User Perceptions of Sound in Simple Linear Regression Diagnostics

Author: Jessica M. Orth* Advisor: Kate Cowles†

Abstract

There are many advanced and sophisticated graphical displays in common use by statisticians today, but there is little in the area of representing graphical displays through sound. We aim to introduce sound to graphical displays in simple linear regression diagnostics. Our goal is to use sound to present another dimension to complement visual tools and at the same time create an alternative to visualization that has a potential to help users with visual disabilities. Using sounds generated from the `audiolyzR` package, we analyze an online survey we designed to test users' perceptions of sound in simple linear regression diagnostics by means of the current software. Based on the analysis of our survey, we can conclude that users are able to correctly identify different diagnostic plots in simple linear regression through the sound medium, though this was more challenging for certain regression concepts than others. Areas of difficulty include identifying a non-constant variance relationship and outliers, based on our survey sample.

Key Words: Auditory Display; Regression Diagnostics; Sonification

1. Introduction and Approach

Statistical graphical displays, no matter how elaborate and involved, are basic images. Applying sound to images is by no means a new concept. With the inauguration of sound to previously silent films in the 1920's film *The Jazz Singer*, images took on a new meaning. Once a person views a film, the images seen on the screen are connected with the soundtrack of the film. Remove the images and retain the sound, and the person is able to visualize the movie in their mind as they hear the soundtrack. Why should a statistical graphical display be any different? Just as a soundtrack tells the story of the film in music, we aim to tell the story of the statistical analyses through sound. This requires the introduction of new concepts and ideas applied to statistical analyses so that the pursuit of eliciting sound in uncertainty is both meaningful and informative to the user of these 'musical reports'.

The main question we are seeking to explore is if users can correctly identify different concepts in simple linear regression through the sound medium. How effective is sound in the detection of outliers, non-constant variance, non-linear trends, and increasing and decreasing relationships? Are users able to connect sound with graphical displays and vice versa? We will highlight the use of the `audioScatter` function in the `audiolyzR` R package to demonstrate the sonification, or the "audio equivalent of visualization" [2], of data in these different settings.

The present paper is by no means a complete user's manual to portraying data with sound, for that remains an innovative field of study the paths of which are yet to be established. We intend to present our ideas on how sound can be introduced to data analyses as well as demonstrate the use of the `audiolyzR` package. An illustration of this package's current functionality is presented and possible modifications for more meaningful results are suggested based on the outcomes of our survey.

*Department of Statistics and Actuarial Science, The University of Iowa, 241 Schaeffer Hall Iowa City, IA 52242-1409

†Department of Statistics and Actuarial Science, The University of Iowa, 241 Schaeffer Hall Iowa City, IA 52242-1409

Section 2 will introduce the audiolyzR package used in the creation of the sound files utilized in the online survey discussed in section 2.2. The analysis of this survey using descriptive statistics, McNemar tests, multiple correspondence analysis, and logistic regression is presented in section 3. Comments and feedback from survey participants are also included in section 3 and section 4 contains our conlusions and plans for future work in this area.

Here is a summary of our approach that will be discussed in this paper.

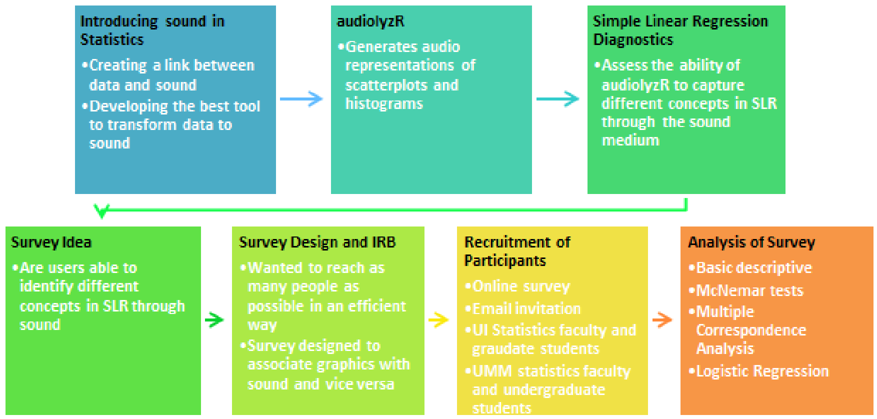

Figure 1: Diagram of our approach

2. Survey Methodology and Procedures

2.1 audiolyzR: give your data a listen

The R package audiolyzR can be used to create audio representations of histograms, scatterplots, and scatterplot matrices. The current version is "intended as a tool for familarization with a dataset, identification of outliers, and further analyses." [11] We will focus on the audioScatter function and analyze these claims based on the results of our survey.

The audiolyzR function generates a static image in the R session and also a JSON file that opens a stand-alone synthesizer that uses Max software, a powerful tool for creating your own audio files and interactive media by connecting objects. We incorporated the use of ggplot in this function to be used as the default plotting package for the static images. The function bins the data into arpeggios and relates the information in an audio graph. It takes as arguments two vectors, x and y, to plot on the respective axes. Once the function is applied to a data set, the user is sent to the Max software where there are several options available for listening to your data.

We selected six different concepts in simple linear regression on which to base our analysis of user perceptions of sound. These include: association (increasing vs. decreasing relationship), variance (non-constant vs. constant), functional (linear vs. non-linear), transformation (of the response variable), outliers (present or absent), and checking the normality assumption in qq-plots. We simulated data for each of these settings and created scatter plots, residual plots, and qq-plots for each using the audioScatter function and integrated these sounds and graphics into our online survey.

2.2 Survey design and implementation

The initial objective of our research was to determine if sound could be introduced to data analyses and provide the same information that a graphical analysis would provide to the user. We wanted to create a link between data and sound and develop the best tool to transform data to sound. Using the `audiolyzR` package we were able to create sound files from data, but we were uncertain of the quality of these sound representations and wanted to test how users perceive sound in simple linear regression diagnostics. Understanding how users who know regression interpret a sound file that characterizes a scatterplot, residual plot, or qq-plot became our main interest in being able to see the potential of data sonification in the future of Statistics.

With the goal of reaching as many people as possible in an efficient way to gain insight into how people who know regression interpret sound files, we designed our survey online using Qualtrics software available from The University of Iowa. Qualtrics is a tool that enables users to create and distribute online surveys easily.

Our survey was designed to test the relationship on how users answer graphics and sound questions and associate sound with graphics and vice versa. Figure 2 summarizes our survey design.

Figure 2: Overview of survey design

The first section of the survey contained background information on the participant including level of academic position (faculty, graduate student, undergraduate student), the level of regression courses taken (basic statistics course at undergraduate level (BU), higher level applied statistics course at undergraduate level (HU), basic statistics course at graduate level (BG), higher level applied statistics course at graduate level (HG), no linear regression in any course taken), and gender.

In the design of our survey, we included graphics questions to be used as a control in our analysis so that we could explore the effectiveness of sound by testing different simple linear regression concepts based on the answers to these control questions. We were interested in determining how users who know regression interpret corresponding sound files of the diganostic plots. Each of these questions asked the respondent to look at a graphical display and select the appropriate answer from a list of textual choices. The control questions from the survey are given below.

How would you classify the relationship between X and Y based on the following scatterplot?

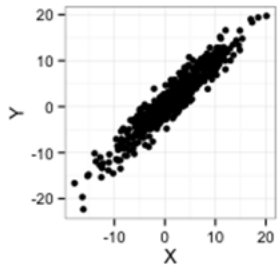

(a) There is an increasing trend

(b) There is a decreasing trend

(c) There is no trend

What would you say about the relationship between X and Y based on the following scatterplot?

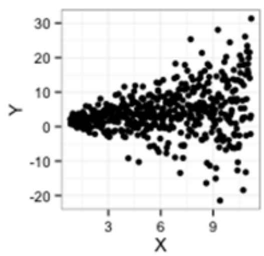

(a) There is a constant variance relationship

(b) There is a non-constant variance relationship

(c) Neither a nor b

The relationship between X and Y in the following scatterplot is best described by a(n)

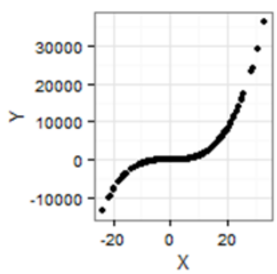

(a) Decreasing linear trend

(b) Increasing linear trend

(c) Increasing nonlinear trend

(d) Decreasing nonlinear trend

Does the following Normal quantile-quantile plot suggest any deviance from the Normality assumption?

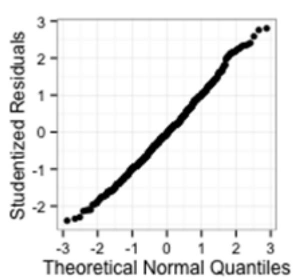

(a) Yes

(b) No

Does the following residual plot suggest a need for a transformation of the response variable to stabilize the variance?

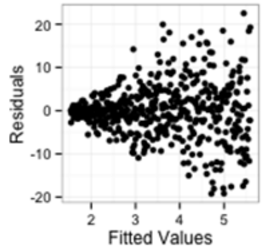

(a) Yes

(b) No

Does the following residual plot reveal an outlier in the data?

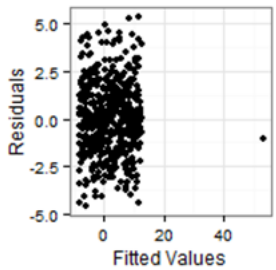

(a) Yes

(b) No

The third part of our survey included sound questions in several forms: sound to graphics, graphics to sound, and one question related with sound and text. The sound to graphics questions asked the respondent to listen to a sound file and match the graphic to the sound they heard. We explored three regression concepts using this technique: association, functional, and variance. These questions are listed below.

Please click on the play button and match the scatterplot to the sound that you hear:

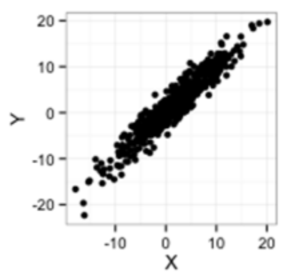

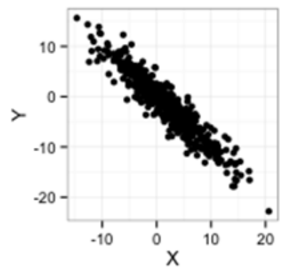

Please click on the play button and match the scatterplot to the sound that you hear:

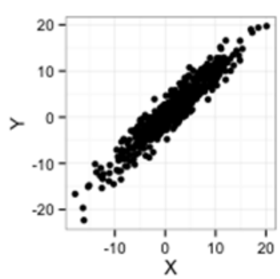

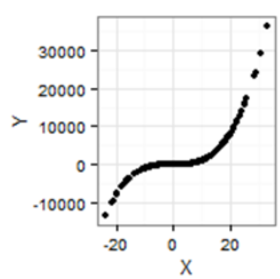

Please click on the play button and match the scatterplot to the sound that you hear:

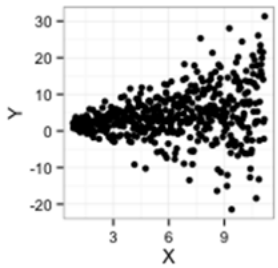

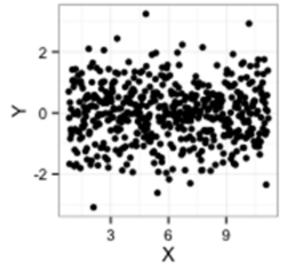

The graphics to sound questions asked the respondent to either look at a graphic and select the appropriate sound representation or read a question and select the correct sound file. For these questions the answers were sound files. We explored two regression concepts using this technique: outlier and transformation. The questions are given below.

Please click on the play button and match the sound to the residual plot that you hear:

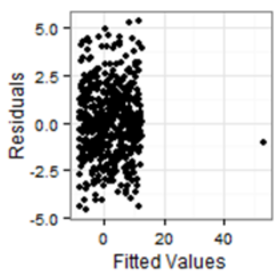

Which of the following sounds suggests a need for a transformation of the response variable to stabilize the variance?

We included one question relating text with sound in our survey. This question depicted two regression concepts at the same time through sound: increasing trend and outlier. Respondents were asked to listen to a sound file and select all that apply in terms of hearing an increasing relationship, a decreasing relationship, an outlier, or a non-constant variance relationship.

Our main point of analysis in designing our survey this way is to explore if users can first of all correctly identify these concepts in graphic form, and then determine the potential of sound representations by analyzing the responses to the corresponding sound questions.

Institutional Review Board approval for conducting this survey was granted on November 26, 2013 and the survey was launched on December 6, 2013. We sent the survey to statistics faculty and students at The University of Iowa and University of Minnesota Morris.

3. Survey Analysis

3.1 Descriptive analysis

We received 101 responses to our survey and 55 completed surveys. We used the 55 completed surveys for our analysis. There were 4 responses in this group that left the last three sound questions unanswered, and we adjusted the proportions for these questions accordingly. The following sections provide the details of our analysis.

Background information

As it is shown in Figure 3, the majority of respondents to the survey were undergraduate students and the highest level of regression course taken for the majority of respondents is a basic statistics course at the undergraduate level. Male and female respondents are approximately equal.

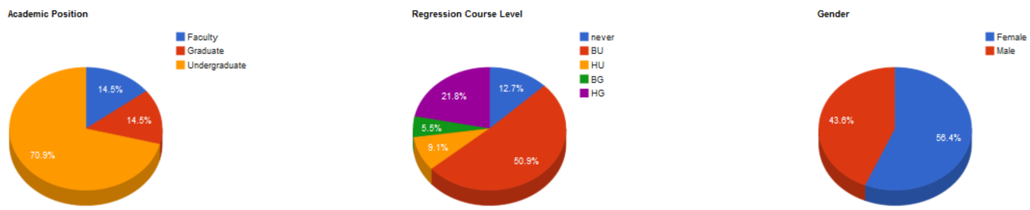

Figure 3: Proportion of respondents for academic position, level of regression course taken, and gender

Analysis of control questions

Table 1 summarizes the proportion of correct responses to each of the control questions. Sample proportions and 95% confidence intervals for the proportion of correct answers overall indicate that the majority of respondents did not have a problem recognizing the different regression concepts, as all proportions and intervals are above 0.50.

Breakdowns by background information are provided in the last three columns of the table. The p-values correspond to a Fisher's Exact Test for Count Data. This test was used because of the small sample size and gives a better approximation than the chi-square test statistic in this small sample setting. The hypothesis being tested in each of these is that there is no association between how respondents answered the control questions with their status of background information. Significant p-values at the 0.05 level are marked with a * and include the level of regression course taken for the variance and outlier questions, and the academic position of the respondent for the transformation and outlier questions. It is important to note that an insignificant p-value implies that our survey data does not provide enough evidence to show a relationship between how respondents answered the question and their academic position, level of regression course taken, or gender. This is possibly due to the small sample size.

Control Question	Overall	Position	Level	Gender
Association	1.00			
Variance	0.84 (0.74, 0.94)	Faculty: 1.00 Graduate: 1.00 Undergraduate: 0.77 (0.1729)	BU: 0.86 HG: 1.00 Other: 0.67 (0.05372)*	Female: 0.87 Male: 0.79 (0.4817)
Functional	0.85 (0.76, 0.94)	Faculty: 1.00 Graduate: 1.00 Undergraduate: 0.79 (0.2115)	BU: 0.86 HG: 1.00 Other: 0.73 (0.1159)	Female: 0.87 Male: 0.83 (0.7176)
Normality	0.65 (0.52, 0.78)	Faculty: 0.88 Graduate: 0.75 Undergraduate: 0.59 (0.3673)	Never: 0.71 BU: 0.57 HG: 0.83 Other: 0.63 (0.5852)	Female: 0.71 Male: 0.58 (0.3973)
Transformation	0.75 (0.64, 0.86)	Faculty: 1.00 Graduate: 1.00 Undergraduate: 0.64 (0.01899)*	Never: 0.71 BU: 0.61 HG: 1.00 Other: 0.88 (0.06413)	Female: 0.71 Male: 0.79 (0.5471)
Outlier	0.87 (0.78, 0.96)	Faculty: 0.63 Graduate: 0.75 Undergraduate: 0.95 (0.02385)*	Never: 1.00 BU: 0.96 HG: 0.58 Other: 0.88 (0.01481)*	Female: 0.94 Male: 0.79 (0.2197)

Table 1: Proportion of correct responses for overall participants and background characteristics for control questions. Numbers correspond to proportion correct (95% confidence) for the overall column and proportion correct (p-value for Fisher's Exact Test for Count Data) for the remaining columns

Analysis of sound questions

Sound to graphics

Table 2 summarizes the proportion of correct responses for each of the sound to graphics questions. The first column corresponds to the responses overall and gives a 95% confidence interval for the proportion of respondents correctly answering the question and breakdowns by background information are provided in the subsequent columns. The 95% confidence intervals for association and variance include 0.50, indicating that these proportions are not significantly different than what would be observed by random guessing based on our small survey sample. None of the breakdowns by background information status presented significant p-values for the hypothesis of no association between how respondents answered the question and their status. This is again possibly due to the small sample size.

Sound to Graphics Question	Overall	Position	Level	Gender
Association	0.56 (0.43, 0.69)*	Faculty: 0.75 Graduate: 0.88 Undergraduate: 0.46 (0.06459)	BU: 0.57 HG: 0.83 Other: 0.33 (0.06014)	Female: 0.55 Male: 0.58 (1.00)
Functional	0.78 (0.67, 0.89)	Faculty: 1.00 Graduate: 0.75 Undergraduate: 0.74 (0.344)		Female: 0.77 Male: 0.79 (1.00)
Variance	0.44 (0.31, 0.57)*		BU: 0.46 HG: 0.45 Other: 0.40 (1.00)	Female: 0.52 Male: 0.35 (0.2738)

Table 2: Proportion of correct responses for overall participants and background characteristics for sound to graphics questions. Numbers correspond to proportion correct (95% confidence) for the overall column and proportion correct (p-value for Fisher's Exact Test for Count Data) for the remaining columns. Empty cells are a consequence of anonymity in a small sample size.

Graphics to sound

Table 3 summarizes the proportion of correct responses for each of the graphics to sound questions. The first column corresponds to the responses overall and gives a 95% confidence interval for the proportion of respondents correctly answering the question and breakdowns by background information are provided in the subsequent columns.

The 95% confidence intervals for both of these questions include 0.50, indicating that these proportions are not significantly different than what would be observed by random guessing based on our small survey sample. Possibly due to the small sample size, none of the breakdowns by background information status presented significant p-values for the hypothesis of no association between how respondents answered the question and their status.

Graphics to Sound Question	Overall	Position	Level	Gender
Outlier	0.61 (0.48, 0.74)*		BU: 0.75 HG: 0.55 Other: 0.40 (0.1125)	Female: 0.65 Male: 0.57 (0.584)
Transformation	0.57 (0.44, 0.70)*		BU: 0.54 HG: 0.67 Other: 0.57 (0.7036)	Female: 0.67 Male: 0.46 (0.1687)

Table 3: Proportion of correct responses for overall participants and background characteristics for graphics to sound questions. Numbers correspond to proportion correct (95% confidence) for the overall column and proportion correct (p-value for Fisher's Exact Test for Count Data) for the remaining columns. Empty cells are a consequence of anonymity in a small sample size.

Sound and text

It is important to note the features of the sound and text question shown in Table 4. Even though the proportion of responses are low, this is not due to random guessing, as 0.35 answered incorrectly and the overall proportion of correct responses is divided among the remaining options. Because the sound file for this question depicted both an increasing trend and an outlier, the best response was both. However, the lowest proportion of selections is for respondents answering ‘both’. Respondents had an easier time hearing the outlier than the increasing trend alone when both scenarios were present in the same sound file.

Question	Overall
Sound and text	Both: 0.13 (0.04, 0.22)
	Increasing: 0.20 (0.09, 0.31)
	Outlier: 0.31 (0.19, 0.43)
	Wrong: 0.35 (0.22, 0.48)

Table 4: Distribution of the proportion of responses for overall participants to the sound and text question. Numbers correspond to proportion correct (95% confidence interval)

3.2 Inferential analysis

McNemar's test

McNemar's test is used to test the same hypotheses of association as a chi-square test but the responses are now paired. This test checks the relationship between discordant pairs of observations. For each regression concept, we paired the graphic and sound questions and used McNemar's test to examine the hypothesis that there is no association between correctly answering the graphics and corresponding sound question. The results are shown in Table 5 along with the conditional probability of correctly answering the sound question based on the answer to the graphics question. A graphical representation of this table is shown in Figure 4.

We observe a significant difference between graphics and sound for association, variance, and outlier, and no significant difference between functional and transformation based on our survey data. The p-value for association is very small compared to the others because all respondents correctly answered the graphics association question. Users faced a challenge trying to interpret the sound representation of association, variance, and outlier, while interpreting the sound representation of functional and transformation proved to be an easier task with the current software based on our sample.

Question Pairs (graphic vs. sound)	McNemar's p-value	Proportion of correct answers to sound question conditional on a correct answer to graphics question	Proportion of correct answers to sound question conditional on an incorrect answer to graphics question
Association	2.668e-06*	0.56	
Variance	0.0002041*	0.44	0.44
Functional	0.3017	0.79	0.71
Transformation	0.08086	0.63	0.43
Outlier	0.000685*	0.67	0.17

Table 5: p-values for McNemar's test of independence in question pairs (graphic vs. sound) and the conditional proportion of correct answers to sound questions given the response to the corresponding graphics question

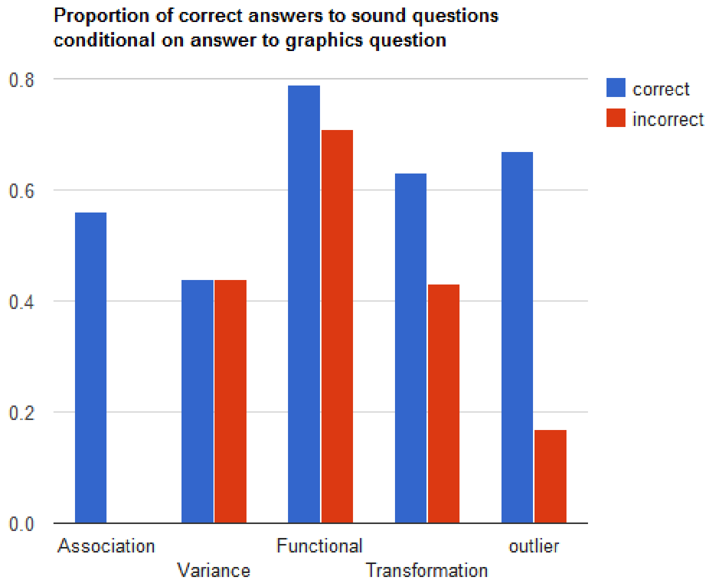

Figure 4: Results of McNemar's Test on Question Pairs

Multiple correspondence analysis

Multiple correspondence analysis (MCA) can be used to demonstrate the relationship between a set of three or more categorical variables in a direct extension of correspondence analysis. MCA seeks to create scales for rows and columns that display the maximum correlation between these pairs and plot the results in a map. Eigenvalues and eigenvectors from the Burt matrix given by

$$B = \begin{pmatrix} Z_1' Z_1 & \cdots & Z_1' Z_k \\ \vdots & \ddots & \vdots \\ Z_k' Z_1 & \cdots & Z_k' Z_k \end{pmatrix} \tag{1}$$

where $Z = [Z_1, \ldots, Z_k]$ denotes the indicator matrix for a k-way contingency table, are used in the creation of the indicator matrix for plotting the MCA. Points that are more associated with each other are plotted closer to each othe on the map and points that are less associated with each other are plotted farther away from each other on the map. The results of MCA on our survey data are shown in Figure 5.

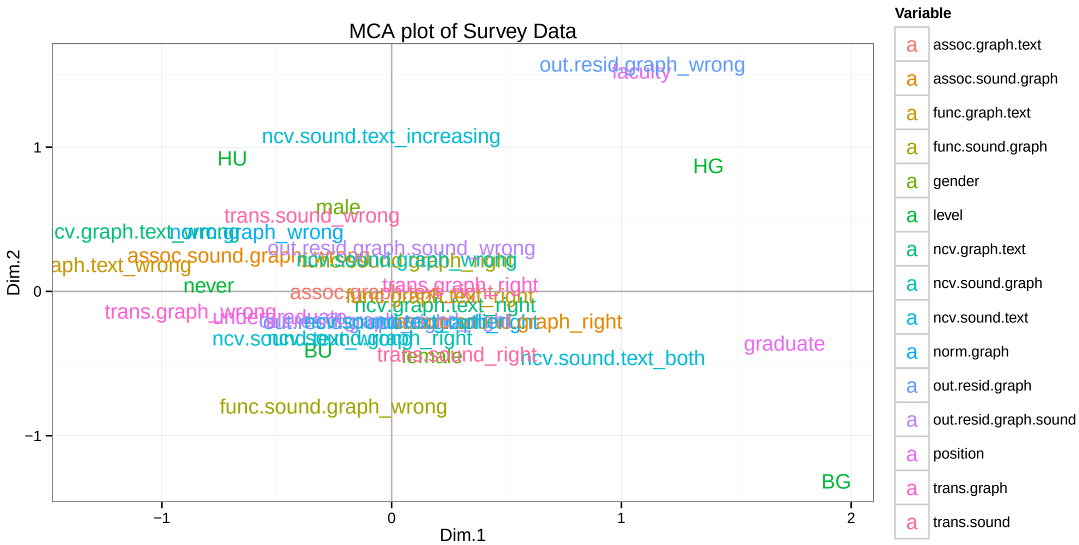

Figure 5: MCA results of survey data

In the upper right corner we can see that the faculty are associated with answering the graphical outlier question incorrectly. The graphical outlier question portrayed an outlier with respect to the explanatory variable only, not an outlier with respect to the relationship of both variables. Faculty had a deeper understanding of this concept and made their selection accordingly. On the left side we can observe that having never seen linear regression in a course taken is associated with incorrectly answering several of the questions including the graphic transformation question and association sound question. Males, located in the upper left quadrant, are associated with incorrectly answering the sound transformation question and are similar to those respondents who have have taken a higher level regression course as an undergraduate student. Females are located in the center of the map in the lower right quadrant and are associated with correctly answering the sound transformation question and several other sound questions as well. More associations can be explored from Figure 5.

Logistic regression

Using the response variable of correctly answering the sound question, we fit all possible logistic models regressing on the corresponding graphic question and background information status. Logistic regression could not be performed on the association questions because all respondents answered the graphics question correctly. In all of these models only one gave a significant term other than the intercept. This is the model for the question on outliers specified below.

$$\begin{aligned} y_i \leftarrow\ Y_i\ indep\ \sim\ binom(n_i, \mu_i), \\ logit\mu_i = \beta_0 + \beta_1 outlier.graphic_i, i = 1, \ldots, 55, \\ where\ \mu_i = P(outlier.sound\ correct | outlier.graphic) \end{aligned}$$

Table 6 provides the estimates, standard errors, and significance for this model.

	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	0.6931	0.3062	2.26	0.0236
outlier.graphic.wrong	-2.3026	1.1374	-2.02	0.0429

Table 6: Estimates for logistic regression model

The estimated regression line for this model is

$$logit(\hat{\mu}(outlier.graphic)) = 0.6931 - 2.3026 * outlier.graphic_{wrong} \tag{2}$$

The probability of correctly answering the sound outlier question for respondents who correctly answered the graphics outlier question is

$$\frac{exp(0.6931)}{1 + exp(0.6931)} = 0.67 \tag{3}$$

The probability of correctly answering the sound outlier question for respondents who incorrectly answered the graphics outlier question is

$$\frac{exp(0.6931 - 2.3026)}{1 + exp(0.6931 - 2.3026)} = 0.17 \tag{4}$$

Based on this model from our survey data, knowledge of traditional graphical diagnostics can be associated with an increase in the probability of correctly answering the corresponding sound question for the detection of outliers.

3.3 Impact of the survey

The survey discussed in this paper was received with much enthusiasm from participants. Here we list serval of many similar evaluations of the survey indicating a strong interest in the idea of introducing sound to Statistics:

1. 'I thought that the study between sound and graphs was very interesting. As I was answering the questions I found it very interesting how one sound can make you think of a certain type of graph or relationship. It would be very cool to see sound/ music used more in math classes and see how that could improve learning.'

2. 'I think that project will be extremely beneficial, in the classroom and in society overall.'

3. 'My thoughts on this survey are that it was a very interesting concept that someone could possibly hear statistics. I personally did not feel as if I could hear the statistical evidence, but that is partially because music is not a talent of mine. However, I think the idea of hearing stats is very interesting based on if the sound increases or decreases. We can apply this knowledge to everyday life.'

4. 'I have never thought of using sound to represent a physical representation of something. If there was a lot of data on a certain spot on the graph, it was easy to tell which sound went with that graph because the sound would be more dense.'

5. 'I thought it was very interesting. The only thing would be that if this idea would go a step farther would maybe be to define or establish what different sounds mean. At least like if a person was solely using the sound. This is because one can assume that a higher pitch may mean a higher point, but for it to be totally effective I feel that the idea behind the pitches would need to be established. This would be to clarify

and avoid misunderstanding. This would of course be an expansion and I am sure it was not originally done because the researcher was trying to see what pitches an individual associated with the different points on the graph.'

6. 'I thought the combination of music and statistics is incredibly interesting and important to develop over time. I was deeply involved in music throughout my entire middle school and high school career, so it has always been easy for me to listen to different melodies and tunes for their intricacies. It was really quite different to see the sounds displayed in a graphical formation because you don't often try to pair something very science and analysis based into a melody or sound effects. I think by doing this comparison, students who learn better by listening rather than seeing or talking would benefit immensely in such a very discussion based course.'

4. Conclusions and Future Work

Based on the analysis of our survey, we can conclude that users are able to correctly identify different diagnostic plots in simple linear regression through the sound medium, though this was more challenging for certain regression concepts than others including the identification of outliers and non-constant variance. Users are able to connect sound with graphical displays and vice versa, and the current software `audiolyzR`, when used in the simple linear regression setting, proved to be effective in using sound to detect outliers, variance, and other trends.

We used only the default settings of `audiolyzR` in terms of chord quality, volume, tempo, starting note, and delay time. The meaning of changing the chord quality and other characteristics related with music remains a point of discussion and interpretation which we intend to address in future research.

In the future we plan to carry out a more in-depth analysis on the implications of sound in data analyses and extend it to statistical concepts beyond the simple linear regression setting. We intend to introduce sound into statistical inference and simulation studies as well. Also, we will search for ways of using sound in Bayesian analyses including the problem of elicitation. In addition, we aim to introduce sound to the use of copulas and directional dependence structures and angular correlation. In all of these applications of sound to Statistics, our goal is to reach a new audience of statisticians and non-statisticians alike by presenting a tool that allows for an innovative way to hear the stories behind the data. In addition, we would like to create a method of teaching statistics that will help those with visual disabilities as they will be able to listen to the trends in the data.

Acknowledgments

The author greatly appreciates the generous comments and suggestions of Professor Cowles throughout this research. This research could not have been started without the numerous responses to the survey, and the author would like to express deep gratitude to all who participated in the survey. A special thanks is extended to the statistics faculty at the University of Minnesota, Morris for distributing the survey to their students. The author extends deep appreciation to Professor Sungur at the University of Minnesota, Morris for introducing the idea of sound in statistics to the author as an undergraduate student and for his incredible support and encouragement. LaTeXassistance provided by Matt Bognar is also greatly appreciated.

References

[1] Herv Abdi and Dominique Valentin. Multiple correspondence analysis. *Encyclopedia of Measurement and Statistics*, pages 651–657, 2007.

[2] Ethan Brown and Nick Bearman. Listening to uncertainty: Information that sings. *Significance*, 9(5):14–17, 2012.

[3] B. Everitt and G. Dunn. *Applied Multivariate Data Analysis*. Hodder Arnold UK, United Kingdom, 2001.

[4] J. Fox. *Applied Regression Analysis and Generalized Linear Models*. SAGE Publications, Inc., USA, 2008.

[5] J. Fox and S. Weisberg. *An R Companion to Applied Regression*. SAGE, Inc., USA, 2011.

[6] Markus Gesmann and Diego de Castillo. googlevis: Interface between r and the google visualisation api. *The R Journal*, 3(2):40–44, December 2011.

[7] Francois Husson, Julie Josse, Sebastien Le, and Jeremy Mazet. *FactoMineR: Multivariate Exploratory Data Analysis and Data Mining with R*, 2013. R package version 1.25.

[8] P. McCullagh and J.A. Nelder. *Generalized Linear Models*. Chapman and Hall, United Kingdom, 1989.

[9] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013.

[10] RStudio and Inc. *shiny: Web Application Framework for R*, 2013. R package version 0.8.0.

[11] Eric Stone and Jesse Garrison. *audiolyzR: Give your data a listen*, 2013. R package version 0.4-9.

[12] W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S*. Springer, New York, fourth edition, 2002. ISBN 0-387-95457-0.