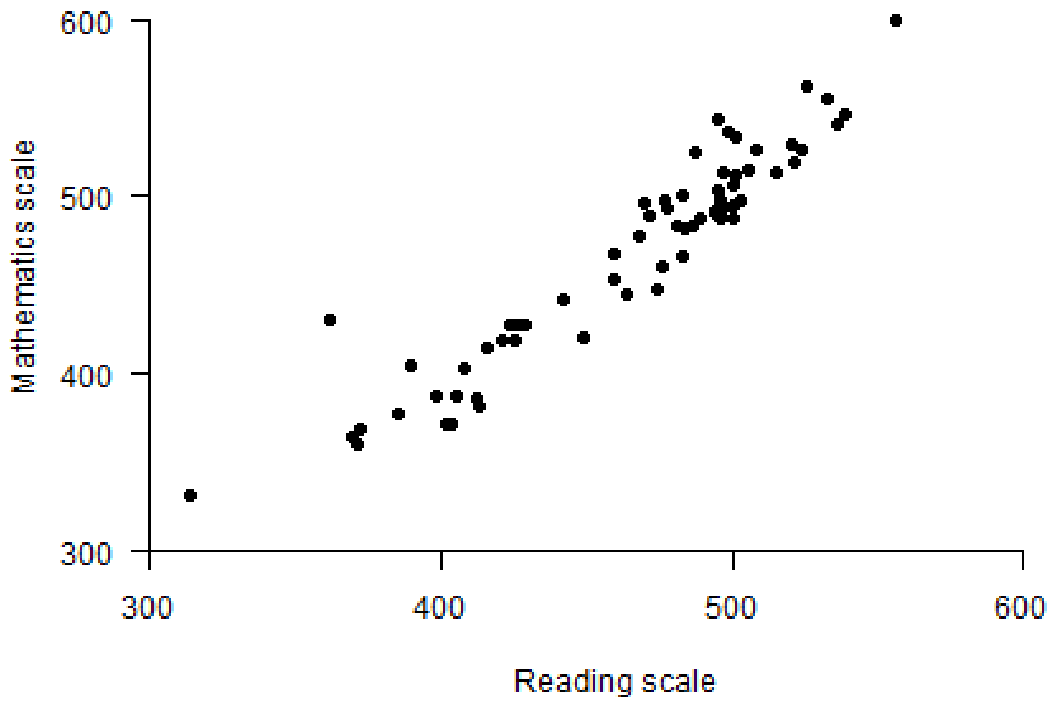

FIGURE 23: Graph for Example 5.

Example 6 - Assessing the fit of lines

Key Idea Which lines fit? That is, summarise the pattern of points?

Introduction In statistics one goal is often to summarise the data. With scatter plots this summarisation is done by drawing a line that represents the general trend in the data.

Of course there are many lines that could be drawn on a scatter plot. Inevitably some will provide better summaries than others. So how can we tell which ones are the better ones? Well, one straightforward way is to just draw the line on the scatter plot and compare it with the points. A good line should go through the middle of the cloud of points - though it may not go through any of the points exactly.

In this example we will consider just one set of data, comprising of a total of 15 points.

Lets listen! First listen to the sonification of the just the points. Allow the student to listen to the sonification of just the points (Example 6 - points only). As with all the sonifications allow the student to listen to this as many times as they like.

engagement questions

What appears to be the general trend in the notes? Do they go up or do they go down? Hopefully the participant will detect that they are going generally up - but its not very strong.

For these data four different straight lines were suggested as reasonable summaries. In each of the following sonifications, a continuous note representing the line is played along with individual notes for each of the points. These sonifications are given in the files Example 6 - Figure A to Example 6 - Figure D.

engagement questions

- Listening to each of the lines, which ones appear to summarise the data well and which ones not so well? Give your reasoning.
- Sketch what you think each of the plots looks like.

Figure Description The figure description for each of the scatter plot with lines are follows.

Line A. The line cuts the negative horizontal axis and the positive vertical axis. It passes through two of the data points that have both coordinates positive. The remaining points lie beneath the line.

Line B. The line is much steeper than Line A. It is shown cutting the positive horizontal axis only, to the left of the point on that axis. However, if extended beyond the scope of the diagram, the line would eventually cut the negative vertical axis. There are nine points to the left of the line and six to the right.

Line C. The line cuts the negative horizontal axis and the positive vertical axis but is less steep than either of the previous two lines. It passes through or is very close to six of the points. There are three points above the line and six points below.

Line D. The line cuts the negative horizontal axis and the positive vertical axis but is less steep than the other three lines. Three points lie close to the line, six lie above it and seven below it.

engagement questions

After reading these descriptions has your opinion about the fit of any of these lines changed? If so how?

Comparison with the graphs Now look at the scatter plots with the different lines on them. (Example 6 - Figure A to Example 6 - Figure D.)

engagement questions

- The scatter plots and lines actually match the scatter plots and lines you listened to. Which scatter plot do you think matches which sonification? Do any seem to not match at all?
- Has your opinion about the fit of any of these lines changed? If so how?

The graphs! The graphs are shown in fig. 24.

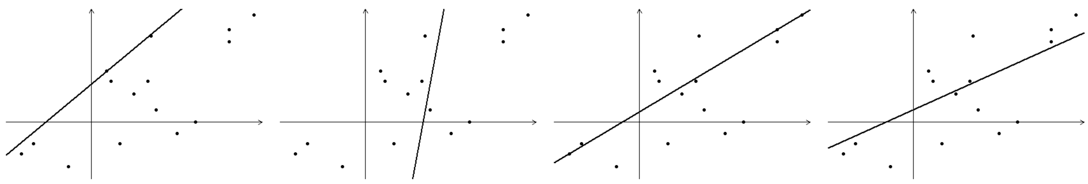

(A) (B) (C) (D)

FIGURE 24: Graphs for Example 6.